

SimEvents

For Use with Simulink®

- Modeling
- Simulation
- Implementation

User's Guide

Version 1



How to Contact The MathWorks:



www.mathworks.com
comp.soft-sys.matlab
www.mathworks.com/contact_TS.html

Web
Newsgroup
Technical Support



suggest@mathworks.com
bugs@mathworks.com
doc@mathworks.com
service@mathworks.com
info@mathworks.com

Product enhancement suggestions
Bug reports
Documentation error reports
Order status, license renewals, passcodes
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.

SimEvents User's Guide

© COPYRIGHT 2005–2006 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, Simulink, Stateflow, Handle Graphics, Real-Time Workshop, and xPC TargetBox are registered trademarks of The MathWorks, Inc.

Other product or brand names are trademarks or registered trademarks of their respective holders.

Patents

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

Revision History

November 2005	Online only	New for Version 1.0 (Release 14SP3+)
March 2006	Online only	Revised for Version 1.1 (Release 2006a)

Working with Entities

1

Generating Entities When Events Occur	1-2
Detecting Sample Time Hits	1-2
Detecting Changes in Signal Values	1-4
Detecting Edges in Trigger Signals	1-5
Detecting Function Calls	1-7
Using Generation Times from a Vector	1-11
Configuring the Block to Generate Entities at Specified Times	1-11
Sample Use Cases	1-13
Setting Attributes of Entities	1-14
Example: Setting Attributes	1-15
When to Use Attributes	1-17
Accessing Attributes of Entities	1-20
Counting Entities	1-21
Counting Departures Across the Simulation	1-21
Counting Departures per Time Instant	1-21
Resetting a Counter Upon an Event	1-23
Associating Each Entity with Its Index	1-24
Replicating Entities on Multiple Paths	1-25
Departure Port Precedence	1-25

Working with Events

2

Supported Events in SimEvents Models	2-2
---	------------

Types of Supported Events	2-2
Signal-Based Events	2-3
Function Calls	2-7
Using the Event Calendar	2-9
Events That Appear on the Event Calendar	2-9
Example: Event Calendar for a Queue-Server Model	2-11
Setting Event Priorities	2-19
Processing Sequence for Simultaneous Events	2-21
Example: Race Conditions at a Switch	2-23
Observing Events	2-30
Example: Observing Service Completions	2-32
Example: Detecting Collisions by Comparing Events	2-34
Generating Function-Call Events	2-36
Generating Events When Other Events Occur	2-36
Generating Events Using Intergeneration Times	2-38
Manipulating Events	2-39
Blocks for Manipulating Events	2-40
Creating a Union of Multiple Events	2-40
Translating Events to Control the Processing Sequence ..	2-43
Conditionalizing Events	2-45

Working with Signals

3

Role of Event-Based Signals in SimEvents Models	3-2
Comparison With Time-Based Signals	3-2
Generating Random Signals	3-4
Generating Random Event-Based Signals	3-4
Examples of Random Event-Based Signals	3-6
Generating Random Time-Based Signals	3-8
Sequence of Updates of Output Signals	3-10

Example: Detecting Changes in the Last-Updated Signal	3-10
Multiple Simultaneous Updates	3-13
Zero-Duration Values of Signals	3-13
Importance of Zero-Duration Values	3-14
Detecting Zero-Duration Values	3-14
Latency in Signal Updates	3-17
Manipulating Signals	3-18
Defining Initial Conditions for Event-Based Signals	3-18
Example: Resampling a Signal Based on Events	3-20
Sending Data to the MATLAB Workspace	3-23
Example: Sending Queue Length to the Workspace	3-23
Using the To Workspace Block with Event-Based Signals	3-24

Modeling Queues and Servers

4

Using a LIFO Queuing Discipline	4-2
Example: Waiting Time in LIFO Queue	4-2
Sorting by Priority	4-4
Example: Serving Preferred Customers First	4-6
Preempting an Entity in a Server	4-9
Criteria for Preemption	4-9
Residual Service Time	4-9
Queuing Disciplines for Preemptive Servers	4-10
Example: Preemption by High-Priority Entities	4-10
Modeling Multiple Servers	4-12
Example: M/M/5 Queuing System	4-12

Modeling the Failure of a Server	4-15
Server States	4-15
Using a Gate to Implement a Failure State	4-15
Using Stateflow to Implement a Failure State	4-16

Routing Techniques

5

Example: Cascaded Switches with Random Selections	5-2
Six-Way Switching Component	5-2
Five-Way Switching Component	5-2
Example: Cascaded Switches with Round-Robin Sequence	5-4
Example: Compound Switching Logic	5-5

Using Logic

6

Role of Logic in SimEvents Models	6-2
Using Embedded MATLAB Function Blocks for Logic	6-3
Example: Choosing the Shortest Queue	6-3
Example: Varying Fluid Flow Rate Based on Batching Logic	6-6
Using Logic Blocks	6-10
Example: Using Servers in Shifts	6-11
Example: Choosing the Shortest Queue Using Logic Blocks	6-16

Regulating Arrivals Using Gates

7

Role of Gates in SimEvents Models	7-2
Accessing Gate Blocks	7-3
Types of Gates	7-3
Keeping a Gate Open Over a Time Interval	7-4
Example: Controlling Joint Availability of Two Servers ...	7-4
Opening a Gate Instantaneously	7-6
Example: Synchronizing Service Start Times with the Clock	7-6
Example: Opening a Gate Upon Entity Departures	7-7
Using Logical Combinations of Gates	7-9
Example: First Entity as a Special Case	7-10

Controlling Timing Using Subsystems

8

Timing Issues in SimEvents Models	8-2
Timing for the End of the Simulation	8-2
Timing for a Statistical Computation	8-3
Timing for Choosing a Port Using a Sequence	8-4
Role of Discrete Event Subsystems in SimEvents Models	8-7
Purpose of Discrete Event Subsystems	8-8
Processing Sequence for Events in Discrete Event Subsystems	8-8
Blocks Inside Discrete Event Subsystems	8-10
Working with Discrete Event Subsystem Blocks	8-11
Setting Up Signal-Based Discrete Event Subsystems	8-11

Signal-Based Events That Control Discrete Event Subsystems	8-14
Examples Using Discrete Event Subsystem Blocks	8-16
Example: Comparing the Lengths of Two Queues	8-16
Example: Normalizing a Statistic to Use for Routing	8-17
Example: Using Event-Based Timing for a Statistical Computation	8-19
Example: Ending the Simulation Upon an Event	8-20
Example: Sending Unrepeated Data to the MATLAB Workspace	8-21
Example: Focusing on Events, Not Values	8-22
Example: Detecting Changes from Empty to Nonempty ..	8-23
Example: Logging Data About the First Entity on a Path	8-24
Creating Entity-Departure Subsystems	8-26
Accessing Blocks for Entity-Departure Subsystems	8-26
Setting Up Entity-Departure Subsystems	8-27
Examples Using Entity-Departure Subsystems	8-29
Example: Using Entity-Based Timing for Choosing a Port	8-29
Example: Performing a Computation on Selected Entity Paths	8-31
Using Function-Call Subsystems	8-32
Use Cases for Function-Call Subsystems	8-32
Setting Up Function-Call Subsystems in SimEvents Models	8-32

Plotting Data

9

Choosing and Configuring Plotting Blocks	9-2
Sources of Data for Plotting	9-2
Inserting and Connecting Scope Blocks	9-3
Connections Among Points in Plots	9-4
Varying Axis Limits Automatically	9-5

Examples Using Scope Blocks	9-6
Plotting Window Features	9-7
Using Plots for Troubleshooting	9-8
Example: Plotting Entity Departures to Verify Timing ...	9-8
Example: Plotting Event Counts to Find Roundoff Error ..	9-12
Comparison with Time-Based Plotting Tools	9-15

Using Statistics

10

Role of Statistics in Discrete-Event Simulation	10-2
Statistics for Data Analysis	10-2
Statistics for Run-Time Control	10-3
Accessing Statistics from SimEvents Blocks	10-4
Accessing Statistics Throughout the Simulation	10-4
Accessing Statistics When Stopping or Pausing Simulation	10-6
Deriving Custom Statistics	10-7
Example: Fraction of Dropped Messages	10-8
Example: Computing a Time Average of a Signal	10-9
Example: Resetting an Average Periodically	10-12
Using Timers	10-18
Basic Example Using Timer Blocks	10-18
Basic Procedure for Using Timer Blocks	10-19
Timing Multiple Entity Paths with One Timer	10-20
Restarting a Timer	10-21
Timing Multiple Processes Independently	10-23
Running a Series of Simulations	10-25
Creating Independent Replications	10-25
Running Simulations from MATLAB	10-27

Troubleshooting Discrete-Event Simulations

11

Viewing the Event Calendar	11-2
Turning Event Logging On	11-2
Logging the Processing of Events	11-3
Logging the Scheduling of Events	11-4
Logging the List of Events	11-4
Example: Event Logging	11-6
Viewing Entity Locations	11-8
Turning Entity Logging On	11-8
Interpreting Entity Logging Messages	11-8
Example: Entity Logging	11-9
Common Problems in SimEvents Models	11-12
Unexpectedly Simultaneous Events	11-13
Unexpectedly Nonsimultaneous Events	11-13
Unexpected Processing Sequence for Simultaneous Events	11-16
Time-Based Block Not Recognizing Certain Trigger Edges	11-17
Incorrect Timing of Signals	11-17
Unexpected Use of Old Value of Signal	11-19
Effect of Initial Condition on Signal Loops	11-23
Loops in Entity Paths Without Storage Blocks	11-26
Unexpected Timing of Random Signal	11-27
Unexpected Correlation of Random Processes	11-29
Configuration Parameters for SimEvents Models	11-30

12

Notifications and Queries Among Blocks	12-2
Querying Whether a Subsequent Block Can Accept an Entity	12-2
Notifying Blocks About Status Changes	12-3
Interleaved Operations of Storage and Nonstorage Blocks	12-4
Example: Sequence of Departures and Statistical Updates	12-5
Zero-Duration Values and Time-Based Blocks	12-10
Example: Using a #n Signal as a Trigger	12-11

Demonstration Models

13

Tutorial Demos	13-2
Attributes: Data Within Entities	13-2
Service Time From Attribute	13-3
Specifying Service Time in Single Server	13-3
Specifying Service Time in Infinite Server Block	13-3
Single Server Block Versus Infinite Server Block	13-3
Start Timer and Read Timer Blocks	13-4
Release Gate: Value Change Versus Trigger	13-4
Input Switching Using Signal	13-4
Output Switching Using Signal	13-4
Path Combiner Versus Input Switch	13-5
Time-Driven and Event-Driven Addition	13-5
Counting Instantaneous Events	13-6
Preload Queue with Entities	13-6
Queuing Systems	13-7
M/M/1 Queuing System	13-7
M/D/1 Queuing System	13-9
G/G/1 Queuing System and Little's Law	13-10
Single Server Versus N-Server	13-13

Single Queue Versus Multiple Queues	13-13
A Queuing System with Discouraged Arrivals	13-13
Prioritized Queuing Policy Comparison	13-15
Preemption Policy Comparison	13-15
Application Demos	13-16
Shared Access Communications Media	13-16
Dynamic Voltage Scaling Using Online Gradient Estimation	13-16
Comparison of Routing Policies	13-19
F-14 Flight Control Over a Network	13-19
Selective-Repeat Automatic Repeat Request	13-19
Tank Filling Station	13-20
Astable Multivibrator Circuit	13-21

Functions — Alphabetical List

14

Blocks — By Category

15

Generators	15-2
Entity Generators	15-2
Event Generators	15-2
Signal Generators	15-2
SimEvents Sinks	15-3
Attributes	15-3
Queues	15-3
Servers	15-4
Routing	15-4

Gates	15-4
SimEvents Ports and Subsystems	15-5
Timing	15-5
Probes	15-5
Event Translation	15-6

Blocks — Alphabetical List

16

Glossary

Examples

A

Attributes of Entities	A-2
Counting Entities	A-2
Working with Events	A-2
Queuing Systems	A-2
Working with Signals	A-2
Server States	A-3

Routing Entities	A-3
Batching	A-3
Gates	A-3
Discrete Event Subsystems	A-3
Troubleshooting	A-4
Statistics	A-4
Timers	A-5

Index

Working with Entities

Generating Entities When Events Occur (p. 1-2)

Using Generation Times from a Vector (p. 1-11)

Setting Attributes of Entities (p. 1-14)

Accessing Attributes of Entities (p. 1-20)

Counting Entities (p. 1-21)

Replicating Entities on Multiple Paths (p. 1-25)

Using events to determine when to generate an entity

Generating entities at explicit values of time

Attaching data to entities

Reading and using data attached to entities

Counting entities per time instant and across the simulation

Creating copies of entities

Generating Entities When Events Occur

The Event-Based Entity Generator block enables you to generate entities in response to events that occur during the simulation. Event times and the time intervals between pairs of successive entities are not necessarily predictable in advance. This section describes the events that can cause entity generation, in these topics:

- “Detecting Sample Time Hits” on page 1-2
- “Detecting Changes in Signal Values” on page 1-4
- “Detecting Edges in Trigger Signals” on page 1-5
- “Detecting Function Calls” on page 1-7

Generating entities when events occur might be appropriate if you want the dynamics of your model to determine when to generate entities. For example, if you want to generate an entity every time a Stateflow® chart transitions from state A to state B, then you configure the Stateflow chart to output a function call upon such a transition and configure the Event-Based Entity Generator block to react to each function call by generating an entity. As another example, if you want to generate an entity every time the length of a queue changes, then you configure the queue to output a signal indicating the queue length and configure the Event-Based Entity Generator block to react to changes in that signal’s value by generating an entity.

Note To specify intergeneration times between pairs of successive entities, use the Time-Based Entity Generator block as described in “Creating Entities Using Intergeneration Times” in the Getting Started documentation.

Detecting Sample Time Hits

You can configure the Event-Based Entity Generator block so that it generates entities in response to updates in a signal. More explicitly, whenever the block producing that signal recomputes and outputs the signal value, the Event-Based Entity Generator block generates an entity. The actual value of the signal and the question of whether recomputing the value yields a different result compared to the previous time step are irrelevant to the entity generation process; only the time of the output is relevant.

Configuring the Block to Detect Sample Time Hits


To use this method for generating entities, set the Event-Based Entity Generator block's **Generate entities upon** parameter to **Sample time hit** from port **ts**. This causes the block to have an input signal port labeled **ts**. During the simulation, the Event-Based Entity Generator block generates entities in response to updates in the signal connected to this **ts** port.

Sample Use Cases

Here are a few scenarios that illustrate this method for generating entities:


- The **ts** signal is the output of a block with an explicit **Sample time** parameter. Regardless of the value of the **ts** signal, the Event-Based Entity Generator block generates an entity periodically, according to the sample time of the driving block.

Compared to using a Time-Based Entity Generator block with **Distribution** set to **Constant**, this event-based approach is a more direct way to synchronize entity generation events with sample time hits and avoid possible roundoff errors. Below is an example.



For other examples that use this entity-generation method to effect desired simultaneity of events, see “Example: Race Conditions at a Switch” on page 2-23 and “Example: Plotting Event Counts to Find Roundoff Error” on page 9-12.

- The **ts** signal is the output of a triggered subsystem whose **Propagate execution context across subsystem boundary** parameter is selected. Whenever Simulink® calls the subsystem and recomputes the output, the Event-Based Entity Generator block generates an entity.



Note Do not put the Event-Based Entity Generator block inside a triggered subsystem. Like other blocks that possess entity ports, the Event-Based Entity Generator block is not valid inside a triggered subsystem. See also “Detecting Edges in Trigger Signals” on page 1-5.

- The **ts** signal is a statistical output signal from a SimEvents block. Whenever the block recomputes and outputs the statistic, the Event-Based Entity Generator block generates an entity.

Detecting Changes in Signal Values

You can configure the Event-Based Entity Generator block so that it generates entities in response to numerical changes in a signal. This can be useful if the changes in the signal’s value have some significance in your simulation; for example, a signal representing the length of a queue changes whenever an entity arrives at or departs from the queue.


Configuring the Block to Detect Value Changes

To use this method for generating entities, set the Event-Based Entity Generator block’s **Generate entities upon** parameter to **Change in signal** from port **vc**. This causes the block to have an input signal port labeled **vc**. Also, set the **Type of value change** parameter to indicate whether the block should generate an entity whenever the signal connected to this **vc** port increases, decreases, or exhibits either type of change.



Sample Use Cases

Typically, the **vc** signal is one that you expect to change values at only a discrete set of times during the simulation. In some applications, the signal is discrete valued. Here are a few scenarios that illustrate this method for generating entities:

- The **vc** signal is an integer-valued statistical output signal from a SimEvents block. For example, the statistic could be the number of entities in a queue (shown below) or the number of entities that have departed from the block. Whenever the statistic changes values, the Event-Based Entity Generator block generates an entity.



- The **vc** signal is the output of a discrete-valued block, such as Relational Operator, Compare to Zero, or Direct Lookup Table (n-D). Whenever the logical value or the lookup table value changes, the Event-Based Entity Generator block generates an entity.



Detecting Edges in Trigger Signals

You can configure the Event-Based Entity Generator block so that it generates entities in response to rising or falling edges in a signal. This can be useful if the signal's zero crossings have some significance in your simulation; for example, a signal representing acceleration crosses zero whenever the velocity reverses direction. A signal whose rising and falling edges are used to invoke a behavior during the simulation is called a trigger signal.

A rising edge is an increase from a negative or zero value to a positive value (or zero if the initial value is negative). A falling edge is a decrease from a positive or a zero value to a negative value (or zero if the initial value is positive).

Configuring the Block to Detect Edges

To use this method for generating entities, set the Event-Based Entity Generator block's **Generate entities upon** parameter to `Trigger` from port `tr`. This causes the block to have an input signal port labeled `tr`. Also, set the **Trigger type** parameter to indicate whether the block should generate an entity whenever the signal connected to this `tr` port has a rising edge, a falling edge, or either type of edge.

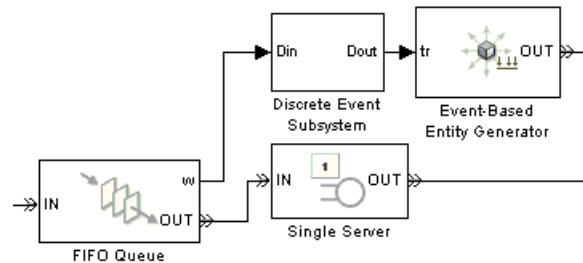
Note Do not put the Event-Based Entity Generator block inside a triggered subsystem, but rather attach the trigger signal directly to the block's `tr` port. Like other blocks that possess entity ports, the Event-Based Entity Generator block is not valid inside a triggered subsystem.

Sample Use Cases

Here are a few scenarios that illustrate this method for generating entities:


- The `tr` signal arises from the time-driven dynamics of your system. Whenever the signal crosses zero, the Event-Based Entity Generator block generates an entity.
- The `tr` signal is a real-valued statistical output signal from a `SimEvents` block, plus a negative constant. Whenever the statistic crosses a threshold that is the absolute value of the constant, the sum crosses zero and the Event-Based Entity Generator block generates an entity.

In the figure below, the Event-Based Entity Generator block generates a new entity each time the queue's average waiting time signal crosses the threshold of 5 seconds.



Top-Level Model

The subsystem between the queue and the entity generator, shown below, adds -5 to the average waiting time value to translate the threshold from 5 to 0. To understand why this computation occurs in a subsystem rather than at the top level of the model hierarchy, see “Timing Issues in SimEvents Models” on page 8-2.



Subsystem Contents

Detecting Function Calls

A function-call signal is a special type of signal that directly defines a time instant and whose typical purpose is to call a subsystem or other functional operation at that instant. A function-call signal can come from an Entity Departure Event to Function-Call Event block, a Signal-Based Event to Function-Call Event block, a Function-Call Generator block, or a Stateflow block.

You can configure the Event-Based Entity Generator block so that it generates entities in response to a function call. This can be useful for generating new entities based on the behavior of existing entities in the simulation,

generating multiple new entities simultaneously, or incorporating Stateflow dynamics into your SimEvents model.

Configuring the Block to Detect Function Calls

To use this method for generating entities, set the Event-Based Entity Generator block's **Generate entities upon** parameter to Function call from port **fcn**. This causes the block to have an input signal port labeled **fcn**. During the simulation, the Event-Based Entity Generator block generates entities in response to function calls in the signal connected to this **fcn** port.

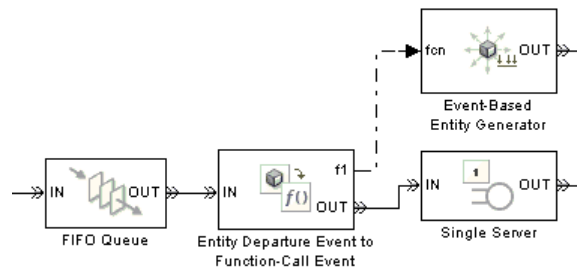
Note Do not put the Event-Based Entity Generator block inside a function-call subsystem, but rather attach the function-call signal directly to the block's **fcn** port. Like other blocks that possess entity ports, the Event-Based Entity Generator block is not valid inside a function-call subsystem.

Sample Use Cases

Here are a few scenarios that illustrate this method for generating entities:


- An Entity Departure Event to Function-Call Event block issues a function call whenever an entity departs from it. Whenever this occurs, the Event-Based Entity Generator block generates an entity.

In the figure below, the Event-Based Entity Generator block generates a new entity each time an entity departs from the queue (or, equivalently, each time an entity arrives at the server).




- A Signal-Based Event to Function-Call Event block with **ts** and **e1** input ports issues a function call whenever the **ts** signal is updated while the **e1** signal is positive. During time intervals when **e1** is positive, the Event-Based Entity Generator block can generate entities, where the specific times depend on updates of the **ts** signal. When **e1** is zero or negative, the Event-Based Entity Generator block generates no entities, even if the **ts** signal has an update. The **e1** signal provides a way to enable and disable entity generation.

In the figure below, the Event-Based Entity Generator block generates entities at integer-valued times when the queue is nonempty.




- A Function-Call Generator block issues one or more function calls periodically during the simulation. Each function call causes the Event-Based Entity Generator block to generate an entity.

In the figure below, the two Function-Call Generator blocks share the same sample time of 1 second, but the top block generates one function call at a time, while the bottom one generates three simultaneous function calls at a time. The icon changes to reflect the multiplicity of function calls. As a result, the top Event-Based Entity Generator block generates one entity each second, while the bottom one generates three entities *simultaneously* each second.



Note If you generate multiple entities simultaneously, then consider the appropriateness of other blocks in the model. For example, if three simultaneously generated entities advance to a single server, then you might want to insert a queue between the generator and the server so that entities (in particular, the second and third entities) have a place to wait for the server to become available.

- A Stateflow block issues a function-call output whenever the state takes a particular transition, as defined in the chart. Each function call causes the Event-Based Entity Generator block to generate an entity.



Using Generation Times from a Vector

If you have an explicit list of unique times at which you want to generate entities, you can configure the Time-Based Entity Generator block so that it generates entities at these times. To do this, create a vector of intergeneration times. Intergeneration times are the differences between pairs of successive time values in your list.

These topics provide instructions and motivation:

- “Configuring the Block to Generate Entities at Specified Times” on page 1-11
- “Sample Use Cases” on page 1-13

Configuring the Block to Generate Entities at Specified Times

To generate entities at specified times, follow this procedure:

- 1** Set the Time-Based Entity Generator block’s **Generate entities with** parameter to Intergeneration time from port **t**. A signal input port labeled **t** appears on the block.
- 2** Depending on whether you want to generate an entity at $T=0$, either select or clear the **Generate entity at simulation start** option in the Time-Based Entity Generator block.
- 3** Create a column vector, `gentimes`, that lists 0 followed by the nonzero times at which you want to create entities, in strictly ascending sequence. You can create this vector by entering the definition in the MATLAB® Command Window, by loading a MAT-file that you previously created, or by manipulating a variable that a To Workspace or Discrete Event Signal to Workspace block previously created.

An example of a column vector listing generation times is below.

```
gentimes = [0; 0.9; 1.7; 3.8; 3.9; 6];
```


- 4 Apply the `diff` function to the vector of generation times, thus creating a vector of intergeneration times.

```
intergentimes = diff(gentimes);
```

- 5 Create a matrix suitable for the From Workspace block. This is a two-column matrix whose first column contains time values and whose second column contains intergeneration times until the next entity generation. Remember that the vector of intergeneration times is shorter than the vector of generation times by one element because of the `diff` operation; this is why the definition of the matrix below ignores the last entry in the vector of generation times.

```
from_workspace_matrix = [gentimes(1:end-1), intergentimes];
```

- 6 Insert a From Workspace block in the model and connect it to the `t` input port of the Time-Based Entity Generator block.



- 7 In the dialog box of the From Workspace block, set **Data** to `from_workspace_matrix`.

Note If your maximum generation time is smaller than the stop time of the simulation, consider setting the **Form output after final data value by** parameter to `Holding final value`. This makes the entity generation process periodic after the simulation reaches the maximum generation time.

Sample Use Cases

Using explicit entity-generation times might be appropriate if you want to

- Recreate an earlier simulation whose intergeneration times you saved using a To Workspace block.
- Study your model's behavior under unusual circumstances and have created a series of entity generation times that you expect to produce unusual circumstances.
- Verify simulation behavior observed elsewhere, such as a result reported in a paper.

Setting Attributes of Entities

You can attach data to an entity using one or more *attributes* of the entity. Each attribute has a name and a numeric scalar value. For example, if your entities represent a message that you are transmitting across a communication network, you might assign data called *length* that indicates the length of each particular message. You can read or change the values of attributes during the simulation.

The Set Attribute block assigns up to four attributes on each arriving entity. Assignments can create new attributes or change the values of existing attributes. Attribute values can come from information you enter in the block's dialog box or from signals. To set more than four attributes on an entity, connect multiple Set Attribute blocks in series.

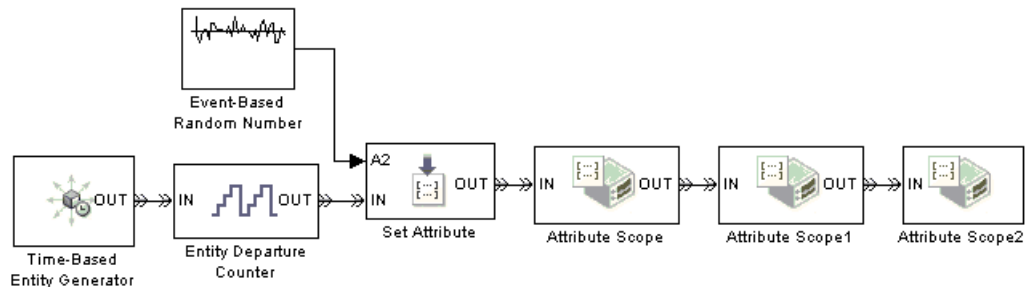
Other blocks can set particular kinds of attributes:

- The Entity Departure Counter block can set an attribute whose value is the entity count.
- The Single Server block can set an attribute whose value is the residual service time for a preempted entity. For more information, see “Preempting an Entity in a Server” on page 4-9.

To learn how to query entities for attribute values, see “Accessing Attributes of Entities” on page 1-20.

Example: Setting Attributes

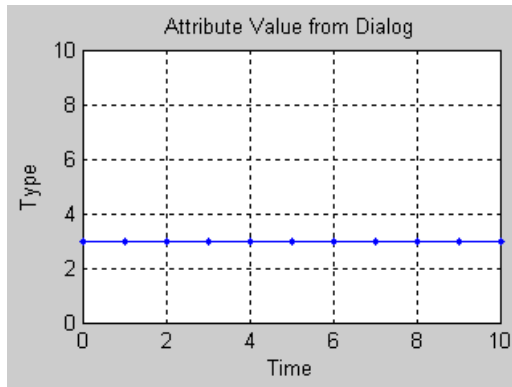

The model below illustrates different ways of assigning attribute values to entities.

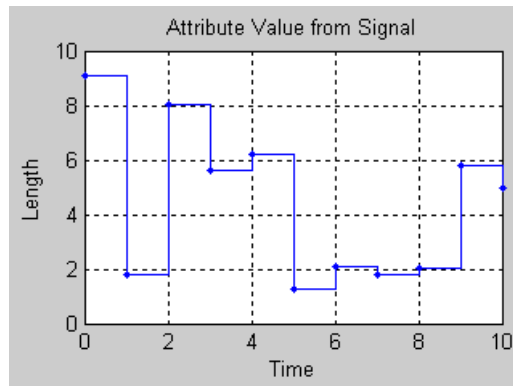


After each entity departs from the Set Attribute block, it possesses the attributes listed in the table.

Attribute Name	Attribute Value	Method for Setting Attribute Value
Count	N, for the Nth entity departing from the Time-Based Entity Generator block	In Entity Departure Counter dialog box: Write count to attribute = On Attribute name = Count Actually, the entity generator creates the Count attribute with a value of 0. The Entity Departure Counter block sets the attribute value according to the entity count.
Type	Constant value of 3	A1 tab of Set Attribute dialog box: Attribute assignment = Specify via dialog Attribute name = Type Value = 3
Length	Value of signal at output of Event-Based Random Number block	Event-Based Random Number block connected to Set Attribute block in which A2 tab of dialog box is configured as follows: Attribute assignment = From signal port A2 Attribute name = Length

In this example, each Attribute Scope block plots values of a different attribute over time. Notice from the images below that the Count plot is a stairstep graph, the Type plot is a constant, and the Length plot looks random.







When to Use Attributes

In some modeling situations, it is important to *attach* data to an entity instead of merely creating the data as the content of a signal. This section discusses the importance of considering not only the topology of your block diagrams, but also the timing of data signals appearing in SimEvents models.

Example Reusing Data

Consider a queue-server model with varying service times, where you want to plot the service time against entity count for each entity that departs from the server. A signal specifies the service time to use for each entity. Although connecting the same signal to the Signal Scope block appears correct topologically, the timing in such an arrangement is incorrect because of the delay at the server. That is, the signal has one value when a given entity arrives at the server and another value when the same entity arrives at the scope.

A correct way to implement such an example involves attaching the service time to each entity using an attribute and retrieving the attribute value from each entity upon its departure from the server. That way, the scope receives the service time associated with each entity, regardless of the delay between arrival times at the server and the scope.

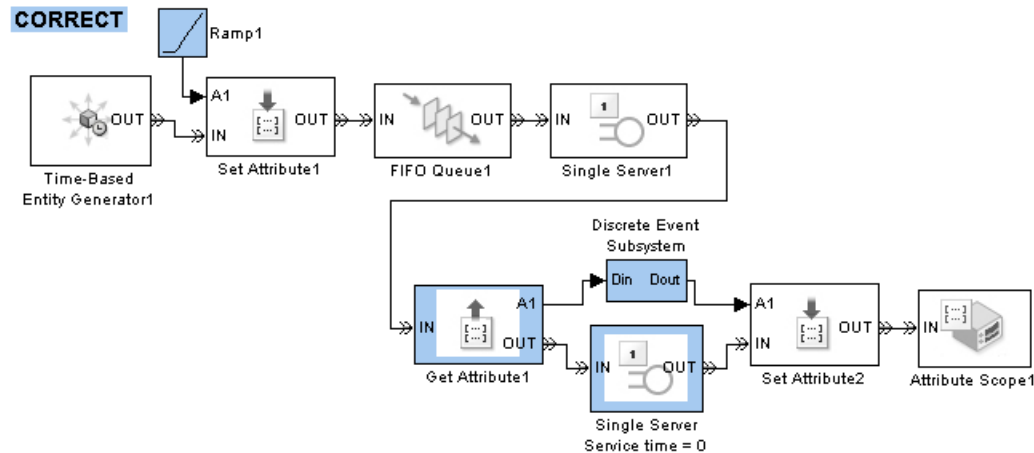
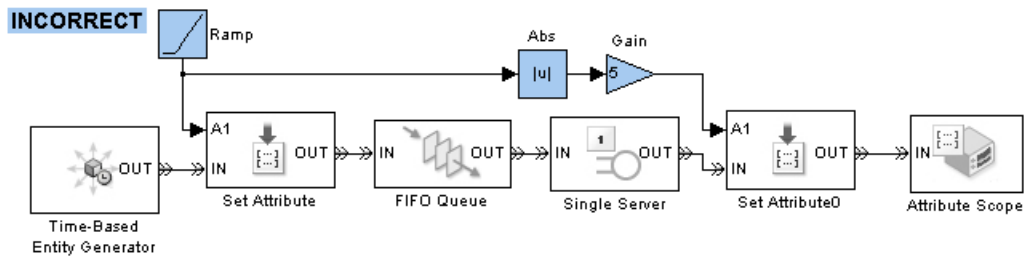


Example Manipulating Data

To manipulate the value of an attribute that you originally set using a signal, follow these rules:

- Query the value and then manipulate it, instead of creating a branch line from the signal and manipulating that. To query an entity for the value of an attribute, use the Get Attribute block.
- Perform the manipulation in a discrete event subsystem, as described in Chapter 8, “Controlling Timing Using Subsystems”. This ensures correct timing of the manipulation of the event-based signal that represents the attribute value.
- Insert a storage block with a delay of zero after the Get Attribute block, if you use the manipulated attribute value in a subsequent block on the same entity path. This ensures that the subsequent block reads the up-to-date results of the manipulation upon the entity’s arrival. For details, see “Interleaved Operations of Storage and Nonstorage Blocks” on page 12-4.

The example below illustrates the use of the Get Attribute block to query an entity for its attribute value, the use of the Discrete Event Subsystem block to contain the manipulation of the attribute value, and the insertion of a Single Server block between the Get Attribute and Set Attribute blocks.



Accessing Attributes of Entities

The section “Setting Attributes of Entities” on page 1-14 described how to use the Set Attribute block to attach data to entities near the time that you generate the entities. You can also use that block to change the value of an attribute at any point in an entity path.

To access data that has been attached to an entity, use one of these methods:

- Send the entity to a Get Attribute block. Each instance of the Get Attribute block can retrieve up to four attributes on each arriving entity and create signals with the attribute values.

For example, see the subsystem of the model described in “Adding Event-Based Behavior” in the Getting Started documentation.

- Send the entity to an Attribute Scope block and set the block’s **Y attribute name** parameter to the name of the attribute. Alternatively, send the entity to an X-Y Attribute Scope block and set the block’s **X attribute name** and **Y attribute name** parameters to the names of two attributes.

For example, see Chapter 9, “Plotting Data”.

- Name the attribute in the dialog box of a block that supports the use of attribute values for block parameters. For example, you can use an attribute value to specify the service time in the Single Server block or the selected entity output port in the Output Switch block.

For example, see “Example: Using an Attribute to Select an Output Port” in the Getting Started documentation.

Counting Entities

Counting entities can be useful for statistical measures and for understanding a simulation. This section describes techniques for counting entities in different ways, in these topics:

- “Counting Departures Across the Simulation” on page 1-21
- “Counting Departures per Time Instant” on page 1-21
- “Resetting a Counter Upon an Event” on page 1-23
- “Associating Each Entity with Its Index” on page 1-24

For troubleshooting purposes, see also “Viewing Entity Locations” on page 11-8.

Counting Departures Across the Simulation

Use the **#d** or **#a** output signal from a block to learn how many entities have departed from (or arrived at) a particular block and when their departures occurred. This method of counting is cumulative throughout the simulation. These examples use the **#d** output signal to count departures:

- “Building a Simple Discrete-Event Model” in the Getting Started documentation
- “Example: First Entity as a Special Case” on page 7-10
- “Stopping Based on Entity Count” on page 10-33

Counting Departures per Time Instant

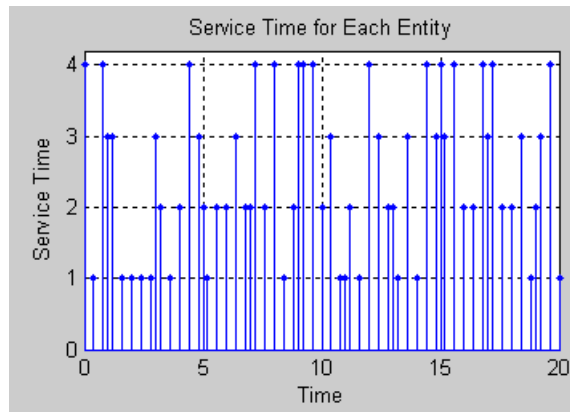
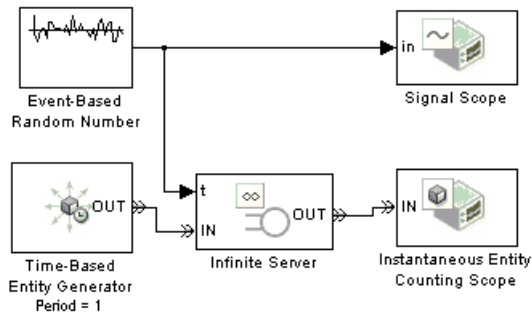
In some cases, you want to visualize how many entities have departed from a particular block and when their departures occurred, but you want to restart the counter at each time instant. This can be useful for


- Detecting simultaneous departures
- Focusing on the departure times without needing to accommodate large counts (for example, in a plot with a large range of axis values)

Use the Instantaneous Entity Counting Scope to plot the number of entities that have arrived at each time instant. The block restarts the count from 1 when the time changes. As a result, the count is cumulative for a given time instant but not cumulative across the entire simulation.

Example: Counting Simultaneous Departures from a Server

In the example below, the Infinite Server block sometimes completes service on multiple entities simultaneously. The Instantaneous Entity Counting Scope indicates how many entities departed from the server at each fixed time instant during the simulation.






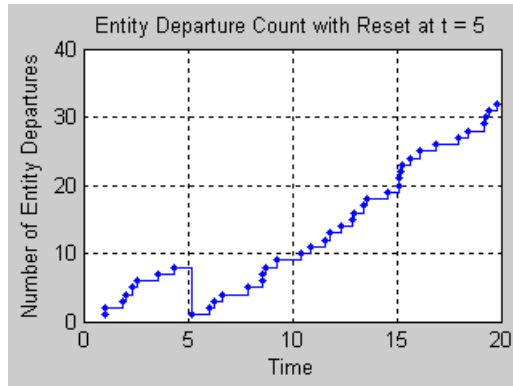
Resetting a Counter Upon an Event

Use the Entity Departure Counter block with **Reset counter upon** set to Change in signal from port `vc` or Trigger from port `tr` to count entity departures via a resettable counter. For details on this feature, see the reference page for the Entity Departure Counter block.

Example: Resetting a Counter After a Transient Period

For example, the model below counts entity departures from a queuing system but resets the counter after an initial transient period.





Associating Each Entity with Its Index


Use the Entity Departure Counter block with **Write count to attribute** set to 0n to associate an entity count with the entities that use a particular entity path. The Nth entity departing from the Entity Departure Counter block has an attribute value of n.

For an example, see “Example: Setting Attributes” on page 1-15.

For an example that illustrates when using the Entity Departure Counter block is more straightforward than storing the #d output signal in an attribute using the Set Attribute block, see “Example: Sequence of Departures and Statistical Updates” on page 12-5.

Replicating Entities on Multiple Paths

You can distribute copies of an entity on multiple entity paths using the Replicate block. Replicating entities might be a requirement of the situation you are modeling or it might be merely a convenient modeling construct. One scenario in which you might replicate entities is when copies of messages in a multicasting communication system advance to multiple transmitters or multiple recipients, as shown in the fragment below.



Alternatively, copies of computer jobs might advance to multiple computers in a cluster so that the jobs can be processed in parallel on different platforms.

Unlike the Output Switch block, the Replicate block has departures at all of its entity output ports that are not blocked, not just a single selected entity output port.

If your model routes the replicates such that they use a common entity path, then be aware that blockages can occur during the replication process. For example, connecting all ports of a Replicate block, Path Combiner block, and Single Server block in that sequence can create a blockage because the server can accommodate at most one of the replicates at a time. The blockage causes fewer than the maximum number of replicates to depart from the block.

Departure Port Precedence

Each time the Replicate block replicates an entity, the copies depart in a sequence whose start is determined by the **Departure port precedence** parameter. Although all copies depart at the same time instant, the sequence might be significant in some modeling situations. For details, see the reference page for the Replicate block.

Working with Events

Supported Events in SimEvents
Models (p. 2-2)

List and discussion of event types

Using the Event Calendar (p. 2-9)

How the event calendar determines
when and how SimEvents blocks act

Observing Events (p. 2-30)

Techniques for determining when
events occur

Generating Function-Call Events
(p. 2-36)

Generating events in an event-based
or time-based manner

Manipulating Events (p. 2-39)

Translating, combining, prioritizing,
and delaying events

Supported Events in SimEvents Models

An event is an instantaneous discrete incident that changes a state variable, an output, and/or the occurrence of other events. This section lists the supported events in SimEvents models and discusses some types of events in greater detail. The topics are

- “Types of Supported Events” on page 2-2
- “Signal-Based Events” on page 2-3
- “Function Calls” on page 2-7

Types of Supported Events

SimEvents supports the events listed below.

Event	Description
Sample time hit	Update in the value of a signal that is connected to a block configured to react to signal updates
Value change	Change in the value of a signal connected to a block that is configured to react to relevant value changes
Trigger	Rising or falling edge of a signal connected to a block that is configured to react to relevant trigger edges
Function call	Discrete invocation request carried from block to block by a special signal called a function-call signal
Entity generation	Creation of an entity
Entity destruction	Arrival of an entity at a block that has no entity output port
Entity advancement	Departure of an entity from one block and arrival at another block
Service completion	Completion of service on an entity in a server
Preemption	Replacement of an entity in a server by a higher-priority entity

Event	Description
Counter reset	Reinitialization of the counter in the Entity Departure Counter block
Gate opening or closing	Change in the state of the gate represented by the Enabled Gate or Release Gate block
Port selection	Selection of an entity input port in the Input Switch block or an entity output port in the Output Switch block
Memory writing	Writing of memory in the Signal Latch block
Memory reading	Reading of memory in the Signal Latch block

Signal-Based Events

Sample time hits, value changes, and triggers are collectively called *signal-based events*. Signal-based events can occur with respect to time-based or event-based signals. Signal-based events provide a mechanism for a block to respond to selected state changes in a signal connected to the block. The kind of state change to which the block responds determines the specific type of signal-based event.


When comparing the types of signal-based events, note that

- The updated value that results in a sample time hit could be the same as or different from the previous value of the signal.
- Event-based signals do not necessarily undergo an update or initialization at the beginning of the simulation.
- Every change in a signal value is also an update in that signal's value. However, the opposite is not true because an update that merely reconfirms the same value is not a change in the value.
- Every rising or falling edge is also a change in the value of the signal. However, the opposite is not true because a change from one positive value to another (or from one negative value to another) is not a rising or falling edge.

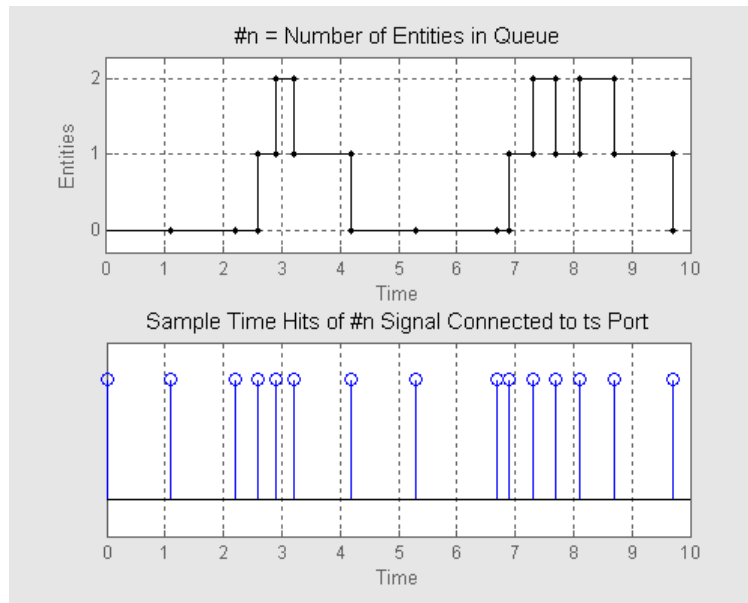
- Triggers and value changes can be rising or falling. You configure a block to determine whether the block considers rising, falling, or either type to be a relevant occurrence.
- Blocks in the Simulink libraries are more restrictive than blocks in the SimEvents libraries regarding trigger edges that rise or fall from zero. Simulink blocks in discrete-time systems do not consider a change from zero to be a trigger edge unless the signal remained at zero for more than one time step; see “Triggered Subsystems”. SimEvents blocks configured with **tr** ports consider any change from zero to a nonzero number to be a trigger edge.

Example: Comparing Types of Signal-Based Events

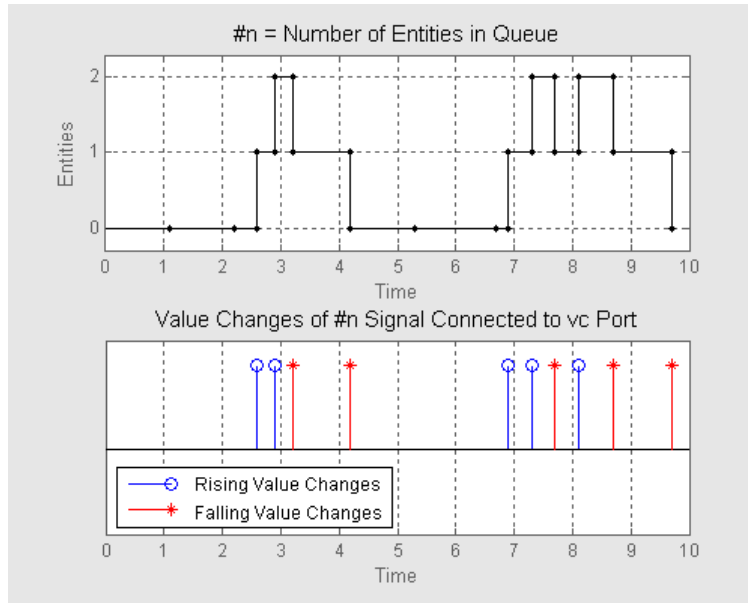
Consider the signal representing the number of entities stored in a FIFO queue. This is the **#n** output signal from the FIFO Queue block in the model below. The **#n** signal is connected to the Event-Based Entity Generator block, which reacts to different types of signal-based events. Parameters in its dialog box determine whether the block has a **ts**, **vc**, or **tr** input port, as well as the types of events to which the block reacts.



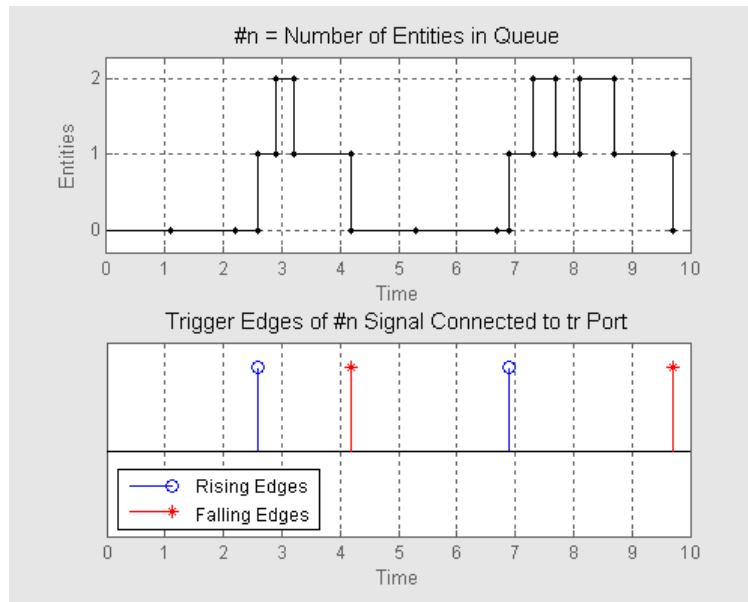
The following figures use a staircase plot to show the values of the **#n** signal and stem plots to indicate when signal-based events occur at different signal input ports (**ts**, **vc**, or **tr**) of the Event-Based Entity Generator block.



Sample Time Hits of #n Signal Connected to ts Port



Value Changes of #n Signal Connected to vc Port



Trigger Edges of #n Signal Connected to tr Port

Function Calls

Function calls are discrete invocation requests carried from block to block by a special signal called a function-call signal. A function-call signal appears as a dashed line, not a solid line. A function-call signal carries a function call at discrete times during the simulation and does not have a defined value at other times. A function-call signal is capable of carrying multiple function calls at the same value of the simulation clock, representing multiple simultaneous events.

In SimEvents models, function calls are the recommended way to make Stateflow blocks and blocks in the Simulink libraries respond to asynchronous state changes.

Function calls and signal-based events are often interchangeable in their ability to elicit reactions from various SimEvents blocks, such as the Event-Based Entity Generator block and the Signal-Based Event to Function-Call Event block.

Function-call signals can be combined, as described in “Creating a Union of Multiple Events” on page 2-40.

Using the Event Calendar

In a discrete-event simulation, state variables change instantaneously at discrete times in response to events. An event is an instantaneous incident that changes a state variable, an output, and/or the occurrence of future events. The *event calendar* lists selected upcoming events. The event calendar provides you with flexibility over the sequencing of events that occur at the same value of the simulation clock.

This section discusses which events appear on the event calendar and the processing sequence of simultaneous events. The topics are

- “Events That Appear on the Event Calendar” on page 2-9
- “Example: Event Calendar for a Queue-Server Model” on page 2-11
- “Setting Event Priorities” on page 2-19
- “Processing Sequence for Simultaneous Events” on page 2-21
- “Example: Race Conditions at a Switch” on page 2-23

To learn how to display event information in the MATLAB Command Window, see “Viewing the Event Calendar” on page 11-2.

Events That Appear on the Event Calendar

The table below indicates which events appear or might appear on the event calendar. In some cases, you have a choice.

Event	On Event Calendar	How to Cause Event to Appear on Event Calendar
Sample time hit	Never	
Value change		
Trigger		
Function call	Sometimes	Select Specify event priority... , if present, in the dialog box of the block that generates the function call. If the dialog box has no such option, then the function call is not on the event calendar.
Entity generation	Sometimes	Use the Time-Based Entity Generator block, or select Specify event priority... in the Event-Based Entity Generator block's dialog box.
Entity destruction	Never	
Entity advancement	Never	
Service completion	Always	
Preemption	Never	
Counter reset	Sometimes	Select Specify event priority... , if present, in the block's dialog box.
Gate opening or closing		
Port selection		
Memory writing	Sometimes	Select Specify event priority... on the Write tab of the block's dialog box.
Memory reading	Sometimes	Select Specify event priority... on the Read tab of the block's dialog box.

If an event does not appear on the event calendar, then the application arbitrarily decides when to process the event relative to other events occurring simultaneously, that is, at the same value of the simulation clock. When multiple unrelated events occur simultaneously, causality considerations alone do not necessarily determine a unique correct sequence.

Choosing Whether to Put Events on the Event Calendar

When you use blocks whose dialog boxes contain an option labeled **Specify event priority...**, your choice determines whether particular events appear on the event calendar. When making these choices, ask these questions about your model:

- How likely are unrelated simultaneous events?

Some models experience many state changes at deterministic intervals, so the probability of unrelated simultaneous events is considerable. Other models, such as high-level abstractions that omit details, experience state changes mostly at random times, so the probability of unrelated simultaneous events is small.

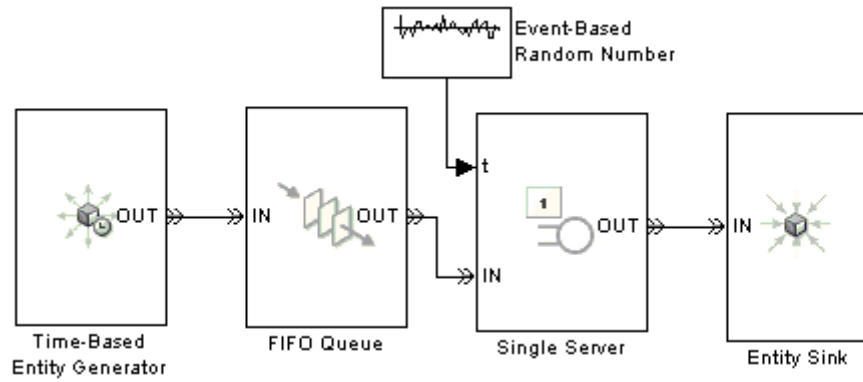
- Does the processing sequence of simultaneous events affect your simulation results, qualitatively or quantitatively?

If unrelated simultaneous events seem likely to occur and if their sequence can alter your simulation results, then you should use the event calendar to control the sequencing of such simultaneous events. Select **Specify event priority...** in those blocks that you expect to produce events simultaneously. Then use event priorities as described in “Setting Event Priorities” on page 2-19 to determine a sequence that meets your simulation needs. Simultaneous events having distinct event priorities are processed in ascending order of the event priority values.

If you do not need such control, then your simulation might run more quickly if you clear the **Specify event priority...** check box.

Example: Event Calendar for a Queue-Server Model

To see how the event calendar drives the simulation of a simple event-based model, consider the queue-server model depicted below.



Assume that the blocks are configured so that

- The Time-Based Entity Generator block generates an entity at $T = 0.9, 1.7, 3.8, 3.9,$ and 6 .
- The queue has infinite capacity.
- The server uses random service times that are uniformly distributed between 0.5 and 2.5 seconds.

The sections below indicate how the event calendar and the system's states change during the simulation.

Start of Simulation

When the simulation starts, the queue and server are empty. The entity generator schedules an event for $T = 0.9$. The event calendar looks like the table below.

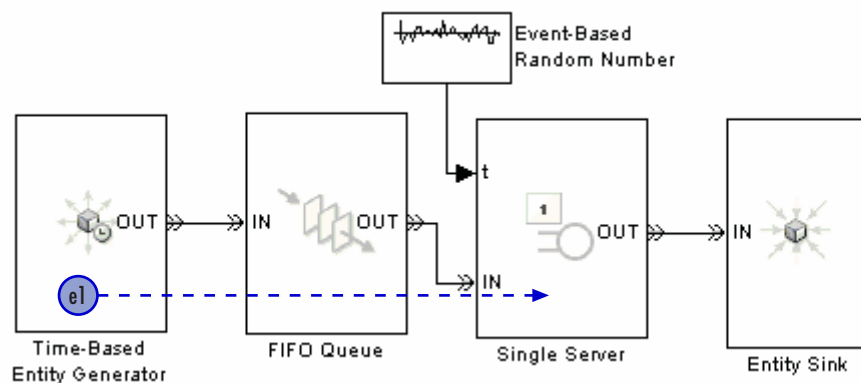
Time of Event (s)	Type of Event
0.9	Time-Based Entity Generator block generates an entity.

Generation of First Entity

At $T = 0.9$,

- The entity generator generates an entity and attempts to output it.
- The queue is empty, so the entity advances from the entity generator to the queue.
- The newly generated entity is the only one in the queue, so the queue attempts to output the entity. It queries the server to determine whether the server can accept the entity.
- The server is empty, so the entity advances from the queue to the server.
- The server's entity input port becomes temporarily unavailable to future entities.
- The server schedules an event that indicates when the entity's service time is completed. Suppose the (stochastic) service time turns out to be 1.3 in this case. This means that the duration of service is 1.3 seconds, so service is complete at $T = 2.2$.
- The entity generator schedules its next entity-generation event, at $T = 1.7$.

In the schematic below, the circled notation “e1” depicts the first entity and the dashed arrow is meant to indicate that this entity advances from the entity generator through the queue to the server.



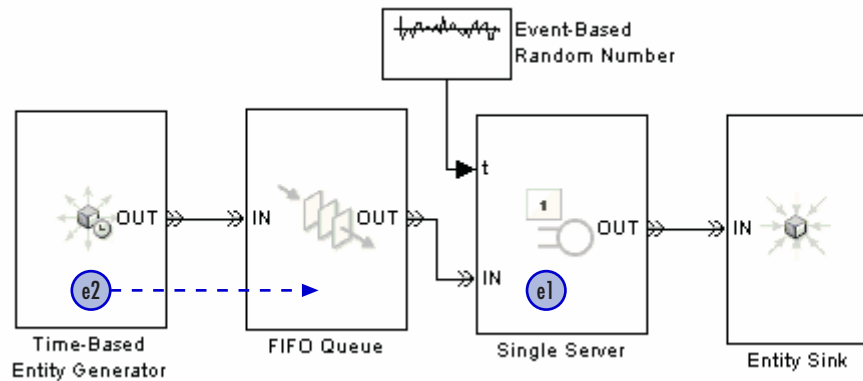
The event calendar looks like this.

Time of Event (s)	Event Description
1.7	Time-Based Entity Generator block generates second entity.
2.2	Single Server block completes service on the first entity.

Generation of Second Entity

At $T = 1.7$,

- The entity generator generates an entity and attempts to output it.
- The queue is empty, so the entity advances from the entity generator to the queue.
- The newly generated entity is the only one in the queue, so the queue attempts to output the entity. However, the server's entity input port is unavailable, so the entity stays in the queue. The queue's entity output port is said to be blocked because an entity has tried and failed to depart via this port.
- The entity generator schedules its next entity-generation event, at $T = 3.8$.



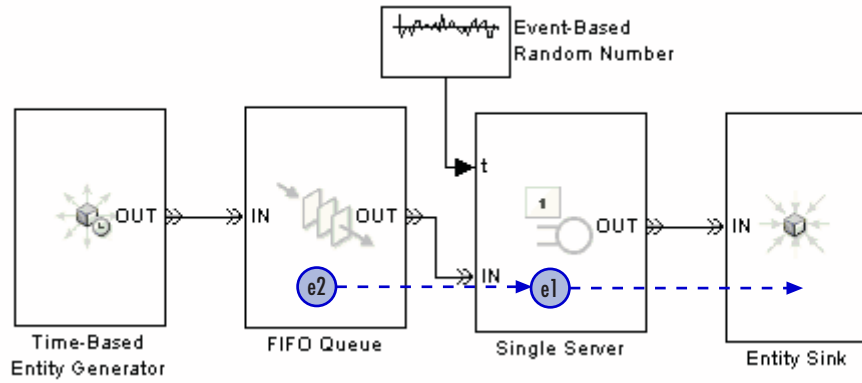
Time of Event (s)	Event Description
2.2	Single Server block completes service on the first entity.
3.8	Time-Based Entity Generator block generates the third entity.

Completion of Service Time

At $T = 2.2$,

- The server finishes serving its entity and attempts to output it. The server queries the next block to determine whether it can accept the entity.
- The sink block accepts all entities by definition, so the entity advances from the server to the sink.
- The server's entity input port becomes available.
- The queue is now able to output the second entity to the server. As a result, the queue becomes empty and the server becomes busy again.
- The server's entity input port becomes temporarily unavailable to future entities.
- The server schedules an event that indicates when the second entity's service time is completed. Suppose the service time turns out to be 2.0 in this case.

Note The server's entity input port started this time instant in the unavailable state, became available (when the first entity departed from the server), and then became unavailable once again (when the second entity arrived at the server). It is not uncommon for a state to change more than once in the same time instant.




Time of Event (s)	Event Description
3.8	Time-Based Entity Generator block generates the third entity.
4.2	Single Server block completes service on the second entity.

Generation of Third Entity

At $T = 3.8$,

- The entity generator generates an entity and attempts to output it.
- The queue is empty, so the entity advances from the entity generator to the queue.
- The newly generated entity is the only one in the queue, so the queue attempts to output the entity. However, the server's entity input port is unavailable, so the entity stays in the queue.
- The entity generator schedules its next entity-generation event, at $T = 3.9$.

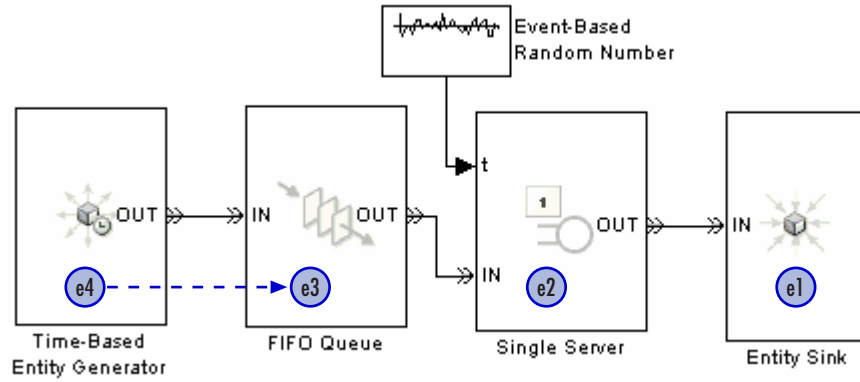


Time of Event (s)	Event Description
3.9	Time-Based Entity Generator block generates the fourth entity.
4.2	Single Server block completes service on the second entity.

Generation of Fourth Entity

At $T = 3.9$,

- The entity generator generates an entity and attempts to output it.
- The queue is not full, so the entity advances from the entity generator to the queue.
- The server's entity input port is still unavailable, so the queue cannot output an entity. The queue length is currently two.
- The entity generator schedules its next entity-generation event, at $T = 6$.



Time of Event (s)	Event Description
4.2	Single Server block completes service on the second entity.
6	Time-Based Entity Generator block generates the fifth entity.

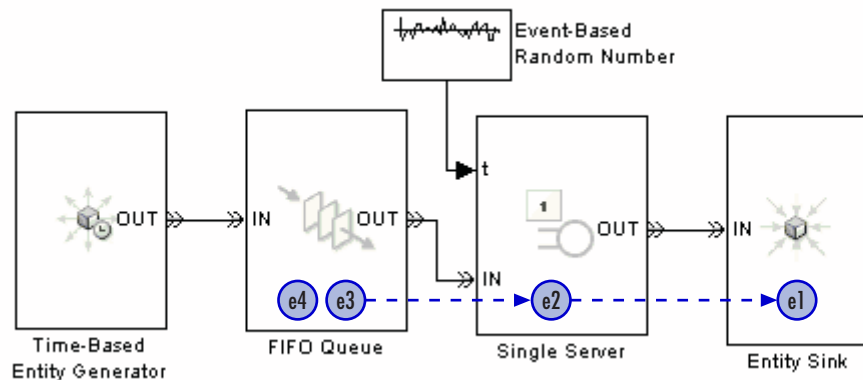
Completion of Service Time

At $T = 4.2$,

- The server finishes serving its entity and attempts to output it.
- The sink block accepts all entities by definition, so the entity advances from the server to the sink.
- The server's entity input port becomes available, so the queue's entity output port becomes unblocked. The queue is now able to output the third entity to the server. As a result, the queue length becomes one, and the server becomes busy.
- The server's entity input port becomes temporarily unavailable to future entities.
- The server schedules an event that indicates when the entity's service time is completed. Suppose the service time turns out to be 0.7 in this case.

- The queue attempts to output the fourth entity. However, the server's entity input port is unavailable, so this entity stays in the queue. The queue's entity output port becomes blocked.

Note The queue's entity output port started this time instant in the blocked state, became unblocked (when it sensed that the server's entity input port became available), and then became blocked once again (when the server began serving the third entity).



Time of Event (s)	Event Description
4.9	Single Server block completes service on the third entity.
6	Time-Based Entity Generator block generates the fifth entity

Setting Event Priorities

If the event calendar contains two or more events that are scheduled for the same time, then you can assign event priorities to influence the processing sequence of the events. Simultaneous events having distinct event priorities are processed in ascending order of the event priority values. To assign event priorities, use this procedure:

- 1 Find the block that produces the event you want to prioritize. For example, it might be an entity generator, a server, a gate, a counter, or a switch.
- 2 If the block's dialog box has an option whose name starts with **Specify event priority**, then select this option. A parameter called **Event priority** appears.
- 3 Set the **Event priority** or similarly named parameter to a positive integer.

Events that are not on the event calendar have no event priority. You cannot vary, or necessarily predict, the processing sequence of unrelated simultaneous events unless all of them appear on the event calendar.

For examples that show the effect of changing event priorities, see “Example: Race Conditions at a Switch” on page 2-23 and the Event Priorities demo.

Events with Equal Priorities

If simultaneous events on the event calendar have equal event priorities, then the application arbitrarily or randomly determines the processing sequence, depending on a modelwide configuration parameter. To access this parameter, use this procedure:

- 1 Select **Simulation > Configuration Parameters** from the model window. This opens the Configuration Parameters dialog box.
- 2 In the left pane, select **SimEvents**.
- 3 In the right pane, set **Execution order of simultaneous events** to either Randomized or Arbitrary.
 - If you select Arbitrary, the application uses an internal algorithm to determine the processing sequence for events on the event calendar that share the same time and same priority.
 - If you select Randomized, the application randomly determines the processing sequence. All possible sequences have equal probability. The **Seed for event randomization** parameter is the initial seed of the random number generator; for a given seed, the generator's output is repeatable.

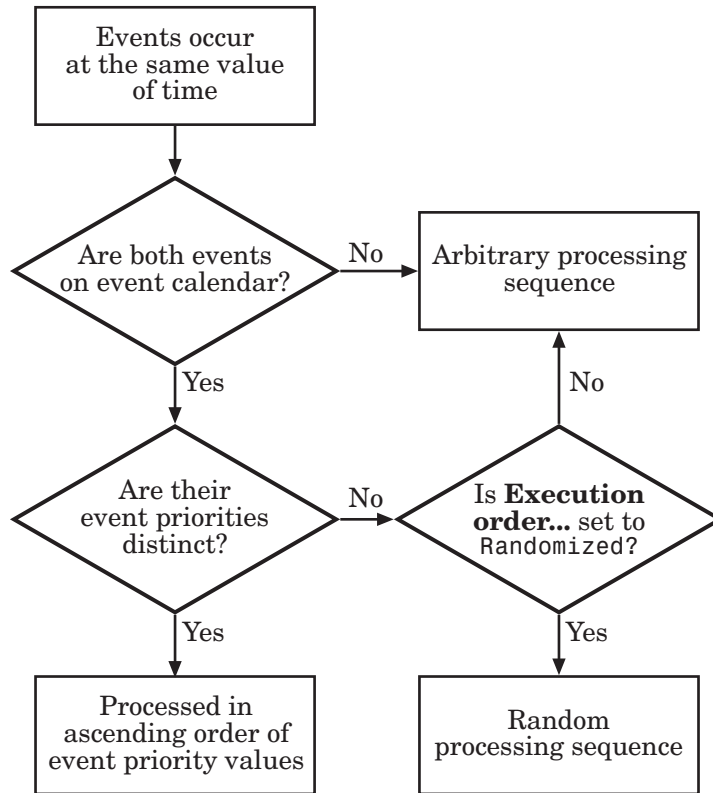
The processing sequence might be different from the sequence in which the events were scheduled on the event calendar.

Processing Sequence for Simultaneous Events

Although simultaneous events occur at the same value of the simulation clock, the application processes them sequentially. This section describes the sequence of processing when the simultaneous events are unrelated, that is, have no causal relationship to each other. Depending on the event types and on how you have designed your model, the application might determine the processing sequence

- Explicitly according to the event priority values you set. Simultaneous events having distinct event priorities are processed in ascending order of the event priority values.
- Randomly, where your choice of a seed for the random number generator provides repeatability
- Arbitrarily, using an internal algorithm

The figure below indicates when each kind of determination is in effect, under the assumption that two events with no causal relationship to each other occur at the same value of the simulation clock. The abbreviation **Execution order...** refers to the **Execution order of simultaneous events** parameter on the **SimEvents** tab of the model's Configuration Parameters dialog box. See "Events with Equal Priorities" on page 2-20 for details.



When the sequence is arbitrary, you should not make any assumptions about the sequence or its repeatability.

Choosing an Approach for Simultaneous Events

Here are some examples of situations in which you might want to use different modeling approaches:

- You want the simulation to run as fast as possible and you know that the specific processing sequence for unrelated simultaneous events has little or no effect on your simulation results. In this case, you don't vary any priority-related parameters from their default values.

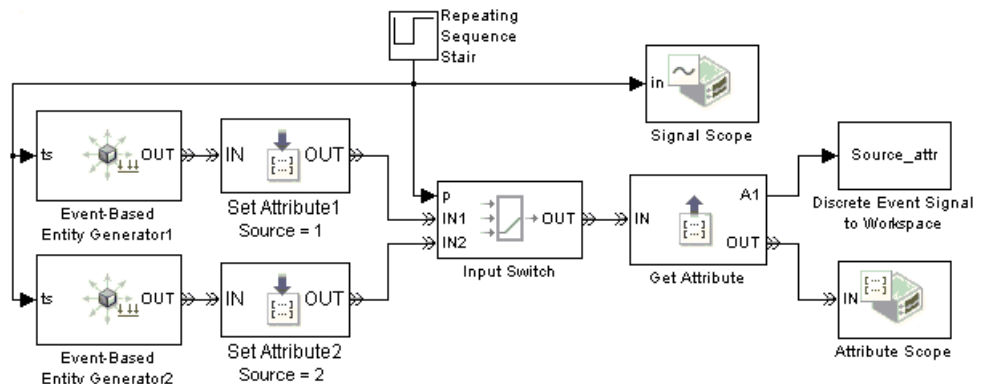
- You need to control the processing sequence for a particular set of events that might occur simultaneously. In this case, you put the events on the event calendar and choose appropriate priority values.

For example, suppose your model includes two single servers, each followed by an enabled gate, and suppose that you want to make a gate-opening event occur before a service completion event whenever the two events occur simultaneously. In this case, configure the Enabled Gate blocks by selecting **Specify event priority for gate opening and closing** and specifying **Event priority** values that are smaller numbers than the **Service completion event priority** values in the Single Server blocks.

- If you also need to have one service completion occur before the other, then make the two **Service completion event priority** values distinct. You would see the effects if the gates were followed by a Path Combiner block and a FIFO Queue block, for example.
- If you want to randomly determine which service completion occurs first if they are scheduled to occur at the same value of the simulation clock, then make the two **Service completion event priority** values equal and set **Execution order of simultaneous events** to Randomized in the model's Configuration Parameters dialog box.

Example: Race Conditions at a Switch

This example shows how you can vary the processing sequence for simultaneous events.



At $T=1, 2, 3, \dots$ the following events occur, not necessarily in this sequence:

- One or both entity generators generate an entity.
- The Repeating Sequence Stair block changes its value from 1 to 2 or vice versa, which causes the Input Switch block to select a different entity input port.

Both entity generators are configured so that if a generated entity cannot depart immediately, the generator holds the entity and temporarily suspends the generation of additional entities.

In the model, the two Set Attribute blocks assign a Source attribute to each entity, where the attribute value is 1 or 2 depending on which entity generator generated the entity. The Attribute Scope block plots the Source attribute values to indicate the source of each entity that departs from the switch.

Selecting a Port First

Suppose the model is configured so that the two entity generators and the switch have the explicit event priorities shown below.

Event Type	Event Priority
Generation event at top entity generator	300
Generation event at bottom entity generator	310
Port selection event at switch	200

At $T=1$,

- 1** The switch selects its **IN2** entity input port.
- 2** The top entity generator generates an entity, which cannot depart because the switch's **IN1** entity input port is unavailable.
- 3** The bottom entity generator generates an entity, which advances from block to block until it reaches the Attribute Scope block.


At T=2,

- 1** The switch selects its **IN1** entity input port. This causes the top entity generator to output the entity it generated 1 second ago. This entity advances from block to block until it reaches the Attribute Scope block.
- 2** The top entity generator generates an entity, which advances from block to block until it reaches the Attribute Scope block. A total of two entities from the top entity generator reach the scope at this time instant.
- 3** The bottom entity generator generates an entity, which cannot depart because the switch's **IN2** entity input port is unavailable.


At T=3,

- 1** The switch selects its **IN2** entity input port. This causes the bottom entity generator to output the entity it generated 1 second ago. This entity advances from block to block until it reaches the Attribute Scope block.
- 2** The top entity generator generates an entity, which cannot depart because the switch's **IN1** entity input port is unavailable.
- 3** The bottom entity generator generates an entity, which advances from block to block until it reaches the Attribute Scope block. A total of two entities from the bottom entity generator reach the scope at this time instant.

The plot of entities' Source attribute values shows an alternating pattern of dots, as does the plot of the port selection signal **p**. The list of times and values of the entities' Source attribute, as recorded in the `Source_attr` variable in the MATLAB workspace, shows that two entities from the same entity generator reach the scope at T=2, 3, 4, etc.



Port Selection Signal



Switch Departures When Port Selection Is Processed First

```
[Source_attr.time, Source_attr.signals.values]
```

```
ans =
```

```
1     2
2     1
2     1
3     2
3     2
4     1
4     1
5     2
5     2
6     1
6     1
7     2
7     2
8     1
8     1
9     2
9     2
10    1
10    1
```

Generating Entities First

Suppose the model is configured so that the two entity generators and the switch have the explicit event priorities shown below.

Event Type	Event Priority
Generation event at top entity generator	300
Generation event at bottom entity generator	310
Port selection event at switch	4000

At the beginning of the simulation, the port selection signal **p** is 1.


At T=1,

- 1 The top entity generator generates an entity, which advances from block to block until it reaches the Attribute Scope block.
- 2 The bottom entity generator generates an entity, which cannot depart because the switch's **IN2** entity input port is unavailable.
- 3 The switch selects its **IN2** entity input port. This causes the bottom entity generator to output the entity it just generated. This entity advances from block to block until it reaches the Attribute Scope block.

At T=2,

- 1 The top entity generator generates an entity, which cannot depart because the switch's **IN1** entity input port is unavailable.
- 2 The bottom entity generator generates an entity, which advances from block to block until it reaches the Attribute Scope block.
- 3 The switch selects its **IN1** entity input port. This causes the top entity generator to output the entity it just generated. This entity advances from block to block until it reaches the Attribute Scope block.

The plot of entities' Source attribute values shows that two entities from different entity generators depart from the switch every second.



```
[Source_attr.time, Source_attr.signals.values]
```

```
ans =
```

```
1      1
1      2
2      2
2      1
3      1
3      2
4      2
4      1
5      1
5      2
6      2
6      1
7      1
7      2
8      2
8      1
9      1
9      2
10     2
10     1
```


Switch Departures When Entity Generations Are Processed First

Randomly Selecting a Sequence

Suppose the model is configured so that the two entity generators and the switch have equal event priorities. By default, the application uses an arbitrary processing sequence for the entity generation events and the port selection events, which might or might not be appropriate in an application. To avoid bias by randomly determining the processing sequence for these

events, set **Execution order of simultaneous events** to Randomized in the model's Configuration Parameters dialog box.

Sample attribute values and the corresponding plot are below, but your results might vary depending on the specific random numbers.



Switch Departures When Processing Sequence is Random

```
[Source_attr.time, Source_attr.signals.values]
```

```
ans =
```

1	2
2	2
2	1
3	2
4	1
4	1
5	1
5	2
5	2
6	1
7	1
7	2
8	2
8	1
9	2
10	1
10	1

Observing Events

The event logging feature described in “Viewing the Event Calendar” on page 11-2 can help you observe events that appear on the event calendar. The table below suggests some ways to observe events that do not appear on the event calendar. Key tools are the Instantaneous Event Counting Scope block, Signal Scope block, and Discrete Event Signal to Workspace block. You can also build a discrete event subsystem that counts events and creates a signal, as illustrated in “Example: Focusing on Events, Not Values” on page 8-22.

Note Do not select **Specify event priority...** options in block dialog boxes for the sole purpose of using the event logging feature to observe those events. When you have a choice about whether an event appears on the event calendar, your decision might affect the processing sequence of simultaneous events and hence the simulation behavior.

Event	Observation Technique
Sample time hit	Use a branch line to connect the signal to an Instantaneous Event Counting Scope block.
Value change	
Trigger	
Function call	If the block issuing the function call provides a #f1 output signal, then observe its increases. Otherwise, configure a Signal-Based Function-Call Event Generator block by enabling the #f1 output port and setting Generate function call only upon to Function call from port fcn; then insert this block between the block issuing the function call and the block reacting to the function call.
Entity generation	Observe values of the entity generator’s pe output signal. Upon an entity generation, pe repeats a previous value of 0 (if the generated entity departs immediately) or increases from 0 to 1 (if the entity cannot depart).
Entity destruction	Observe increases in the #a output signal. The Instantaneous Entity Counting Scope block provides a plot in place of a #a signal.

Event	Observation Technique
Entity advancement	Observe increases in the #d output signal of the block from which the entity departs. Alternatively, use the entity logging feature described in “Viewing Entity Locations” on page 11-8.
Service completion	Observe values of the server’s #pe (or pe in the case of the Single Server block) output signal. Upon a service completion, the signal repeats a previous value of 0 (if the entity departs immediately) or increases (if the entity cannot depart). Because simultaneous repeated values would not be easy to observe on a plot, it is best to observe the signal value in the MATLAB workspace if a departure and a new service completion might occur simultaneously in your model.
Preemption	Observe increases in the #p output signal of the server block.
Counter reset	Observe falling edges in the counter block’s #d output signal. Alternatively, use a branch line to connect the counter block’s input signal to an Instantaneous Event Counting Scope block.
Gate opening or closing	Use a branch line to connect the gate block’s input signal to an Instantaneous Event Counting Scope block. In the case of an enabled gate, rising trigger edges of the input signal indicate gate opening events, while falling trigger edges of the input signal indicate gate closing events.
Port selection	If the block has a p input signal, use a branch line to connect the p signal to an Instantaneous Event Counting Scope block, configured to plot value changes. Otherwise, observe the block’s last output signal.
Memory writing	Observe sample time hits in the Signal Latch block’s mem output signal.
Memory reading	Observe sample time hits in the Signal Latch block’s out output signal.


For examples that use one or more of these techniques, see

- “Example: Plotting Event Counts to Find Roundoff Error” on page 9-12
- “Example: Entity Logging” on page 11-9
- “Example: Observing Service Completions” on page 2-32
- “Example: Focusing on Events, Not Values” on page 8-22

Also, “Example: Detecting Collisions by Comparing Events” on page 2-34 shows how to use a Signal Latch block to observe which of two types of events occurred more recently.

Example: Observing Service Completions

The model below writes an N-Server block’s **#pe** signal to a structure called **pe** in the MATLAB workspace. In the structure, **pe.time** indicates when each value is attained and **pe.signals.values** indicates the corresponding value.



An entity completes its service at precisely those times when the **#pe** signal repeats its previous value or increases to a larger value. After the simulation is over, you can form a vector of service completion times using the code below.

```
% Output #pe times and values.
pe_matrix = [pe.time, pe.signals.values]

% Determine when #pe changes its value.
dpe = [0; diff(pe.signals.values)];
```



```
% Service completions occur when #pe does not decrease.  
t_svcp = pe.time(dpe >= 0)
```

Sample output, which depends on the entity generation, service completion, gate opening, and gate closing times in the model, is below. Notice the rows of `pe_matrix` in which the first column is 2.0000 and the second column has decreasing values. These correspond to decreases in the **#pe** signal value at $T=2$, which result from departures of entities that previously completed their service. The value $T=2$ does not appear in the `t_svcp` vector because this is not a time at which service completion events occur.

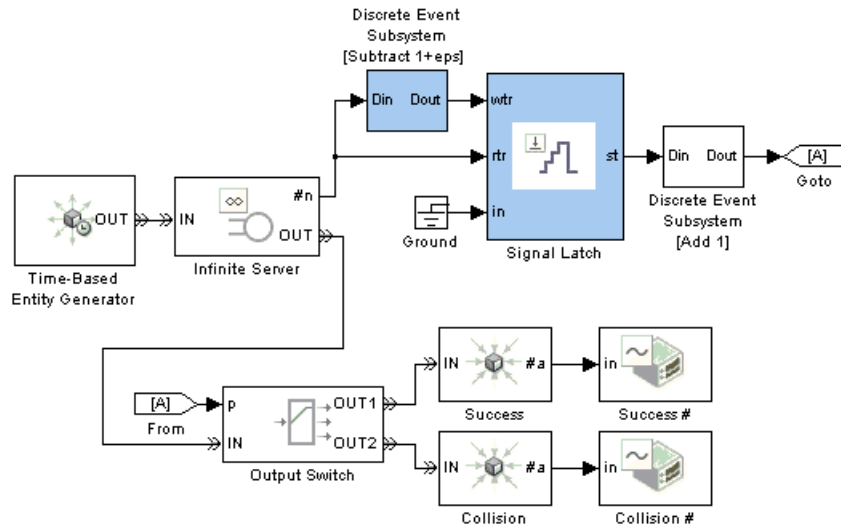
```
pe_matrix =  
  
    0.9077    1.0000  
    1.0849    2.0000  
    1.8895    3.0000  
    2.0000    2.0000  
    2.0000    1.0000  
    2.0000         0  
    2.4487         0  
    3.0707         0  
    3.1958         0  
    3.4068         0  
    3.5856         0  
    3.7906         0  
    4.3674         0  
    4.8672    1.0000  
    4.9545    2.0000
```

```
t_svcp =  
  
    0.9077  
    1.0849  
    1.8895  
    2.4487  
    3.0707  
    3.1958  
    3.4068
```

- 3.5856
- 3.7906
- 4.3674
- 4.8672
- 4.9545

Example: Detecting Collisions by Comparing Events

The model below aims to determine whether an entity is the entity in an infinite server for the entire duration of service. The model uses the Signal Latch block to compare the times of two kinds of events and report which kind occurred more recently. This usage of the Signal Latch block relies on the block's status output signal, **st**, rather than the default **in** and **out** ports.



In the model, entities arrive at an infinite server, whose **#n** output signal indicates how many entities are in the server. The Signal Latch block responds to these signal-based events involving the integer-valued **#n** signal:

- If **#n** increases from 0 to a larger integer, then
 - **rtr** has a rising edge.
 - The Signal Latch block processes a read event.
 - The Signal Latch block's **st** output signal becomes 0.
- If **#n** increases from 1 to a larger integer, then
 - **wtr** has a rising edge.
 - The Signal Latch block processes a write event.
 - The Signal Latch block's **st** output signal becomes 1.
- If **#n** increases from 0 to 2 at the same value of the simulation clock, then it also assumes the value 1 as a zero-duration value. As a result,
 - **rtr** and **wtr** both have rising edges, in that sequence.
 - The Signal Latch block processes a read event followed by a write event.
 - The Signal Latch block's **st** output signal becomes 1.

By the time the entity departs from the Infinite Server block, the Signal Latch block's **st** signal is 0 if and only if that entity has been the only entity in the server block for the entire duration of service. This outcome is considered a success for that entity. Other outcomes are considered collisions between that entity and one or more other entities.

Generating Function-Call Events

You can generate an event and use it to

- Invoke a discrete event subsystem or a Stateflow block
- Cause certain events, such as the opening of a gate or the reading of memory in a Signal Latch block
- Generate an entity

For most purposes, a function call is an appropriate type of event to generate.

Note While you can invoke triggered subsystems and Stateflow blocks upon trigger edges, this method has limitations in discrete-event simulations. In particular, you should use function calls instead of trigger edges if you want the invocations to occur asynchronously, to be prioritized among other simultaneous events, or to occur more than once in a fixed time instant.

These topics describe how to generate function calls in an event-based or time-based manner:

- “Generating Events When Other Events Occur” on page 2-36
- “Generating Events Using Intergeneration Times” on page 2-38

Generating Events When Other Events Occur

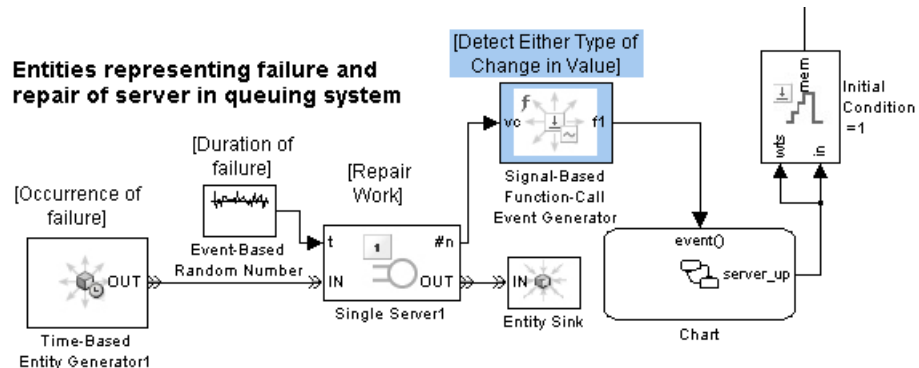
The table below indicates which blocks generate function calls when other events occur.

Event Upon Which to Generate Another Event	Block
Entity advancement	Entity-Based Function-Call Event Generator
Signal-based event	Signal-Based Function-Call Event Generator
Function call	Signal-Based Function-Call Event Generator

Example: Calling a Stateflow Block Upon Changes in Server Contents

The fragment below, which is part of an example in “Using Stateflow to Implement a Failure State” on page 4-16, uses entities to represent failures and repairs of a server elsewhere in the model:

- A failure of the server is modeled as an entity’s arrival at the block labeled Repair Work. When the Repair Work block’s **#n** signal increases to reflect the entity arrival, the Signal-Based Function-Call Event Generator block generates a function call that calls the Stateflow block to change the state of the server from up to down.
- A completed repair of the server is modeled as an entity’s departure from the Repair Work block. When the Repair Work block’s **#n** signal decreases to reflect the entity departure, the Signal-Based Function-Call Event Generator block generates a function call that calls the Stateflow block to change the state of the server from down to up.



One reason to use function calls rather than trigger signals to call a Stateflow block in discrete-event simulations is that an event-based signal can experience a trigger edge due to a zero-duration value that a time-based block would not recognize. The Signal-Based Function-Call Event Generator can detect signal-based events that involve zero-duration values.

Generating Events Using Intergeneration Times

To generate events using intergeneration times from a signal or a statistical distribution, use this procedure:


- 1 Use the signal or statistical distribution with the Time-Based Entity Generator block to generate entities.
- 2 Use the Entity-Based Function-Call Event Generator block to generate an event associated with each entity.
- 3 Terminate the entity path with an Entity Sink block.

In the special case when the intergeneration time is constant, a simpler alternative is to use the Function-Call Generator block in the Simulink Ports & Subsystems library.

Example: Opening a Gate Upon Random Events

The model below uses the top entity generator to generate entities whose sole purpose is to cause the generation of events with intergeneration times from a statistical distribution. The bottom entity generator generates entities that enter a gated queuing system.

Generating Events with Random Intergeneration Times



Manipulating Events

You can manipulate an event to accomplish any of these goals:

- To invoke a function-call subsystem or Stateflow block upon entity departures or signal-based events.

Note You can invoke triggered subsystems and Stateflow blocks upon trigger edges, which are a type of signal-based event. However, you will need to translate the trigger edges into function calls if you want the invocations to occur asynchronously, to be prioritized among other simultaneous events, or to occur more than once in a fixed time instant.

- To create a union of events from multiple sources. See “Creating a Union of Multiple Events” on page 2-40.
- To prioritize the reaction to an event relative to simultaneous events. See “Translating Events to Control the Processing Sequence” on page 2-43.
- To delay the reaction to an event. See the **Function-call time delay** parameter on the Signal-Based Event to Function-Call Event block’s reference page.
- To conditionalize the reaction to an event. See “Conditionalizing Events” on page 2-45.

The term *event translation* refers to the conversion of one event into another. The result of the translation is often a function call, but can be another type of event. The result of the translation can occur at the same time as, or a later time than, the original event.

Blocks for Manipulating Events

The table below lists blocks that are useful for manipulating events.

Event to Manipulate	Block
Entity advancement	Entity Departure Event to Function-Call Event
Signal-based event	Signal-Based Event to Function-Call Event
Function call	Signal-Based Event to Function-Call Event
	Mux

If you connect the Entity Departure Counter block's **#d** output port to a block that detects sample time hits or rising value changes, then you can view the counter as a mechanism for converting an entity advancement event into a signal-based event. Corresponding to each entity departure from the block is an increase in the value of the **#d** signal.

Creating a Union of Multiple Events

To generate a function-call signal that represents the union (logical OR) of multiple events, use this procedure:

- 1** Generate a function call for each event that is not already a function call. Use blocks in the Event Generators or Event Translation library.
- 2** Use the Mux block to combine the function-call signals.

The multiplexed signal carries a function call when any of the individual function-call signals carries a function call. If two individual signals carry a function call at the same time instant, then the multiplexed signal carries two function calls at that time instant.

Examples are in “Example: Performing a Computation on Selected Entity Paths” on page 8-31 and below.


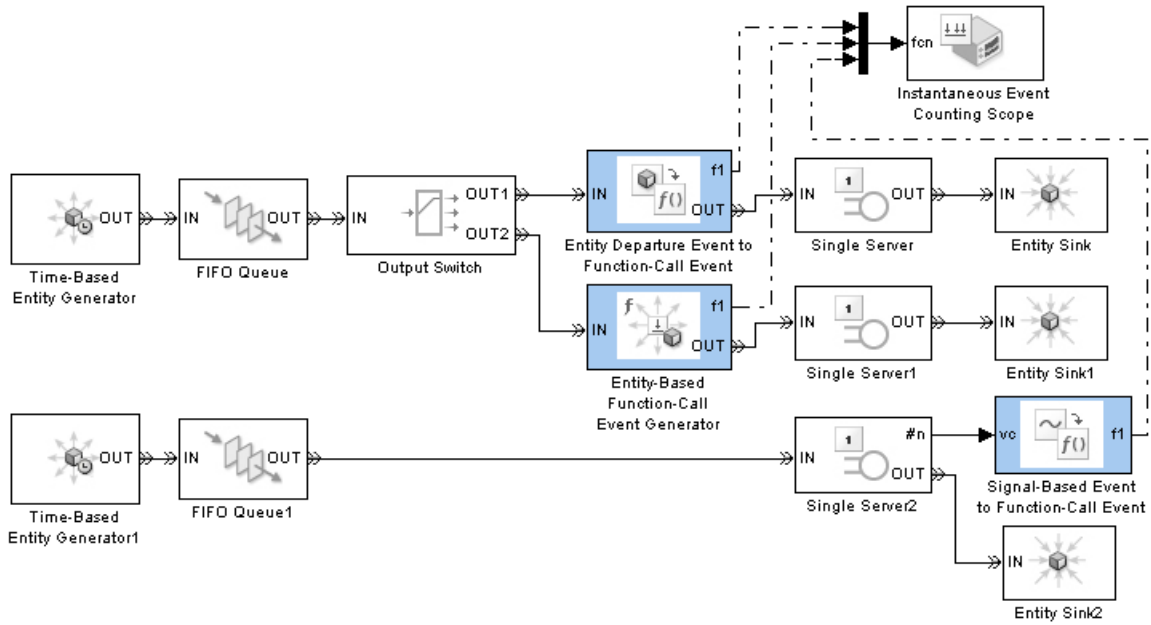
Example: Counting Events from Multiple Sources

The example below illustrates different approaches to event translation and event generation. This example varies the approach for illustrative purposes; in your own models, you might decide to use a single approach that you find most intuitive.

The goal of the example is to plot the number of arrivals at a bank of three servers at each value of time. Entities advance to the servers via one or two FIFO Queue blocks. To count arrivals and create the plot, the model translates each arrival at a server into a function call; the Mux block combines the three function-call signals to create an input to the Instantaneous Event Counting Scope block.

The three server paths use these methods for translating an entity arrival into a function call:

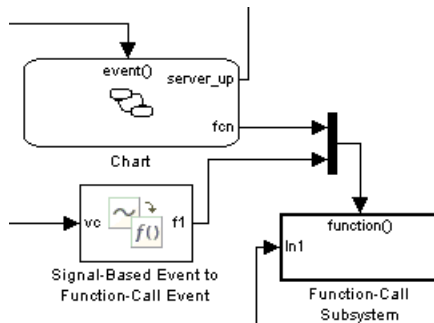
- One path uses the Entity Departure Event to Function-Call Event block, treating the problem as one of event translation.
- One path uses the Entity-Based Event Generator block, treating the problem as one of event generation. This is similar to the approach above.
- One path uses the Signal-Based Event to Function-Call Event block to translate an increase in the value of the server block's **#n** signal into a function call. This approach uses the fact that each arrival at the server block causes a simultaneous increase in the block's **#n** signal.



Example: Executing a Subsystem Based on Multiple Types of Events

You can configure a Discrete Event Subsystem block to detect signal-based events from one or more sources, and you can configure a Function-Call Subsystem block to detect function calls from one or more sources. Using an event translation block to convert a signal-based event into a function call, the fragment below effectively creates a subsystem that detects a function call from a Stateflow block and a signal-based event from another source. The subsystem is executed when either the Stateflow block generates a function call or the signal connected to the **vc** port of the Signal-Based Event to Function-Call Event block changes. If both events occur simultaneously, then the subsystem executes twice.

“Block execution” in this documentation is shorthand for “block methods execution.” Methods are functions that Simulink uses to solve a model. Blocks are made up of multiple methods. For details, see “Block Methods” in the Simulink documentation.




Another similar example is in “Example: Performing a Computation on Selected Entity Paths” on page 8-31.

Translating Events to Control the Processing Sequence

In some situations, event translation blocks can help you prescribe the processing sequence for simultaneous events. The examples below illustrate how to do this by taking advantage of the sequence in which an event translation block issues two function calls, and by converting an unprioritized function call into a function call having an event priority.

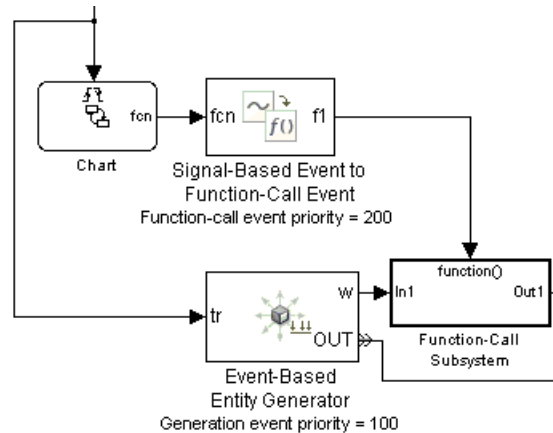
Example: Issuing Two Function Calls in Sequence

In the model below, entity generation and the execution of a function-call subsystem can occur simultaneously. At such times, it is important that the entity generation occur first, so that the entity generator updates the value of the **w** signal before the function-call subsystem uses **w** in its computation. This model ensures a correct processing sequence by using the same Signal-Based Event to Function-Call Event block to issue both function calls and by relying on the fact that the block always issues the **f1** function call before the **f2** function call.



Example: Generating a Function Call with an Event Priority

The model below uses an event translation block to prioritize the execution of a function-call subsystem correctly on the event calendar, relative to a simultaneous event. In the model, a Stateflow block and an entity generator respond to edges of the same trigger signal. The Stateflow block calls an event translation block, which in turn calls a function-call subsystem. The subsystem performs a computation using the **w** output signal from the entity generator.



As in the earlier example, it is important that the entity generator update the value of the w signal before the function-call subsystem uses w in its computation. To ensure a correct processing sequence, the Signal-Based Event to Function-Call Event block replaces the original function call, which is not on the event calendar, with a new function call that appears on the event calendar with a priority of 200. The Event-Based Entity Generator block creates an entity generation event on the event calendar with a priority of 100. As a result of the event translation and the relative event priorities, the entity generator generates the entity before the event translator issues the function call to the function-call subsystem whenever these events occur at the same value of the simulation clock.



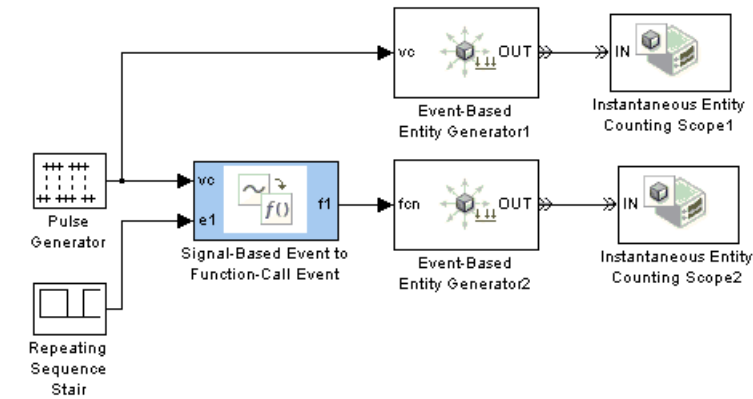
Conditionalizing Events

The Entity Departure Event to Function-Call Event and Signal-Based Event to Function-Call Event blocks provide a way to suppress the output function call based on a control signal. If the control signal is zero or negative when the block is about to issue the function call, then the block suppresses the function call. You can use this feature to

- Prevent simulation problems. The example in “Example: Detecting Changes in the Last-Updated Signal” on page 3-10 uses conditional function calls to prevent division-by-zero warnings.
- Model an inoperative state of a component of your system. See the next example.

Example: Modeling Periodic Shutdown of an Entity Generator

The model below uses Event-Based Entity Generator blocks to generate entities when a pulse signal changes its value. The top entity generator generates an entity upon each such event. The bottom entity generator responds to a function call issued by an event translation block that detects changes in the pulse signal's value. However, the event translation block issues a function call only upon value changes that occur while the **e1** input signal is positive. In this model, a nonpositive value of the **e1** signal corresponds to a failure or resting period of the entity generator.



Working with Signals

Role of Event-Based Signals in SimEvents Models (p. 3-2)	Overview of event-based signals and issues involving them
Generating Random Signals (p. 3-4)	Producing random numbers in an event-based or time-based manner
Sequence of Updates of Output Signals (p. 3-10)	Signal updates relative to each other
Multiple Simultaneous Updates (p. 3-13)	Working with zero-duration values
Latency in Signal Updates (p. 3-17)	Delays in signal updates or reactions to updates
Manipulating Signals (p. 3-18)	Using the Signal Latch block to delay or resample signals
Sending Data to the MATLAB Workspace (p. 3-23)	Collecting data from event-based signals for manipulation in MATLAB

Role of Event-Based Signals in SimEvents Models

Discrete-event simulations often involve signals that change when events occur; for example, the number of entities in a server is a statistical output signal from a server block and the signal value changes when an entity arrives at or departs from the server. An event-based signal is a signal that can change in response to discrete events. Most output signals from SimEvents blocks are event-based signals.

Comparison With Time-Based Signals

Unlike time-based signals, event-based signals

- Do not have a true sample time. (These are *not* continuous signals, even though the sample time coloration feature makes the signal connection line black or gray, and a Probe block reports a sample time of zero.)
- Might be updated at time instants that do not correspond to time steps determined by time-based dynamics.
- Might undergo multiple updates in a single time instant.
- Might have the first update upon the first relevant event, not necessarily at the beginning of the simulation.

For example, consider a signal representing the number of entities in a server. Computing this value at fixed intervals is wasteful if no entities arrive or depart for long periods. Computing the value at fixed intervals is inaccurate if entities arrive or depart in the middle of an interval, because the computation misses those events. Simultaneous events can make the signal multivalued; for example, if an entity completes its service and departs, which permits another entity to arrive at the same time instant, then the count at that time equals both 0 and 1 at that time instant. Furthermore, if an updated value of the count signal causes an event, then the processing of the signal update relative to other operations at that time instant can affect the processing sequence of simultaneous events and change the behavior of the simulation.

When you use output signals from SimEvents blocks to examine the detailed behavior of your system, you should understand when the blocks update the signals, including the possibility of multiple simultaneous updates. When you use event-based signals for controlling the dynamics of the simulation,

understanding when blocks update the signals and when other blocks react to the updated values is even more important.

Note Blocks in the SimEvents libraries process signals whose data type is double. To convert between data types, use the Data Type Conversion block in the Simulink Signal Attributes library.

Generating Random Signals

Discrete-event simulations often use random numbers for entity intergeneration times, service times, routing, and other purposes. An important block for generating random signals is the Event-Based Random Number block. These topics describe how to use this block to produce random signals:

- “Generating Random Event-Based Signals” on page 3-4
- “Examples of Random Event-Based Signals” on page 3-6
- “Generating Random Time-Based Signals” on page 3-8

For details on the connectivity restrictions of the Event-Based Random Number block, see its reference page.

Generating Random Event-Based Signals

The Event-Based Random Number block is designed to create event-based signals using a variety of distributions. The block generates a new random number from the distribution upon notifications from a port of a subsequent block. For example, when connected to the **t** input port of a Single Server block, the Event-Based Random Number block generates a new random number each time it receives notification that an entity has arrived at the server.

Implied Timing of Random Number Generation

Notifying ports, listed in the table below, notify the preceding block when a certain event has occurred. When the preceding block is the Event-Based Random Number block, it responds to the notification by generating a new random number. The Event-Based Random Number block must be connected to exactly one notifying port.

Notifying Ports

Signal Input Port	Block	Generate New Random Number Upon
A1, A2, A3, A4	Set Attribute	Entity arrival
in	Signal Latch	Write event
e1, e2	Entity Departure Event to Function-Call Event	Entity arrival
	Signal-Based Event to Function-Call Event	Relevant signal-based event, depending on configuration of block
t	Infinite Server	Entity arrival
	N-Server	Entity arrival
	Single Server	Entity arrival
t	Time-Based Entity Generator	Simulation start and subsequent entity departures
x	X-Y Signal Scope	Sample time hit at in signal input port

To observe signal values, you can optionally use a branch line to connect the Event-Based Random Number block to one or more monitoring ports. Monitoring ports are listed on the reference page for the Event-Based Random Number block. Connections from the Event-Based Random Number block to ports other than notifying ports and monitoring ports are not supported.

Generating Random Signals Based on Arbitrary Events

When generating random event-based signals other than the cases covered in “Implied Timing of Random Number Generation” on page 3-4, you should use the Signal Latch block to indicate explicitly which events cause the Event-Based Random Number block to generate a new random number. Use this procedure:

- 1 Insert an Event-Based Random Number block into your model and configure it to indicate the distribution and parameters you want to use.

- 2 Insert a Signal Latch block and set **Read from memory upon** to Write to memory event. The block no longer has an **rvc** signal input port.
- 3 Determine which events should result in the generation of a new random number, and set the Signal Latch block's **Write to memory upon** accordingly.
- 4 Connect the signal whose events you identified in the previous step to the write-event port (**wts**, **wvc**, **wtr**, or **wfcn**) of the Signal Latch block. Connect the Event-Based Random Number block to the **in** port of the Signal Latch block.

The **out** port of the Signal Latch block is the desired random event-based signal.

Examples of Random Event-Based Signals

For examples using the Event-Based Random Number block in some of the contexts described in “Implied Timing of Random Number Generation” on page 3-4, see

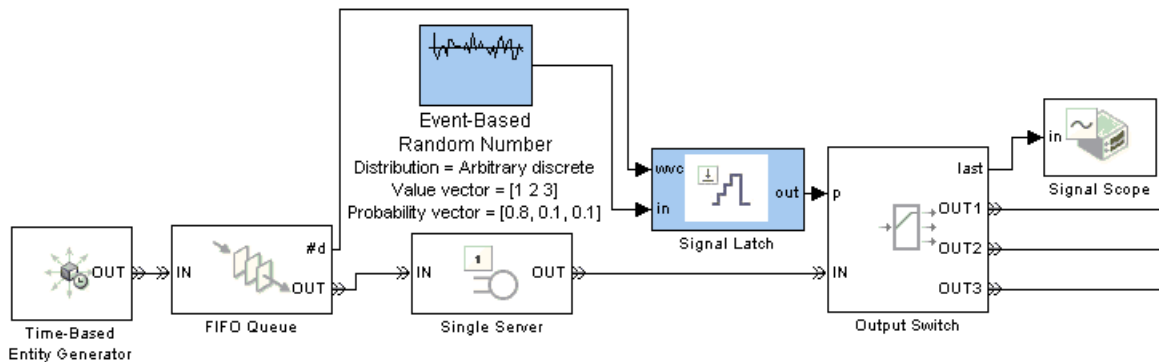
- “Example: Using an Arbitrary Discrete Distribution as Intergeneration Time” in the Getting Started documentation
- “Example: Setting Attributes” on page 1-15
- “Generating Packets” in the Getting Started documentation
- “Example: Using Random Service Times in a Queuing System” in the Getting Started documentation
- “Example: Event Calendar for a Queue-Server Model” on page 2-11
- “Example: M/M/5 Queuing System” on page 4-12
- “Example: Compound Switching Logic” on page 5-5


The model in “Example: Compound Switching Logic” on page 5-5 also illustrates how to use the Signal Latch block as described in “Generating Random Signals Based on Arbitrary Events” on page 3-5, to generate a random number upon each departure from an Input Switch block.

The models in “Example: Invalid Connection of Event-Based Random Number Generator” on page 11-28 illustrate how to follow the connection rules for the Event-Based Random Number block.

Example: Creating a Random Signal for Switching

The model below, similar to the one in “Example: Using Entity-Based Timing for Choosing a Port” on page 8-29, implements random output switching with a skewed distribution. The Signal Latch block causes the Event-Based Random Number block to generate a new random number upon each increase in the FIFO Queue block’s #d output signal, that is, each time an entity advances from the queue to the server. The random number becomes the switching criterion for the Output Switch block that follows the server. The plot reflects the skewed probability defined in the Event-Based Random Number block, which strongly favors 1 instead of 2 or 3.





Generating Random Time-Based Signals



The Random Number and Uniform Random Number blocks in the Simulink Sources library create time-based random signals with Gaussian and uniform distributions, respectively. The Event-Based Random Number block supports other distributions but is designed to create event-based signals. To generate time-based random signals using the Event-Based Random Number block, use this procedure:

- 1 Insert an Event-Based Random Number block into your model and configure it to indicate the distribution and parameters you want to use.
- 2 Insert and configure a Signal Latch block:
 - a Set **Write to memory upon** to Sample time hit from port **wts**.
 - b Set **Read from memory upon** to Write to memory event.

The block now has input ports **wts** and **in**, but not **wvc** or **rvc**.

- 3 Insert a Step block (or another time-based source block) and set **Sample time** to the desired sample time of the time-based signal you want to create.
- 4 Connect the Step block to the **wts** port of the Signal Latch block. Connect the Event-Based Random Number block to the **in** port of the Signal Latch block.

The **out** port of the Signal Latch block is a time-based signal whose sample time is the one specified in the Step block and whose values come from the Event-Based Random Number block. An example is below.



Sequence of Updates of Output Signals

When a block produces more than one output signal in response to events, the simulation behavior might depend on the sequence of signal updates relative to each other. This is especially likely if you use one of the signals to influence a behavior or computation that also depends on another one of the signals, as in “Example: Detecting Changes in the Last-Updated Signal” on page 3-10 and “Example: Detecting Changes from Empty to Nonempty” on page 8-23.

When you turn on more than one output signal from a SimEvents block’s dialog box (typically, from the **Statistics** tab), the block updates each of the signals in a sequence. See the Signal Output Ports table on the block’s reference page to learn about the update order:

- In some cases, a block’s reference page specifies the sequence explicitly using unique numbers in the Order of Update column.

For example, the reference page for the N-Server block indicates that upon entity departures, the **w** signal is updated before the **#n** signal. The Order of Update column in the Signal Output Ports table lists different numbers for the **w** and **#n** signals.

- In some cases, a block’s reference page lists two or more signals without specifying their sequence relative to each other. Such signals are updated in an arbitrary sequence relative to each other and you should not rely on a specific sequence for your simulation results.


For example, the reference page for the N-Server block indicates that the **w** and **util** signals are updated in an arbitrary sequence relative to each other. The Order of Update column in the Signal Output Ports table lists the same number for both the **w** and **util** signals.

- When a block offers fewer than two signal output ports, the sequence of updates does not need explanation on the block’s reference page. For example, the reference page for the Enabled Gate block does not indicate an update sequence because the block can output only one signal.

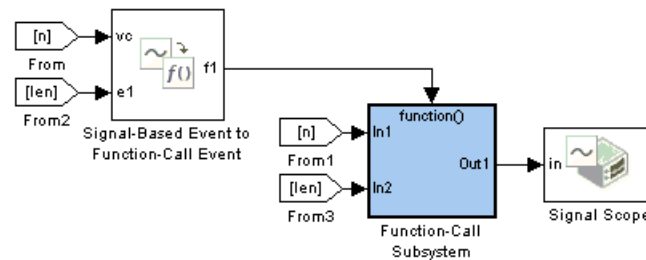
Example: Detecting Changes in the Last-Updated Signal

The model below plots the ratio of the queue’s current length to the time average of the queue length. The FIFO Queue block produces **#n** and **len**


signals representing the current and average lengths, respectively. The computation of the ratio occurs in a function-call subsystem that is called when the Signal-Based Event to Function-Call Event block detects a change in **#n** (as long as **len** is positive, to avoid division-by-zero warnings). Because the FIFO Queue block updates the **len** signal before updating the **#n** signal, both signals are up to date when the value change occurs in the **#n** signal.



Detect Changes in #n Signal



Top-Level Model



Subsystem Contents

If you instead connect the **len** signal to the Signal-Based Event to Function-Call Event block's **vc** input port, then the block issues a function

call upon detecting a change in the **len** signal. At that point, the **#n** value is left over from the block's previous arrival or departure, so the computed ratio is incorrect.

Multiple Simultaneous Updates

Simultaneous events, which might be causally related to each other, are common in discrete-event simulation. This section describes how they affect output signals from SimEvents blocks. Common scenarios involving simultaneous events include the following:

- An entity completes its service and departs from a server, which permits an entity to arrive at the same time instant from a queue that precedes the server.
- An entity arrives at an empty queue, finds that the subsequent server block is also empty, and advances immediately to the server.
- An Enabled Gate block between a queue and an Entity Sink block changes from the closed state to the open state, which permits all entities in the queue to depart simultaneously.

This section defines zero-duration values and illustrates how you can detect them in your simulation. The topics are as follows:

- “Zero-Duration Values of Signals” on page 3-13
- “Importance of Zero-Duration Values” on page 3-14
- “Detecting Zero-Duration Values” on page 3-14

Zero-Duration Values of Signals

Some output signals from SimEvents blocks produce a new output value for each departure from the block. When multiple departures occur in a single time instant, the result is a multivalued signal. That is, at a fixed instant in time, the signal assumes multiple values in sequence. The sequence of values corresponds to the sequence of departures. Although the departures and values have a well-defined sequence, no time elapses between adjacent events.

Scenario: Server Departure and New Arrival

For example, consider the scenario in which an entity departs from a server at time T and, consequently, permits another entity to arrive from a queue that precedes the server. The statistic representing the number of entities in the server is 1 just before time T because the first entity has not completed

its service. The statistic is 1 just after time T because the second entity has begun its service. At time T, the statistic is 0 before it becomes 1 again. The value of 0 corresponds to the server's empty state after the first entity has departed and before the second entity has arrived. Like this empty state, the value of 0 does not persist for a positive duration.

Scenario: Status of Pending Entities in a Queue

Another example of zero-duration values is in “Plotting the Pending-Entity Signal”, which discusses a signal that indicates when the entity at the head of a queue is unable to depart. This signal becomes 0 if the entity at the head of the queue, previously unable to depart, finally departs. If the queue is left with other entities that cannot depart at this time, then the signal becomes 1 again. That is, the value of 0 does not persist for a positive duration.

Importance of Zero-Duration Values

The values of signals, even values that do not persist for a positive duration, can help you understand or debug your simulations. In the example described in “Scenario: Server Departure and New Arrival” on page 3-13, the zero-duration value of 0 in the signal tells you that the server experienced a departure. If the signal assumed only the value 1 at time T (because 1 is the final value at time T), then the constant values before, at, and after time T would fail to indicate the departure. While you could use a departure count signal to detect departures specifically, the zero-duration value in the number-in-block signal provides you with more information in a single signal.

Detecting Zero-Duration Values


These topics describe ways to detect and examine zero-duration values:

- “Plotting Signals that Exhibit Zero-Duration Values” on page 3-14
- “Plotting the Number of Signal Changes Per Time Instant” on page 3-16
- “Viewing Zero-Duration Values in the MATLAB Workspace” on page 3-16

Plotting Signals that Exhibit Zero-Duration Values

One way to visualize event-based signals, including signal values that do not persist for a positive duration, is to use the Signal Scope or X-Y Signal Scope

block. Either of these blocks can produce a plot that includes a marker for each signal value (or each signal-based event, in the case of the event counting scope). For example, the figure below uses a plot to illustrate the situation described in “Scenario: Server Departure and New Arrival” on page 3-13.



When multiple plotting markers occur along the same vertical line, it means that the signal assumes multiple values at a single time instant. The callouts in the figure describe the server states that correspond to a few key points of the plot.

By contrast, some of the vertical lines have exactly one marker, meaning that the signal value at that time instant is unique. For example, at time 85, an arrival at the previously empty server is the only server-related activity at that time instant.

The preceding figure was created by modifying the mm1 demo model: in particular, raising the arrival rate somewhat, enabling the **Number of entities in block** signal in the Single Server block, inserting a Signal Scope block, and then attaching the scope to the server’s #n signal output port.

Note Unlike the Signal Scope block and X-Y Signal Scope blocks, the Scope block in the Simulink Sinks library does not detect zero-duration values. For more information, see “Comparison with Time-Based Plotting Tools” on page 9-15.

Plotting the Number of Signal Changes Per Time Instant

To detect the presence of zero-duration values, but not the values themselves, use the Instantaneous Event Counting Scope block with the **Type of value change** parameter set to `Either`. When a stem in the plot has a height of two or greater, the input signal has assumed multiple values at that instant of time.

For an example using this block, see “Example: Plotting Event Counts to Find Roundoff Error” on page 9-12.

Viewing Zero-Duration Values in the MATLAB Workspace

If an event-based signal assumes many values at one time instant and you cannot guess the sequence from a plot of the signal versus time, then you can get more information by examining the signal in the MATLAB workspace. By creating a variable that contains each time and signal value, you can recover the exact sequence in which the signal assumed each value during the simulation.

See “Sending Data to the MATLAB Workspace” on page 3-23 for instructions and an example.

Latency in Signal Updates

In some cases, the updating of an output signal or the reaction of a block to updates in the signal can experience a delay:

- When you use an event-based signal as an input to a time-based block that is not in a discrete event subsystem, the block might not react to changes in the input at exact event times but instead might react at the next time-based sample time hit for that block.

To make time-based blocks react to changes immediately when an event occurs in another block, use a discrete event subsystem. For details and examples, see Chapter 8, “Controlling Timing Using Subsystems”.

- The update of an output signal in one block might occur after other operations occur at that value of time, in the same block or in other blocks. This latency does not last a positive length of time, but might affect your simulation results. For details and an example, see “Interleaved Operations of Storage and Nonstorage Blocks” on page 12-4.
- When the definition of a statistical signal suggests that its value can vary *continuously* as simulation time elapses, the block increases efficiency by updating the signal value only at key moments during the simulation. As a result, the signal has a somewhat outdated “approximate” value between such key moments, but corrects the value later.

The primary examples of this phenomenon are the signals that represent time averages, such as a server’s utilization percentage. The definitions of time averages involve the current time, but simulation performance would suffer drastically if the block recomputed the percentage at each time-based simulation step. Instead, the block recomputes the percentage only upon the arrival or departure of an entity, when the simulation ends, and when you pause the simulation. For an example, see the reference page for the Single Server block.

When plotting statistics that, by definition, vary continuously as simulation time elapses, consider using a continuous-style plot. For example, set **Plot type** to Continuous in the Signal Scope block.

Manipulating Signals

The Signal Latch is a versatile block for manipulating event-based signals. You can use it to delay or resample signals based on events, not time. You can also use it to add an initial condition to event-based signals. The topics here are as follows:

- “Defining Initial Conditions for Event-Based Signals” on page 3-18
- “Example: Resampling a Signal Based on Events” on page 3-20

In addition, see these examples:

- “Generating Random Event-Based Signals” on page 3-4
- “Generating Random Time-Based Signals” on page 3-8
- “Example: Detecting Collisions by Comparing Events” on page 2-34
- “Example: Compound Switching Logic” on page 5-5

Defining Initial Conditions for Event-Based Signals

Many event-based signals have the first update upon the first relevant event, not necessarily at the beginning of the simulation. If you need to initialize the signal before the first relevant event occurs, you can define an initial condition for the signal using the Signal Latch block. An initial condition is useful when a feedback loop involves an event-based signal, as in “Example: Controlling Joint Availability of Two Servers” on page 7-4.

To define an initial condition, use this procedure:

- 1 Set these parameters in the Signal Latch block:
 - **Initial memory value** = your desired initial condition
 - **Write to memory upon** = Sample time hit from port wts
 - **Read from memory upon** = Write to memory event
 - Select **Report memory value upon write event, mem**
 - Clear **Report memory value upon read event, out**

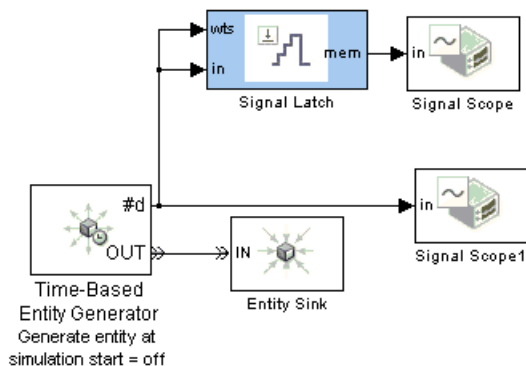
The block now has signal input ports **wts** and **in**, and a signal output port **mem**.

- 2 Connect the signal whose initial condition you want to define to both the **in** and **wts** ports of the Signal Latch block.


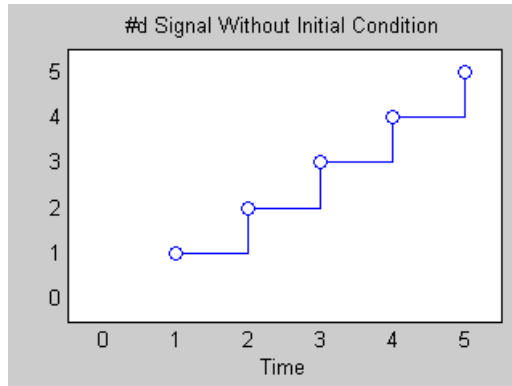
The Signal Latch block's **mem** output signal uses your initial condition until your original signal has its first update. Afterward, the **mem** signal and your original signal are identical.

Example: Defining #d Before the First Entity Departure

The model below illustrates the procedure above by defining an initial condition for the Time-Based Entity Generator block's **#d** output signal. The initial condition is in effect until the first entity departs from the generator block.

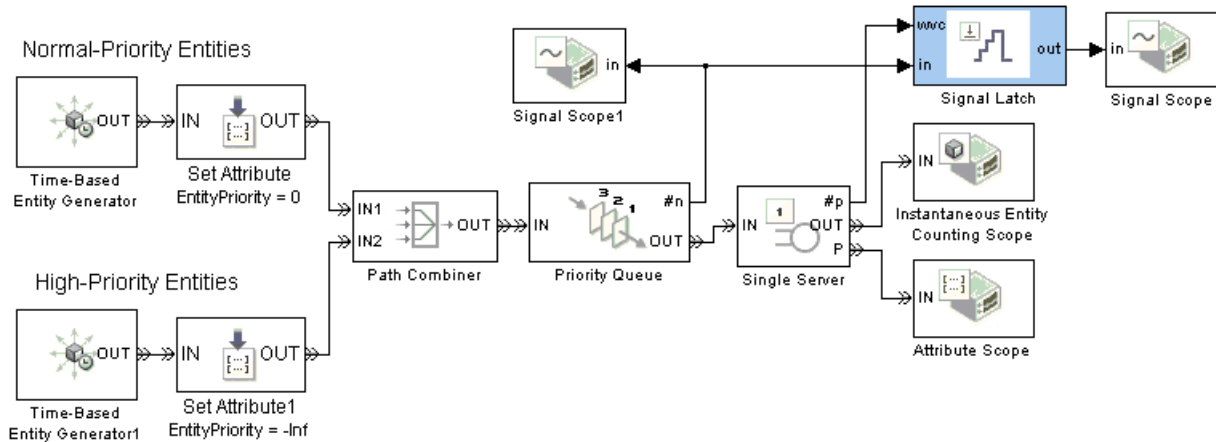


Two plots show the **#d** signal as produced by the generator block and the **mem** signal produced by the Signal Latch block. Notice that before the first entity's departure from the generator block, the plot of the **#d** signal is blank, while the plot of the **mem** signal reflects an initial condition of zero.




Example: Resampling a Signal Based on Events


The model below contains a server that supports preemption of normal-priority entities by high-priority entities. This is similar to “Example: Preemption by High-Priority Entities” on page 4-10. Suppose that a preemption and the subsequent service of a high-priority entity represents a time interval during which the server is inoperable. The goal of this example is to find out how many entities are in the queue when the breakdown begins.



A plot of the Priority Queue block's **#n** output signal indicates how many entities are in the queue at all times during the simulation.



The Signal Latch block resamples the **#n** signal, focusing only on the values that **#n** assumes when a high-priority queue preempts an entity already in the server. The Signal Latch block outputs a sample from the **#n** signal whenever the Single Server block's **#p** output signal increases, where **#p** is the number of entities that have been preempted from the server. Between pairs of successive preemption events, the Signal Latch block does not update its output signal, ignoring changes in **#n**. A plot of the output from the Signal Latch block makes it easier to see how many entities are in the queue when the breakdown begins, compared to the plot of the entire **#n** signal.




Sending Data to the MATLAB Workspace

The Discrete Event Signal to Workspace block writes event-based signals to the MATLAB workspace when the simulation stops or pauses. When inside a discrete event subsystem, the To Workspace block can also be useful for writing event-based signals to the MATLAB workspace.

Example: Sending Queue Length to the Workspace

The model below shows one way to write the times and values of an event-based signal to the MATLAB workspace. In this case, the signal is the **#n** output from a FIFO Queue block, which indicates how many entities the queue holds.



After you run this simulation, you can use the following code to create a two-column matrix containing the times values in the first column and the signal values in the second column.

```
times_values = [num_in_queue.time, num_in_queue.signals.values]
```

The output below reflects the Time-Based Entity Generator block's constant intergeneration time of 0.8 second and the Single Server block's constant service time of 1.1 second.

```
times_values =
    0         0
  0.8000    1.0000
  1.1000         0
  1.6000    1.0000
```

2.2000	0
2.4000	1.0000
3.2000	2.0000
3.3000	1.0000
4.0000	2.0000
4.4000	1.0000
4.8000	2.0000
5.5000	1.0000
5.6000	2.0000
6.4000	3.0000
6.6000	2.0000
7.2000	3.0000
7.7000	2.0000
8.0000	3.0000
8.8000	4.0000
8.8000	3.0000
9.6000	4.0000
9.9000	3.0000

From the output, you can see that the number of entities in the queue increases at times that are a multiple of 0.8, and decreases at times that are a multiple of 1.1. At $T=8.8$, a departure from the server and an entity generation occur simultaneously; both events influence the number of entities in the queue. The output below shows two values corresponding to $T=8.8$, enabling you to see the zero-duration value that the signal assumes at this time.

Using the To Workspace Block with Event-Based Signals

The To Workspace block in the Simulink Sinks library can be useful for working with event-based signals in special ways, such as

- Omitting repeated values of the signal and focusing on changes in the signal's value. For an example, see "Example: Sending Unrepeated Data to the MATLAB Workspace" on page 8-21.
- Recording values of multiple signals when any *one* of the signals has an update. To accomplish this, place multiple To Workspace blocks in a discrete event subsystem that has multiple input ports.

If you use the To Workspace block in the Simulink Sinks library to write event-based signals to the MATLAB workspace, you should

- 1** Set the block's **Save format** parameter to Structure With Time, which causes the block to record time values, not just signal values.
- 2** Place the To Workspace block in a discrete event subsystem to ensure that the workspace variable records data at appropriate times during the simulation.

For more details about discrete event subsystems, see “Role of Discrete Event Subsystems in SimEvents Models” on page 8-7.

Modeling Queues and Servers

The topics below supplement the discussion in “Basic Queues and Servers” in the Getting Started documentation.

Using a LIFO Queuing Discipline
(p. 4-2)

Sorting by Priority (p. 4-4)

Preempting an Entity in a Server
(p. 4-9)

Modeling Multiple Servers (p. 4-12)

Modeling the Failure of a Server
(p. 4-15)

Comparing LIFO and FIFO queues

Using attribute values to control the queue discipline

Enabling an entity to replace a lower-priority entity in a server

Modeling a bank of servers

Using Stateflow to model the behavior of a server that might require maintenance

Using a LIFO Queuing Discipline

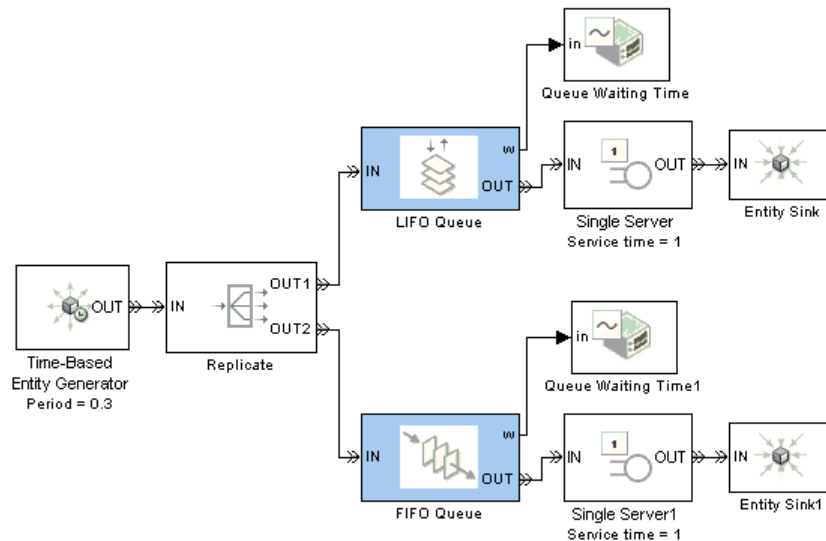
The LIFO Queue block supports the last-in, first-out (LIFO) queuing discipline. The entity that departs from the queue at a given time is the most recent arrival. You can interpret a LIFO queue as a stack.



Some ways to see the difference between FIFO and LIFO queuing disciplines are to

- Attach data to entities to distinguish entities from each other. For more information about using entities to carry data, see “Setting Attributes of Entities” on page 1-14.
- View simulation statistics that you expect the queuing discipline to influence. One such statistic is the average waiting time in the queue; to compute the waiting time of each entity, the block must know which entity is departing at a given departure time.

Example: Waiting Time in LIFO Queue

As an example, compare the FIFO and LIFO disciplines in a D/D/1 queuing system with an intergeneration time of 0.3 and a service time of 1.



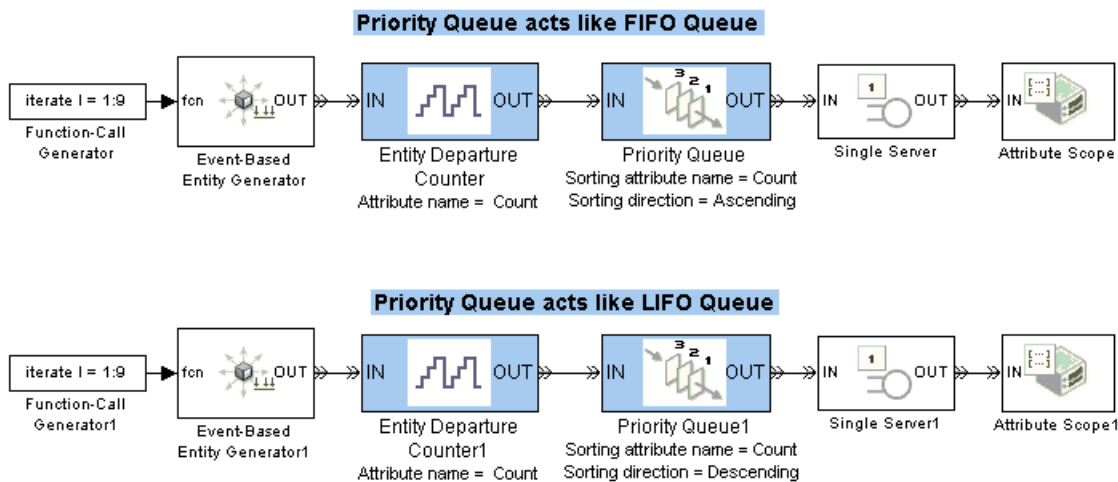


Sorting by Priority

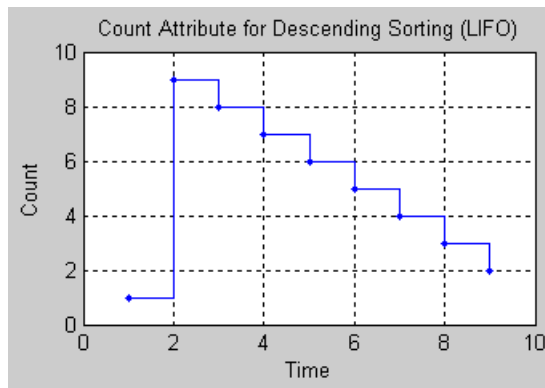

The Priority Queue block supports queuing in which entities' positions in the queue are based primarily on their attribute values. Arrival times are relevant only when attribute values are equal. You specify the attribute and the sorting direction using the **Sorting attribute name** and **Sorting direction** parameters in the block's dialog box. To assign values of the attribute for each entity, you can use the Set Attribute block as described in "Setting Attributes of Entities" on page 1-14.

Note While you can view the value of the sorting attribute as an entity priority, this value has nothing to do with event priorities or block priorities.

Two familiar cases are shown below, in which a priority queue acts like a FIFO or LIFO queue. At the start of the simulation, the FIFO and LIFO sections of the model each generate nine entities, the first of which advances immediately to a server. The remaining entities stay in the queues until the server becomes available. The sorting attribute is Count, whose values are the entities' arrival sequence at the queue block. In this example, the servers do not permit preemption; preemptive servers would behave differently.

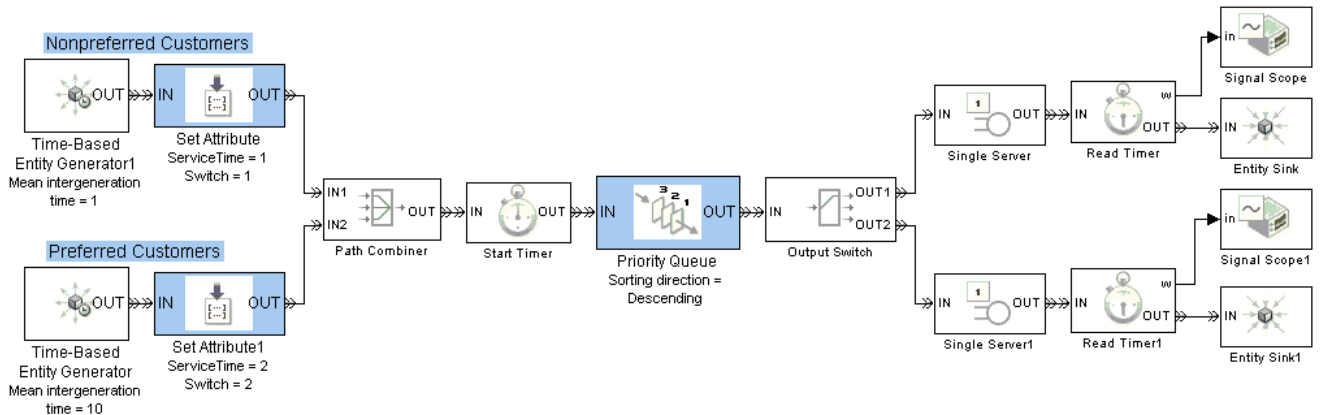


The FIFO plot reflects an increasing sequence of Count values. The LIFO plot reflects a descending sequence of Count values, except for the Count=1 entity that advances to the server before the queue has any other entities to sort.




Example: Serving Preferred Customers First


In the model below, two types of customers enter a queuing system. One type, considered to be preferred customers, are less common but require longer service. The priority queue places preferred customers ahead of nonpreferred customers. The model plots the average system time for the set of preferred customers and separately for the set of nonpreferred customers.



You can see from the plots that despite the shorter service time, the average system time for the nonpreferred customers is much longer than the average system time for the preferred customers.




Average System Time for Nonpreferred Customers Sorted by Priority




Average System Time for Preferred Customers Sorted by Priority

Comparison with Unsorted Behavior

If the queue used a FIFO discipline for all customers instead of a priority sorting, then the average system time would decrease slightly for the nonpreferred customers and increase markedly for the preferred customers.



Average System Time for Nonpreferred Customers Unsorted



Average System Time for Preferred Customers Unsorted

Preempting an Entity in a Server

The Single Server block supports preemption, which is the replacement of an entity in a server block by an entity that satisfies certain criteria. The preempted entity departs from the block via the **P** entity output port instead of the usual **OUT** port.

Criteria for Preemption

Whether preemption occurs depends on attribute values of the entity in the server and of the entity attempting to arrive at the server. You specify the attribute and the orientation of the comparison using the **Sorting attribute name** and **Sorting direction** parameters in the Single Server block's dialog box. (These parameters are available after you select **Permit preemption based on attribute**.) To assign values of the sorting attribute for each entity, you can use the Set Attribute block as described in "Setting Attributes of Entities" on page 1-14. Valid values for the sorting attribute are any real numbers, Inf, and -Inf.

If the attribute values are equal, no preemption occurs.

When preemption is supposed to occur, the **P** port must not be blocked. Consider connecting the **P** port to a queue or server with infinite capacity, to prevent a blockage during the simulation.

Note While you can view the value of the sorting attribute as an entity priority, this value has nothing to do with event priorities or block priorities.

Residual Service Time

A preempted entity might or might not have completed its service time. The remaining service time the entity would have required if it had not been preempted is called the entity's *residual* service time. If you select **Write residual service time to attribute** in the Single Server block, then the block records the residual service time of each preempted entity in an attribute of that entity. If the entity completes its service time before preemption occurs, then the residual service time is zero.

For entities that depart from the block's **OUT** entity output port (that is, entities that are not preempted), the block records a residual service time only if the entity already has an attribute whose name matches the **Residual service time attribute name** parameter value. In this case, the block sets that attribute to zero when the entity departs from the **OUT** port.

Queuing Disciplines for Preemptive Servers

When you permit preemption in a Single Server block preceded by a queue, only the entity at the head of the queue can preempt an entity in the server.

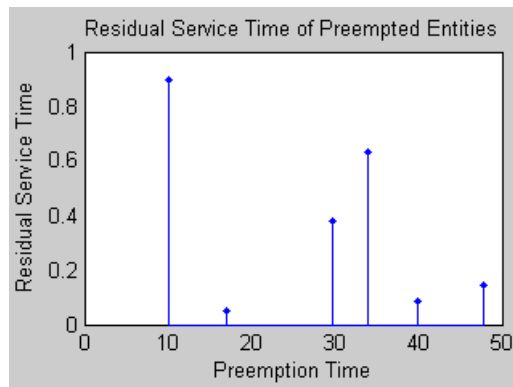
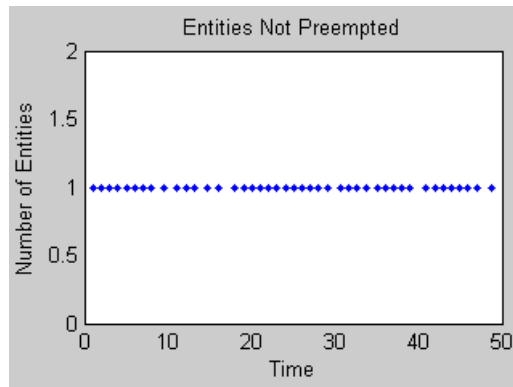
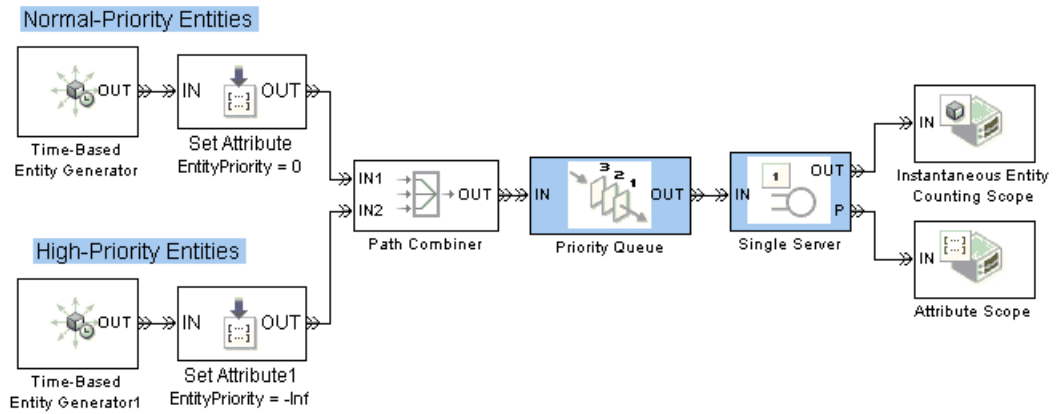
The Priority Queue block is particularly appropriate for use with the preemption feature of the Single Server block. When an entity with sufficiently high priority arrives at the Priority Queue block, the entity goes to the head of the queue and immediately advances to the server.

When using the Single Server and Priority Queue blocks together, you typically set the **Sorting attribute name** and **Sorting direction** parameters to the same values in both blocks.

Example: Preemption by High-Priority Entities

The model below generates two classes of entities, most with an EntityPriority attribute value of 0 and some with an EntityPriority attribute value of -Inf. The sorting direction in the Priority Queue and Single Server blocks is Ascending, so entities with sorting attribute values of -Inf go to the head of the priority queue and immediately preempt any entity in the server except another entity whose sorting attribute value is -Inf.

One plot shows when nonpreemptive departures occur, while another plot indicates the residual service time whenever preemptive departures occur.



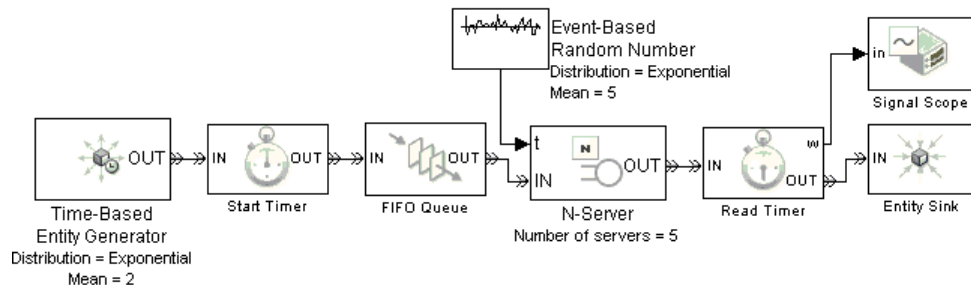
Modeling Multiple Servers

You can use the N-Server and Infinite Server blocks to model a bank of identical servers operating in parallel. The N-Server block lets you specify the number of servers using a parameter, while the Infinite Server block models a bank of infinitely many servers.


To model multiple servers that are not identical to each other, you must use multiple blocks. For example, to model a pair of servers whose service times do not share the same distribution, use a pair of Single Server blocks rather than a single N-Server block. The example in “Example: Selecting the First Available Server” in the Getting Started documentation illustrates the use of multiple Single Server blocks with a switch.

Example: M/M/5 Queuing System

The model below shows a system with infinite storage capacity and five identical servers. In the notation, the M stands for Markovian; M/M/5 means that the system has exponentially distributed interarrival and service times, and five servers.



The plot below shows the waiting time in the queuing system.




You can compare the empirical values shown in the plot with the theoretical value, $E[S]$, of the mean system time for an M/M/m queuing system with an arrival rate of $\lambda=1/2$ and a service rate of $\mu=1/5$. Using expressions in [2], the computation is as follows.

$$\rho = \frac{\lambda}{m\mu} = \frac{(1/2)}{5(1/5)} = \frac{1}{2}$$

$$\pi_0 = \left[1 + \sum_{n=1}^{m-1} \frac{(m\rho)^n}{n!} + \frac{(m\rho)^m}{m!} \frac{1}{1-\rho} \right]^{-1} \approx 0.0801$$

$$E[S] = \frac{1}{\mu} + \frac{1}{\mu} \frac{(m\rho)^m}{m!} \frac{\pi_0}{m(1-\rho)^2} \approx 5.26$$

Zooming in the plot shows that the empirical value is close to 5.26.



Modeling the Failure of a Server

In some applications, it is useful to model situations in which a server fails. For example, a machine might break down and later be repaired, or a network connection might fail and later be restored. This section explores ways to model failure of a server, as well as server states. The topics are as follows:

- “Server States” on page 4-15
- “Using a Gate to Implement a Failure State” on page 4-15
- “Using Stateflow to Implement a Failure State” on page 4-16


Server States

The server blocks in SimEvents do not have built-in states, so you can design states in any way that is appropriate for your application. Some examples of possible server states are in the table below.

Server as Communication Channel	Server as Machine	Server as Human Processor
Transmitting message	Processing part	Working
Connected but idle	Waiting for new part to arrive	Waiting for work
Unconnected	Off	Off duty
Holding message (pending availability of destination)	Holding part (pending availability of next operator)	Waiting for resource
Establishing connection	Warming up	Preparing to begin work

Using a Gate to Implement a Failure State

For any state that represents a server’s inability or refusal to accept entity arrivals even though the server is not necessarily full, a common implementation involves an Enabled Gate block preceding the server.



The gate prevents entity access to the server whenever the gate’s control signal at the **en** input port is zero or negative. The logic that creates the **en** signal determines whether or not the server is in a failure state. You can implement such logic using the techniques described in Chapter 6, “Using Logic” or using Stateflow to transition among a finite number of server states.

For an example in which an Enabled Gate block precedes a server, see “Example: Controlling Joint Availability of Two Servers” on page 7-4. The example is not specifically about a failure state, but the idea of controlling access to a server is similar. Also, you can interpret the Signal Latch block with the **st** output signal enabled as a two-state machine that changes state when read and write events occur.

Note A gate prevents new entities from arriving at the server but does not prevent the current entity from completing its service. If you want to eject the current entity from the server upon a failure occurrence, then you can use the preemption feature of the server to replace the current entity with a high-priority “placeholder” entity.

Using Stateflow to Implement a Failure State

Stateflow is a suitable tool for implementing transitions among a finite number of server states. If you need to support more than just two states, then a Stateflow block might be more natural than a combination of Enabled Gate and logic blocks.


When modeling interactions between the state chart and discrete-event aspects of the model, note that a function call is the recommended way to make Stateflow blocks respond to asynchronous state changes. You can use blocks in the Event Generators and Event Translation libraries to produce a function call upon signal-based events or entity departures; the function call can invoke a Stateflow block. Conversely, a Stateflow block can output a

function call that can cause a gate to open, an entity counter to reset, or an entity generator to generate a new entity.

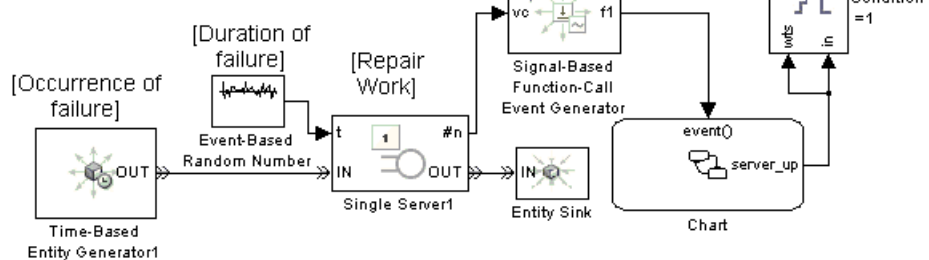
Example: Failure and Repair of a Server

The model below uses a Stateflow block to describe a two-state machine. A server is either down (failed) or up (operable). The state of the server is an output signal from the Stateflow block and is used to create the enabling signal for an Enabled Gate block that precedes a server in a queuing system.

Entities representing customers in queuing system



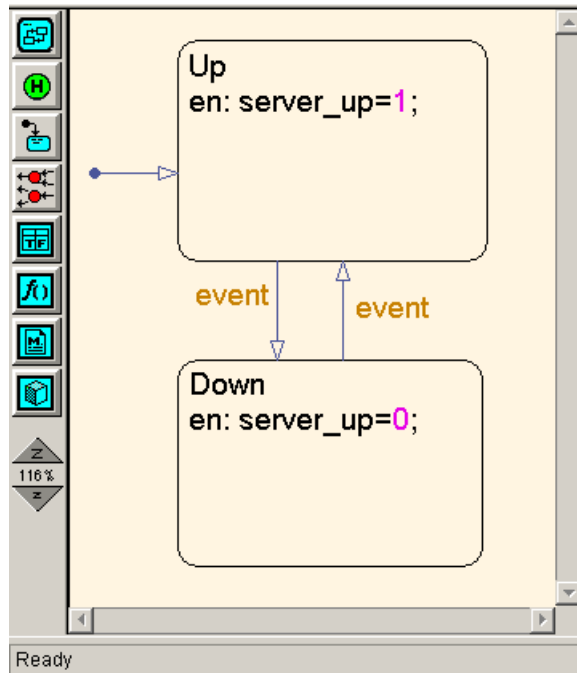
Entities representing failure and repair of server in queuing system




The lower portion of the model contains a parallel queuing system. The entities in the lower queuing system represent failures, not customers. Generation of a failure entity represents a failure occurrence in the upper queuing system. Service of a failure entity represents the time during which the server in the upper queuing system is down. Completion of service of a failure entity represents a return to operability of the upper queuing system.

When the lower queuing system generates an entity, changes in its server's #n signal invoke the Stateflow block that determines the state of the upper

queuing system. Increases in the #n signal cause the server to go down, while decreases cause the server to become operable again.



While this simulation runs, Stateflow alternately highlights the up and down states. The plot showing entity departures from the upper queuing system shows gaps, during which the server is down.




Although this two-state machine could be modeled more concisely with a Signal Latch block instead of a Stateflow block, the Stateflow chart scales more easily to include additional states or other complexity.

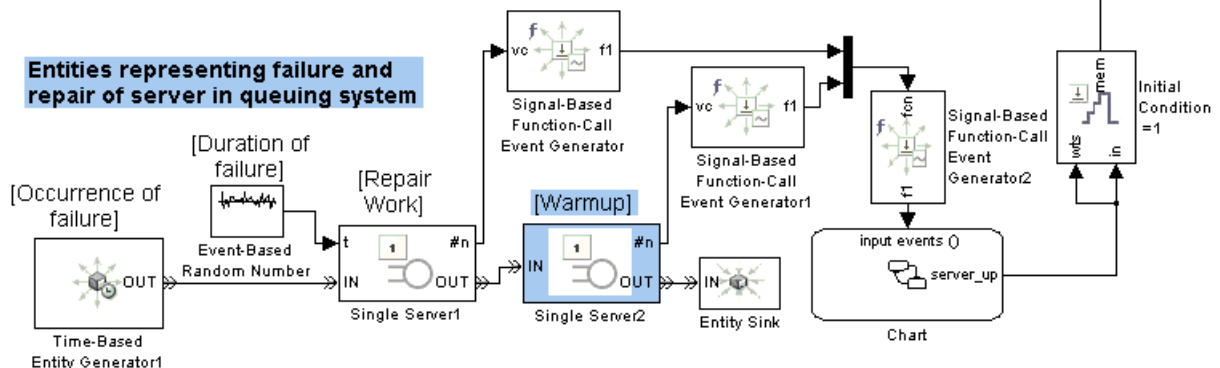
Example: Adding a Warmup Phase

The model below modifies the one in “Example: Failure and Repair of a Server” on page 4-17 by adding a warmup phase after the repair is complete. The Enabled Gate block in the upper queuing system does not open until the repair and the warmup phase are complete. In the lower queuing system, an additional Single Server block represents the duration of the warmup phase.

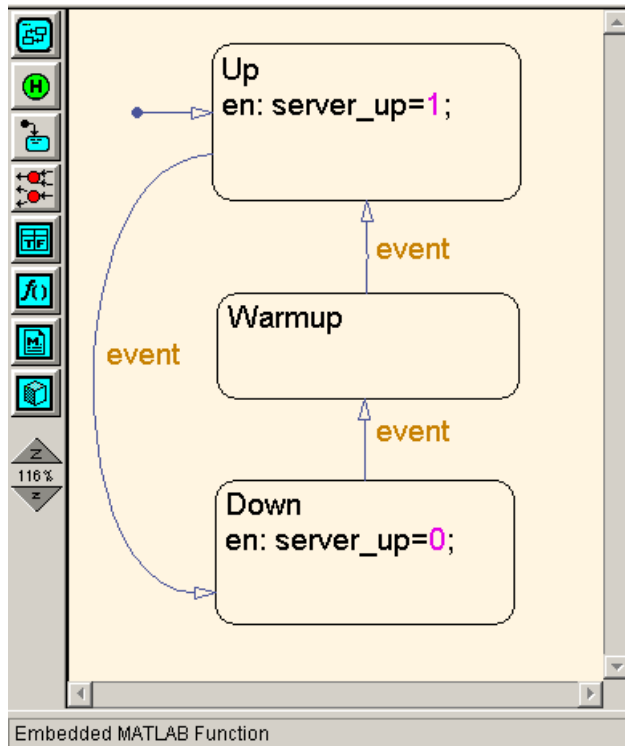
Entities representing customers in queuing system




Entities representing failure and repair of server in queuing system



In the Stateflow block, the input function calls controls when the repair operation starts, when it ends, and when the warmup is complete. The result of the function-call event depends on the state of the chart when the event occurs. A rising edge of the Repair Work block’s #n signal starts the repair operation, a falling edge of the same signal ends the repair operation, and a falling edge of the Warmup block’s #n signal completes the warmup.



While this simulation runs, Stateflow alternates among the three states. The plot showing entity departures from the upper queuing system shows gaps, during which the server is either under repair or warming up. By comparing the plot to the one in “Example: Failure and Repair of a Server” on page 4-17, you can see that the gaps in the server’s operation last slightly longer. This is because of the warmup phase.



Routing Techniques

The topics below supplement the discussion in “Designing Paths for Entities” in the Getting Started documentation by providing examples of cascaded switch blocks.

Example: Cascaded Switches with Random Selections (p. 5-2)

Equiprobable switching using cascaded switch blocks

Example: Cascaded Switches with Round-Robin Sequence (p. 5-4)

Round-robin switching using cascaded switch blocks

Example: Compound Switching Logic (p. 5-5)

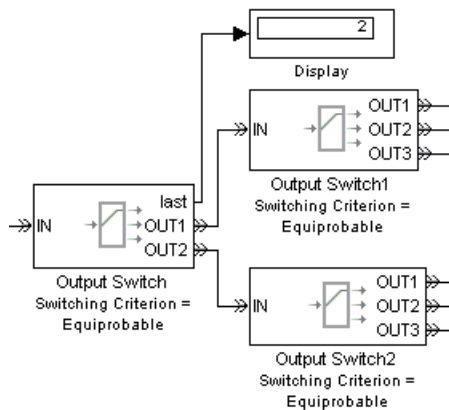
Combination of round-robin and weighed random switching

Example: Cascaded Switches with Random Selections

The Input Switch and Output Switch blocks support random selection of ports via the Equiprobable switching criterion. You can cascade multiple switch blocks to form a composite switching component with equiprobable outcomes.

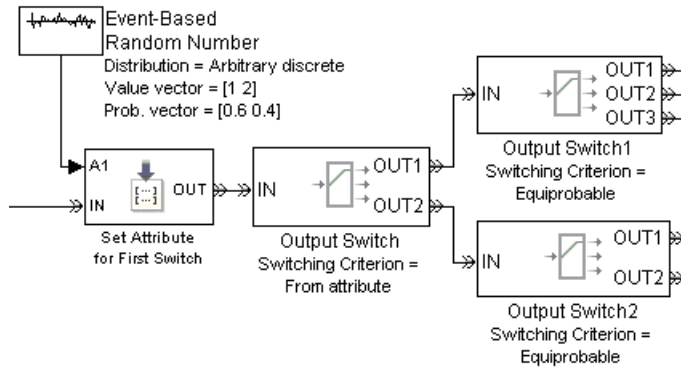
Six-Way Switching Component

Suppose you want to form a six-way switching component and that it is useful to know whether each entity uses one of the first three entity paths as opposed to one of the last three entity paths. The example below uses three Output Switch blocks to form a six-way switching component. Each Output Switch block has its **Switching criterion** set to Equiprobable. The **last** signal output port on the first switch block provides a way to find out whether each entity uses one of the first three or last three entity paths.




Five-Way Switching Component

Suppose you want to form a five-way switching component in which each of the five entity paths has equal probability of being selected. You cannot imitate the technique used in “Six-Way Switching Component” on page 5-2 because five is a prime number. Instead, use a skewed probability distribution in the first Output Switch block’s switching criterion to compensate for the unequal numbers of entity output ports in the subsequent two Output Switch blocks.



Example: Cascaded Switches with Round-Robin Sequence

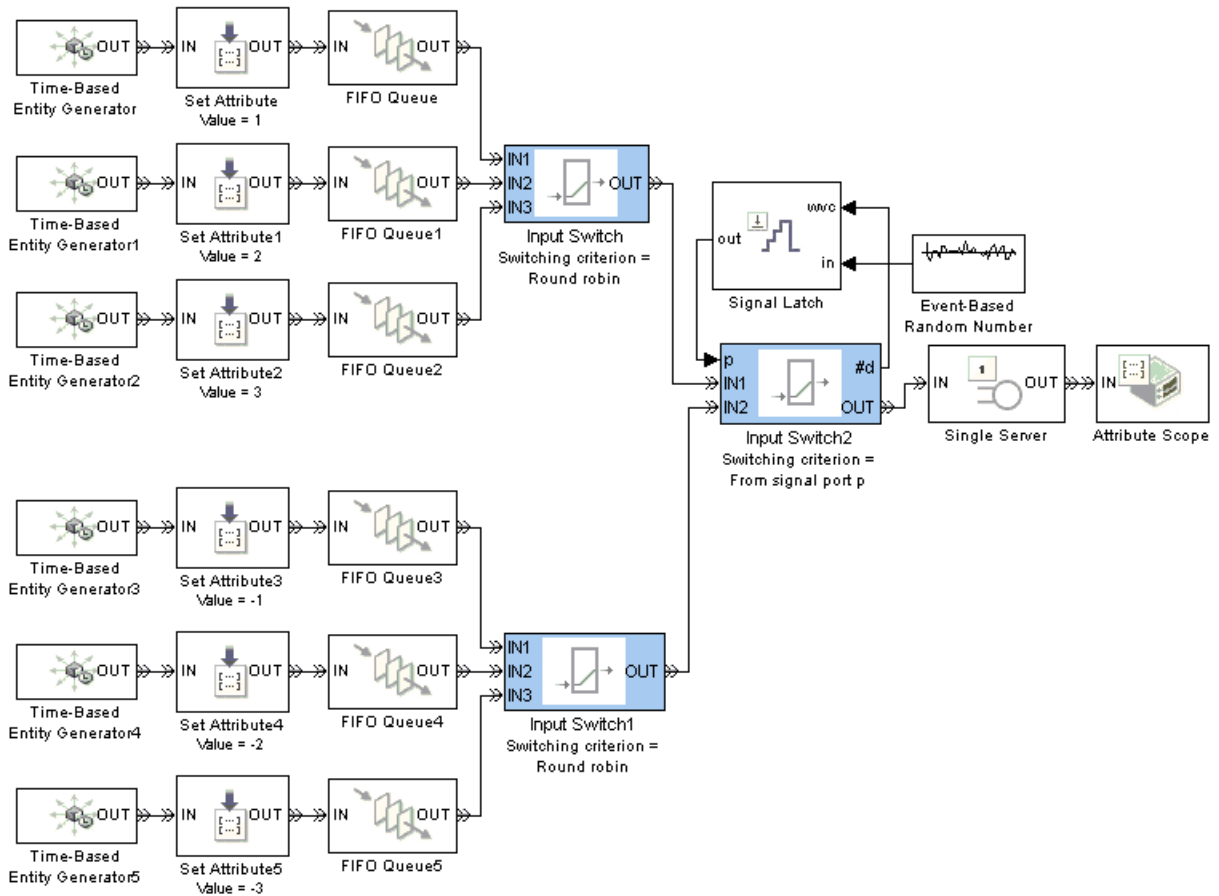
Suppose you want to form an eight-way switching component using a round-robin switching criterion. For simplicity, the example below uses a constant intergeneration time between entities so that counting time steps is equivalent to counting entities. The first Output Switch block changes its selected port every four time steps, while the subsequent Output Switch blocks change their selected ports with every entity. The result of this cascade is a round-robin sequence among the entity output ports of the subsequent Output Switch blocks.




Example: Compound Switching Logic

Suppose a single server processes entities from two groups each consisting of three sources. The switching component between the entity sources and the server determines which entities proceed to the server whenever it is available. The switching component uses a distribution that is skewed toward entities from the first group. Within each group, the switching component uses a round-robin approach.

The example below shows how to implement this design using three Input Switch blocks. The first two Input Switch blocks have their **Switching criterion** parameter set to Round robin to represent the processing of entities within each group of entity sources. The last Input Switch block uses a random signal with a skewed probability distribution to choose between the two groups. The Signal Latch block causes the random number generator to draw a new random number after each departure from the last Input Switch block.



For tracking purposes, the model assigns an attribute to each entity based on its source. The attribute values are 1, 2, and 3 for entities in the first group and -1, -2, and -3 for entities in the second group. You can see from the plot below that negative values occur less frequently than positive values, reflecting the skewed probability distribution. You can also see that the positive values reflect a round-robin approach among servers in the top group, while negative values reflect a round-robin approach among servers in the bottom group.



Using Logic

Role of Logic in SimEvents Models
(p. 6-2)

Typical situations in which logic
affects simulation behavior

Using Embedded MATLAB Function
Blocks for Logic (p. 6-3)

Specifying logic using MATLAB code

Using Logic Blocks (p. 6-10)

Specifying logic using a block
diagram

Role of Logic in SimEvents Models

Logic can be an important component in a discrete-event simulation, for specifying

- Normal but potentially complex routing or gating behavior.

For example, you might want to model a multiple-queue system in which entities advance to the shortest queue. Such a model must also indicate what happens if the minimum length is not unique.

- Handling of overflows, blockages, and other special cases.

For example, a communication system might drop packets that overflow a queue, while a manufacturing assembly line might pause processing at one machine if it releases parts that overflow a second machine.

Using Embedded MATLAB Function Blocks for Logic

If your logic algorithm is easier to express in MATLAB code than in a block diagram, then you can use the Embedded MATLAB Function block to implement the logic. Details about how to use this block are in “Using the Embedded MATLAB Function Block” in the Simulink documentation. This section provides examples that are particularly relevant for logic in SimEvents models:

- “Example: Choosing the Shortest Queue” on page 6-3
- “Example: Varying Fluid Flow Rate Based on Batching Logic” on page 6-6

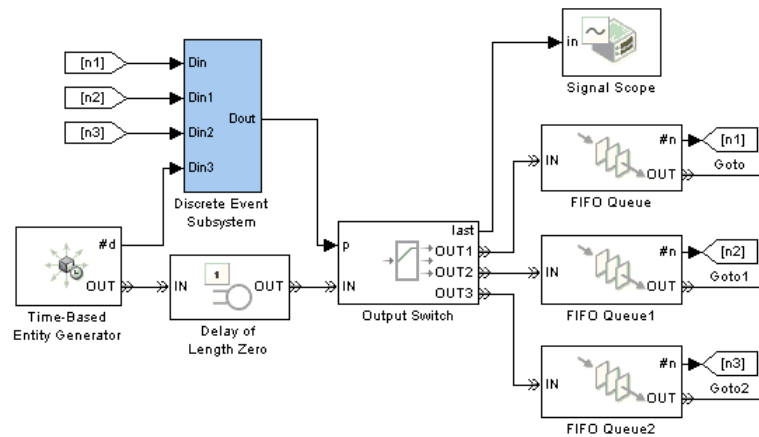
If your logic algorithm requires data from an earlier call to the function, then you can use persistent variables to retain data between calls. For examples of this technique, see

- The Switching Logic subsystems in the Astable Multivibrator Circuit demo.
- “Example: Computing a Time Average of a Signal” on page 10-9

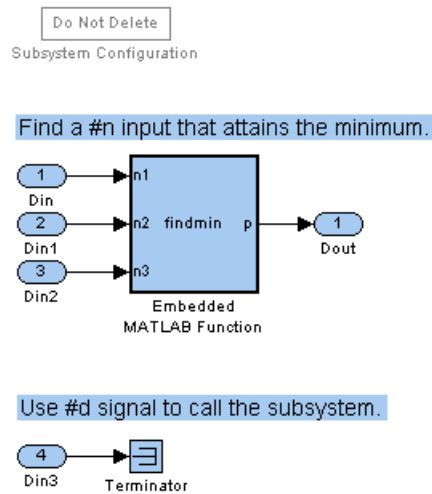
Note If you put an Embedded MATLAB Function block in a Discrete Event Subsystem block, use the Ports and Data Manager instead of Model Explorer to view or change properties such as the size or source of an argument. Model Explorer does not show the contents of Discrete Event Subsystem blocks.

Example: Choosing the Shortest Queue

The model below directs entities to the shortest of three queues. It uses an Output Switch block to create the paths to the different queues. To implement the choice of the shortest queue, a discrete event subsystem queries each queue for its current length, determines which queue or queues achieve the minimum length, and provides that information to the Output Switch block. For simplicity, the model omits any further processing of the entities after they leave their respective queues.



Although the block diagram shows signals at the $\#n$ signal output ports from the queue blocks and another signal at the **p** signal input port of the switch block, the block diagram does not indicate how to compute **p** from the set of $\#n$ values. That computation is performed inside a discrete event subsystem that contains an Embedded MATLAB Function block.



Subsystem Contents


If you double-click an Embedded MATLAB Function block in a model, an editor window shows the MATLAB function that specifies the block. In this example, the following MATLAB function computes the index of a queue having the shortest length, where the individual queue lengths are $n1$, $n2$, and $n3$. If more than one queue achieves the minimum, then the computation returns the smallest index among the queues that minimize the length.

```
function p = findmin(n1, n2, n3)
```

```
% p is the index of a queue having the shortest length.
[minlength,p] = min([n1 n2 n3]);
```

Note For visual simplicity, the model uses Goto and From blocks to connect the #n signals to the computation.

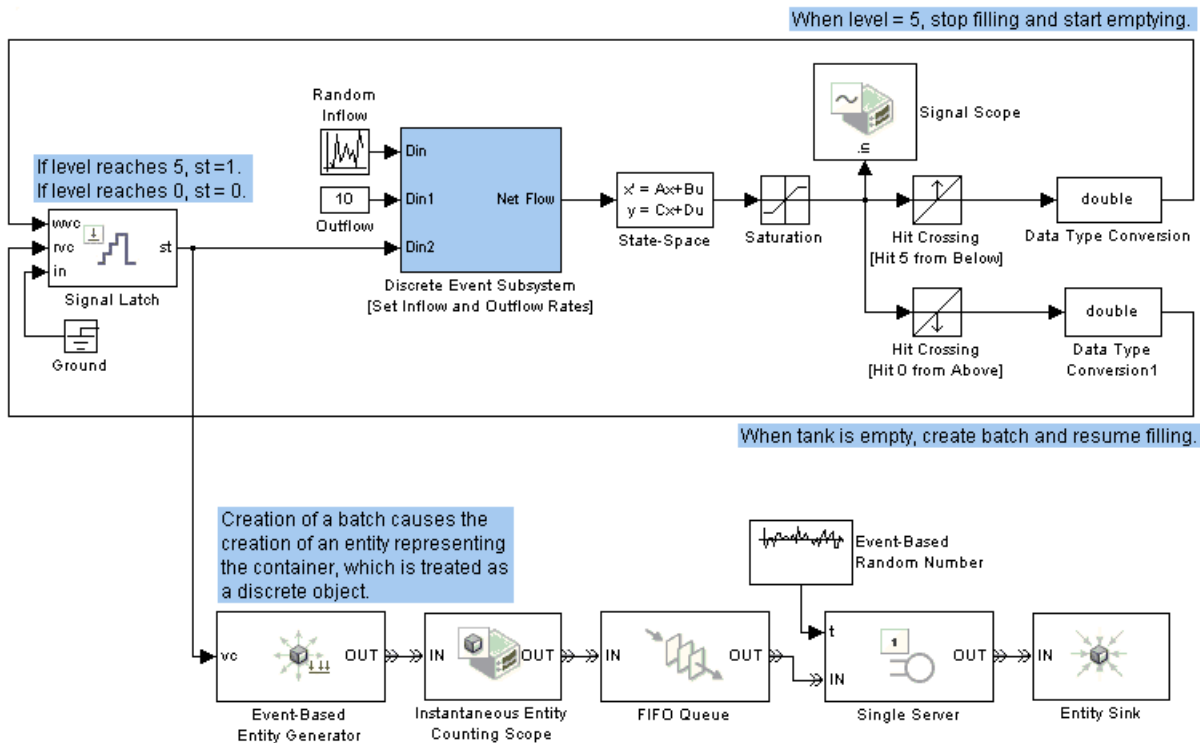
The figure below shows a sample plot. Each stem corresponds to an entity departing from the switch block via one of the three entity output ports.



For a variation on this model that uses logic blocks instead of the Embedded MATLAB Function block, see “Example: Choosing the Shortest Queue Using Logic Blocks” on page 6-16.


Example: Varying Fluid Flow Rate Based on Batching Logic

The model below represents a batching process in which a tank accumulates fluid up to a desired level of 5. When the level reaches 5, the tank switches to an emptying mode, modeling the creation of a batch. When the tank is empty, an entity generator creates an entity that represents the container for the batch of fluid. Batching logic determines whether the tank is filling or emptying.



Top-Level Model

Within a discrete event subsystem, an Embedded MATLAB Function block uses a logical if-then structure to set the inflow and outflow rates. The MATLAB code also computes the net flow rate, which forms the block's output signal. The subsystem and code are below.



Subsystem Contents

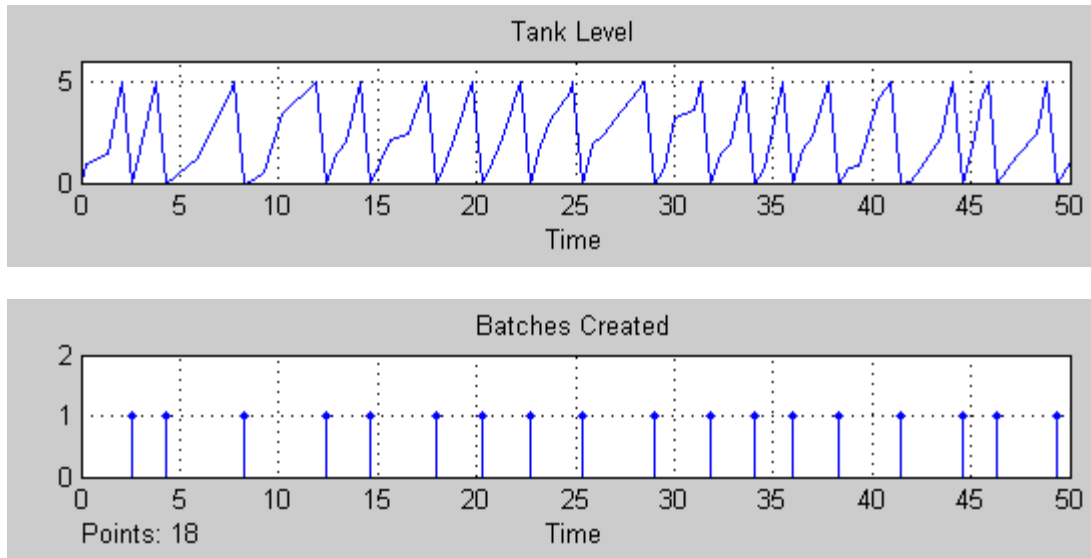
```
function net = netflowrate(ran,maxout,state)

% Compute the inflow and outflow rates.
if (state == 1)
    % Empty the tank.
    in = 0;
    out = maxout;
else
    % Fill the tank.
    in = ran;
    out = 0;
end

% Compute the net flow rate, which forms the output.
net = in - out;
```

While you could alternatively use a Switch block to compute the net flow rate, the choice of MATLAB code or a block-diagram representation might depend on which approach you find more intuitive.

The model plots the continuously changing level of fluid in the tank. A stem plot shows when each batch is created.



Using Logic Blocks

The following blocks can be useful for modeling logic because they return a 0 or 1:

- Relational Operator
- Compare To Constant and Compare To Zero
- Interval Test and Interval Test Dynamic
- Detect Change, Detect Decrease, Detect Increase, etc.
- Signal Latch

Note Some blocks return a 0 or 1 of a Boolean or integer data type. Blocks in the SimEvents libraries process signals whose data type is double. To convert between data types, use the Data Type Conversion block in the Simulink Signal Attributes library.

For switching, you might need to compute an integer that indicates a port number. Here are some useful blocks for this situation:

- Switch
- Lookup Table
- Bias
- Rounding Function, if an earlier computation returns a noninteger
- Other blocks in the Math Operations library

See these examples:

- “Example: Using Servers in Shifts” on page 6-11
- “Example: Choosing the Shortest Queue Using Logic Blocks” on page 6-16
- The logic diagrams depicted in “Stopping Upon Reaching a Particular State” on page 10-35

- The gate examples, especially “Example: Controlling Joint Availability of Two Servers” on page 7-4

Example: Using Servers in Shifts

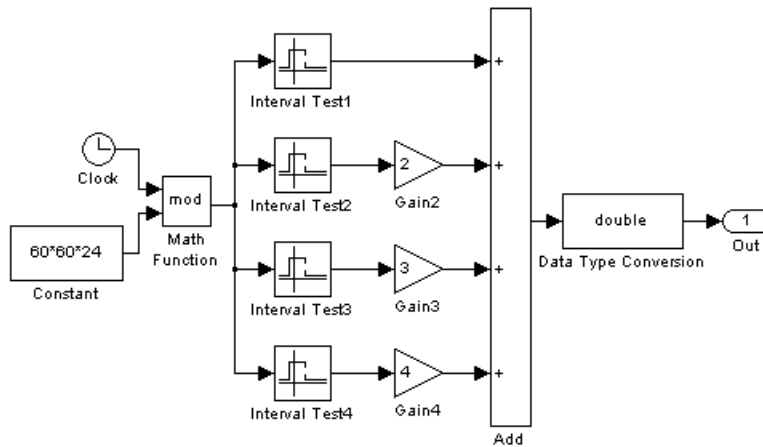
Suppose you have four servers that operate in mutually exclusive shifts of unequal lengths, and you want to direct each entity to the server that is currently on duty. Suppose that the server on duty has the following index between 1 and 4:

$$\text{Index} = \begin{cases} 1 & \text{between midnight and 4 A.M.} \\ 2 & \text{between 4 A.M. and noon} \\ 3 & \text{between noon and 8 P.M.} \\ 4 & \text{between 8 P.M. and midnight} \end{cases}$$

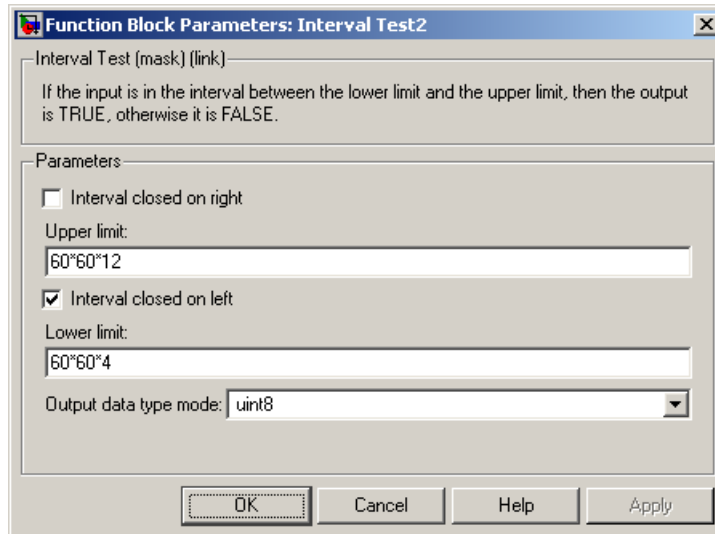
Below are two methods of computing this index in a subsystem. One method uses both logical and numerical blocks, while the other method is purely numerical.

Index Computation 1

You can compute the index of the server on duty using a subsystem like the one shown below, where the Interval Test blocks use **Lower limit** and **Upper limit** parameter values that represent the start and end of each shift in each 24-hour day.



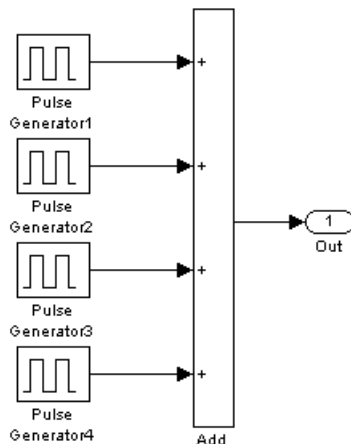
For example, the second shift is represented by the second Interval Test block, whose dialog box is shown below.



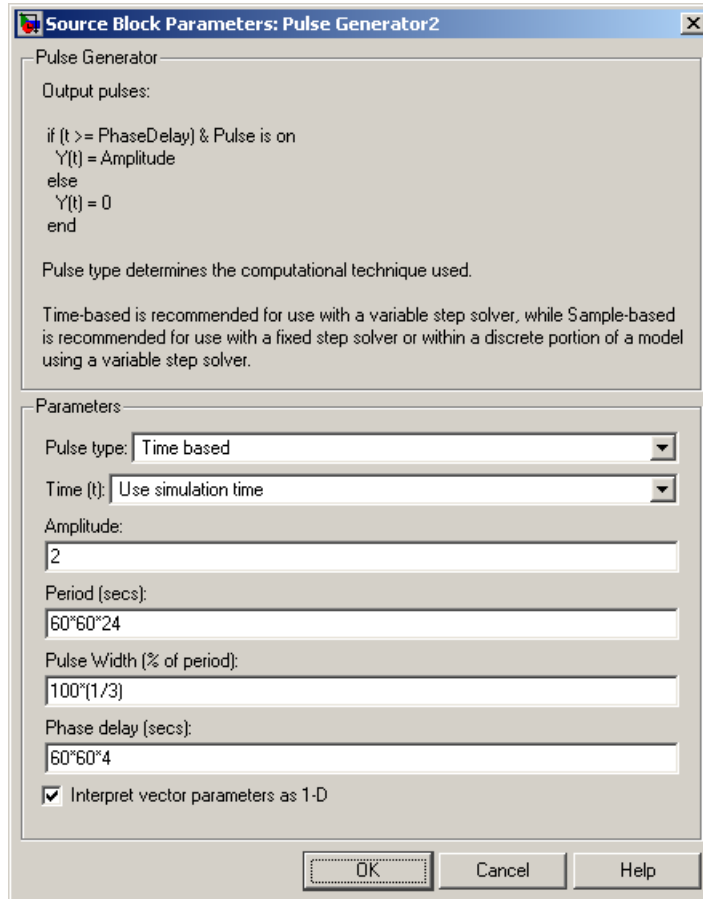
Index Computation 2

Alternatively, you can compute the index of the server on duty using a subsystem like the one shown below, where each Pulse Generator block assumes the corresponding index value when that server is on duty, and assumes the value 0 at other times. In particular,

- **Period** = $60 \times 60 \times 24$ for all Pulse Generator blocks, to represent a daily cycle in seconds
- **Amplitude** gives the index for each server, between 1 and 4
- **Pulse width** gives the length of each server's shift
- **Phase delay** gives the starting time of each server's shift

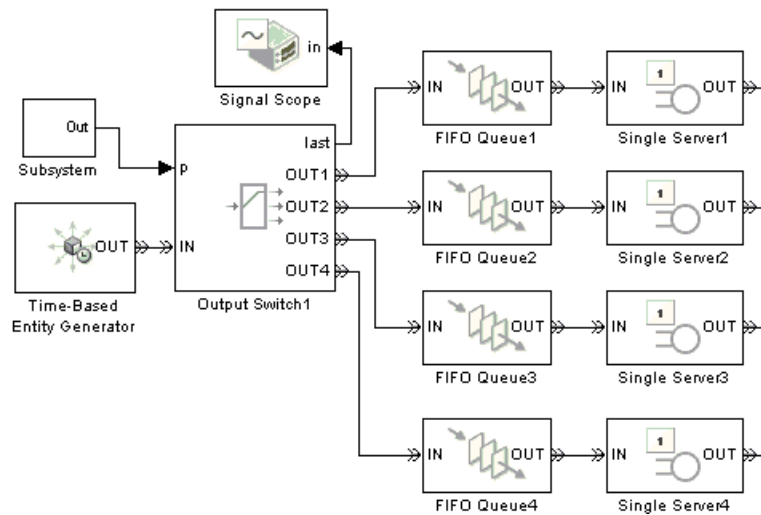


For example, the second shift is represented by the second Pulse Generator block, whose dialog box is shown below.



Top-Level Model

The figure below shows how you can integrate either kind of index computation, contained in a subsystem, into a larger model. It is similar to the example in “Example: Choosing the Shortest Queue Using Logic Blocks” on page 6-16 except that this example uses different switching logic that does not depend on feedback from the queues. The subsystem in this model is a virtual subsystem used for visual simplicity, not a discrete event subsystem.

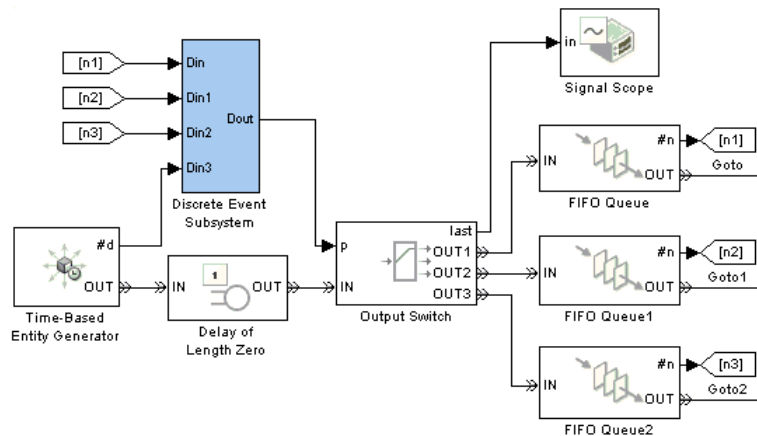


The sample plot below reflects the use of shifts. Each plotting marker corresponds to an entity departing from the switch block via one of the four entity output ports.



Example: Choosing the Shortest Queue Using Logic Blocks

This example, a variation on the model in “Example: Choosing the Shortest Queue” on page 6-3, directs entities to the shortest of three queues. The discrete event subsystem computes the index of the shortest queue using logic blocks. If more than one queue achieves the minimum, then the computation returns the smallest index among the queues that minimize the length.

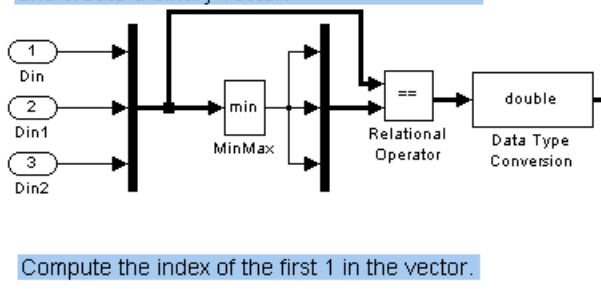


Top-Level Model


Do Not Delete

Subsystem Configuration


Determine which #n inputs attain the minimum and create a binary vector.



Compute the index of the first 1 in the vector.



Use #d signal to call the subsystem.



Subsystem Contents

Regulating Arrivals Using Gates

Role of Gates in SimEvents Models
(p. 7-2)

What gates represent in various models

Keeping a Gate Open Over a Time Interval (p. 7-4)

Admitting entities during a time interval of arbitrary length

Opening a Gate Instantaneously
(p. 7-6)

Admitting an entity when an event occurs

Using Logical Combinations of Gates
(p. 7-9)

Implementing compound logic for entity admission

Role of Gates in SimEvents Models

By design, certain blocks change their availability to arriving entities depending on the circumstances. For example,

- A queue or server accepts arriving entities as long as it is not already full to capacity.
- An input switch accepts an arriving entity through a single selected entity input port but forbids arrivals through other entity input ports.

Some applications require more control over whether and when entities advance from one block to the next. A gate provides flexible control via its changing status as either open or closed: by definition, an open gate permits entity arrivals as long as the entities would be able to advance immediately to the next block, while a closed gate forbids entity arrivals. You configure the gate so that it opens and closes under circumstances that are meaningful in your model.

For example, you might use a gate

- To create periods of unavailability of a server. For example, you might be simulating a manufacturing scenario over a monthlong period, where a server represents a machine that runs only 10 hours per day. An enabled gate can precede the server, to make the server's availability contingent upon the time.

To learn about enabled gates, which can remain open for a time interval of nonzero length, see “Keeping a Gate Open Over a Time Interval” on page 7-4.

- To make departures from one queue contingent upon departures from a second queue. A release gate can follow the first queue. The gate's control signal determines when the gate opens, based on decreases in the number of entities in the second queue.

To learn about release gates, which open and then close in the same time instant, see “Opening a Gate Instantaneously” on page 7-6.

- With the `First` port that is not blocked mode of the Output Switch block. Suppose each entity output port of the switch block is followed by a

gate block. An entity attempts to advance via the first gate; if it is closed, then the entity attempts to advance via the second gate, and so on.

This arrangement is explored in “Using Logical Combinations of Gates” on page 7-9.

To learn about implementing logic that determines when a gate is open or closed, see Chapter 6, “Using Logic”.

Accessing Gate Blocks

The gate blocks reside in the Gates library of SimEvents.

A gate block forbids or permits entities to advance from the block before the gate to the block after the gate. For example, if you want to control advancement from a queue to a server, then place the gate block after the queue and before the server. Many models follow a gate with a storage block, such as a queue or server.

Types of Gates

The Gates library offers these fundamentally different kinds of gate blocks:

- The Enabled Gate block, which uses a control signal to determine time intervals over which the gate is open or closed. For more information, see “Keeping a Gate Open Over a Time Interval” on page 7-4.
- The Release Gate block, which uses a control signal to determine a discrete set of times at which the gate is instantaneously open. The gate is closed at all other times during the simulation. For more information, see “Opening a Gate Instantaneously” on page 7-6.

Keeping a Gate Open Over a Time Interval

The Enabled Gate block uses a control signal at the input port labeled **en** to determine when the gate is open or closed:

- When the **en** signal is positive, the gate is open and an entity can arrive as long as it would be able to advance immediately to the next block.
- When the **en** signal is zero or negative, the gate is closed and no entity can arrive.

Because the **en** signal can remain positive for a time interval of arbitrary length, an enabled gate can remain open for a time interval of arbitrary length. The length can be zero or a positive number.

Depending on your application, the **en** signal can arise from time-driven dynamics, state-driven dynamics, a SimEvents block's statistical output signal, or a computation involving various types of signals.

Example: Controlling Joint Availability of Two Servers


Suppose that each entity undergoes two processes, one at a time, and that the first process does not start if the second process is still in progress for the previous entity. Assume for this example that it is preferable to model the two processes using two Single Server blocks in series rather than one Single Server block whose service time is the sum of the two individual processing times; for example, you might find a two-block solution more intuitive or you might want to access the two Single Server blocks' utilization output signals independently in another part of the model.

If you connect a queue, a server, and another server in series, then the first server can start serving a new entity while the second server is still serving the previous entity. This does not accomplish the stated goal. The model needs a gate to prevent the first server from accepting an entity too soon, that is, while the second server still holds the previous entity.

One way to implement this is to precede the first Single Server block with an Enabled Gate block that is configured so that the gate is closed when an entity is in either server. In particular, the gate

- Is open from the beginning of the simulation until the first entity's departure from the gate
- Closes whenever an entity advances from the gate to the first server, that is, when the gate block's **#d** output signal increases
- Reopens whenever that entity departs from the second server, that is, when the second server block's **#d** output signal increases

This arrangement is shown below.



The Signal Latch block's **st** output signal becomes 0 when the block's **rvc** input signal increases and becomes 1 when the **wvc** input signal increases. That is, the **st** signal becomes 0 when an entity departs from the gate and becomes 1 when an entity departs from the second server. The block labeled Open Gate at Start of Simulation is another Signal Latch block; its purpose is to modify the **st** signal only by defining an initial condition of 1 (using the technique described in "Defining Initial Conditions for Event-Based Signals" on page 3-18). In summary, the entity at the head of the queue advances to the first Single Server block if and only if both servers are empty.


Opening a Gate Instantaneously

The Release Gate block opens instantaneously at a discrete set of times during the simulation and is closed at all other times. The gate opens when a signal-based event or a function call occurs. By definition, the gate's opening permits one pending entity to arrive if able to advance immediately to the next block. No simulation time passes between the opening and subsequent closing of the gate; that is, the gate opens and then closes in the same time instant. If no entity is already pending when the gate opens, then the gate closes without processing any entities. It is possible for the gate to open multiple times in a fixed time instant, if multiple gate-opening events occur in that time instant.



An entity passing through a gate must already be pending before the gate-opening event occurs. Suppose a Release Gate block follows a Single Server block and a gate-opening event is scheduled simultaneously with a service completion event. If the gate-opening event is processed first, then the gate opens and closes before the entity completes its service, so the entity does not pass through the gate at that time instant. If the service completion is processed first, then the entity is already pending before the gate-opening event is processed, so the entity passes through the gate at that time instant. To learn more about the processing sequence for simultaneous events, see "Setting Event Priorities" on page 2-19.

Example: Synchronizing Service Start Times with the Clock

In the example below, a Release Gate block with an input signal from a Pulse Generator block ensures that entities begin their service only at fixed time steps of 1 second, even though the entities arrive asynchronously. In this example, the Release Gate block has **Open gate upon** set to `Change in signal from port vc` and **Type of change in signal value** set to `Rising`, while the Pulse Generator block has **Period** set to 1. (Alternatively, you could set **Open gate upon** to `Trigger from port tr` and **Trigger type** to `Rising`.)

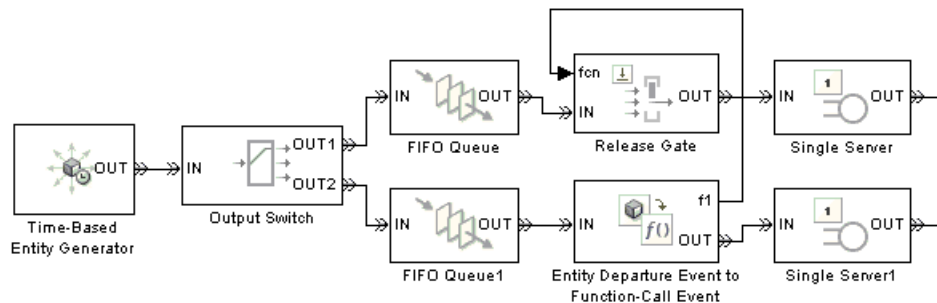


The plots below show that the entity generation times can be noninteger values, but the service beginning times are always integers.



Example: Opening a Gate Upon Entity Departures

In the model below, two queue-server pairs operate in parallel and an entity departs from the top queue only in response to a departure from the bottom queue. In particular, departures from the bottom queue block cause the Entity Departure Event to Function-Call Event block to issue a function call, which in turn causes the gate to open. The Release Gate block in this model has the **Open gate upon** parameter set to Function call from port fcn.



If the top queue in the model is empty when the bottom queue has a departure, then the gate opens but no entity arrives there.

When configuring a gate to open based on entity departures, be sure the logic matches your intentions. For example, when looking at the model shown above, you might assume that entities advance through the queue-server pairs during the simulation. However, if the Output Switch block is configured to select the first entity output port that is not blocked, and if the top queue has a large capacity relative to the number of entities generated during the simulation duration, then you might find that all entities advance to the top queue, not the bottom queue. As a result, no entities depart from the bottom queue and the gate never opens to permit entities to depart from the top queue. By contrast, if the Output Switch block is configured to select randomly between the two entity output ports, then it is likely that some entities reach the servers as expected.

Alternative Using Value Change Events

An alternative to opening the gate upon departures from the bottom queue is to open the gate upon changes in the value of the **#d** signal output from that queue block. The **#d** signal represents the number of entities that have departed from that block, so changes in the value are equivalent to entity departures. To implement this approach, set the Release Gate block's **Open gate upon** parameter to Change in signal from port **vc** and connect the **vc** port to the queue block's **#d** output signal.


Using Logical Combinations of Gates

You can use multiple gate blocks in combination with each other:

- Using a Release Gate block and/or one or more Enabled Gate blocks in series is equivalent to a logical AND of their gate-opening criteria. For an entity to pass through the gates, they must all be open at the same time.

Note You should not connect two Release Gate blocks in series. No entities would ever pass through such a series of gates because each gate closes before the other gate opens, even if the gate-opening events occur at the same value of the simulation clock.

- Using multiple gate blocks in parallel, you can implement a logical OR of their gate-opening criteria. Use the Output Switch and Path Combiner blocks as in the figure below and set the Output Switch block's **Switching criterion** parameter to First port that is not blocked.




Each entity attempts to arrive at the first gate; if it is closed, the entity attempts to arrive at the second gate, and so on. If all gates are closed, then the Output Switch block's entity input port is unavailable and the entity must stay in a preceding block (such as a queue or server preceding the switch).

Note The figure above uses two Release Gate blocks and one Enabled Gate block, but you can use whatever combination is suitable for the logic of your application and whatever sequence you prefer. Also, the figure above omits the control signals (**vc** and **en**) for visual clarity but in your model these ports must be connected.

If you expect gate-opening or gate-closing events for different gates to occur at the same time, then consider the sequence of events. For example, if you connect an Enabled Gate block to a Release Gate block in series and the enabled gate closes at the same time instant that the release gate opens, then the sequence of events matters. If the gate-closing event is processed first, then a pending entity cannot pass through the gates at that time; if the gate-opening event is processed first, then a pending entity can pass through the gates before the gate-closing event is processed. To control the sequence, select the **Specify event priority...** parameters in the gate blocks and specify appropriate **Event priority** parameters. For details, see “Setting Event Priorities” on page 2-19.

Example: First Entity as a Special Case

This example illustrates the use of a Release Gate block and an Enabled Gate block connected in parallel. The Release Gate block permits the arrival of the first entity of the simulation, which receives special treatment, while the Enabled Gate block permits entity arrivals during the rest of the simulation. In this example, a warmup period at the beginning of the simulation precedes normal processing.




The Release Gate block is open precisely when the **#d** output signal from the Time-Based Entity Generator block rises from 0 to 1. That is, the gate is open for the first entity of the simulation and no other entities. The first entity arrives at an Infinite Server block, which represents the warmup period.

Subsequent entities find the Release Gate block's entity input port unavailable, so they attempt to arrive at the Enabled Gate block. The Enabled Gate block is open during the entire simulation, except when the first entity has not yet departed from the Infinite Server block. This logic is necessary to prevent the second entity from jumping ahead of the first entity before the warmup period is over.

The Path Combiner block merges the two entity paths, removing the distinction between them. Subsequent processing depends on your application; this model merely uses a queue-server pair as an example.

The plot below shows which path each entity takes during the simulation. The plot shows that the first entity advances from the first (path=1) entity output port of the Output Switch block to the Release Gate block, while subsequent entities advance from the second (path=2) entity output port of the Output Switch block to the Enabled Gate block.



Controlling Timing Using Subsystems

Timing Issues in SimEvents Models (p. 8-2)	Understanding timing in situations where time-based and event-based behavior interact
Role of Discrete Event Subsystems in SimEvents Models (p. 8-7)	Overview of discrete event subsystems as a way to ensure appropriate simulation timing
Blocks Inside Discrete Event Subsystems (p. 8-10)	Blocks that are suitable for use in a discrete event subsystem
Working with Discrete Event Subsystem Blocks (p. 8-11)	Setting up and configuring Discrete Event Subsystem blocks
Examples Using Discrete Event Subsystem Blocks (p. 8-16)	Examples of situations in which discrete event subsystems are useful
Creating Entity-Departure Subsystems (p. 8-26)	How to call a subsystem when an entity departs from a block
Examples Using Entity-Departure Subsystems (p. 8-29)	Examples of situations in which entity-departure subsystems are useful
Using Function-Call Subsystems (p. 8-32)	Using the Function-Call Subsystem block in SimEvents models

Timing Issues in SimEvents Models

Most SimEvents models contain one or more inherently time-based blocks from the Simulink libraries, in addition to inherently event-based blocks from the SimEvents libraries. When time-based and event-based behavior combine in one model, it is important to use correct modeling techniques to ensure correct timing in the simulation.


When you combine time-based and event-based blocks, consider whether you want the operation of a time-based block to depend *only* on time-oriented considerations or whether you want the time-based block to respond to events. The following example scenarios illustrate the difference:

- “Timing for the End of the Simulation” on page 8-2
- “Timing for a Statistical Computation” on page 8-3
- “Timing for Choosing a Port Using a Sequence” on page 8-4

The section “Role of Discrete Event Subsystems in SimEvents Models” on page 8-7 introduces the discrete event subsystem, an important modeling construct that you can use to make time-based blocks respond to events.

Timing for the End of the Simulation

Consider a queuing model in which you want to end the simulation precisely when the number of entities in the queue first equals or exceeds some threshold value. One way to model the constraint is to use the Compare To Constant block to check whether the queue length is greater than or equal to the threshold value and the Stop Simulation block to end the simulation whenever the comparison is true.



The queue length is an optional output from the FIFO Queue block from the signal output port labeled **#n**.



Time-Based Default Behavior

Connecting the **#n** signal directly to the signal input port of the Compare To Constant block looks correct topologically. However, that would *not* cause the simulation to stop at the exact moment when the **#n** signal reaches or exceeds the threshold for the first time. Instead, it would cause the simulation to stop upon the next time step for the time-based Compare To Constant and Stop Simulation blocks, which could be at a later time. The **#n** signal experiences a change in value based on an event, at a time that is unrelated to the time-based simulation clock, whereas Simulink defaults to time-based behavior for most blocks.

Achieving Correct Event-Based Behavior

The correct way to cause the simulation to stop at the exact moment when the **#n** signal reaches or exceeds the threshold for the first time is to construct the model so that the Compare To Constant and Stop Simulation blocks respond to events, not the time-based simulation clock. Put these blocks inside a discrete event subsystem that is executed at exactly those times when the FIFO Queue block's **#n** signal increases from one integer to another.

For details on this solution, see “Example: Ending the Simulation Upon an Event” on page 8-20.

Timing for a Statistical Computation

Suppose that you want to compute the total length of time for which the queue length equals or exceeds a threshold value during the simulation. A symbolic expression for this computation is

$$f(T) = \int_0^T I_{\#n \geq 3}(t) dt$$

where T is the total length of the simulation, $\#n$ is the instantaneous queue length, the threshold is three entities, and I is an indicator function defined by

$$I_{\#n \geq 3}(t) = \begin{cases} 1 & \text{if } \#n(t) \geq 3 \\ 0 & \text{otherwise} \end{cases}$$

You can use the optional **#n** signal output port from the FIFO Queue block to produce the queue length signal, and the Compare To Constant block to implement the indicator function. You can connect the resulting signal to an integrator block.

Time-Based Integration

The Integrator and Discrete-Time Integrator blocks in the Simulink libraries are inherently time-based. In this example, you are integrating a signal that experiences asynchronous discontinuities. A discrete-time integrator with an explicit sample time yields incorrect results because a discontinuity might occur between successive sample time values. A continuous-time integrator can make the simulation clock adjust to detected discontinuities in the signal, yet it is still possible for asynchronous discontinuities to go undetected by the time-based simulation clock.

Some applications require time-based integration, but this application requires that the integration respond to events that change the value of the **#n** signal.


Achieving Correct Event-Based Behavior

To make the Compare To Constant and integrator blocks respond to events, not the time-based simulation clock, you can put these blocks inside a discrete event subsystem that is executed at exactly those times when the FIFO Queue block's **#n** signal increases from one integer to another. This solution causes the simulation to compute the indicator function correctly, which results in a correct value for the statistic.

For details on this solution, see “Example: Using Event-Based Timing for a Statistical Computation” on page 8-19.

Timing for Choosing a Port Using a Sequence

Consider an Output Switch block that directs each arriving entity to one of three entity output ports, depending on the value of an input signal.



Suppose a Repeating Sequence Stair block generates the input signal by cycling through the values 3, 2, and 1 throughout the simulation.



So far, this description does not indicate *when* the Repeating Sequence Stair block changes its value. Here are some possibilities:

- It chooses a new value from the repeating sequence at regular *time* intervals, which might be appropriate if the switch is intended to select among three operators who work in mutually exclusive shifts in time.

You can specify the time interval in the Repeating Sequence Stair block's **Sample time** parameter.

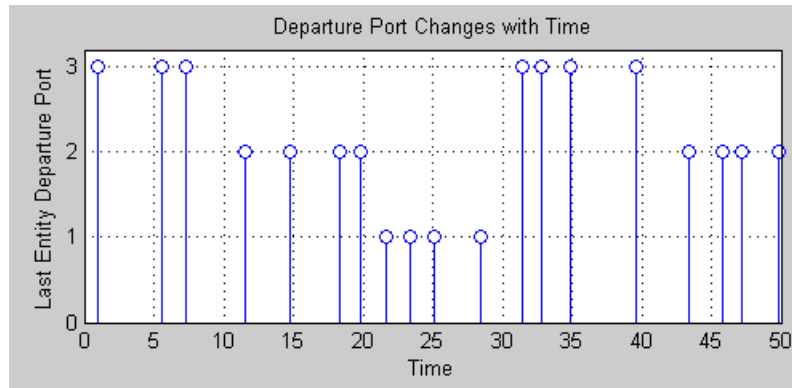
- It chooses a new value from the repeating sequence whenever an entity arrives at the Output Switch block, which might be appropriate if the switch is intended to distribute a load fairly among three operators who work in parallel.

To make the Repeating Sequence Stair block respond to entity advancement events, not the time-based simulation clock, you can put these blocks inside an appropriately configured function-call subsystem, as discussed in “Creating Entity-Departure Subsystems” on page 8-26.


For details on this approach, see “Example: Using Entity-Based Timing for Choosing a Port” on page 8-29.

These possibilities correspond to qualitatively different interpretations of the model as well as quantitatively different results. If the Output Switch block reports the index of its last selected entity output port (that is, the entity output port through which the most recent entity departure occurred), then a

plot of this statistic against time looks quite different depending on the timing of the Repeating Sequence Stair block's operation. Sample plots are below.



Departure Port Changes with Time



Departure Port Changes with Each Entity

Role of Discrete Event Subsystems in SimEvents Models

Given the questions raised in “Timing Issues in SimEvents Models” on page 8-2 about the response of time-based blocks to events, this section gives an overview of discrete event subsystems and describes how you can use them to ensure appropriate simulation timing. A discrete event subsystem

- Contains time-based blocks. Examples include
 - Constant with a sample time of `inf`
 - Sum or Relational Operator with a sample time of `-1`
 - Stop Simulation
- Cannot contain blocks from the SimEvents libraries, except the Discrete Event Inport, Discrete Event Outport, and Subsystem Configuration blocks

Note If you want to put blocks that have entity ports into a subsystem as a way to group related blocks or to make a large model more hierarchical, then use an ordinary Subsystem block from the Simulink Ports & Subsystems library. Alternatively, select one or more blocks and use the **Edit > Create Subsystem** menu option. In either case, the use of a subsystem does not affect the timing of the simulation but is merely a graphical construct.

- Has two basic forms: a Discrete Event Subsystem block and an appropriately configured Function-Call Subsystem block.
- Is executed in response to signal-based events that you specify in the Discrete Event Inport blocks inside the Discrete Event Subsystem window, or in response to function calls in the case of a function-call subsystem.

“Block execution” in this documentation is shorthand for “block methods execution.” Methods are functions that Simulink uses to solve a model. Blocks are made up of multiple methods. For details, see “Block Methods” in the Simulink documentation.

Purpose of Discrete Event Subsystems


The purpose of a discrete event subsystem is to call the blocks in the subsystem at the exact time of each qualifying event and not at times suggested by the time-based simulation clock. This is an important change in the semantics of the model, not merely an optimization.

Processing Sequence for Events in Discrete Event Subsystems

When creating a discrete event subsystem, you might need to confirm or manipulate the processing sequence for two or more events, such as

- Signal-based events that execute a Discrete Event Subsystem block
- Entity departures that execute an entity-departure subsystem
- Function calls that execute a function-call subsystem
- Updates in the values of signals that serve as inputs to any kind of discrete event subsystem

Consider the schematic below involving a discrete event subsystem. Suppose an entity departure from Block A, an entity arrival at Block C, and updates in all of the signals occur at a given value of the simulation clock.



Typically, the goal is to execute the subsystem

- After the entity departure from Block A, which produces a signal that is an input to the subsystem.
- After both Block A and Block B update their output signals at that value of the simulation clock.

Be especially aware of this if you clear the **Execute subsystem upon signal-based events** option in a Discrete Event Inport block. Because the subsystem uses the most recent value of the signal, you should make sure that value is up to date, rather than a value from a previous call to the block that creates the signal.

See “Example: Detecting Changes from Empty to Nonempty” on page 8-23 for an example in which the last-updated signal executes the subsystem. See “Sequence of Updates of Output Signals” on page 3-10 to learn more about when blocks update their output signals.

- Before the entity arrival at Block C, which uses an output signal from the subsystem. “Example: Normalizing a Statistic to Use for Routing” on page 8-17 shows how an extra server block whose service time is zero can help produce the correct processing sequence.

For details on processing sequences, see “Interleaved Operations of Storage and Nonstorage Blocks” on page 12-4 and “Processing Sequence for Simultaneous Events” on page 2-21.

Blocks Inside Discrete Event Subsystems

The only blocks that are suitable for use in a discrete event subsystem are

- Blocks having a **Sample time** parameter of -1, which indicates that the sample time is inherited from the driving block.
- Blocks that always inherit a sample time from the driving block, such as the Bias block or the Unary Minus block. To determine whether a block in one of the Simulink libraries inherits its sample time from the driving block, see the “Characteristics” table near the end of the block’s online reference page.
- Blocks whose outputs cannot change from their initial values during a simulation. For more information, see “Constant Sample Time” in the Simulink documentation.

Types of blocks that are *not* suitable for use in a discrete event subsystem include

- Continuous-time blocks
- Discrete-time blocks with a **Sample time** parameter value that is positive and finite
- Blocks from the SimEvents libraries, except the Discrete Event Inport, Discrete Event Outport, and Subsystem Configuration blocks. In particular, a discrete event subsystem cannot contain blocks that possess entity ports or nested Discrete Event Subsystem blocks.

In some cases, you can work around these restrictions by entering a **Sample time** parameter value of -1 and/or by finding a discrete-time analogue of a continuous-time block. For example, instead of using the continuous-time Clock block, use the discrete-time Digital Clock block with a **Sample time** parameter value of -1.

Working with Discrete Event Subsystem Blocks

Building on the conceptual information in “Role of Discrete Event Subsystems in SimEvents Models” on page 8-7, this section provides some practical information to help you incorporate Discrete Event Subsystem blocks into your SimEvents models. The topics are as follows:

- “Setting Up Signal-Based Discrete Event Subsystems” on page 8-11
- “Signal-Based Events That Control Discrete Event Subsystems” on page 8-14


For discrete event subsystems that respond to entity departures rather than signal-based events, see “Creating Entity-Departure Subsystems” on page 8-26.

Setting Up Signal-Based Discrete Event Subsystems


To create discrete event subsystems that respond to signal-based events, follow the procedure below using blocks in the SimEvents Ports and Subsystems library.

Note You cannot create a signal-based discrete event subsystem by selecting blocks and using **Edit > Create Subsystem** or by converting a time-based subsystem into a discrete event subsystem. If your model already contains the blocks you want to put into a discrete event subsystem, then you can copy them into the subsystem window of a Discrete Event Subsystem block while following the procedure below.


- 1 Drag the Discrete Event Subsystem block from the SimEvents Ports and Subsystems library into your model. Initially, it shows one signal input port **Din** and one output signal port **Dout**. Note that these are signal ports, not entity ports, because the subsystem is designed to process signals rather than entities. Furthermore, the signal input ports carry only real scalar signals of data type double.




- 2** In the model window, double-click the Discrete Event Subsystem block to open the subsystem it represents. Initially, the subsystem contains an inport connected to an outport. Note that these are Discrete Event Inport and Discrete Event Outport blocks, which are not the same as the Inport and Outport blocks in the Simulink Ports & Subsystems library. The subsystem also contains a Subsystem Configuration block, which you should not delete.




- 3** A discrete event subsystem must have at least one input that determines when the subsystem executes. To change the number of inputs or outputs to the subsystem, change the number of inport and outport blocks in the subsystem window:
- If your subsystem requires an additional input or output, then copy and paste an inport block or outport block in the subsystem window. (Copying and pasting is different from duplicating the inport, which is not supported.) The figure below shows a pasted inport block.




As a result, the subsystem block at the upper level of your model shows the additional port as appropriate. The figure below shows an additional input port on the subsystem block.



- If your subsystem needs no outputs, then select and delete the output block in the subsystem window. The figure below shows the absence of output blocks.

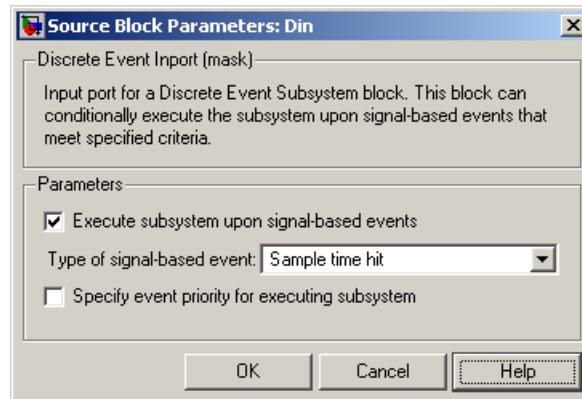


As a result, the subsystem block at the upper level of your model omits the output port. The figure below shows the absence of output ports on the subsystem block.



- 4 Drag other blocks into the subsystem window as appropriate to build the subsystem. This step is similar to the process of building the top level of your model, except that only certain types of blocks are suitable for use inside the subsystem. See “Blocks Inside Discrete Event Subsystems” on page 8-10 for details.

- 5 Configure each of the Discrete Event Inport blocks to indicate when the subsystem should be executed. Each inport block independently determines criteria for executing the subsystem:
- To execute the subsystem when the signal corresponding to that inport exhibits a qualifying signal-based event, select **Execute subsystem upon signal-based events** and use additional parameters in the dialog box to describe the signal-based event.
 - To have the subsystem use the most recent value of the signal corresponding to that inport without responding to signal-based events in that signal, clear the **Execute subsystem upon signal-based events** option.



Signal-Based Events That Control Discrete Event Subsystems

Blocks in a Discrete Event Subsystem are called in response to signal-based events. Using the dialog box of the Discrete Event Inport blocks inside the subsystem, you configure the subsystem so that it is executed in response to

- An updated value in an input signal, even if the updated value is the same as the previous value
- A change in the value of an input signal
- A rising edge or falling edge in an input signal known as a trigger signal

Note To call a subsystem upon an entity departure or upon a function call, see “Creating Entity-Departure Subsystems” on page 8-26 or “Using Function-Call Subsystems” on page 8-32, respectively.

In a discrete event subsystem containing multiple Discrete Event Inport blocks, the subsystem is executed when at least one of the inports detects a qualifying event. Two qualifying events at the same time cause two calls to the subsystem. If you want one of the inports to provide an input signal without affecting the times at which the subsystem is executed, then clear the **Execute subsystem upon signal-based event** check box on that inport block; however, always select **Execute subsystem upon signal-based event** for at least one inport of the subsystem or else the subsystem will never be executed.

Comparison of Event Types for Discrete Event Subsystems

Here are some points to keep in mind when deciding which type of signal-based event should call your discrete event subsystem:

- Value changes are similar to sample time hits, except that a sample time hit might cause a signal to be updated with the same value. If you expect that calling the subsystem for repeated values of an input signal would produce incorrect numerical results or would be wasteful, then execute the subsystem upon changes in the signal value rather than upon every sample time hit.
- The Discrete Event Subsystem block is similar to the ordinary Triggered Subsystem block in the Simulink Ports & Subsystems library. However, the Discrete Event Subsystem block can detect zero-duration values in the input signal, which are signal values that do not persist for a positive duration. (See “Multiple Simultaneous Updates” on page 3-13 for details on zero-duration values.) Unlike the Triggered Subsystem block, the Discrete Event Subsystem block can detect and respond to a trigger edge formed by a zero-duration value, as well as multiple edges in a single instant of time.

For more information about signal-based events, see “Types of Supported Events” on page 2-2.

Examples Using Discrete Event Subsystem Blocks

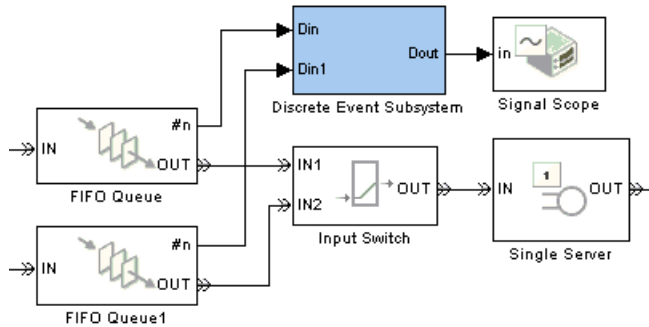
The topics listed below illustrate the use of the Discrete Event Subsystem block.

Topic	Feature of Example
“Example: Comparing the Lengths of Two Queues” on page 8-16	Manipulating two event-based signals having different timing
“Example: Normalizing a Statistic to Use for Routing” on page 8-17	Performing event-oriented computation for use in a subsequent block
“Example: Using Event-Based Timing for a Statistical Computation” on page 8-19	Integrating a signal having asynchronous discontinuities
“Example: Ending the Simulation Upon an Event” on page 8-20	Responding immediately to asynchronous discontinuities
“Example: Sending Unrepeated Data to the MATLAB Workspace” on page 8-21	Logging data only when a signal changes
“Example: Focusing on Events, Not Values” on page 8-22	Counting changes in a signal’s value
“Example: Detecting Changes from Empty to Nonempty” on page 8-23	Manipulating one signal when another signal crosses zero
“Example: Logging Data About the First Entity on a Path” on page 8-24	Focusing on the first entity that uses a path

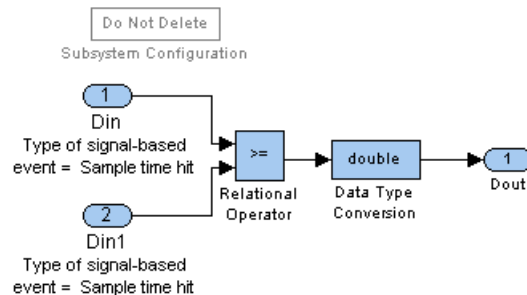
Example: Comparing the Lengths of Two Queues

In a model containing two queues, a logical comparison of the lengths of the queues changes when either queue has an arrival or departure. A queue block’s **#n** output signal is updated after each arrival if the queue is nonempty, and after each departure. By contrast, the Relational Operator block is a time-based block. The model below performs the comparison inside a discrete event subsystem whose Discrete Event Inport blocks both have **Type of signal-based event** set to **Sample time hit**. This way, the comparison occurs whenever either **#n** signal is updated. If both queues update their

#n values at the same time on the simulation clock, then the discrete event subsystem is called twice at that time.



Top-Level Model



Subsystem Contents

Example: Normalizing a Statistic to Use for Routing

Suppose you want to make a routing decision based on an output signal from a SimEvents block, but you must manipulate or normalize the statistic so that the routing block receives a value in the correct range. In the model shown below, the **util** output signal from the Single Server block assumes real values between 0 and 1, while the **p** input signal to the Output Switch block expects values of 1 or 2. The discrete event subsystem adds 1 to the rounded utilization value.

The Discrete Event Inport block inside the subsystem has **Type of signal-based event** set to **Sample time hit** so that the computation occurs


whenever the server block updates the value of the utilization signal. The subsystem's configuration and the presence of the Delay of Length Zero block ensure that the routing decision always uses the most up-to-date value of the **util** signal.

At a time instant at which an entity departs from the Single Server block, the sequence of operations is as follows:


- 1 The entity advances from the Single Server block to the block labeled Delay of Length Zero, which is merely a single server whose service time is zero.
- 2 The Single Server block updates its **util** output signal.

The fact that this occurs after the entity has already departed from the Single Server block and arrived at a storage block is the reason for the Delay of Length Zero block. See “Interleaved Operations of Storage and Nonstorage Blocks” on page 12-4 for details.

- 3 The subsystem reacts to the updated value of **util** by computing an updated value at the **p** input port of the Output Switch block.
- 4 The entity advances from the Delay of Length Zero block to the Output Switch, which uses an up-to-date **p** signal to determine the routing behavior.



Top-Level Model




Subsystem Contents


Example: Using Event-Based Timing for a Statistical Computation

This example performs the computation described in “Timing for a Statistical Computation” on page 8-3, to find the total length of time during the simulation that a queue’s length equals or exceeds a threshold value.

In the model shown below, the threshold is three entities. The Compare To Constant block produces an indicator function for that threshold. The Discrete Time Integrator block, configured to have an inherited sample time rather than an explicit discrete sample time, computes the area under the curve of the indicator function, that is, the total amount of time that the input to the subsystem exceeds the threshold. The Discrete Event Inport block inside the subsystem has **Type of signal-based event** set to Sample time hit so that the computation occurs whenever the queue block updates the value of the queue length signal. The subsystem’s configuration ensures that the comparison block does not miss any asynchronous discontinuities in the queue length signal and that the integrator processes the correct indicator function.



Top-Level Model




Subsystem Contents


Note In this example, it is important to use a discrete-time integrator with an inherited sample time instead of a continuous-time integrator. See “Blocks Inside Discrete Event Subsystems” on page 8-10 for more information.

Example: Ending the Simulation Upon an Event

This example ends the simulation as described in “Timing for the End of the Simulation” on page 8-2, precisely when the number of entities in a queue first equals or exceeds a threshold. In the model shown below, the Compare To Constant and Stop Simulation blocks are in a discrete event subsystem. The Discrete Event Inport block inside the subsystem has **Type of signal-based event** set to Sample time hit so that Simulink calls the subsystem at exactly those times when the FIFO Queue block updates the value of the queue length signal.



Top-Level Model

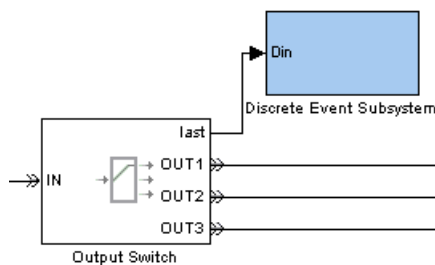


Subsystem Contents


Example: Sending Unrepeated Data to the MATLAB Workspace

Suppose you want to log statistics to the MATLAB workspace, but you want to save simulation time and memory by capturing only values that are relevant to you. You might want to suppress repeated values, for example, or capture only values that represent an increase from the previous value.

In the model shown below, the discrete event subsystem contains a To Workspace block whose **Save format** parameter is set to Structure With Time. The Discrete Event Inport block inside the subsystem has **Type of signal-based event** set to Change in signal and **Type of change in signal value** set to Either, so that the MATLAB workspace variable tells you when the Output Switch block selects an entity output port that differs from the previously selected one. If, for example, the switch is configured to select the first port that is not blocked, then a change in the port selection indicates a change in the state of the simulation (that is, a previously blocked port has become unblocked, or a port becomes blocked that previously was not).



Top-Level Model



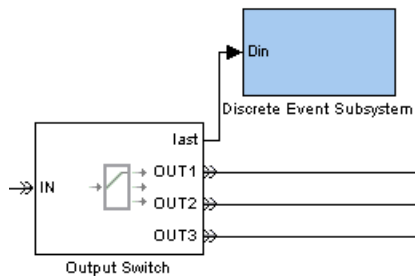
Subsystem Contents

Example: Focusing on Events, Not Values

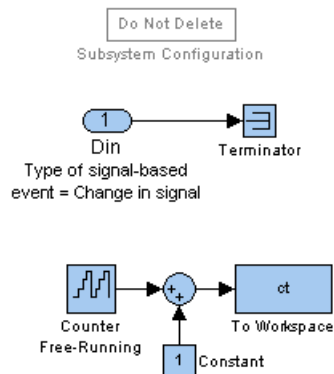
The example below counts the number of times a signal changes its value, ignoring times at which the signal might be updated with the same value. The discrete event subsystem contains a Counter Free-Running block with an inherited sample time. Because the Counter Free-Running counts starting from 0, the subsystem also adds 1 to the counter output.

The Discrete Event Inport block inside the subsystem has **Type of signal-based event** set to **Change in signal** and **Type of change in signal value** set to **Either**, so that the subsystem is executed each time the input signal changes its value. In contrast to other subsystem examples, this subsystem does not use the signal's specific values for computations; the input signal is connected to a Terminator block inside the subsystem. The counter's value is what the subsystem sends to the MATLAB workspace.

In this example, avoiding extraneous calls to the subsystem is not merely a time-saver or memory-saver, but rather a strategy for producing the correct results.



Top-Level Model




Subsystem Contents

Example: Detecting Changes from Empty to Nonempty

The example below executes a subsystem only when an N-server changes from empty to nonempty, or vice versa, but not when the number of entities in the server remains constant or changes between two nonzero values. In the model, the N-Server block produces a $\#n$ signal that indicates the number of entities in the server. The server is empty if and only if the $\#n$ signal is zero. Connected to the $\#n$ signal is a Discrete Event Import block inside the subsystem that has **Type of signal-based event** set to Trigger and **Trigger type** set to Either, so that the subsystem is executed each time the $\#n$ signal changes from zero to one or from one to zero. Connected to the w signal is another Discrete Event Import block inside the subsystem; this

block has **Execute subsystem upon signal-based events** cleared so that this signal does not cause additional calls to the subsystem; the subsystem merely uses the most recent value of the **w** signal whenever the **#n** signal exhibits a trigger edge.

Note Because the N-Server block updates the **w** signal before updating the **#n** signal, both signals are up to date when the trigger edge occurs.



If the server changes instantaneously from empty to nonempty and back to empty, then the subsystem is called exactly twice in the same time instant, once for the rising edge and once for the subsequent falling edge. The Triggered Subsystem block might not detect the edges that the zero-duration value of 1 creates, and thus might not call the subsystem at that time instant. This is why the Discrete Event Subsystem block is more appropriate for this application.

Example: Logging Data About the First Entity on a Path


Suppose your model includes a particular entity path that entities rarely use, and you want to record certain attributes of the *first* entity that takes that path during the simulation. You can send the attribute values to the MATLAB workspace by using a To Workspace block inside a discrete event subsystem.

In the example below, the **#d** output signal from the Get Attribute block indicates how many entities have departed from the block. The other outputs from that block are the attribute values that you want to send to the MATLAB workspace. Connected to the **#d** signal is a Discrete Event Inport block inside the subsystem that has **Type of signal-based event** set to Trigger, so that the subsystem is executed each time the **#d** signal changes from zero to one. Connected to the **A1** and **A2** signals are additional Discrete Event Inport blocks inside the subsystem. These blocks have **Execute subsystem**

upon signal-based events cleared so that the attribute signals do not cause additional calls to the subsystem; the subsystem merely uses the most recent value of the **A1** and **A2** signals whenever the **#d** signal exhibits a trigger edge.

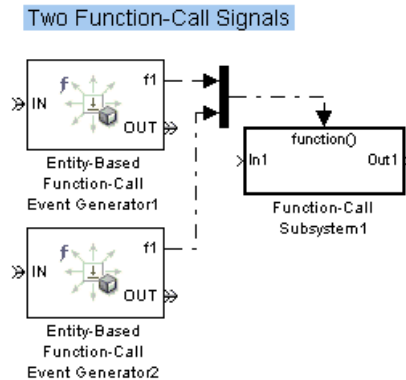

The To Workspace block inside the subsystem does not actually create the variables in the MATLAB workspace until the simulation ends, but the variable contents are correct because the timing of the subsystem corresponds to the time of the **#d** signal's first positive value.

Note Because the Get Attribute block updates the **A1** and **A2** signals before updating the **#d** signal, all signals are up to date when the trigger edge occurs.



Creating Entity-Departure Subsystems

You can create a subsystem that Simulink calls only when an entity departs from a particular block in your model. The figure below shows a prototype, although most ports are not yet connected. The prototype uses the Entity-Based Function-Call Event Generator block to generate a function call when an entity departs. The function call executes the subsystem.



Prototype of Entity-Departure Subsystems

Accessing Blocks for Entity-Departure Subsystems

To create discrete event subsystems that respond to entity departures, use some or all of the blocks listed below.

Block	Library Location	Purpose
Entity-Based Function-Call Event Generator	Event Generators library in SimEvents	Issues a function call corresponding to each entity departure
Entity Departure Event to Function-Call Event	Event Translation library in SimEvents	
Mux	Signal Routing library in Simulink	Combines multiple function-call signals into a single function-call signal, if needed
Function-Call Subsystem	Ports & Subsystems library in Simulink	Contains blocks to execute upon each function call. You must configure the subsystem to propagate its execution context, as described in “Creating Entity-Departure Subsystems” on page 8-26.
Inport	Ports & Subsystems library in Simulink	Links a subsystem to its parent system
Outport		

Setting Up Entity-Departure Subsystems

To create subsystems that respond to entity departures, follow the procedure below.

- 1 Insert and configure the Function-Call Subsystem block as described in “Setting Up Function-Call Subsystems in SimEvents Models” on page 8-32.
- 2 Insert one or more of these blocks into your model. The first is easier to use but less flexible.
 - Entity-Based Function-Call Event Generator
 - Entity Departure Event to Function-Call Event

Note You can configure these blocks to issue a function call either before or after the entity departs. In most situations, the After entity departure option is more appropriate. The Before entity departure option can be problematic if a subsystem is executed too soon, but is an appropriate choice in some situations.

- 3** Connect the newly inserted blocks to indicate which entity departures should call the subsystem. If entity departures from multiple blocks should call the subsystem, then combine multiple function-call signals using a Mux block.

Examples Using Entity-Departure Subsystems

The topics listed below illustrate the use of entity-departure subsystems.


Topic	Feature of Example
“Example: Using Entity-Based Timing for Choosing a Port” on page 8-29	Calling a signal source when an entity departs from a block
“Example: Performing a Computation on Selected Entity Paths” on page 8-31	Detecting departures from multiple blocks

Example: Using Entity-Based Timing for Choosing a Port


This example performs the entity-based routing described in “Timing for Choosing a Port Using a Sequence” on page 8-4. The example routes entities by establishing a sequence of paths and then choosing a number from that sequence for each entity that arrives at the routing block. This is the situation shown in the figure Departure Port Changes with Each Entity on page 8-6.

In the model shown below, the Function-Call Subsystem block contains a Repeating Sequence Stair block whose **Sample time** parameter is set to -1 (inherited). Any entity that arrives at the Output Switch block previously departed from the Entity Departure Event to Function-Call Event block. The function-call output from that block caused the subsystem to produce a number that indicates which entity output port the entity uses when it departs from the Output Switch block.

If you used the Repeating Sequence Stair block with an explicit sample time and not inside a subsystem, then the routing behavior would depend on the clock, as shown in the figure Departure Port Changes with Time on page 8-6, rather than on entity departures.



Top-Level Model




Subsystem Contents

The Entity Departure Event to Function-Call Event block, which issues function calls after entity departures, appears before the Single Server block instead of between the Single Server block and the Output Switch block. This placement ensures that when the function call executes the subsystem, the entity has not yet arrived at the switch but rather is stored in the server. Possible alternative approaches are to

- Place the Entity Departure Event to Function-Call Event block, followed by an extra server, between the Single Server block and the switch. This approach is similar to the use of the Delay of Length Zero block in “Example: Normalizing a Statistic to Use for Routing” on page 8-17.
- Configure the Entity Departure Event to Function-Call Event block to issue function calls before entity departures and place the block between the server and the switch. Issuing function calls before departures does not lead to causality problems in this particular example.

Example: Performing a Computation on Selected Entity Paths

The model below performs a computation whenever an entity arrives at the **IN2** or **IN3** entity input port of a Path Combiner block, but not when an entity arrives at the **IN1** port of the Path Combiner block. The computation occurs inside the Function-Call Subsystem block. When an entity departs from specific blocks that precede the Path Combiner block, the corresponding Entity-Based Function-Call Event Generator block issues a function call. A Mux block combines the two function-call signals, creating a function-call signal that calls the subsystem. If both event generators issue a function call at the same value of the simulation clock, then the subsystem is called twice at that time.



Using Function-Call Subsystems

The most general kind of discrete event subsystem is an appropriately configured Function-Call Subsystem block, where the appropriate configuration requires selecting the **Propagate execution context across subsystem boundary** option as a subsystem property. You can execute such a subsystem at the exact time of an input function call, whether the function call comes from a Stateflow block, a block in the Event Generators library, a block in the Event Translation library, or the Function-Call Generator block.

Use Cases for Function-Call Subsystems

The Discrete Event Subsystem block and the entity-departure subsystems discussed in “Working with Discrete Event Subsystem Blocks” on page 8-11 and “Creating Entity-Departure Subsystems” on page 8-26, respectively, are special cases of the Function-Call Subsystem block configured as a discrete event subsystem. You might require the additional flexibility provided by the Function-Call Subsystem approach if you want to execute the subsystem upon

- The logical OR of multiple event occurrences, where the events can come from any combination of a Stateflow block, another source of function calls, or a signal-based event. To do this, use the Mux block to combine multiple function-call signals into one function-call signal.

For an example, see “Example: Executing a Subsystem Based on Multiple Types of Events” on page 2-43.

- The logical AND of an event occurrence and some underlying condition. To do this, use blocks in the Event Translation library and select **Suppress function call f1 if enable signal e1 is not positive** (or the similar option for the **f2** and **e2** ports).

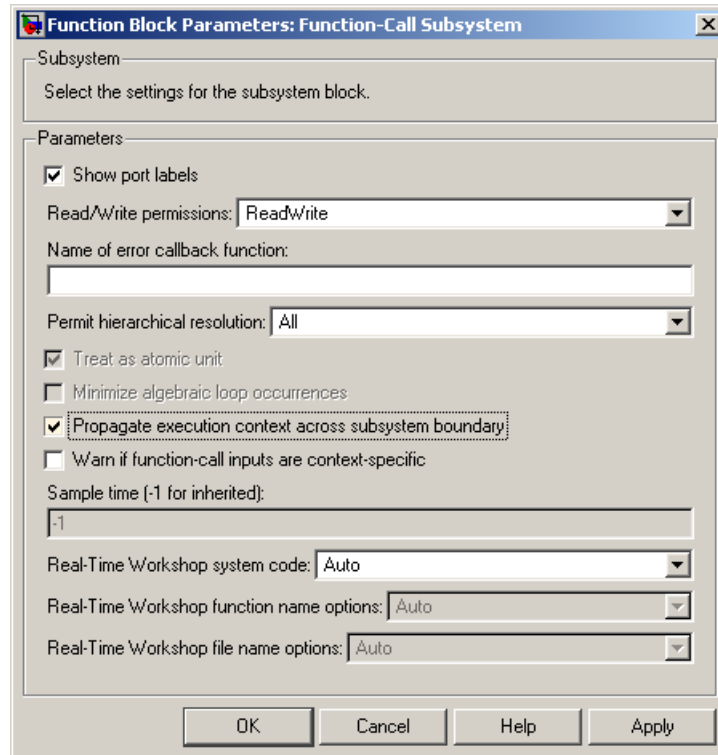
For an example, see “Example: Detecting Changes in the Last-Updated Signal” on page 3-10.

Setting Up Function-Call Subsystems in SimEvents Models

To use a Function-Call Subsystem block in a model that uses event-based blocks or event-based signals, follow the procedure below.

Note The selection of the **Propagate execution context across subsystem boundary** option is particularly important for ensuring that the subsystem executes at the correct times.

- 1** Drag the Function-Call Subsystem block from the Simulink Ports & Subsystems library into your model. Initially, it shows one signal input port **In1**, one output signal port **Out1**, and a control port **function()**. Note that these are signal ports, not entity ports, because the subsystem is designed to process signals rather than entities.
- 2** Select the subsystem block and choose **Edit > SubSystem Parameters** from the model window's menu bar.
- 3** In the dialog box that opens, select **Propagate execution context across subsystem boundary** and click **OK**.



- 4 In the model window, double-click the Function-Call Subsystem block to open the subsystem it represents. Initially, the subsystem contains an Inport block connected to an Outport block. Note that these are from the Simulink Ports & Subsystems library, and are not the same as the Discrete Event Inport and Discrete Event Outport blocks in the SimEvents Ports and Subsystems library. The subsystem also contains a block labeled “function.”
- 5 To change the number of inputs or outputs to the subsystem, change the number of Inport and Outport blocks in the subsystem window. You can copy, paste, and delete these blocks.
- 6 Drag other blocks into the subsystem window as appropriate to build the subsystem. This step is similar to the process of building the top level of hierarchy in your model, except that only certain types of blocks are

suitable for use inside the subsystem. See “Blocks Inside Discrete Event Subsystems” on page 8-10 for details.

Plotting Data

Choosing and Configuring Plotting Blocks (p. 9-2)

Plotting Window Features (p. 9-7)

Using Plots for Troubleshooting (p. 9-8)

Comparison with Time-Based Plotting Tools (p. 9-15)

Available types of plots and how to access them

Customizing plots

Visualizing a simulation to understand and debug it

Notes about plots designed for time-based simulations

Choosing and Configuring Plotting Blocks

SimEvents provides scope blocks that help you visualize data from your simulation. When you insert a scope block into your model, it creates a plot as the simulation runs. This section describes capabilities of the scope blocks in these topics:

- “Sources of Data for Plotting” on page 9-2
- “Inserting and Connecting Scope Blocks” on page 9-3
- “Connections Among Points in Plots” on page 9-4
- “Varying Axis Limits Automatically” on page 9-5
- “Examples Using Scope Blocks” on page 9-6

Sources of Data for Plotting

The table below indicates the kinds of data you can plot using various combinations of blocks and parameter values.

Data	Block	Parameter
Scalar signal vs. time	Signal Scope	X value from = Event time
Scalar signal values without regard to time	Signal Scope	X value from = Index
Two scalar signals (X-Y plot)	X-Y Signal Scope	
Attribute vs. time	Attribute Scope	X value from = Event time
Attribute values without regard to time	Attribute Scope	X value from = Index
Two attributes of same entity (X-Y plot)	X-Y Attribute Scope	
Attribute vs. scalar signal	Get Attribute block to assign the attribute value to a signal; followed by X-Y Signal Scope	
Scalar signal vs. attribute		

Data	Block	Parameter
Number of entity arrivals per time instant	Instantaneous Entity Counting Scope	
Number of events per time instant	Instantaneous Event Counting Scope	

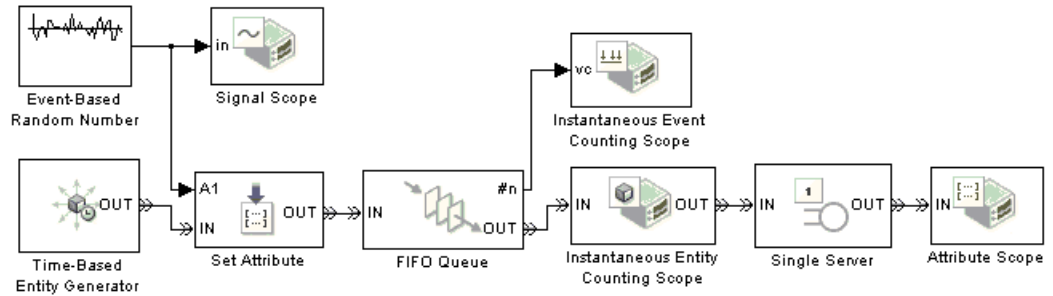
The Signal Scope and X-Y Signal Scope blocks are particularly appropriate for data arising from discrete-event simulations because the plots can include zero-duration values. That is, the plots can include multiple values of the signal at a given time. By contrast, the Scope block in the Simulink Sources library plots at most one value of the signal at each time.

Inserting and Connecting Scope Blocks

The scope blocks reside in the SimEvents Sinks library. The table below indicates the input ports on each scope block.

Block	Input Port(s)	Port Description
Signal Scope	One signal input port	Signal representing the data to plot
X-Y Signal Scope	Two signal input ports	Signals representing the data to plot
Attribute Scope	One entity input port	Entities containing the attribute value to plot
X-Y Attribute Scope	One entity input port	Entities containing the attribute values to plot
Instantaneous Entity Counting Scope	One entity input port	Entities whose arrivals the block counts
Instantaneous Event Counting Scope	One signal input port	Signal whose signal-based events or function calls the block counts

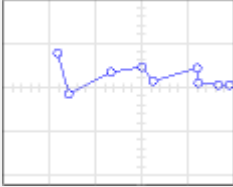
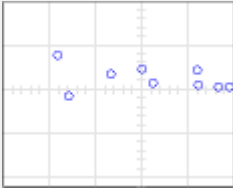
The figure below shows some typical arrangements of scope blocks in a model. Notice that the blocks that have entity input ports can have optional entity output ports, and that signal lines can branch whereas entity connection lines cannot.



Connections Among Points in Plots

You can configure certain scope blocks to determine if and how it connects the points that it plots. The table below indicates the options.

Connection Characteristics	Setting	Sample Plot
Stairstep across, then up or down. Also known as a zero-order hold.	Plot type = Stair in the block's dialog box	
Vertical line from horizontal axis to point. No connection with previous or next plotted point. Also known as a stem plot.	Plot type = Stem in the block's dialog box	

Connection Characteristics	Setting	Sample Plot
Single line segment from point to point. Also known as a first-order hold.	Plot type = Continuous in the block's dialog box	
No connection with other points or with axis. Also known as a scatter plot.	Style > Line > None in the plot window	

Note If a signal does not have an initial condition at $T=0$, then the plotting window does not show a point at $T=0$ and does not connect the first plotted point to the $T=0$ edge of the plot.

Varying Axis Limits Automatically

Using parameters on the **Axes** tab of the dialog box scope blocks, you set the initial limits for the horizontal and vertical axes of the plot. Also, the **If X value is beyond limit** and **If Y value is beyond limit** parameters let you choose how the block responds if a point does not fit within the current axis limits. Choices are in the table below.

Option	Description
Stretch axis limits	Maintain one limit while doubling the size of the displayed interval
Keep axis limits unchanged	Maintain both limits, which means that points outside the limits do not appear
Shift axis limits	Maintain the size of the displayed interval while changing both limits

Examples Using Scope Blocks

The following examples use scope blocks to create different kinds of plots:

Example	Description
“Plotting the Pending-Entity Signal” and “Observations from Plots” in the Getting Started documentation	Stairstep and continuous plots of statistical signals
“Example: Round-Robin Approach to Choosing Inputs” in the Getting Started documentation	Stem plot of data from an attribute
“Example: Using Servers in Shifts” on page 6-11	Unconnected plot of a signal using dots
“Example: Setting Attributes” on page 1-15	Stairstep plots of data from attributes using Attribute Scope blocks as sinks
“Example: Synchronizing Service Start Times with the Clock” on page 7-6	Stem plots that count entities using Instantaneous Entity Counting Scope blocks with entity output ports
X-Y Signal Scope reference page	Continuous plot of two signals
X-Y Attribute Scope reference page	Unconnected plot of two attributes using x’s as plotting markers

Plotting Window Features

After a scope block opens its plotting window, you can modify several aspects of the plot by using the plotting window's menu and toolbar:

- **Autoscale** resizes the axes to fit the range of the data and replaces the values of the upper and lower limit parameters on the **Axes** tab of the block's dialog box.
- The Zoom In and Zoom Out toolbar buttons change the axes as described in the MATLAB documentation about zooming in 2-D views.
- The **Style** menu lets you change the line type, marker type, and color of the plot. Your changes are saved when you save the model.
- **Save position** updates the **Position** parameter on the **Figure** tab of the block's dialog box to reflect the window's current position and size.

Note When a menu option duplicates the behavior of a parameter in the block's dialog box, selecting the menu option *replaces* the corresponding parameter value in the dialog. You can still edit the parameter values in the dialog manually. An example of this is the **Show grid** menu option and dialog box parameter.

The Save Figure toolbar button lets you save the current state of the plot to a FIG-file. You can reload the file in a different MATLAB session. Reloading the file creates a plot that is *not* associated with the original scope block and that does not offer the same menu and toolbar options as in the original plotting window.

Using Plots for Troubleshooting

Typical ways to use plotting blocks in the SimEvents Sinks library to troubleshoot problems are described in the table below.

Technique	Example
Plot a block's #d output signal to check when an entity departs from the block.	"Example: Plotting Entity Departures to Verify Timing" on page 9-8
Plot statistical output signals such as pe or last , if applicable, to check whether operations such as service completion or routing are occurring as you expect.	"Example: Using Servers in Shifts" on page 6-11 and "Timing for Choosing a Port Using a Sequence" on page 8-4
Plot input signals such as port selection, service time, or intergeneration time, and compare the values with observations of how the corresponding blocks use those signals.	"Example: Race Conditions at a Switch" on page 2-23
Plot the output of a Read Timer block to check how long entities spend in a region of the model.	"Example: M/M/5 Queuing System" on page 4-12
Use the Instantaneous Entity Counting Scope or Instantaneous Event Counting Scope block to check whether events you expect to be simultaneous are, in fact, simultaneous.	"Example: Plotting Event Counts to Find Roundoff Error" on page 9-12

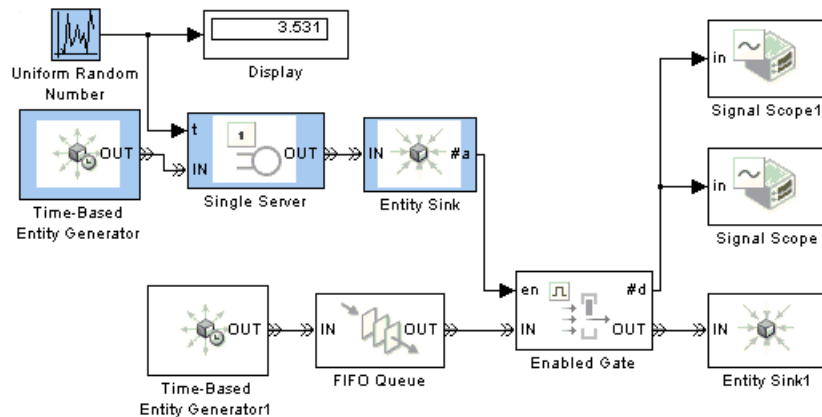
Example: Plotting Entity Departures to Verify Timing

The example below compares two methods of opening a gate at a random time and leaving the gate open for the rest of the simulation. The Signal Scope block lets you see when the gate opens, to check whether the timing is what you expected. One method exhibits latency in the gate's opening. For a



nonvisual way to determine when entities depart from the gate, see “Viewing Entity Locations” on page 11-8.

Model Exhibiting Correct Timing

The model below views the random opening of the gate as a discrete event, and models it via an entity departure from a server at a random time. The Time-Based Entity Generator block generates exactly one entity, at $T=0$. The Single Server block delays the entity for the amount of time indicated by the Uniform Random Number block, 3.531 seconds in this case. At $T=3.531$, the entity arrives at the Entity Sink block. This time is exactly when the sink block's **#a** signal changes from 0 to 1, which in turn causes the gate to open.

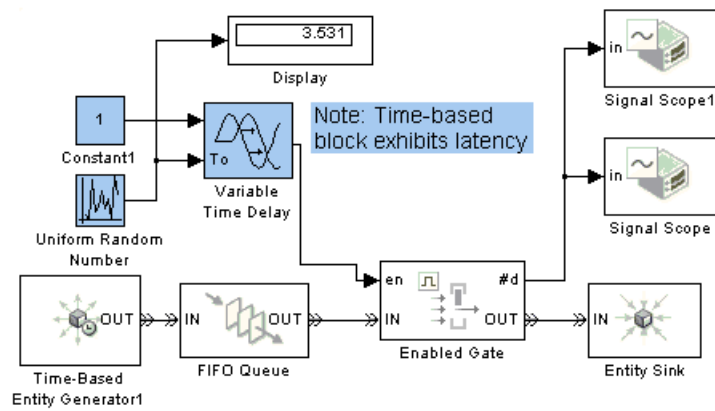


By using the zoom feature of the scope, you can compare the time at which entities depart from the Enabled Gate block with the random time shown on the Display block in the model.

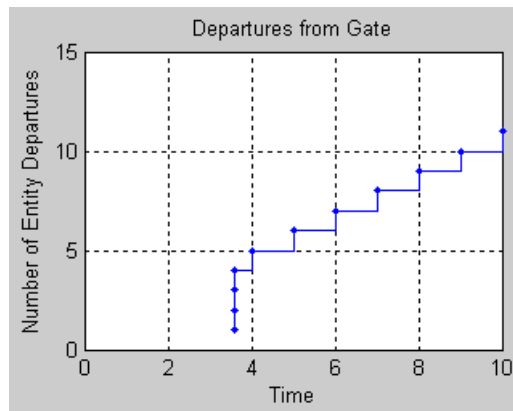



Model Exhibiting Latency

The model below uses the Variable Time Delay block to create a step signal that is intended to change from 0 to 1 at a random time. However, because the Variable Time Delay block is time-based, it updates its output signal at times dictated by the time-based simulation clock. The step signal does not actually change from 0 to 1 until the next sample time hit after the time indicated by the random number. This is the intentional documented behavior of this time-based block.



By using the zoom feature of the scope, you can see that entities depart from the Enabled Gate block later than the random time shown on the Display block in the model.







Example: Plotting Event Counts to Find Roundoff Error

The example below suggests how to use the Instantaneous Event Counting Scope block to determine whether events you want to be simultaneous are truly simultaneous.

Suppose you want two entity generators with periods of 1 and $1/3$ to create simultaneous entity departures every second, so that event priorities determine which entity arrives at the queue first. By counting events at each value of time and checking when the count is 2, you can see whether two entity generation events are truly simultaneous.

Model Exhibiting Simultaneous Events


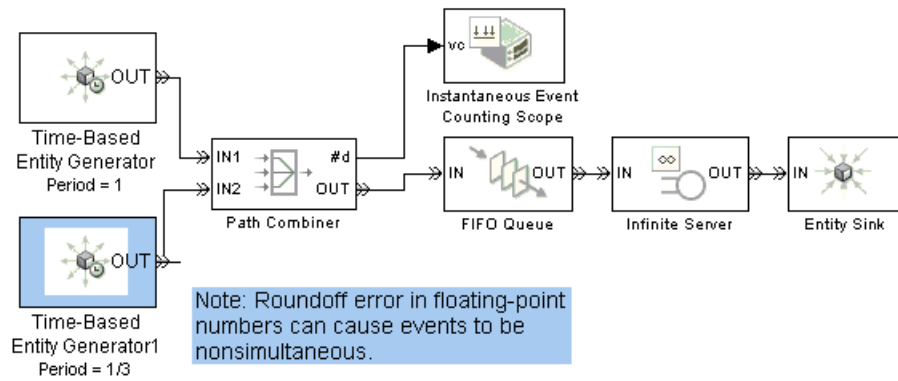
The model below uses two Event-Based Entity Generator blocks receiving the same input signal. You can see from the plot that simultaneous events occur every second, as desired.



Although this example uses the Instantaneous Event Counting Scope to plot a **#d** signal, you can alternatively use the Instantaneous Entity Counting Scope to count entities departing from the Path Combiner block.

Model Exhibiting Nonsimultaneous Events

Using the **Period** parameter of the Time-Based Entity Generator block shows how roundoff error in the computer's representation of the floating-point number $1/3$ can cause events to be close in time, but not simultaneous. In general, comparing floating-point numbers within a tolerance is a better practice than testing for equality. Computers' use of floating-point arithmetic involves a finite set of numbers with finite precision.



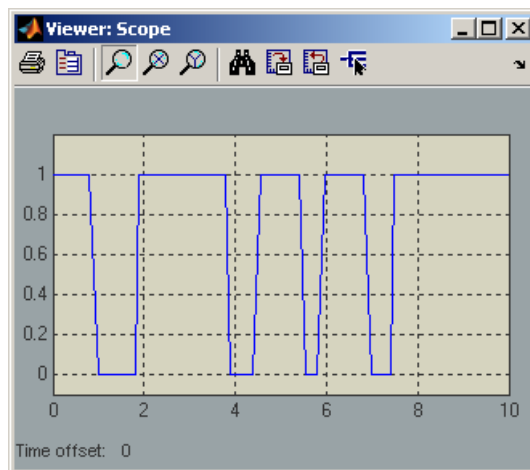
Comparison with Time-Based Plotting Tools

Simulink offers plotting tools designed for signals in time-based simulations. To learn more, see these parts of the Simulink documentation:

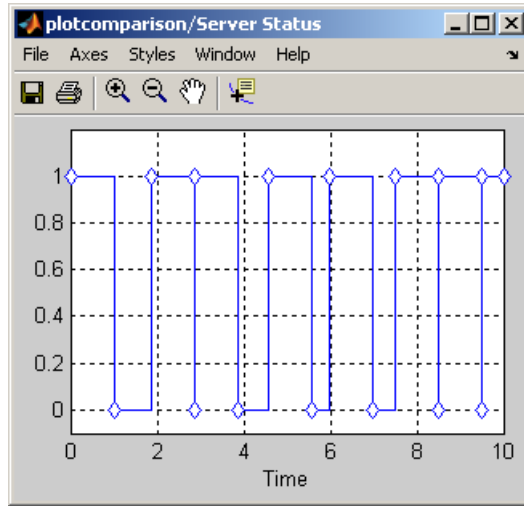
- “The Signal & Scope Manager”
- Scope block reference page
- XY Graph block reference page
- `simplot` function reference page

If you use these time-based tools instead of event-based tools, note that they typically join each pair of successive points with a single line segment. To create a stairstep plot, use the blocks in the SimEvents Sinks library. Also, time-based plotting tools plot at most one value at each time, whereas blocks in the SimEvents Sinks library can include zero-duration values.

Compare the two figures below, which depict the same data, when you consider which plotting tools are more appropriate in your event-based simulations.



Time-Based Plot of Number of Entities in a Server



Discrete-Event Plot of Number of Entities in a Server

Using Statistics

Role of Statistics in Discrete-Event Simulation (p. 10-2)	How well-designed simulation statistics can help you learn or make decisions about your system
Accessing Statistics from SimEvents Blocks (p. 10-4)	Computing statistics while the simulation runs
Deriving Custom Statistics (p. 10-7)	Creating your own statistics by manipulating outputs from SimEvents blocks
Using Timers (p. 10-18)	Recording how long entities take to reach a block
Running a Series of Simulations (p. 10-25)	Running multiple simulations and regulating the simulation length

Role of Statistics in Discrete-Event Simulation

Most SimEvents blocks are capable of producing one or more statistical output signals. You can use these signals to gather data from the simulation or to influence the dynamics of the simulation. This section gives an overview of both purposes, in these topics:

- “Statistics for Data Analysis” on page 10-2
- “Statistics for Run-Time Control” on page 10-3

The rest of this chapter illustrates some modeling and analysis techniques that you can use with SimEvents. However, a detailed treatment of statistical analysis is well beyond the scope of this User’s Guide; see the works listed in “Selected Bibliography” for more information.

Statistics for Data Analysis

Often, the purpose of creating a discrete-event simulation is to improve understanding of the underlying system being modeled or to use simulation results to help make decisions about the underlying system. Numerical results gathered during simulation can be important tools.

For example, if you simulate the operation and maintenance of equipment on an assembly line, then you might use the computed production and defect rates to help decide whether to change your maintenance schedule. As another example, if you simulate a communication bus under varying bus loads, then you might use computed average delays in high- or low-priority messages to help determine whether a proposed architecture is viable.

Just as you decide how to design a simulation model that adequately describes the underlying system, you decide how to design the statistical measures that you will use to learn about the system. Some questions to consider are

- Which statistics are meaningful for your investigation or decision? For example, if you are trying to maximize efficiency, then what is an appropriate measure of efficiency in your system? As another example, does a mean give the best performance measure for your system, or is it also worthwhile to consider the proportion of samples in a given interval?

- How can you compute the desired statistics? For example, do you need to ignore any transient effects, does the choice of initial conditions matter, and what stopping criteria are appropriate for the simulation?
- To ensure sufficient confidence in the result, how many replications of the simulation do you need? One simulation run, no matter how long, is still a single sample and probably inadequate for valid statistical analysis.

For details concerning statistical analysis and variance reduction techniques, see the works [7], [4], [1], and [2] listed in “Selected Bibliography” in the “Getting Started with SimEvents” book.

Statistics for Run-Time Control

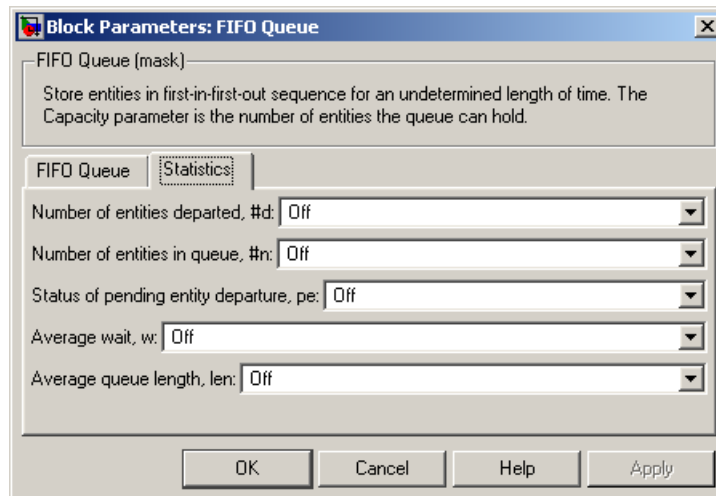
Some systems rely on statistics to influence the dynamics. For example, a queuing system with discouraged arrivals has a feedback loop that adjusts the arrival rate throughout the simulation based on statistics reported by the queue and server, as illustrated in the `discouragearrival` demo model.

When you create simulations that use statistical signals to control the dynamics, you must have access to the current values of the statistics at key times throughout the simulation, not just at the end of the simulation. Some questions to consider while designing your model are

- Which statistics are meaningful, and how should they influence the dynamics of the system?
- How can you compute the desired statistics at the right times during the simulation? It is important to understand when SimEvents blocks update each of their statistical outputs and when other blocks can access the updated values. This topic is discussed in Chapter 3, “Working with Signals”.
- Do you need to account for initial conditions or extreme values in any special way? For example, if your control logic involves the number of entities in a queue, then be sure that the logic is sound even when the queue is empty or full.
- Will small perturbations result in large changes in the system’s behavior? When using statistics to control the model, you might want to monitor those statistics or other statistics to check whether the system is undesirably sensitive to perturbations.

Accessing Statistics from SimEvents Blocks

Most SimEvents blocks can produce one or more statistical outputs. To see which statistics are available, open the block's dialog box. In most cases, the list of available statistics appears on the **Statistics** tab of the dialog box. For example, the figure below shows the **Statistics** tab of the FIFO Queue block's dialog box.



In cases where the dialog box has no **Statistics** tab, such as the Entity Sink block, the dialog box has so few parameters that the statistical options are straightforward to locate.


To use one or more statistics, see the options described in these sections:

- “Accessing Statistics Throughout the Simulation” on page 10-4
- “Accessing Statistics When Stopping or Pausing Simulation” on page 10-6

Accessing Statistics Throughout the Simulation

To configure a block so that it outputs an available statistic throughout the simulation, set the corresponding parameter in the dialog box to On. After you apply the change, the block has a signal output port corresponding to that

statistic. The figure below shows the dialog box and icon for the FIFO Queue block after the average wait statistic is enabled.



You can connect this signal output port to the signal input port of any block. For example, you can connect the statistical signal output port to

- A Signal Scope or X-Y Signal Scope block, to create a plot using the statistic.
- A Display block, which shows the statistic on the block icon throughout the simulation.
- A Discrete Event Signal to Workspace block, which writes the entire data set to the MATLAB workspace when the simulation stops or pauses. To learn more, see “Sending Data to the MATLAB Workspace” on page 3-23.
- A custom subsystem or computational block, for further data processing. See “Deriving Custom Statistics” on page 10-7 for some specific examples.

For more information about when SimEvents blocks update their statistical signals and when other blocks react to the updated values, see Chapter 3, “Working with Signals”.

Accessing Statistics When Stopping or Pausing Simulation

In some cases, you can configure a block to report a statistic only when the simulation stops or pauses. To do this, find the dialog box parameter that corresponds to the given statistic and set it to `Upon stop or pause`, if available. After you apply the change, the block has a signal output port corresponding to that statistic.

Because the statistic is not reported throughout the simulation, not all uses of the signal are appropriate. One appropriate use is in conjunction with a Discrete Event Signal to Workspace block with **Save format** set to `Structure With Time`.

Deriving Custom Statistics

You can use the built-in statistical signals from SimEvents blocks to derive more specialized or complex statistics that are meaningful in your model. One approach is to compute statistics during the simulation in discrete event subsystems. Inside the subsystems, you can implement your computations using

- A graphical block-diagram approach. See the Math Operations library in Simulink and the Statistics library in the Signal Processing Blockset. For examples using Simulink blocks, see
 - “Example: Using Event-Based Timing for a Statistical Computation” on page 8-19, which computes the length of time during the simulation that a queue length is above a threshold
 - “Example: Fraction of Dropped Messages” on page 10-8
 - “Example: Detecting Changes in the Last-Updated Signal” on page 3-10, which computes the ratio of an instantaneous queue length to its long-term average
 - The function-call subsystem within the DVS Optimizer subsystem in the Dynamic Voltage Scaling Using Online Gradient Estimation demo
- A nongraphical coded approach. See the blocks in the User-Defined Functions library in Simulink. For examples using the Embedded MATLAB Function block, see
 - “Example: Computing a Time Average of a Signal” on page 10-9
 - “Example: Resetting an Average Periodically” on page 10-12

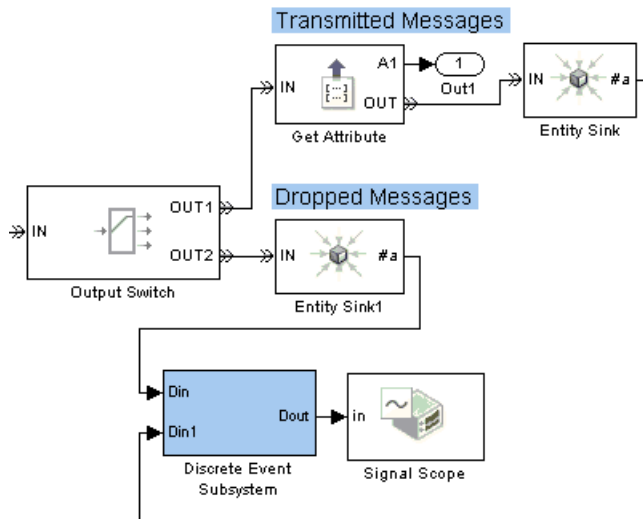
Note If you put an Embedded MATLAB Function block in a Discrete Event Subsystem block, use the Ports and Data Manager instead of Model Explorer to view or change properties such as the size or source of an argument. Model Explorer does not show the contents of Discrete Event Subsystem blocks.

Alternatively, you can use the Discrete Event Signal to Workspace block to log data to the MATLAB workspace and compute statistics after the simulation is complete. For an example of post-simulation analysis, see


- “Example: Observing Service Completions” on page 2-32, which computes the times at which a signal value either increases or remains the same
- “Example: Running a Simulation and Varying a Parameter” on page 10-29

Example: Fraction of Dropped Messages

The example below shows how to compute a ratio of event-based signals in a discrete event subsystem. The Output Switch block either transmits or drops the message corresponding to each entity. The goal is to compute the fraction of dropped messages, that is, the fraction of entities that depart via **OUT2** as opposed to **OUT1** of the Output Switch block.




Upper-Level System



Subsystem Contents

Two Entity Sink blocks produce **#a** signals that indicate how many messages the communication link transmits or drops, respectively. The discrete event subsystem divides the number of dropped messages by the sum of the two **#a** signals. Because the discrete event subsystem performs the division only when one of the **#a** signals increases, no division-by-zero instances occur.

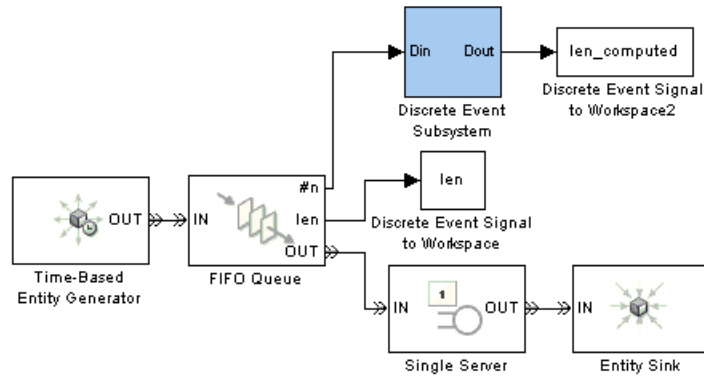


Example: Computing a Time Average of a Signal

This example illustrates how to compute a time average of a signal using the Embedded MATLAB Function block, and especially how to make the block retain data between calls to the function.

The model below implements a simple queuing system in which the FIFO Queue produces the output signals

- **#n**, the instantaneous length of the queue.
- **len**, the time average of the queue length. This is the time average of **#n**.



Top-Level Model

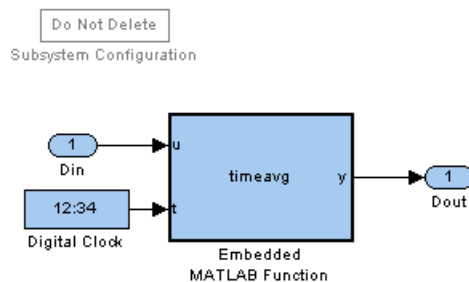
The discrete event subsystem uses **#n** to compute the time average. In this case, the time average should equal **len**. You can use a similar subsystem in your own models to compute time averages of other signals.

Computation of the Time Average

In the example, the discrete event subsystem performs computations each time a customer arrives at or departs from the queue. Within the subsystem, an Embedded MATLAB Function block keeps a running weighted sum of the **#n** values that form the input, where the weighting is based on the length of time over which the signal assumes each value.

The block uses persistent variables for quantities whose values it must retain from one invocation to the next, namely, the running weighted sum and the previous values of the inputs.

Below are the subsystem contents and the function that the Embedded MATLAB Function block represents.



Subsystem Contents

```
function y = timeavg(u,t)
%TIMEAVG Compute time average of input signal U
% Y = TIMEAVG(U,T) computes the time average of U,
% where T is the current simulation time.

% Declare variables that must retain values between iterations.
persistent running_weighted_sum last_u last_t;

% Initialize persistent variables in the first iteration.
if isempty(last_t)
    running_weighted_sum = 0;
    last_u = 0;
    last_t = 0;
end

% Update the persistent variables.
running_weighted_sum = running_weighted_sum + last_u*(t-last_t);
last_u = u;
last_t = t;

% Compute the outputs.
if t > 0
    y = running_weighted_sum/t;
end
```

```
else
    y = 0;
end
```

Verifying the Result

After running the simulation, you can verify that the computed time average of **#n** is equal to **len**.

```
isequal([len.time, len.signals.values],...
        [len_computed.time, len_computed.signals.values])
```

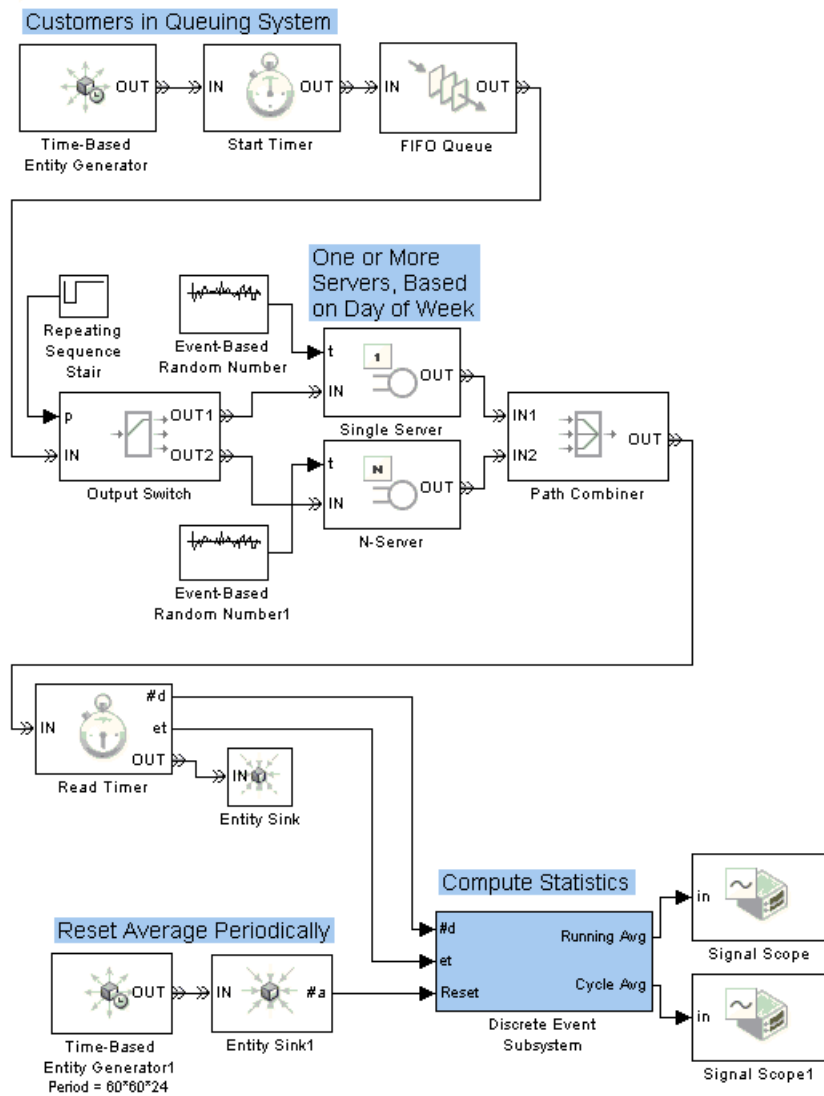
The output indicates that the comparison is true.

```
ans =
     1
```

Example: Resetting an Average Periodically

This example illustrates how to compute a sample mean over each of a series of contiguous time intervals of fixed length, rather than the mean over the entire duration of the simulation. The example simulates a queuing system for four weeks' worth of simulation time, where customers have access to one server during the first two days of the week and five servers on the other days of the week. The average of customers' waiting times over a daily cycle depends on how many servers are operational that day. However, you might expect the averages taken over weekly cycles to be stable from one week to the next.

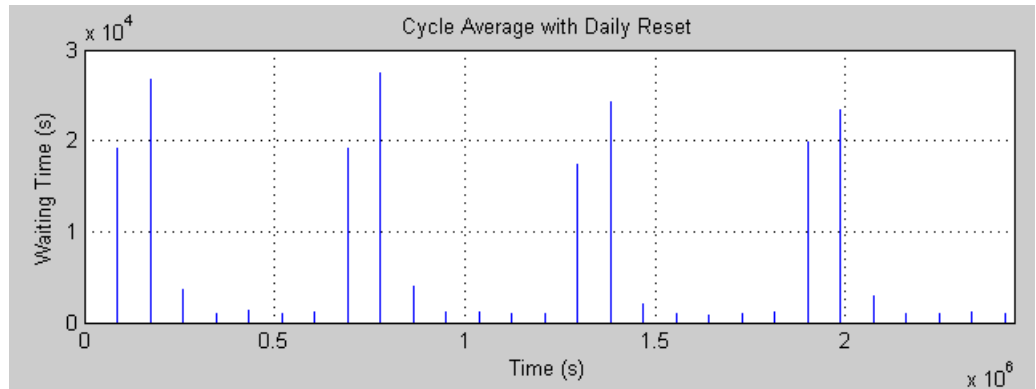

The model below uses a time-based Repeating Sequence Stair block to determine whether entities advance to a Single Server or N-Server block, thus creating variations in the number of operational servers. The Start Timer and Read Timer blocks compute each entity's waiting time in the queuing system. A computational discrete event subsystem processes the waiting time by computing a running sample mean over a daily or weekly cycle, as well as the final sample mean for each cycle. Details about this subsystem are in "Computation of the Cycle Average" on page 10-15.



Top-Level Model



Performance of Daily Averages

When considering daily cycles, you can see that the cycle averages do not stabilize at a single value.



Performance of Weekly Averages

When considering weekly cycles, you can see less variation in the cycle averages because each cycle contains the same pattern of changing service levels. To compute the cycle average over a weekly cycle, change the **Period** parameter in the Time-Based Entity Generator1 block at the bottom of the model to $60 \times 60 \times 24 \times 7$, which is the number of seconds in a week.



Computation of the Cycle Average

In the example, the discrete event subsystem performs computations each time a customer departs from the queuing system and at each boundary of a daily or weekly cycle. Within the subsystem, an Embedded MATLAB Function block counts the number of customers and the total waiting time among all customers so far. The block resets these quantities to zero at each boundary of a cycle.

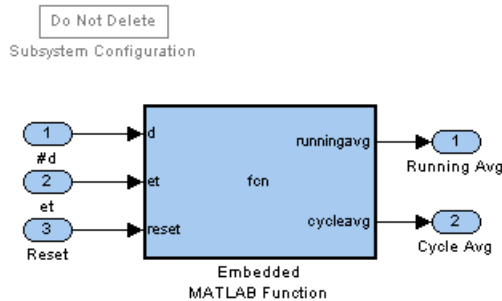
The block uses persistent variables for quantities whose values it must retain from one invocation to the next. The number of customers and total waiting time are important to retain for the computation of an average over time rather than an instantaneous statistic. Previous values of inputs are

important to retain for comparison, so the function can determine whether it needs to update or reset its statistics.

The outputs of the Embedded MATLAB Function block are

- runningavg, the running sample mean of the input waiting times
- cycleavg, a signal that, at reset times, represents the sample mean over the cycle that just ended

Below are the subsystem contents and the function that the Embedded MATLAB Function block represents.



Subsystem Contents

```
function [runningavg, cycleavg] = fcn(d,et,reset)
%FCN    Compute average of ET, resetting at each update of RESET
% [RUNNINGAVG,CYCLEAVG] = FCN(D,ET,RESET) computes the average
% of ET over contiguous intervals. D is the number of samples
% of ET since the start of the simulation. Increases in
% RESET indicate when to reset the average.
%
% Assume this function is invoked when either D or RESET
% (but not both) increases. This is consistent with the
% behavior of the Discrete Event Subsystem block that contains
% this block in this example.
%
% RUNNINGAVG is the average since the start of the interval.
%
```

```
% At reset times, CYCLEAVG is the average over the interval
% that just ended; at other times, CYCLEAVG is 0.

% Declare variables that must retain values between iterations.
persistent total customers last_reset last_d;

% Initialize outputs.
cycleavg = 0;
runningavg = 0;

% Initialize persistent variables in the first iteration.
if isempty(total)
    total = 0;
    customers = 0;
    last_reset = 0;
    last_d = 0;
end

% If RESET increased, compute outputs and reset the statistics.
if (reset > last_reset)
    cycleavg = total / customers; % Average over last interval.
    runningavg = cycleavg; % Maintain running average.
    total = 0; % Reset total.
    customers = 0; % Reset number of customers.
    last_reset = reset;
end

% If D increased, then update the statistics.
if (d > last_d)
    total = total + et;
    customers = customers + 1;
    last_d = d;
    runningavg = total / customers;
end
```

Using Timers

Suppose you want to determine how long each entity takes to advance from one block to another, or how much time each entity spends in a particular region of your model. To compute these durations, you can attach a timer to each entity that reaches a particular spot in the model. Then you can

- Start the timer. The block that attaches the timer also starts it.
- Read the value of the timer whenever the entity reaches a spot in the model that you designate.
- Restart the timer, if desired, whenever the entity reaches a spot in the model that you designate.

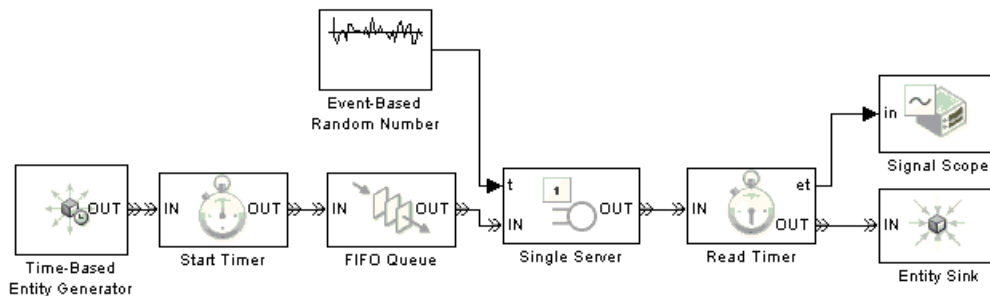
The following sections describe how to arrange the Start Timer and Read Timer blocks to accomplish several common timing goals:

- “Basic Example Using Timer Blocks” on page 10-18
- “Basic Procedure for Using Timer Blocks” on page 10-19
- “Timing Multiple Entity Paths with One Timer” on page 10-20
- “Restarting a Timer” on page 10-21
- “Timing Multiple Processes Independently” on page 10-23

Note Timers measure durations, or relative time. By contrast, clocks measure absolute time. For details about implementing clocks, see the descriptions of the Clock and Digital Clock blocks in the Simulink documentation.

Basic Example Using Timer Blocks

A typical block diagram for determining how long each entity spends in a region of the model is in the figure below. The Start Timer and Read Timer blocks jointly perform the timing computation.



The model shown above measures the time each entity takes between arriving at the queue and departing from the server. The Start Timer block attaches, or associates, a timer to each entity that arrives at the block. Each entity has its own timer. Each entity's timer starts timing when the entity departs from the Start Timer, or equivalently, when the entity arrives at the FIFO Queue block. Upon departing from the Single Server block, each entity arrives at a Read Timer block. The Read Timer block reads data from the arriving entity's timer and produces a signal at the **et** port whose value is the instantaneous elapsed time for that entity. For example, if the arriving entity spent 12 seconds in the queue-server pair, then the **et** signal assumes the value 12.

Basic Procedure for Using Timer Blocks

A typical procedure for setting up timer blocks is as follows:

- 1 Locate the spots in the model where you want to begin timing and to access the value of the timer.
- 2 Insert a Start Timer block in the model at the spot where you want to begin timing.
- 3 In the Start Timer block's dialog box, enter a name for the timer in the **Timer tag** field. This timer tag distinguishes the timer from other independent timers that might already be associated with the same entity.

When an entity arrives at the Start Timer block, the block attaches a named timer to the entity and begins timing.

- 4 Insert a Read Timer block in the model at the spot where you want to access the value of the timer.

- 5 In the Read Timer block's dialog box, enter the same **Timer tag** value that you used in the corresponding Start Timer block.

When an entity having a timer with the specified timer tag arrives at the block, the block reads the time from that entity's timer. Using the **Statistics** tab of the Read Timer block's dialog box, you can configure the block to report this instantaneous time or the average of such values among all entities that have arrived at the block.

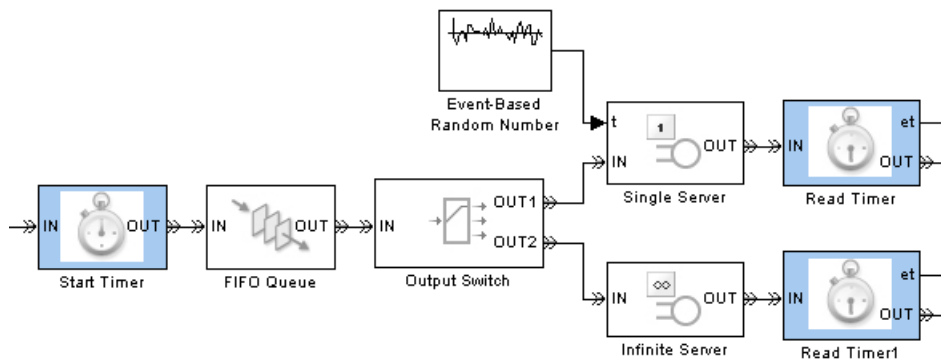
If you need multiple independent timers per entity (for example, to time an entity's progress through two possibly overlapping regions of the model), then follow the procedure above for each of the independent timers. For more information, see "Timing Multiple Processes Independently" on page 10-23.

Timing Multiple Entity Paths with One Timer

If your model includes routing blocks, then different entities might use different entity paths. To have a timer cover multiple entity paths, you can include multiple Start Timer or multiple Read Timer blocks in a model, using the same **Timer tag** parameter value in all timer blocks.

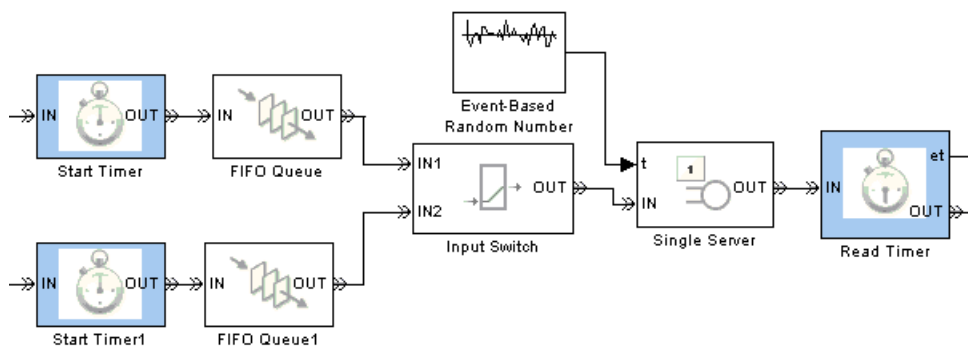
Output Switch Example

In the figure below, each entity advances along one of two different entity paths via the Output Switch block. The timer continues timing, regardless of the selected path. Finally, each entity advances to one of the two Read Timer blocks, which reads the value of the timer.



Input Switch Example

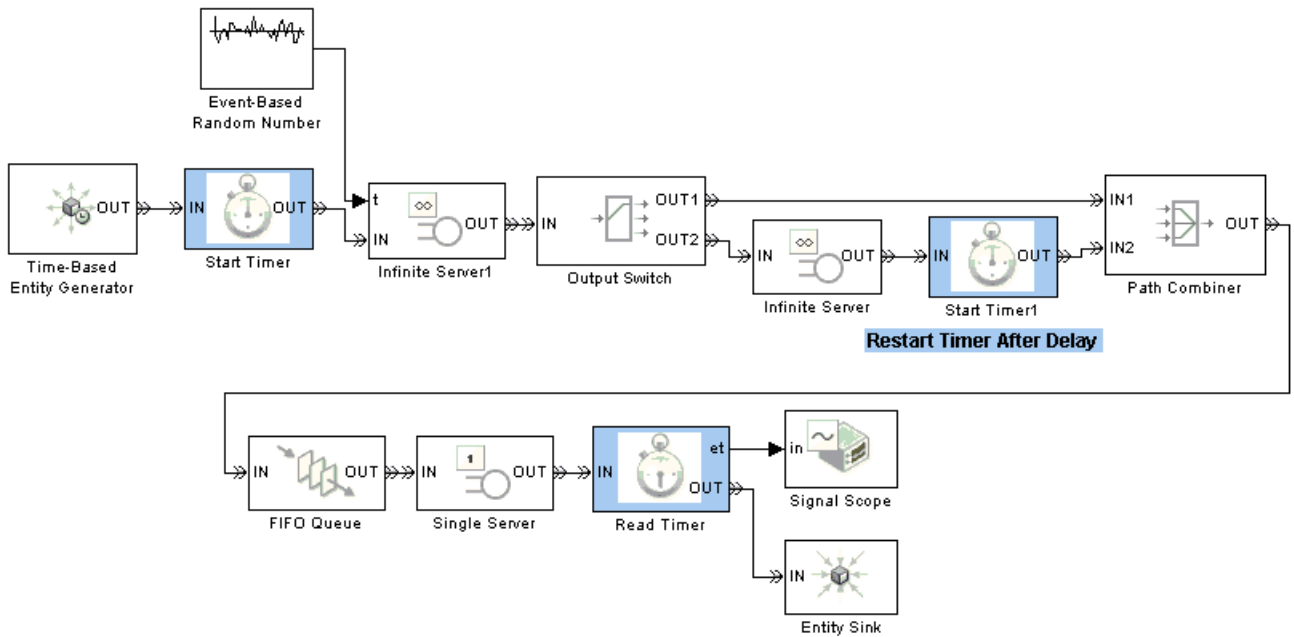
In the figure below, entities wait in two different queues before advancing to a single server. The timer blocks measure the time each entity spends in its respective queue-server pair. Two Start Timer blocks, configured with the same **Timer tag** parameter value, ensure that all entities possess a timer regardless of the path they take before reaching the server.



Restarting a Timer

You can restart an entity's timer, that is, reset its value to zero, whenever the entity reaches a spot in the model that you designate. To do this, insert a Start Timer block in the model where you want to restart the timer. Then set the block's **If timer has already started** parameter to Restart.

The figure below shows an example of restarting a timer.



All timer blocks share the same **Timer tag** parameter value. All entities that arrive at the first **Start Timer** block acquire a timer, which starts timing immediately. All entities incur an initial delay, modeled by an **Infinite Server** block. When entities reach the **Output Switch** block, they depart via one of the two entity output ports and receive different treatment:

- Entities that depart via the **OUT1** port advance to the queue with no further delay, and the timer continues timing.
- Entities that depart via the **OUT2** port incur an additional delay, modeled by another **Infinite Server** block. After the delay, the timer restarts from zero and the entity advances to the queue.

When entities finally advance from the server to the Read Timer block, the elapsed time is one of these quantities:

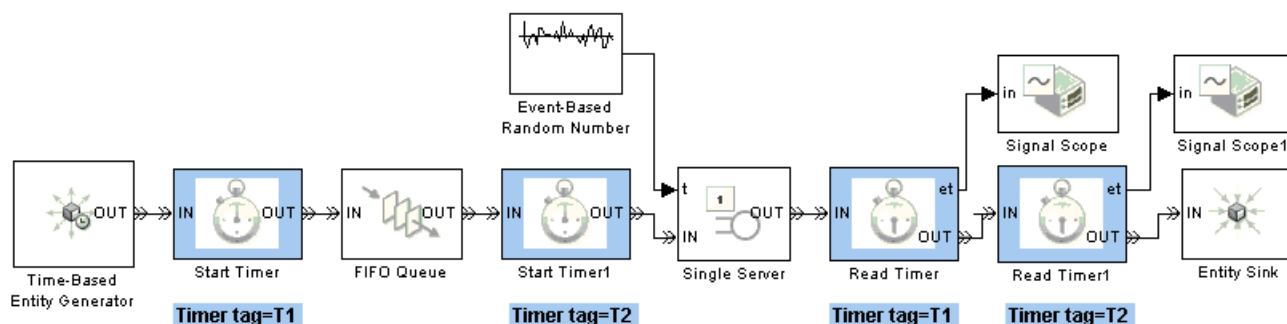
- The initial delay plus the time in the queue plus the time in the server, for entities that departed from the Output Switch block's **OUT1** port
- The time in the queue plus the time in the server, for entities that departed from the Output Switch block's **OUT2** port

Timing Multiple Processes Independently

You can measure multiple independent durations using the Start Timer and Read Timer blocks. To do this, create a unique **Timer tag** parameter for each independent timer. For clarity in your model, consider adding an annotation or changing the block names to reflect the **Timer tag** parameter in each timer block.

The figure below shows how to measure these quantities independently:

- The time each entity spends in the queue-server pair, using a timer with tag T1
- The time each entity spends in the server, using a timer with tag T2



The annotations beneath the blocks in the figure indicate the values of the **Timer tag** parameters. Notice that the T1 timer starts at the time when entities arrive at the queue, while the T2 timer starts at the time when entities depart from the queue (equivalently, at the time when entities arrive

at the server). The two Read Timer blocks read both timers when entities depart from the server. The sequence of the Read Timer blocks relative to each other is not relevant in this example because no time elapses while an entity is in a Read Timer block.

Running a Series of Simulations

This section describes some techniques that can help you gather statistical data from a series of simulations of your model. The topics are as follows:

- “Creating Independent Replications” on page 10-25
- “Running Simulations from MATLAB” on page 10-27
- “Regulating the Simulation Length” on page 10-32

Creating Independent Replications

When you run a simulation multiple times to gather statistics, you can create independent replications by using a different stream of random numbers in each replication. To vary the stream of random numbers, vary the *initial seed* on which the stream of random numbers is based. **Initial seed** is a parameter in the following blocks:

- Time-Based Entity Generator
- Event-Based Random Number
- Uniform Random Number
- Random Number
- Blocks in the Routing library

Also, if your simulation is configured to randomize the sequence of certain simultaneous events, then the Configuration Parameters dialog box has a parameter called **Seed for event randomization**, which indicates the initial seed for that stream of random numbers.

Choosing Values for Initial Seed

Here are some recommendations for choosing appropriate values for **Initial seed** parameters:

- Choose a large odd number, such as a five-digit odd number.
- To obtain the same stream of random numbers the next time you run the same simulation, set **Initial seed** to an unchanging value.

- To obtain a different stream of random numbers the next time you run the same simulation, either change the value of **Initial seed** or set it to a varying expression such as `ceil(cputime*99999)*2+1`. See the `cputime` function for more details.
- If **Initial seed** parameters appear in multiple places in your model, then choose different values or different expressions for each **Initial seed** parameter.

Setting Values for Initial Seed


When running a simulation interactively, you can simply enter your chosen value for the initial seed in the dialog box for a given block.

When you run a simulation multiple times to gather statistics, you might want to use a MATLAB variable to make it easier to assign a different seed each time you run the simulation. Follow this procedure for each **Initial seed** parameter that you need to set:

- 1 Assign a value to the MATLAB variable. For example, the code below defines a variable called `initseed`.

```
initseed = 68521;
```

- 2 In the **Initial seed** field of the dialog box, enter the name of the MATLAB variable. An example is below.



- 3** Run the simulation using **Simulation > Start** or the `sim` function. During the simulation, the value that the block uses for the initial seed is the one that you assigned to the MATLAB variable.
- 4** Repeat steps 1 and 3 if appropriate, using a different value each time.

If you are using the above procedure within an M-file that runs a series of simulations in a loop, put the variable assignment and the simulation command (`sim`) in the M-file. For an example of this approach, see “Example: Running a Simulation Repeatedly to Gather Results” on page 10-27.

Running Simulations from MATLAB

When you analyze simulation statistics, you typically need to run the simulation many times. One simulation run, no matter how long, is still a single sample and probably inadequate for valid statistical analysis. The `sim` function enables you to run simulations unattended, while the Discrete Event Signal to Workspace block writes signal contents to the MATLAB workspace for subsequent storage or analysis. To learn about this function and block, see “Running a Simulation Programmatically” in the Simulink documentation and the reference page for the Discrete Event Signal to Workspace block. This section provides examples and tips that focus on SimEvents simulations, in these topics:

- “Example: Running a Simulation Repeatedly to Gather Results” on page 10-27
- “Example: Running a Simulation and Varying a Parameter” on page 10-29

Example: Running a Simulation Repeatedly to Gather Results

Suppose you want to run the M/D/1 Queuing System demo model many times to check whether the ensemble average for the waiting time is close to the value predicted by queuing theory. You can do this by modifying the model to make it suitable for statistical analysis, simulating the modified model multiple times, and then analyzing the results.

- 1** Open the model by entering `md1` in the MATLAB Command Window.
- 2** Save the model as `md1_stats.mdl` in either the current directory or a writable directory on the MATLAB path.

- 3 From the SimEvents Sinks library, drag the Discrete Event Signal to Workspace block into the model.
- 4 In the Discrete Event Signal to Workspace block's dialog box,
 - Set **Limit data points to last** to 1 because only the final value of the statistic is relevant for this example.
 - Set **Save format** to Array.
- 5 Create a branch line from the output of the subsystem labeled Waiting Time Evaluation and connect the branch line to the Discrete Event Signal to Workspace block.
- 6 In the Exponential Interarrival Time Distribution block's dialog box, set **Initial seed** to seedvalue. The use of a variable facilitates using a different set of random numbers in each run. However, do not run the simulation yet because seedvalue is not defined.
- 7 Resave the model, which is now suitable for statistical analysis of the waiting time.
- 8 To run the simulation repeatedly and capture the statistic from each run, execute the following code in MATLAB.

```
% Set up.
load_system('md1_stats'); % Load system if not already loaded.
nruns = 100; % Number of simulation runs
w = zeros(nruns,1); % Preallocate space for results
opts = simset('SrcWorkspace','Current','DstWorkspace','Current');
h = waitbar(0,'Running simulations. Please wait...');
seedarray = ceil(rand(nruns,1)*999999)*2+1;

% Main simulation loop
for k = 1:nruns
    waitbar(k/nruns,h); % Update progress indicator.
    seedvalue = seedarray(k);
    sim('md1_stats',[],opts); % Run simulation.
    w(k) = simout(2); % Store empirical value of w.
end

w_theor = simout(1); % Store theoretical value.
```

```
close(h); % Close progress indicator window.
```

- 9 To display some information about the collected statistics, enter the following code in the MATLAB Command Window.

```
format long; % Show more digits in Command Window output.
disp('Theoretical value and ensemble average are:');
disp([w_theor, mean(w)]);
disp('Standard deviation for empirical values is:');
disp(std(w));
disp('Relative error in ensemble average is:');
disp([num2str(100*abs(w_theor - mean(w))/w_theor), '%']);
```

Sample output is below. Your results might vary because the initial seed for the random number generator is itself random.

```
Theoretical value and ensemble average are:
0.333333333333333 0.33163393290150
```

```
Standard deviation for empirical values is:
0.01961357797243
```

```
Relative error in ensemble average is:
0.50982%
```

Example: Running a Simulation and Varying a Parameter

Suppose you want to run the M/D/1 Queuing System demo model with different values of the mean arrival rate. You can do this by modifying the model to make it suitable for statistical analysis with a varying parameter, simulating the modified model multiple times with different values of the parameter, and then analyzing the results. Part of the procedure below is similar to the one in “Example: Running a Simulation Repeatedly to Gather Results” on page 10-27; however, the filename, MATLAB code, and instructions regarding the Constant and Arrival Rate Gain blocks are different.

- 1 Open the model by entering `md1` in the MATLAB Command Window.
- 2 Save the model as `md1_varymean.mdl` in either the current directory or a writable directory on the MATLAB path.

- 3 From the SimEvents Sinks library, drag the Discrete Event Signal to Workspace block into the model.
- 4 In the Discrete Event Signal to Workspace block's dialog box,
 - Set **Limit data points to last** to 1 because only the final value of the statistic is relevant for this example.
 - Set **Save format** to Array.
- 5 Create a branch line from the output of the subsystem labeled Waiting Time Evaluation and connect the branch line to the Discrete Event Signal to Workspace block.
- 6 In the Exponential Interarrival Time Distribution block's dialog box, set **Initial seed** to `ceil(cputime*99999)*2+1`. This causes the simulation to use a different set of random numbers in each run, although the results are not repeatable.
- 7 Remove the Arrival Rate Gain block and close the connection gap between the Constant block (labeled Maximal Arrival Rate, but now signifying the mean arrival rate) and the subsystem labeled Exponential Interarrival Time Distribution.
- 8 Open the dialog box of the Constant block labeled Maximal Arrival Rate and set **Constant value** to `m`, which is a MATLAB variable to be defined later.
- 9 Remove the model's callbacks that are not relevant for this modified version, by executing the following in MATLAB.

```
set_param('md1_varymean', 'PostLoadFcn', '')  
set_param('md1_varymean', 'CloseFcn', '')
```

- 10 Resave the model, which is now suitable for statistical analysis of the waiting time with varying mean arrival rates.
- 11 To run the simulation repeatedly while varying the mean arrival rate, and capture the statistic from each run, execute the following code in MATLAB. Note that it takes some time to run.

```
% Set up.  
load_system('md1_varymean'); % Load system if not yet loaded.  
nruns = 100; % Number of simulation runs
```

```

mvec = (0.2 : 0.1 : 0.4); % Values of mean arrival rate
nm = length(mvec); % Number of values in mvec
opts = simset('SrcWorkspace','Current','DstWorkspace','Current');

% Preallocate space for results.
w = zeros(nruns,1);
wavg = zeros(nm,1);
w_theor = zeros(nm,1);

% Vary the mean arrival rate.
for midx = 1:nm
    % m is a parameter in the Constant block, so changing m
    % changes the mean arrival rate in the simulation.
    m = mvec(midx);
    disp(['Simulating with mean arrival rate=' num2str(m)]);

    % Replicate for each value of m
    for k = 1:nruns
        sim('mdl1_varymean',[],opts); % Run simulation.
        w(k) = simout(2); % Store empirical value of w.
    end

    wavg(midx) = mean(w); % Average for fixed value of m
    w_theor(midx) = simout(1); % Theoretical value
end
end

```

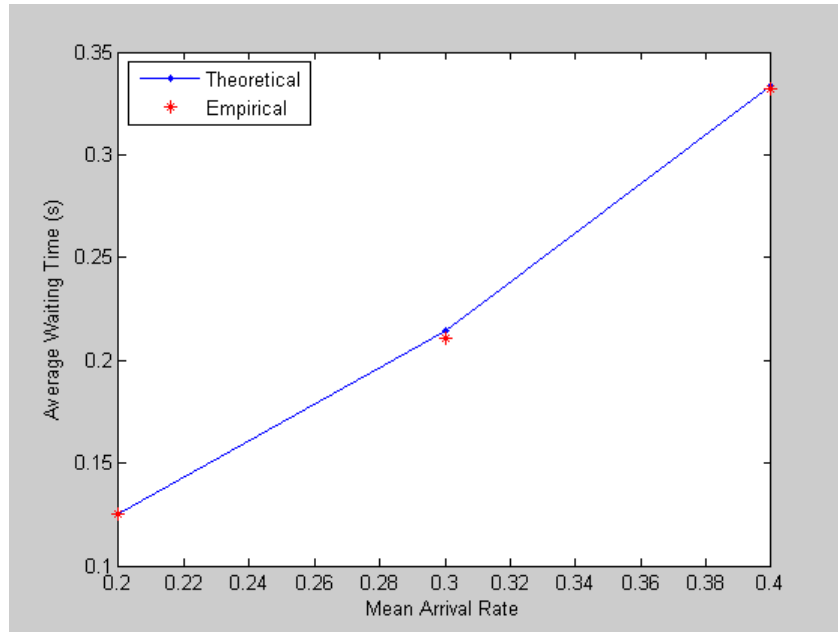
- 12** To plot the average waiting time (averaged over multiple simulations for each fixed value of mean arrival rate) against mean arrival rate, execute the following code in MATLAB.

```

figure; % Create new figure window.
plot(mvec,w_theor,'b.-'); % Plot theoretical curve.
hold on;
plot(mvec,wavg,'r*'); % Plot empirical values.
legend('Theoretical','Empirical','Location','NorthWest')
xlabel('Mean Arrival Rate')
ylabel('Average Waiting Time (s)')
hold off

```

A sample plot is below.



Regulating the Simulation Length

When you run a simulation interactively to observe behavior qualitatively, the stop time of the simulation might not matter. However, when you need to gather statistics from a simulation, knowing when to end the simulation is more important. Typical criteria for ending a discrete-event simulation include the following:

- A fixed amount of time passes
- The cumulative number of entity departures from a particular block crosses a fixed threshold. This might be analogous to processing a fixed number of packets, parts, or customers, for example.
- The simulation achieves a particular state, such as an overflow or a machine failure

Setting a Fixed Stop Time

To run a simulation interactively with a fixed stop time, do the following:

- 1 Open the **Configuration Parameters** dialog box by choosing **Simulation > Configuration Parameters** in the menu of the model window.
- 2 In the dialog box, set **Stop time** to the desired stop time.
- 3 Run the simulation by choosing **Simulation > Start**.

To fix the stop time when running a simulation from MATLAB, use syntax like

```
sim('model',timespan)
```

where `model` is the name of the model and `timespan` is the desired stop time.

Stopping Based on Entity Count


The basic procedure for stopping a simulation based on the total number of entity departures is as follows:

- 1 Find the block's parameter that enables the departure counter as a signal output, usually called **Number of entities departed**. Exceptions are the **Write count to signal port #d** parameter of the Entity Departure Counter block and the **Number of entities arrived** parameter of the Replicate and Entity Sink blocks.
- 2 Set the parameter to On. This causes the block to have a signal output port corresponding to the entity count.
- 3 Connect the new signal output port to a Discrete Event Subsystem block, from the SimEvents Ports and Subsystems library.
- 4 Double-click the Discrete Event Subsystem block to open the subsystem it represents.
- 5 Delete the Discrete Event Outport block labeled `Dout`.
- 6 Connect the Discrete Event Inport block labeled `Din` to a Compare To Constant block, from the Logic and Bit Operations library in Simulink.


7 In the Compare To Constant block,

- Set **Operator** to \geq .
- Set **Constant value** to the desired number of entity departures.
- Set **Output data type mode** to boolean.

8 Connect the Compare To Constant block to a Stop Simulation block, from the Sinks library in Simulink. The result should look like the following, except that your SimEvents block might be a block other than Entity Departure Counter.



Top-Level Model



Subsystem Contents


See the considerations discussed in “Tips for Using State-Based Stopping Conditions” on page 10-38 below. They are relevant if you are stopping the simulation based on an entity count, where “desired state” means the entity-count threshold.

Stopping Upon Reaching a Particular State


Suppose you want the simulation to end when it achieves a particular state, such as an overflow or a machine failure. The state might be the only criterion for ending the simulation, or the state might be one of multiple criteria, each of which is sufficient reason to end the simulation. An example that uses multiple criteria is a military simulation that ends when all identified targets are destroyed or all resources (ammunition, aircraft, etc.) are depleted, whichever occurs first.

Once you have identified a state that is relevant for ending the simulation, you typically create a Boolean signal that queries the state and connect the signal to a Stop Simulation block. Typical ways to create a Boolean signal that queries a state include the following:

- Connect a signal to a logic block to determine whether the signal satisfies some condition. See the blocks in the Logic and Bit Operations library in Simulink. The figure below illustrates one possibility.




Top-Level Model




Subsystem Contents

- Use a Get Attribute block to query an attribute and a logic block to determine whether the attribute value satisfies some condition. The next figure illustrates one possibility.




Top-Level Model




Subsystem Contents

- To end the simulation whenever an entity reaches a particular entity path, you can end that path with an Entity Sink block, enable that block's output signal to count entities, and check whether the output signal is greater than zero.

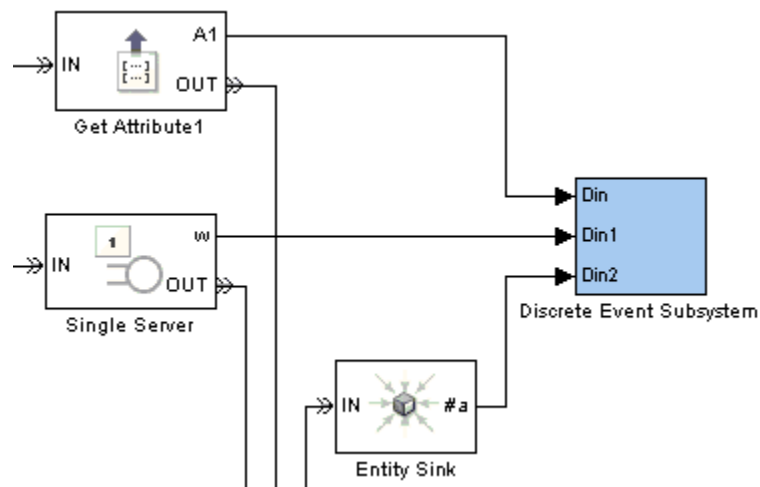


Top-Level Model

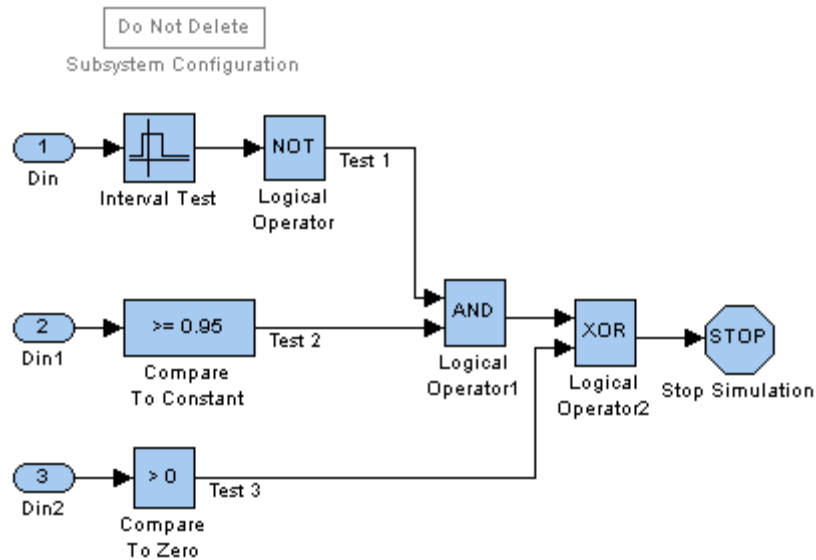


Subsystem Contents

- Logically combine multiple tests using logic blocks to build the final Boolean signal that connects to a Stop Simulation block. (A logical OR operation is implied if your model contains an independent Stop Simulation block for each of the multiple tests, meaning that the simulation ends when the first such block processes an input signal whose value is true.) The figure below illustrates one possibility using the exclusive-OR of two tests, one of which is in turn the logical AND of two tests.



Top-Level Model



Subsystem Contents

Tips for Using State-Based Stopping Conditions. When using a state rather than a time to determine when the simulation ends, keep in mind the following considerations:

- If the model has a finite stop time, then the simulation might end before reaching the desired state. Depending on your needs, this might be a desirable or undesirable outcome. If it is important that the simulation not stop too early, then you can follow the instructions in “Setting a Fixed Stop Time” on page 10-33 and use Inf as the **Stop time** parameter.
- If you set the **Stop time** parameter to Inf, then you should ensure that the simulation actually stops. For example, if you want to stop based on an entity count but the simulation either reaches a deadlock or sends most entities on a path not involving the block whose departure count is the stopping criterion, then the simulation might not end.
- Checking for the desired state throughout the simulation might make the simulation run more slowly than if you used a fixed stop time.

Troubleshooting Discrete-Event Simulations

Viewing the Event Calendar (p. 11-2)	Logging information about events
Viewing Entity Locations (p. 11-8)	Logging information about entities advancing from block to block
Common Problems in SimEvents Models (p. 11-12)	Some modeling errors and ways to avoid them
Configuration Parameters for SimEvents Models (p. 11-30)	SimEvents parameters in the Configuration Parameters dialog box

Additional troubleshooting techniques are in “Using Plots for Troubleshooting” on page 9-8.

Viewing the Event Calendar

Knowing which events are on the event calendar at relevant times during the simulation can help you learn and troubleshoot. SimEvents models offer an option to have the MATLAB Command Window tell you

- When each event is placed on the event calendar
- When each event is processed
- The list of events on the event calendar

This section describes how to enable this option and interpret the information. The topics are as follows:

- “Turning Event Logging On” on page 11-2
- “Logging the Processing of Events” on page 11-3
- “Logging the Scheduling of Events” on page 11-4
- “Logging the List of Events” on page 11-4
- “Example: Event Logging” on page 11-6

To create a file containing messages that appear in the MATLAB Command Window, use the `diary` function.

Turning Event Logging On

To enable event logging for a particular model that contains one or more blocks from the SimEvents libraries, use this procedure:

- 1** Select **Simulation > Configuration Parameters** from the model window’s menu to open the Configuration Parameters dialog box.
- 2** Click **SimEvents** on the left side of the Configuration Parameters dialog box.
- 3** Select one or more of the following options on the right side of the dialog box:
 - **Display events in event calendar**
 - **Log events when executed**

- **Log events when scheduled**

Selecting **Display events in event calendar** disables the other two options because displaying the events in the event calendar automatically includes logging events when scheduled and executed. For details, see “Logging the List of Events” on page 11-4.

The next time you run the simulation, the MATLAB Command Window displays information about the event calendar.

Logging the Processing of Events

When you select **Log events when executed** as described in “Turning Event Logging On” on page 11-2 and run the simulation, the MATLAB Command Window displays a message like the following each time an event is processed:

```
SimEvents: Executing @ 1.0000000000000000 (T=1.0000000000000000
P=300 B='mymodel/Time-Based Entity Generator' N='EntityGeneration')
```

This indicates that the Time-Based Entity Generator block in a model called mymodel generates an entity at time 1 and that this event has event priority 300.

The table below lists the pieces of information contained in messages like this.

Portion of Message	Description
Executing	Distinguishes this message from a scheduling message described in “Logging the Scheduling of Events” on page 11-4
T= followed by a number	Simulation time at which the event is processed
P= followed by a number	Event priority of the event
B= followed by a block pathname	Block that processes the event
N= followed by a string	Name of the event, such as EntityGeneration or ServiceCompletion

Logging the Scheduling of Events

When you select **Log events when scheduled** as described in “Turning Event Logging On” on page 11-2 and run the simulation, the MATLAB Command Window displays a message like the following each time a new event appears on the event calendar:

```
SimEvents: Scheduling @ 0.0000000000000000 (T=1.0000000000000000
P=300 B='mymodel/Time-Based Entity Generator' N='EntityGeneration')
```

This indicates that at time 0, the Time-Based Entity Generator block in a model called mymodel schedules an entity generation to occur at time 1 and that this event has event priority 300.

The table below lists the pieces of information contained in messages like this.

Portion of Message	Description
Scheduling	Distinguishes this message from an execution message described in “Logging the Processing of Events” on page 11-3
@ followed by a number	Simulation time at which the scheduling occurs
T= followed by a number	Simulation time at which the event occurs
P= followed by a number	Event priority of the event
B= followed by a block pathname	Block that schedules the event
N= followed by a string	Name of the event, such as EntityGeneration or ServiceCompletion

The processing sequence and the scheduling sequence might differ for simultaneous events having equal priority.

Logging the List of Events

When you select **Display events in event calendar** as described in “Turning Event Logging On” on page 11-2 and run the simulation, the MATLAB Command Window displays a message like the following each time a new event appears on the event calendar or is processed:


```

SimEvents: Scheduling @ 2.000000000000000 (T=3.000000000000000
P=300 B='mymodel/Time-Based Entity Generator' N='EntityGeneration')
%BEGIN list event in calendar @ 2.000000000000000
SimEvents: Event in calendar @ 2.000000000000000 (T=2.300000000000000
P=500 B='mymodel/Single Server' N='ServiceCompletion')
SimEvents: Event in calendar @ 2.000000000000000 (T=3.000000000000000
P=300 B='mymodel/Time-Based Entity Generator' N='EntityGeneration')
%END list event in calendar @ 2.000000000000000

```

The message has two key components:

- A line that begins with `SimEvents: Scheduling` is a scheduling message as described in “Logging the Scheduling of Events” on page 11-4. The excerpt above indicates that at time 2, the Time-Based Entity Generator block in a model called `mymodel` schedules an entity generation to occur at time 3 and that this event has event priority 300.

Alternatively, a line that begins with `SimEvents: Executing` is an event processing message as described in “Logging the Processing of Events” on page 11-3.

- Lines between `%BEGIN` and `%END` form a list of all events on the event calendar at a given time, after the scheduling or execution mentioned in the earlier part of the message has occurred.

The table below describes the pieces of information contained in the list of all events.

Portion of Message	Description
@ followed by a number	Simulation time at which the event calendar is being displayed
T= followed by a number	Simulation time at which the event occurs
P= followed by a number	Event priority of the event
B= followed by a block pathname	Block that processes the event
N= followed by a string	Name of the event, such as <code>EntityGeneration</code> or <code>ServiceCompletion</code>

Example: Event Logging

You can view the event calendar from the first few seconds of simulation of the M/M/1 Queuing System demo using this procedure:

- 1 Open the demo by entering `mm1` in the MATLAB Command Window or by using the MATLAB Help browser.
- 2 Select **Simulation > Configuration Parameters** from the model window's menu to open the Configuration Parameters dialog box.
- 3 Click **SimEvents** on the left side of the Configuration Parameters dialog box.
- 4 Select **Display events in event calendar** on the right side of the dialog box and click **OK**.
- 5 Run the simulation for 3 seconds by entering the following in the MATLAB Command Window:

```
sim('mm1',3);
```

Interpreting the Event Logging Messages

The resulting messages in the MATLAB Command Window show you the state of the event calendar during the simulation:

- The first entity generation event is scheduled at time 0, to occur at time 0. This is because the entity generator has the **Generate entity at simulation start** option selected. The event calendar contains the entity generation event.

```
SimEvents: Scheduling @ 0.000000000000000 (T=0.000000000000000  
P=1 B='mm1/Time-Based Entity Generator' N='EntityGeneration')  
%BEGIN list event in calendar @ 0.000000000000000  
SimEvents: Event in calendar @ 0.000000000000000 (T=0.000000000000000  
P=1 B='mm1/Time-Based Entity Generator')  
%END list event in calendar @ 0.000000000000000
```

- The entity generation event is processed, leaving the event calendar empty.

```
SimEvents: Executing @ 0.0000000000000000 (T=0.0000000000000000
P=1 B='mm1/Time-Based Entity Generator' N='EntityGeneration')
%BEGIN list event in calendar @ 0.0000000000000000
%END list event in calendar @ 0.0000000000000000
```

- The entity advances immediately to the server, which schedules a service completion event. The event calendar contains the service completion event.

```
SimEvents: Scheduling @ 0.0000000000000000 (T=2.991406386946900
P=1 B='mm1/Single Server' N='ServiceCompletion')
%BEGIN list event in calendar @ 0.0000000000000000
SimEvents: Event in calendar @ 0.0000000000000000 (T=2.991406386946900
P=1 B='mm1/Single Server')
%END list event in calendar @ 0.0000000000000000
```

- The second entity generation event is scheduled for a future time. The event calendar contains two events: the service completion event and the entity generation event.

```
SimEvents: Scheduling @ 0.0000000000000000 (T=3.184988194595833
P=1 B='mm1/Time-Based Entity Generator' N='EntityGeneration')
%BEGIN list event in calendar @ 0.0000000000000000
SimEvents: Event in calendar @ 0.0000000000000000 (T=2.991406386946900
P=1 B='mm1/Single Server')
SimEvents: Event in calendar @ 0.0000000000000000 (T=3.184988194595833
P=1 B='mm1/Time-Based Entity Generator')
%END list event in calendar @ 0.0000000000000000
```

- The service completion event is processed, leaving only the entity generation event on the event calendar.

```
SimEvents: Executing @ 2.991406386946900 (T=2.991406386946900
P=1 B='mm1/Single Server' N='ServiceCompletion')
%BEGIN list event in calendar @ 2.991406386946900
SimEvents: Event in calendar @ 2.991406386946900 (T=3.184988194595833
P=1 B='mm1/Time-Based Entity Generator')
%END list event in calendar @ 2.991406386946900
```

Viewing Entity Locations

Knowing when an entity departs from one block and arrives at another block during the simulation can help you learn and troubleshoot. SimEvents models offer an option to have the MATLAB Command Window display information about entity locations. This section describes how to enable this option and interpret the information. The topics are as follows:

- “Turning Entity Logging On” on page 11-8
- “Interpreting Entity Logging Messages” on page 11-8
- “Example: Entity Logging” on page 11-9

To create a file containing messages that appear in the MATLAB Command Window, use the `diary` function.

Turning Entity Logging On

To enable entity logging for a particular model that contains one or more blocks from the SimEvents libraries, use this procedure:

- 1** Select **Simulation > Configuration Parameters** from the model window’s menu to open the Configuration Parameters dialog box.
- 2** Click **SimEvents** on the left side of the Configuration Parameters dialog box.
- 3** Select **Log entities advancing from block to block** on the right side of the dialog box:

The next time you run the simulation, the MATLAB Command Window displays information about entities advancing from block to block.

Interpreting Entity Logging Messages

When you select **Log entities advancing from block to block** as described in “Turning Entity Logging On” on page 11-8 and run the simulation, the MATLAB Command Window displays a message like the following each time an entity departs from one block and advances to another block:

```
SimEvents: Entity advancing @ 0.0000000000000000
(From='mymodel/Time-Based Entity Generator' To='mymodel/FIFO Queue')
```

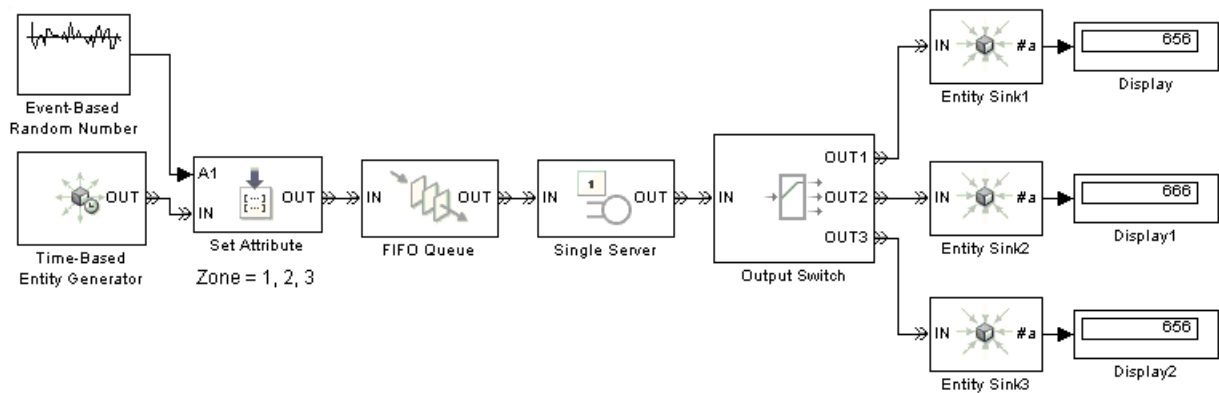
This indicates that at time 0, an entity departs from the Time-Based Entity Generator block and arrives at the FIFO Queue block, in a model called mymodel.

The table below lists the pieces of information contained in messages like this.

Portion of Message	Description
@ followed by a number	Simulation time at which the entity advances
From= followed by a block pathname	Block from which the entity departs
To= followed by a block pathname	Block at which the entity arrives

Example: Entity Logging

By viewing entity locations, you can get information about the model described in “Example: Using an Attribute to Select an Output Port” in the Getting Started documentation.



- At the beginning of the simulation, the first entity advances from the entity generator to the server.

```
SimEvents: Entity Advancing @ 0.0000000000000000
(From='doc_outsw_attr/Time-Based Entity Generator' To='doc_outsw_attr/Set Attribute')
SimEvents: Entity Advancing @ 0.0000000000000000
(From='doc_outsw_attr/Set Attribute' To='doc_outsw_attr/FIFO Queue')
SimEvents: Entity Advancing @ 0.0000000000000000
(From='doc_outsw_attr/FIFO Queue' To='doc_outsw_attr/Single Server')
```

- After completing its service, the first entity departs from the server and is routed to Entity Sink1.

```
SimEvents: Entity Advancing @ 1.0000000000000000
(From='doc_outsw_attr/Single Server' To='doc_outsw_attr/Output Switch')
SimEvents: Entity Advancing @ 1.0000000000000000
(From='doc_outsw_attr/Output Switch' To='doc_outsw_attr/Entity Sink1')
```

- The second entity advances from the entity generator to the server.

```
SimEvents: Entity Advancing @ 4.898639080694728
(From='doc_outsw_attr/Time-Based Entity Generator' To='doc_outsw_attr/Set Attribute')
SimEvents: Entity Advancing @ 4.898639080694728
(From='doc_outsw_attr/Set Attribute' To='doc_outsw_attr/FIFO Queue')
SimEvents: Entity Advancing @ 4.898639080694728
(From='doc_outsw_attr/FIFO Queue' To='doc_outsw_attr/Single Server')
```

- The third entity advances from the entity generator to the queue. (The third entity cannot advance to the server because the server is busy serving the second entity.)

```
SimEvents: Entity Advancing @ 5.234759100998515
(From='doc_outsw_attr/Time-Based Entity Generator' To='doc_outsw_attr/Set Attribute')
SimEvents: Entity Advancing @ 5.234759100998515
(From='doc_outsw_attr/Set Attribute' To='doc_outsw_attr/FIFO Queue')
```

- At the end of its service time, the second entity departs from the server and is routed to Entity Sink2. As a result, the third entity advances from the queue to the server.

```
SimEvents: Entity Advancing @ 5.898639080694728
(From='doc_outsw_attr/Single Server' To='doc_outsw_attr/Output Switch')
```

```
SimEvents: Entity Advancing @ 5.898639080694728  
(From='doc_outsw_attr/Output Switch' To='doc_outsw_attr/Entity Sink2')  
SimEvents: Entity Advancing @ 5.898639080694728  
(From='doc_outsw_attr/FIFO Queue' To='doc_outsw_attr/Single Server')
```

The entity logging messages do not count or otherwise identify which entity is advancing. The descriptions above indicate when the first, second, or third entity is the one that advances because such inferences are straightforward in this example.

Common Problems in SimEvents Models

Troubleshooting a discrete-event simulation can be challenging because blocks that form an entity path operate in coupled ways. The block whose behavior surprises you might not be the source of a mistake. For example, after troubleshooting a surprising set of values in the **#d** output signal from a server, you might find that the problem is not in the server itself but rather in the configuration of a gate or switch block in another part of the model.

Some common problems relate to the simultaneity of events and the sequence in which the events are processed. When events occur at the same value of the simulation clock, it is because they have a causal relationship to each other or because their occurrence times happen to match. Because the value of the simulation clock is a floating-point number, testing two event times for equality is subject to the inherent limitations of comparing floating-point numbers using a computer.

This section describes some common problems. Specific symptoms and fixes are difficult to generalize, but this section offers examples or tips where feasible.

The problems are presented in these topics:

- “Unexpectedly Simultaneous Events” on page 11-13
- “Unexpectedly Nonsimultaneous Events” on page 11-13
- “Unexpected Processing Sequence for Simultaneous Events” on page 11-16
- “Time-Based Block Not Recognizing Certain Trigger Edges” on page 11-17
- “Incorrect Timing of Signals” on page 11-17
- “Unexpected Use of Old Value of Signal” on page 11-19
- “Effect of Initial Condition on Signal Loops” on page 11-23
- “Loops in Entity Paths Without Storage Blocks” on page 11-26
- “Unexpected Timing of Random Signal” on page 11-27
- “Unexpected Correlation of Random Processes” on page 11-29

Unexpectedly Simultaneous Events

An unexpected simultaneity of events can result from roundoff error in event times or other floating-point quantities, and might cause the processing sequence to differ from your expectation about when each event should occur. Computers' use of floating-point arithmetic involves a finite set of numbers with finite precision.

If you have a guess about which events' processing is suspect, then adjusting event priorities or using the Instantaneous Event Counting Scope block can help diagnose the problem. For examples involving event priorities, see "Example: Race Conditions at a Switch" on page 2-23 and the Event Priorities demo. For examples using the Instantaneous Event Counting Scope block, see "Example: Plotting Event Counts to Find Roundoff Error" on page 9-12 and "Example: Counting Events from Multiple Sources" on page 2-41.

Unexpectedly Nonsimultaneous Events

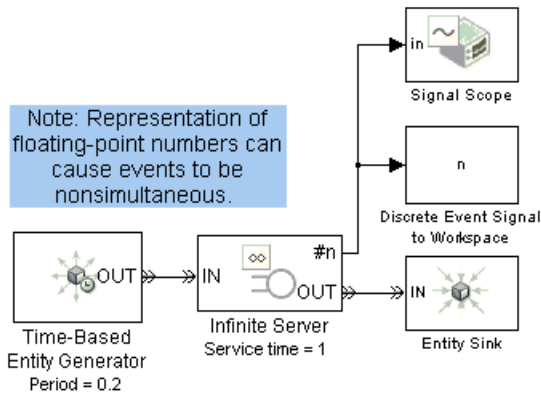
An unexpected lack of simultaneity can result from roundoff error in event times or other floating-point quantities. Computers' use of floating-point arithmetic involves a finite set of numbers with finite precision.

If roundoff error is very small, then the event logging feature and scope blocks might not reveal enough precision to confirm whether events are simultaneous or only close. An alternative technique is to use the Discrete Event Signal to Workspace block to collect data in MATLAB as in the example below.

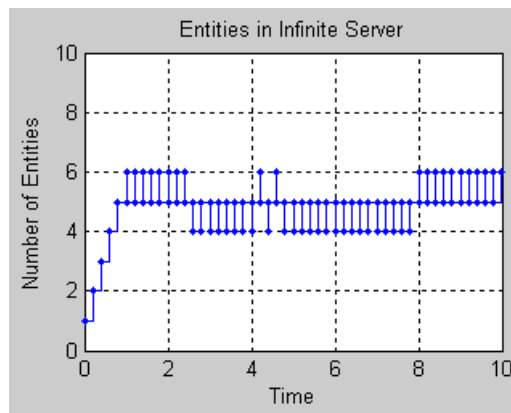
If your model requires that certain events be simultaneous, then use modeling techniques aimed at effecting simultaneity. For examples, see "Example: Plotting Event Counts to Find Roundoff Error" on page 9-12 and "Example: Race Conditions at a Switch" on page 2-23.

Example: Discrepancies in Event Times

In the model below, the Time-Based Entity Generator block has a constant service time of 0.2 and the Infinite Server block has a constant service time of 1.



After the first entity completes its service, you might expect the server to have a simultaneous arrival and departure every 0.2 seconds. However, roundoff error sometimes prevents the generation of a new entity from being exactly simultaneous with the service completion of a previous entity. The plot suggests that the generation events are sometimes, but not always, processed before the service completion events. This might be unexpected if the generation events and service completion events have distinct event priority values.



The Discrete Event Signal to Workspace block records the server's $\#n$ signal as the workspace variable n . By examining n , you can see detail that is not apparent from the plot.

```
times_values = [n.time, n.signals.values];
format short
times_values(90:91,:)
```

The first column of the output below suggests that two events occur simultaneously at T=9.4.

```
ans =

    9.4000    6.0000
    9.4000    5.0000
```

To find out whether the events are actually simultaneous, use the hexadecimal format to view the same data.

```
format hex
times_values(90:91,:)
```

The output below shows that the event times in the first column are not equal. That is, the events are not simultaneous but only nearly simultaneous. As a result, their event priorities are not relevant. Because roundoff error can make entity generation events occur slightly before or slightly after service completion events, the plot above looks irregular.

```
ans =

4022cccccccccccc 4018000000000000
4022cccccccccccc 4014000000000000
```

To find the magnitude of the difference, subtract.



```
format short
times_values(91,1)-times_values(90,1)
```

The output is

```
ans =

3.5527e-015
```

An alternative model would use the Event-Based Entity Generator block with a **ts** input signal having a sample time of 0.2 seconds, instead of the Time-Based Entity Generator block.



Unexpected Processing Sequence for Simultaneous Events

An unexpected sequence for simultaneous events could result from the arbitrary or random handling of events having equal priorities, mentioned in “Processing Sequence for Simultaneous Events” on page 2-21. The sequence might even change when you run the simulation again. When the sequence

is arbitrary, you should not make any assumptions about the sequence or its repeatability.

If you copy and paste blocks that have an event priority parameter, the parameter values do not change unless you manually change them.

Time-Based Block Not Recognizing Certain Trigger Edges

Time-based blocks have a slightly different definition of a trigger edge compared to event-based blocks. If you use event-based signals with Triggered Subsystem blocks or Stateflow blocks with trigger inputs, then the blocks might not run when you expect them to. For more information, suggestions, and an example, see “Zero-Duration Values and Time-Based Blocks” on page 12-10.

Incorrect Timing of Signals

If you use a time-based block to process event-based signals, then the output signal might be a time-based signal. Depending on your model, you might notice that

- The output signal assumes a new value at a later time than the event that caused the last update of the event-based signal.
- The output signal assumes incorrect values.
- An event-based block that uses the output signal, such as an Event-Based Entity Generator block, operates with incorrect timing.


You can avoid these problems by putting the time-based block in a discrete event subsystem, as described in Chapter 8, “Controlling Timing Using Subsystems”. If your time-based block is in a Function-Call Subsystem, then be sure to select **Propagate execution context across subsystem boundary** as described in “Setting Up Function-Call Subsystems in SimEvents Models” on page 8-32.

Example: Time-Based Addition of Event-Based Signals


The model below adds the lengths of two queues. The queue lengths are event-based signals, while the Add block is a time-based block. It is important

that the Add block use up-to-date values of its input signals each time the queue length changes and that the output signal's updates correspond to updates in one of the queue length signals.


**Modeling error
in use of time-based block**



If you build this model without having used `simeventsstartup` previously, or without using `simeventsconfig` later, then you might see the plot below. The incorrect timing is evident because the sum signal has updates at regular intervals that are smaller than the minimum intergeneration time of the entity generators.



If you correct the simulation parameters by using `simeventsconfig` on this model (with either the 'des' or 'hybrid' input argument), then the plot reveals correct update times but incorrect values. To check the values, you can connect the inputs and outputs of the Add block to separate Discrete Event Signal to Workspace blocks and examine the data in the MATLAB workspace.



A better model uses the discrete event subsystem illustrated in the Time-Driven and Event-Driven Addition demo.

Unexpected Use of Old Value of Signal

During a discrete-event simulation, multiple events or signal updates can occur at a fixed value of the simulation clock. If these events and signal updates are not processed in the sequence that you expect, then you might notice that a computation or other operation uses a signal value from a previous time instead of from the current time. Some common situations occur when

- A block defers the update of an output signal until a departing entity has finished advancing to a subsequent storage block, but an intermediate nonstorage block in the sequence uses that signal in a computation or to control an operation. Such deferral of updates applies to most SimEvents blocks that have both an entity output port and a signal output port.

For details and an example, see “Interleaved Operations of Storage and Nonstorage Blocks” on page 12-4.

- A computation involving multiple signals is performed before all of the signals have been updated.

For details and an example, see “Sequence of Updates of Output Signals” on page 3-10.

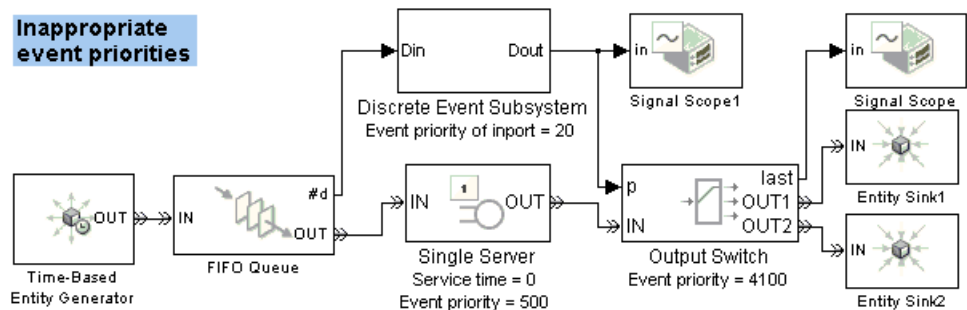
- An inappropriate processing sequence for simultaneous events causes a signal update to occur after a block uses that signal in a computation or to control an operation. See the example below.
- A time-based block’s use of a value of an event-based signal persists until the next time step of the time-based simulation clock, even if the block producing the event-based signal has already updated the value. In many cases, this is the correct behavior of the time-based block.

For an example, see “Example: Plotting Entity Departures to Verify Timing” on page 9-8.

If you need a time-based block to respond to events, consider using a discrete event subsystem as described in Chapter 8, “Controlling Timing Using Subsystems”.

Example: Incorrect Sequence of Event Priorities

In the model below, an inappropriate set of event priorities causes the update of the **Dout** signal to occur after the Output Switch has already performed an operation using its input signal. The goal is to have the third entity in the simulation use the **OUT2** entity output port from the Output Switch block, while all other entities use the **OUT1** port. The Discrete Event Subsystem block uses the queue block’s **#d** signal to determine whether three entities have departed from the queue and output the appropriate port index. The **#d** signal is updated only after the entity departing from the queue has arrived at a subsequent storage block, which is why the model contains a Single Server block whose service time is 0.





The desired sequence of operations at queue departure times is as follows:

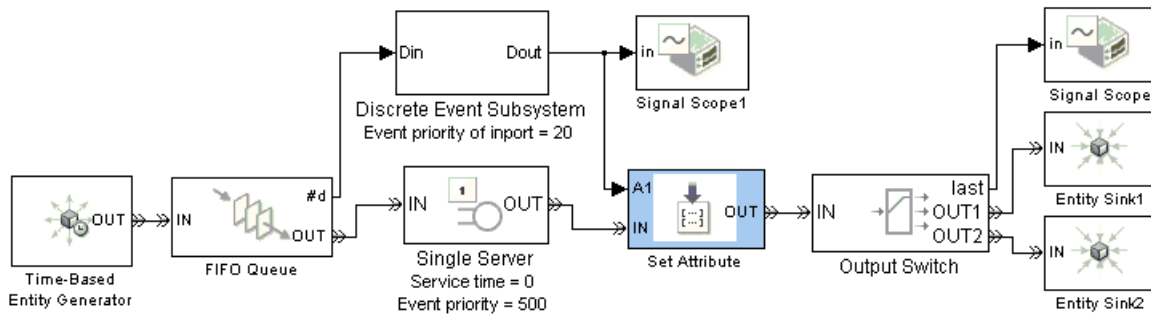
- 1 The entity advances from the queue to the server.
- 2 The queue block updates its **#d** output signal. This automatically follows from the previous step.
- 3 The subsystem responds to the update in the **#d** signal by computing the appropriate port index.
- 4 If the port index has changed, then the switch block selects a different output port.
- 5 The entity advances from the server to the switch block and then to one of the entity sinks.

However, the event priority parameters of the Discrete Event Inport block (inside the Discrete Event Subsystem), Single Server block, and Output Switch block can either support or contradict the desired sequence listed above. If the subsystem is executed after the server has completed its service, then the entity arrives at the switch while the **Dout** signal still retains the value corresponding to a previous entity departure. Also, if the switch schedules its port selection after the server has completed its service, then the entity arrives at the switch while the switch is still configured for a previous entity. In summary, the subsystem's execution and the switch's port selection must precede the server's service completion. The relative values of event priorities for the subsystem and the switch are not relevant because the events have a causal relationship; that is, a change in the port selection depends on, and thus follows, a change in the output signal from the subsystem.

The event priorities shown in the figure above cause the third entity to advance to the switch before the switch has responded to the change in the **Dout** signal. As a result, the entity that uses the **OUT2** entity output port is the fourth entity, not the third entity. You can see this by comparing the **Dout** and **last** plots, which show the desired port index and the actual port index, respectively.



A better model would set the Output Switch block's **Switching criterion** parameter to From attribute and connect **Dout** to the **A1** input port of a Set Attribute block. An attribute-based approach is appropriate here because the switching criterion is closely related to the series of entities. While the Output Switch block monitors changes in its **p** input signal and processes a port selection event after detecting a change, the Set Attribute block reads its **A1** input signal upon each entity arrival. As before, the subsystem's execution must precede the server's service completion; however, two priorities are easier to understand than three.




Effect of Initial Condition on Signal Loops

When you create a loop in a signal connection, consider the effect of initial conditions and be aware that event-based signals might have the first update upon the first relevant event, not necessarily at the beginning of the simulation. If you need to specify initial conditions for event-based signals, see “Defining Initial Conditions for Event-Based Signals” on page 3-18.


Example: Intergeneration Time of Zero at Simulation Start

The model below is problematic at $T=0$ because the initial reading of the t input signal representing the intergeneration time is 0. This signal does not assume a positive value until the first entity departs from the Read Timer block, which occurs after the first completion of service at $T=1$.

Modeling error at $T = 0$




A better model would use the technique described in “Defining Initial Conditions for Event-Based Signals” on page 3-18 to specify a nonzero initial condition for the w output signal from the Read Timer block.

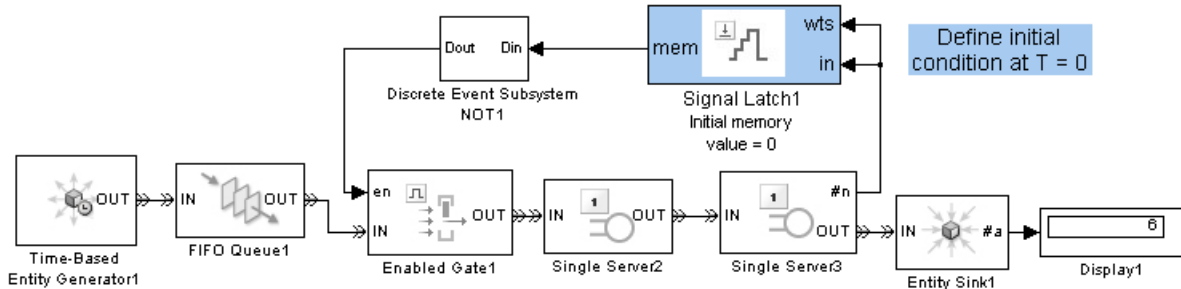
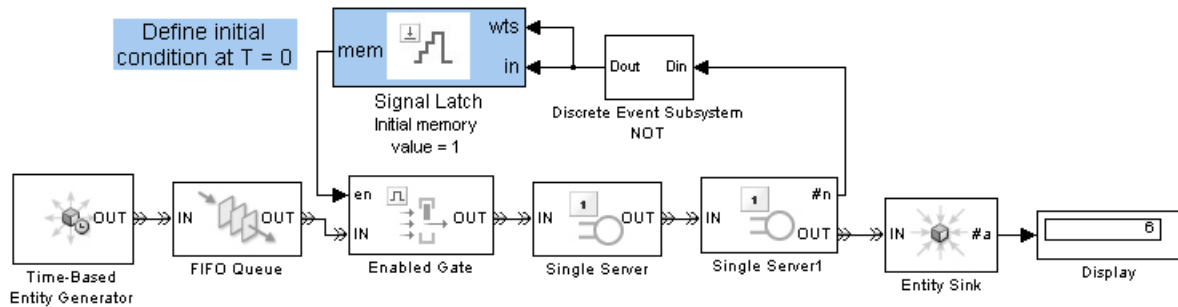


Example: Absence of Sample Time Hit at Simulation Start

In the model below, the second server's #n signal has no updates before the first entity arrival there. As a result, the discrete event subsystem, whose role is to perform a NOT operation on the #n signal, does not execute before the first entity arrival at the server. However, no entity can arrive at the server until the gate opens. This logic causes entities to accumulate in the queue instead of advancing past the gate and to the servers.




A better model would use the technique described in “Defining Initial Conditions for Event-Based Signals” on page 3-18 to define either a positive initial condition for the **en** input signal to the gate, or a zero initial condition for the **#n** output signal from the server.



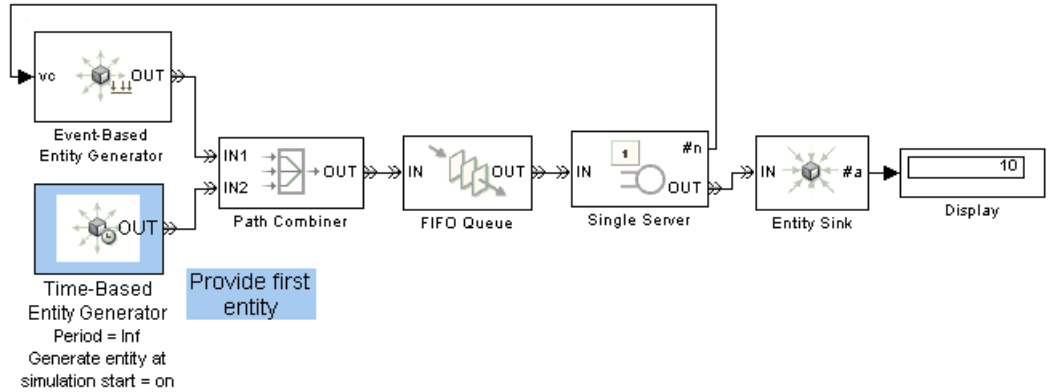
Example: Faulty Logic in Feedback Loop

The model below generates no entities because the logic is circular. The entity generator is waiting for a change in its input signal, but the server's output signal never changes until an entity arrives or departs at the server.

Modeling error at T = 0



A better model would provide the first entity in a separate path. In the revised model below, the Time-Based Entity Generator block generates exactly one entity during the simulation, at T=0.




Loops in Entity Paths Without Storage Blocks

An entity path that forms a loop should contain a storage block. Storage blocks include queues and servers; for a list of storage blocks, see “Interleaved Operations of Storage and Nonstorage Blocks” on page 12-4. The example below illustrates how the storage block can prevent a deadlock.


Example: Deadlock Resulting from Loop in Entity Path

The model below contains a loop in the entity path from the Output Switch block to the Path Combiner block. The problem occurs when the switch selects the entity output port **OUT2**. The entity attempting to depart from the server looks for a subsequent storage block where it can reside, and it cannot reside in a routing block. Until the entity confirms that it can advance to a storage block, the entity cannot depart. However, until it departs, the server is not available to accept a new arrival. The result is a deadlock.

Modeling error when switch uses OUT2



A better model would include a server with a service time of 0 in the looped entity path. This storage block provides a place for an entity to reside after it departs from the Output Switch block. After the service completion event is processed, the entity advances to the Path Combiner block and back to the Single Server block. Notice also that the looped entity path connects to the Path Combiner block's **IN1** entity input port, not **IN2**. This ensures that entities on the looped path, not new entities from the queue, arrive back at the Single Server block.

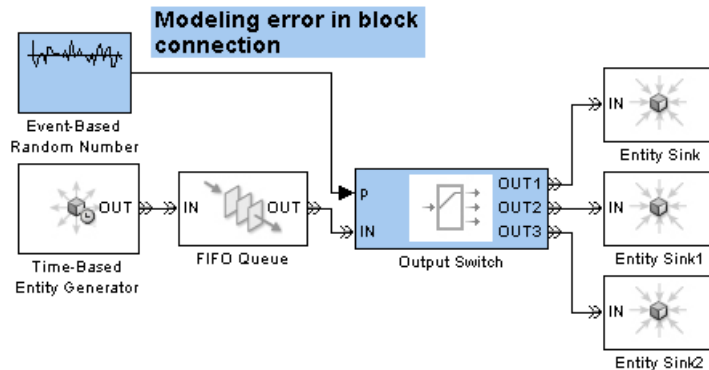


Unexpected Timing of Random Signal

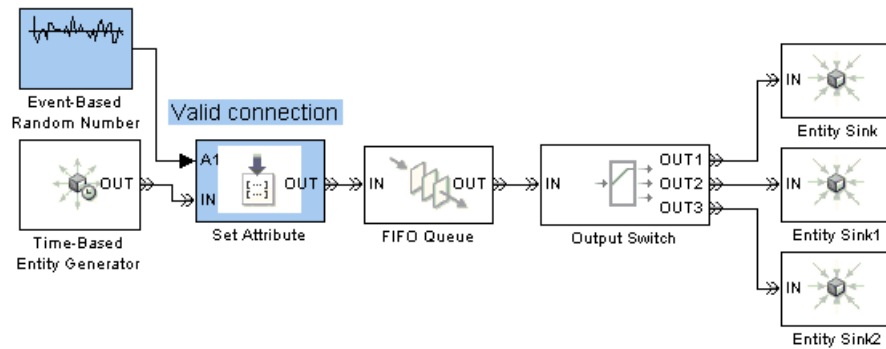
When you use the Event-Based Random Number block to produce a random event-based signal, the block infers from a subsequent block the events upon which to generate a new random number from the distribution. The sequence of times at which the block generates a new random number depends on the port to which the block is connected and on events occurring in the simulation. To learn how to use this block, see “Generating Random Signals” on page 3-4.

Example: Invalid Connection of Event-Based Random Number Generator

The model below is incorrect because the Event-Based Random Number block cannot infer from the **p** input port of an Output Switch block when to generate a new random number. The Output Switch block is designed to listen for changes in its **p** input signal and respond when a change occurs; that is, the Output Switch cannot cause changes in the input signal value or tell the random number generator when to generate a new random number. The **p** input port of the Output Switch block is called a reactive port and it is not valid to connect a reactive signal input port to the Event-Based Random Number block.



If your goal is to generate a new random number corresponding to each entity that arrives at the switch, then a better model connects the Event-Based Random Number block to a Set Attribute block and sets the Output Switch block's **Switching criterion** parameter to **From** attribute. The random number generator then generates a new random number upon each entity arrival at the Set Attribute block. The connection of the Event-Based Random Number block to the **A1** input port of the Set Attribute block is a supported connection because the **A2** port is a notifying port. To learn more about reactive ports and notifying ports, see the reference page for the Event-Based Random Number block.




Unexpected Correlation of Random Processes

An unexpected correlation between random processes could result from equal initial seeds in different dialog boxes. If you copy and paste blocks that have an **Initial seed** parameter, the parameter values do not change unless you manually change them. Such blocks include

- Time-Based Entity Generator
- Event-Based Random Number
- Uniform Random Number
- Random Number
- Blocks in the Routing library

Configuration Parameters for SimEvents Models

When a model contains at least one SimEvents block, the model's Configuration Parameters dialog box has a tab with parameters specific to discrete-event simulation.



Execution order of simultaneous events

If you select **Arbitrary**, an internal algorithm determines the sequence for processing simultaneous events having equal priorities. If you select **Randomized**, all possible sequences have equal probability. In either case, the processing sequence might be different from the sequence in which the events were scheduled on the event calendar. For more information, see “Events with Equal Priorities” on page 2-20.

Seed for event randomization

The initial seed of the random number generator used to determine the sequence for processing simultaneous events having equal priorities. For a given seed, the generator's output is repeatable. This field appears only if you set **Execution order of simultaneous events** to **Randomized**.

Display events in event calendar

If you select this option, the MATLAB Command Window displays a message and the list of events in the event calendar, each time an event is either scheduled or processed. For more information, see “Logging the List of Events” on page 11-4.

Log events when executed

If you select this option, the MATLAB Command Window displays a message each time an event is processed. For more information, see “Logging the Processing of Events” on page 11-3.

Log events when scheduled

If you select this option, the MATLAB Command Window displays a message each time an event is scheduled on the event calendar. For more information, see “Logging the Scheduling of Events” on page 11-4.

Log entities advancing from block to block

If you select this option, the MATLAB Command Window displays information about entities advancing from block to block. For more information, see “Viewing Entity Locations” on page 11-8.

How SimEvents Works

Complementing the information in “How Simulink Works” and “Simulating Dynamic Systems” in the Simulink documentation, this section describes some aspects that are different for models that involve both time-based and event-based processing.

Notifications and Queries Among
Blocks (p. 12-2)

When and why SimEvents blocks
interact with each other, and the
impact on simulation behavior

Interleaved Operations of Storage
and Nonstorage Blocks (p. 12-4)

Nonlinear sequence of block
operations and the impact on
simulation behavior

Zero-Duration Values and
Time-Based Blocks (p. 12-10)

Caveats and techniques for working
with multivalued signals

Notifications and Queries Among Blocks

In a variety of situations, a SimEvents block notifies other blocks about changes in its status or queries other blocks about their status. These interactions among blocks are essential to the proper functioning of a discrete-event simulation. The interactions occur automatically without being reported to you explicitly.

This section gives examples of several types of notifications and queries. The topics are

- “Querying Whether a Subsequent Block Can Accept an Entity” on page 12-2
- “Notifying Blocks About Status Changes” on page 12-3

Querying Whether a Subsequent Block Can Accept an Entity

Before a SimEvents block outputs an entity, it queries the next block to determine whether that block can accept the entity. For example,

- When an entity arrives at an empty FIFO Queue block, the queue queries the next block. If that block can accept an entity, the queue outputs the entity at the head of the queue; otherwise, the queue holds the entity.
- While a Single Server block is busy serving, it does not query the next block. Upon completion of the service time, the server queries the next block. If that block can accept an entity, the server outputs the entity that has completed its service; otherwise, the server holds the entity.
- When an entity attempts to arrive at a Replicate block, the block queries each of the blocks connected to its entity output ports. If all of them can accept an entity, then the Replicate block copies its arriving entity and outputs the copies; otherwise, the block does not permit the entity to arrive there and the entity must stay in a preceding block.
- When a Time-Based Entity Generator block generates a new entity, it queries the next block. If that block can accept an entity, then the generator outputs the new entity; otherwise, the behavior of the Time-Based Entity Generator block depends on the value of its **Response when blocked** parameter.

- When a block (for example, a Single Server block) attempts to advance an entity to the Input Switch block, the server uses a query to check whether it is connected to the currently selected entity input port of the Input Switch block. If so, the Input Switch queries the next block to determine whether it can accept the entity because the Input Switch block cannot hold an entity for a nonzero duration.

Notifying Blocks About Status Changes

When a SimEvents block undergoes certain kinds of status changes, it notifies other blocks of the change. This notification might cause the other blocks to change their behavior or status in some way, depending on the circumstances. For example,

- When an entity departs from a Single Server block, it notifies the preceding block that the server's entity input port has changed from unavailable to available.
- When an entity departs from a queue that was full to capacity, the queue notifies the preceding block that the queue's entity input port has changed from unavailable to available.
- When a Path Combiner block receives notification that the next block's entity input port has changed from unavailable to available, the Path Combiner block's entity input ports also become available. The block notifies preceding blocks that its entity input ports are available.

This case is subtle because the Path Combiner block usually has more than one block to notify, and the sequence of notifications can be significant. See the block's reference page for more information about the options.

- When an entity arrives at a Single Server block that has a **t** signal input port representing the service time, that port notifies the preceding block of the need for a new service time value. If the preceding block is the Event-Based Random Number block, then it responds by generating a new random number that becomes the service time for the arriving entity.

Interleaved Operations of Storage and Nonstorage Blocks

At all simulation times from an entity's generation to destruction, the entity resides in a block. Blocks capable of holding an entity for a nonzero duration are called storage blocks. Blocks that permit an entity arrival but must output the entity at the same value of the simulation clock are called nonstorage blocks. During a simulation, whenever an entity departs from a block, the application processes enough queries, departures, arrivals, and other block operations until the entity arrives at a subsequent storage block. Some block operations, including the updates of statistical output signals that are intended to be updated after the entity's departure, are deferred until after the entity's arrival at a subsequent storage block.

An important consequence of the deferral of block operations is that certain operations of storage and nonstorage blocks might become interleaved. You might want to insert storage blocks in key locations along entity paths in your model to change the sequence of block operations, as illustrated in "Example: Sequence of Departures and Statistical Updates" on page 12-5.

Storage Blocks

- Blocks in Queues library (However, these can act like nonstorage blocks in some circumstances; see the note below.)
- Blocks in Servers library
- Blocks in Entity Generators library
- Entity Sink block
- Attribute Scope, X-Y Attribute Scope, and Instantaneous Entity Counting Scope blocks when configured as a sink, that is, without an entity output port


Note In the special case of an entity arriving at an empty queue whose entity output port is not blocked, the queue acts like a nonstorage block in that block operations are deferred until the entity's arrival at a storage block subsequent to the queue.

Nonstorage Blocks

- Blocks in Routing library
- Blocks in Gates library
- Blocks in Timing library
- Blocks in Probes library
- Blocks in Attributes library
- Attribute Scope, X-Y Attribute Scope, and Instantaneous Entity Counting Scope blocks when configured with an entity output port
- Entity Departure Event to Function-Call Event block
- Entity-Based Function-Call Event Generator block

Example: Sequence of Departures and Statistical Updates

Consider the sequence of operations in the Time-Based Entity Generator, Set Attribute, and Attribute Scope blocks shown below.




At each time $T = 1, 2, 3, \dots, 10$, Simulink processes the following operations in the order listed:

Order	Operation	Block
1	Entity generation	Time-Based Entity Generator
2	Entity departure	Time-Based Entity Generator
3	Arrival at nonstorage block	Set Attribute
4	Assignment of attribute using value at A1 signal input port	Set Attribute
5	Entity departure	Set Attribute


Order	Operation	Block
6	Arrival at storage block	Attribute Scope
7	Update of plot	Attribute Scope
8	Update of signal at #d signal output port	Time-Based Entity Generator

The final operation of the Time-Based Entity Generator block is deliberately processed *after* operations of subsequent blocks in the entity path are processed. This explains why the plot shows a value of 0, not 1, at T=1.

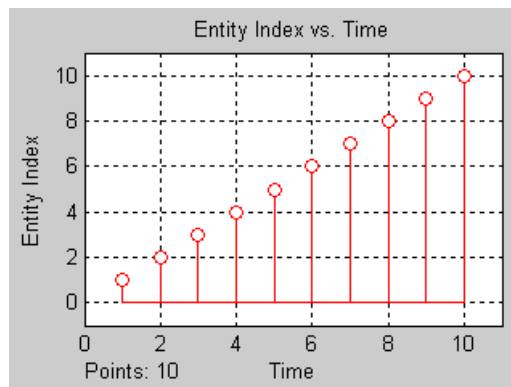


Altering the Processing Sequence

If you want to be sure that the Set Attribute block reads the value at the **A1** signal input port after the Time-Based Entity Generator block has updated its #d output signal, then insert a storage block between the two blocks. In this simple model, you can use a Single Server block with a **Service time** parameter of 0. The model, table, and plot are below.





Order	Operation	Block
1	Entity generation	Time-Based Entity Generator
2	Entity departure	Time-Based Entity Generator
3	Arrival at storage block	Single Server
4	Update of signal at #d signal output port	Time-Based Entity Generator
5	Service completion	Single Server
6	Entity departure	Single Server
7	Arrival at nonstorage block	Set Attribute
8	Assignment of attribute using value at A1 signal input port	Set Attribute
9	Entity departure	Set Attribute
10	Arrival at storage block	Attribute Scope
11	Update of plot	Attribute Scope



Consequences of Inserting a Storage Block


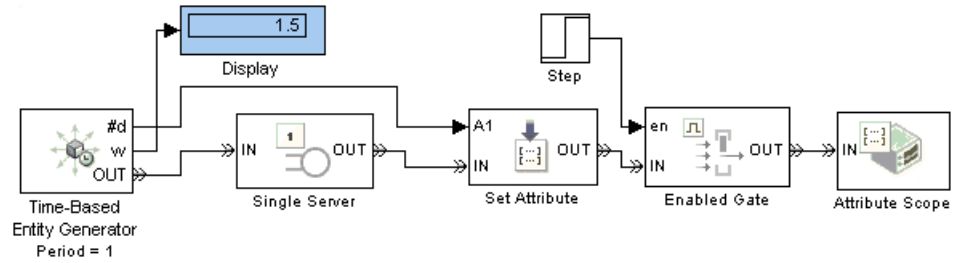
If the storage block you have inserted to alter the processing sequence holds the entity longer than you expect (beyond the zero-duration service time, for example), be aware that your simulation might change in other ways. You should consider the impact of either inserting or not inserting the storage block.

For example, suppose you add a gate block to the preceding example and view the average intergeneration time, w , of the entity generator block. When the gate is closed, a newly generated entity cannot advance immediately to the scope block. Whether this entity stays in the entity generator or a subsequent server block affects the w signal, as shown in the figures below.



Model with Gate and Without Storage Block

When a storage block is present, the first pending entity stays there instead of in the entity generator. The earlier departure of the first entity from the entity generator increases the value of the w signal.



Model with Gate and Storage Block

Zero-Duration Values and Time-Based Blocks

Because time-based simulations involve signals that assume a unique value at each value of the simulation clock, some blocks designed for time-based simulations recognize only one value of a signal per time instant. Because zero-duration values commonly occur in discrete-event simulations (for example, statistical output signals from SimEvents blocks), you should be aware of techniques for working with zero-duration values. The table below lists examples of time-based blocks that recognize one signal value per time instant, along with similar blocks or techniques that recognize multivalued signals.


Time-Based Block	Block or Technique for Working with Multivalued Signals
Scope	Signal Scope in the SimEvents Sinks library
To Workspace	Discrete Event Signal to Workspace in the SimEvents Sinks library. Alternatively, put the To Workspace block in a discrete event subsystem.
Triggered Subsystem	Discrete Event Subsystem in the SimEvents Ports and Subsystems library, with the Discrete Event Inport block configured to execute the subsystem upon trigger edges. Alternatively, use the Signal-Based Event to Function-Call Event block in the Event Translation library to convert the trigger signal to a function call, and then use a Function-Call Subsystem instead of a Triggered Subsystem.
Stateflow with a trigger input signal	Use the Signal-Based Event to Function-Call Event block in the Event Translation library to convert the trigger signal to a function call, then call the Stateflow block with a function-call signal instead of a trigger signal.

For an example comparing the Scope viewer with the Signal Scope block, see “Comparison with Time-Based Plotting Tools” on page 9-15.


Example: Using a #n Signal as a Trigger

Suppose you want to call a subsystem each time the #n signal from a Single Server block rises from 0 to 1. This signal is 0 when the server is not storing an entity and 1 when the server is storing an entity. It is common for an entity to arrive at a server at the same time instant that the previous entity departed from the server. In this case, the #n signal changes from 1 to 0 and back to 1 in the same time instant. A time-based block that recognizes only one value of a signal per time instant might not recognize a rising edge that occurs in a time interval of length zero.

This example uses a Counter Free-Running block inside a subsystem to count the number of times the subsystem is called. (Be aware that the Counter Free-Running block starts counting from zero, not one.)

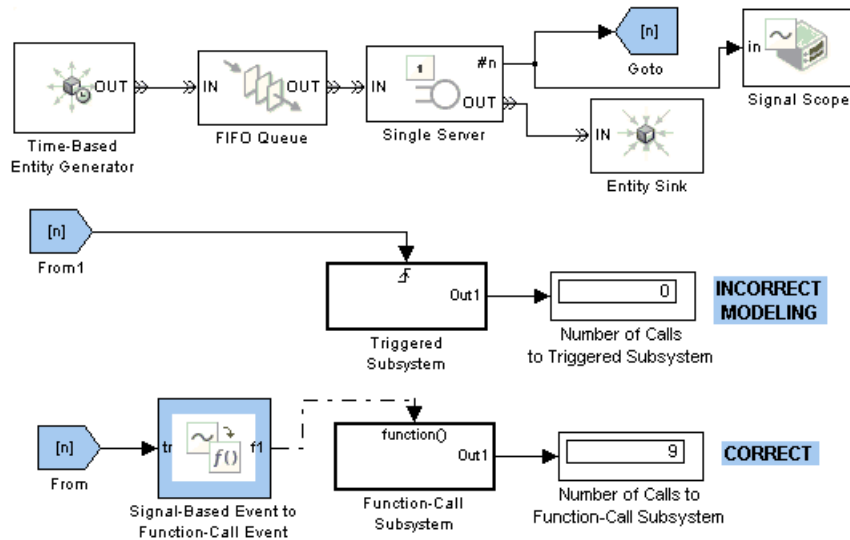


The discrete-event portion of the simulation involves a D/D/1 queuing system in which the server is never idle for a nonzero period of time. As a result, the #n signal exhibits many zero-duration values, shown in the plot below.



The example uses two approaches to try to call the subsystem each time the server's #n signal rises from 0 to 1:

- The approach using a Triggered Subsystem is unsuitable because it does not count changes that occur in a time interval of length zero. You can see from the Display block that the triggered subsystem is never called.
- The approach using function calls is appropriate because the Signal-Based Event to Function-Call Event block recognizes rising edges of #n even when they involve zero-duration values. The block converts these rising edges into function calls to which the Function-Call Subsystem responds.



Demonstration Models

Tutorial Demos (p. 13-2)

Models illustrating how to use
SimEvents blocks

Queuing Systems (p. 13-7)

Models involving queue-server
systems

Application Demos (p. 13-16)

Models for specific discrete-event
applications

Tutorial Demos

Tutorial demos include

- “Attributes: Data Within Entities” on page 13-2
- “Service Time From Attribute” on page 13-3
- “Specifying Service Time in Single Server” on page 13-3
- “Specifying Service Time in Infinite Server Block” on page 13-3
- “Single Server Block Versus Infinite Server Block” on page 13-3
- “Start Timer and Read Timer Blocks” on page 13-4
- “Release Gate: Value Change Versus Trigger” on page 13-4
- “Input Switching Using Signal” on page 13-4
- “Output Switching Using Signal” on page 13-4
- “Path Combiner Versus Input Switch” on page 13-5
- “Time-Driven and Event-Driven Addition” on page 13-5
- “Counting Instantaneous Events” on page 13-6
- “Preload Queue with Entities” on page 13-6

Attributes: Data Within Entities

This model demonstrates how entities carry data from block to block. A piece of such data is called an attribute. The attribute can be read by a subsequent block.

The Set Attribute block assigns an `EntityIndex` attribute in every entity that arrives at its input port. The value for this attribute comes from the Event-Based Random Number block connected to the **A1** port of the Set Attribute block.

The Get Attribute block reads the value of the `EntityIndex` attribute in each arriving entity and writes this value to the **A1** signal output port.

Service Time From Attribute

This model demonstrates one way to use attributes of entities.

The Set Attribute block attaches a `ServiceTime` attribute to each arriving entity. The value for this attribute is the current value of the Repeating Sequence Stair block. The Single Server block, which is configured to read service time from an attribute, delays each arriving entity by the value of that entity's `ServiceTime` attribute. The Departing Entity Service Time plot shows the times at which each entity completes service and leaves the Single Server block.

Specifying Service Time in Single Server

On the left side of this model, the service time for the single server is specified via the `myServiceTime` attribute of every entity arriving at this block. The value of `myServiceTime` is specified as 1.5 seconds. The plot labeled Number of entities departed1 shows that the entities depart from the Single Server block every 1.5 seconds.

On the right side, the service time for the single server is specified via the signal from the Constant block. The value of the signal is 1.5. The plot labeled Number of entities departed2 shows that the entities depart from the Single Server block every 1.5 seconds.

Specifying Service Time in Infinite Server Block

On the left side of this model, the service time for the Infinite Server block is specified via the `myServiceTime` attribute of every entity arriving at this block. The value of `myServiceTime` is specified as 1.5 seconds. The plot labeled Number of entities departed1 shows that every entity is delayed 1.5 seconds.

On the right side, the service time for the Infinite Server block is specified via the signal from the Constant block. The value of this signal is 1.5 seconds. The plot labeled Number of entities departed2 shows that every entity is delayed 1.5 seconds.

Single Server Block Versus Infinite Server Block

On the left side of this model, the Time-Based Entity Generator block generates entities faster than the Single Server block processes them. The IN

port of the Single Server block is unavailable while the block serves an entity because the Single Server block can store at most one entity at a time.

On the right side, the Time-Based Entity Generator block generates entities faster than the Infinite Server block processes them. The **IN** port of the Infinite Server block is always available because this block can store infinitely many entities at a time.

Start Timer and Read Timer Blocks

This model shows how to use the Start Timer and Read Timer blocks to measure entity-based delays.

The Start Timer block starts a specified timer (equivalent to time-stamping an entity) for each arriving entity.

The Read Timer block reads this timer for each entity and computes the delay that the entity incurred between the two timer blocks.

Release Gate: Value Change Versus Trigger

This model compares the change-in-signal and Trigger modes of the Release Gate block. The Gating Signal scope shows the signal used to open the gates.

The change-in-signal gate opens whenever the gating signal changes. The triggered gate opens whenever the gating signal changes from zero to positive, or vice versa.

Input Switching Using Signal

In this model, the Input Switch block determines which input port to accept entities from by reading values from the **p** input port. These values are shown in the plots. The markers on the lines in these plots indicate times at which entities pass through the Input Switch block.

Output Switching Using Signal

In this model, the Set Attribute block sets the **Index** attribute for the entities passing through. The Output Switch block chooses the entity path as specified

by the value on the **p** input port. The scopes plot the Index attribute to show which entity traveled via each path.

Path Combiner Versus Input Switch

This model contrasts the Path Combiner block and the Input Switch block. The Path Combiner block accepts entities as they arrive at any port. The Input Switch accepts entities only from the port that is selected, where the selection changes in a round-robin fashion with the passage of each entity.

In the upper part of the model, entities from the high-priority entity generator arrive at the Path Combiner block before simultaneously generated entities having a lower priority. The plots from the Path Combiner block reflect the prioritized order in which the entities arrive at and pass through the Path Combiner block.

In the lower part of the model, the Input Switch block permits arrival only at the selected port. The prioritization of entity generation events does not influence the switch block. The plots from the Input Switch block reflect the ascending (round-robin) order in which the entities are selected by the Input Switch block.

Time-Driven and Event-Driven Addition

This model demonstrates the difference between time-based blocks and the event-based operations inside a discrete event subsystem.

The two FIFO Queue blocks independently update their **#n** signals when entities arrive or depart. The model adds the **#n** signals in two ways:

- A Sum block at the top level of the model adds the **#n** signals at times determined by the time-based solver. This block's output is temporarily inaccurate between a change in **#n** and the next time step. Furthermore, this block adds the signals at all time steps, even when neither of the **#n** signals has been updated.
- A Sum block inside a discrete event subsystem adds the **#n** signals precisely upon each update of either signal. The subsystem's awareness of events makes the sum accurate at all times.

Counting Instantaneous Events

In this model, the Instantaneous Event Counting Scope block counts instantaneous signal-based events in a statistical output signal from a gate block, while the Instantaneous Entity Counting Scope block counts instantaneous entity departures from a server block.

Preload Queue with Entities

This model loads a queue with entities at the beginning of the simulation. In the Function-Call Generator block,

- **Sample time** exceeds the simulation stop time.
- **Number of iterations** is the number of entities to create at $T=0$.

Queuing Systems

Queuing system demos include

- “M/M/1 Queuing System” on page 13-7
- “M/D/1 Queuing System” on page 13-9
- “G/G/1 Queuing System and Little’s Law” on page 13-10
- “Single Server Versus N-Server” on page 13-13
- “Single Queue Versus Multiple Queues” on page 13-13
- “A Queuing System with Discouraged Arrivals” on page 13-13
- “Prioritized Queuing Policy Comparison” on page 13-15
- “Preemption Policy Comparison” on page 13-15

M/M/1 Queuing System

This model contains a single-queue single-server system having a single traffic source and an infinite storage capacity. In the notation, the M stands for Markovian; M/M/1 means that the system has a Poisson arrival process, an exponential service time distribution, and one server. Queuing theory provides exact theoretical results for some performance measures of an M/M/1 queuing system and this model makes it easy to compare empirical results with the corresponding theoretical results.

Structure of the Model

The model includes the components listed below.

Component	Description
Time-Based Entity Generator block	Models a Poisson arrival process by generating entities (also known as “customers” in queuing theory).
Exponential Interarrival Time Distribution subsystem	Creates a signal representing the interarrival times for the generated entities. The interarrival time of a Poisson arrival process is an exponential random variable.

Component	Description
FIFO Queue block	Stores entities that have yet to be served.
Single Server block	Models a server whose service time has an exponential distribution.

Results and Displays

The model includes these visual ways to understand its performance:

- Display blocks that show the waiting time in the queue and the server utilization
- A scope showing the number of entities (customers) in the queue at any given time
- A scope showing the theoretical and empirical values of the waiting time in the queue, on a single set of axes. You can use this plot to see how the empirical values evolve during the simulation and compare them with the theoretical value.

Theoretical Results

Queuing theory provides the following theoretical results for an M/M/1 queue with an arrival rate of λ and a service rate of μ :

- Mean waiting time in the queue =

$$\frac{1}{\mu - \lambda} - \frac{1}{\mu}$$

The first term is the mean total waiting time in the combined queue-server system and the second term is the mean service time.

- Utilization of the server = λ/μ

Experimenting with the Model

Move the Arrival Rate Gain slider during the simulation and observe the change in the queue content, shown in the Q Content Scope.

References

[1] Kleinrock, Leonard, *Queueing Systems, Volume I: Theory*, New York, Wiley, 1975.

M/D/1 Queuing System

This model contains a single-queue single-server system having a Poisson arrival process and a server with constant service time. The queue has an infinite storage capacity. In the notation, the M stands for Markovian; M/D/1 means that the system has a Poisson arrival process, a deterministic service time distribution, and one server.

Structure of the Model

The model includes the components listed below.

Component	Description
Time-Based Entity Generator block	Models a Poisson arrival process by generating entities (also known as “customers” in queuing theory).
Exponential Interarrival Time Distribution subsystem	Creates a signal representing the interarrival times for the generated entities. The interarrival time of a Poisson arrival process is an exponential random variable.
FIFO Queue block	Stores entities that have yet to be served.
Single Server block	Models a server having a constant service time.

This model is similar to the M/M/1 Queuing System model, except that the service time in this model is constant.

Results and Displays

The model includes these visual ways to understand its performance:

- A Display block that shows the waiting time in the queue
- A scope showing the number of entities (customers) in the queue at any given time

- A scope showing the theoretical and empirical values of the mean waiting time in the queue, on a single set of axes. You can use this plot to see how the empirical values evolve during the simulation and compare them with the theoretical value.

Theoretical Results

According to queuing theory, the mean waiting time in the queue equals

$$\frac{1}{2(\mu - \lambda)} - \frac{1}{2\mu}$$

where λ is the arrival rate and μ is the service rate. This duration is half the theoretical mean waiting time in the queue for the M/M/1 queuing system with the same arrival rate and service rate.

Experimenting with the Model

Move the Arrival Rate Gain slider during the simulation and observe the change in the mean waiting time using the Display block labeled Queue Waiting Time.

References

[1] Kleinrock, Leonard, *Queueing Systems, Volume I: Theory*, New York, Wiley, 1975.

G/G/1 Queuing System and Little's Law

This model contains a single-queue single-server system in which the interarrival time and the service time are uniformly distributed with fixed means of 1.1 and 1, respectively. The queue has an infinite storage capacity. In the notation, the G stands for a general distribution with a known mean and variance; G/G/1 means that the system's interarrival and service times are governed by such a general distribution, and that the system has one server. You can change the variances of the uniform distributions. You can use this model to examine Little's law.

Structure of the Model

The model includes the components listed below.

Component	Description
Time-Based Entity Generator block	Source of entities (also known as “customers” in queuing theory).
Uniform Distribution for Interarrival Time subsystem	Creates a signal representing the interarrival times for the generated entities. After you set the distribution’s variance using the Arrival Process Variance block, the subsystem computes a uniform random variate with the chosen variance and mean 1 . 1. The variance must be at most $(1 . 1^2)/3$ in order for the subsystem to return only positive values; for details, select the Uniform Distribution for Interarrival Time subsystem and choose Edit > Look Under Mask .
FIFO Queue block	Stores entities that have yet to be served.
Single Server block	Models a server whose service time has a uniform distribution.
Uniform Distribution for Service Time subsystem	Creates a signal representing the service time for entities in the server. This is similar to the subsystem that computes interarrival times, except that the mean service time is 1 and hence the maximum variance is 1/3.

Results and Displays

The model includes these visual ways to understand its performance:

- Display blocks that show the queue workload, average waiting time in the queue, average service time, and server utilization.
- A scope showing the number of entities (customers) in the queue at any given time
- A scope comparing empirical and theoretical ratios, as described in “Little’s Law” on page 13-12

Little's Law

You can use this model to verify Little's law, which states the linear relationship between average queue length and average waiting time in the queue. In particular, the expected relationship is as follows:

$$\text{Average queue length} = (\text{Mean arrival rate})(\text{Average waiting time in queue})$$

The FIFO Queue block computes the current queue length and average waiting time in the queue. The subsystem called Little's Law Evaluation computes the ratio of average queue length (derived from the instantaneous queue length via integration) to average waiting time, as well as the ratio of mean service time to mean arrival time. The two ratios appear on the plot labeled Arrival Rate: Theoretical vs. Simulation Results.

Another way to interpret the equation above is that, given a normalized mean service time of 1, you can use the average waiting time and average queue length to derive the system's arrival rate.

Little's Law Applied to the Server. You can also use this model to verify the linear relationship that Little's law predicts between the server utilization and the average service time. The Single Server block computes the server utilization and average waiting time in the server. Because each entity can depart from the server immediately upon completing service, waiting time is equivalent to service time for the server in this model.

Experimenting with the Model

Move the Arrival Process Variance slider or the Service Process Variance slider during the simulation and observe how the queue content changes. When traffic intensity is high, the average waiting time in the queue is approximately linear in the variances of the interarrival time and service time. The larger the variances are, the longer an entity has to wait, and the more entities are waiting in the system.

References

[1] Kleinrock, Leonard, *Queueing Systems, Volume I: Theory*, New York, Wiley, 1975.

Single Server Versus N-Server

This model compares a single-server queuing system with an N-server system. The single server can serve three times as fast as each of the three servers in the N-Server block.

The displays show that the faster single server achieves smaller delays than the set of three slower servers. The throughput of both systems is the same.

Single Queue Versus Multiple Queues

This model compares a single-queue, multiple-server system with a multiple-queue, multiple-server system. Both systems share the same average service time and achieve the same throughput. However, the single-queue system exhibits shorter delays than the multiple-queue system.

The comparison of these two systems is the classic “queuing for multiple bank tellers” problem and demonstrates that forming a single line for multiple tellers reduces the average delay, compared to forming one line for each teller.

A Queuing System with Discouraged Arrivals

This model contains a queuing system in which feedback influences the arrival rate. As more entities accumulate in the queue-server pair, the entity generator slows its arrival rate. This model illustrates how to use statistical signals from SimEvents blocks not only to report data but also to control the dynamics of the simulation.

Structure of the Model

The model includes the components listed below.

Component	Description
Time-Based Entity Generator block	Source of entities (also known as “customers” in queuing theory).
FIFO Queue block	Stores entities that have yet to be served.
Single Server block	Models a server whose service time has an exponential distribution.

Component	Description
Interarrival Time Distribution with Discouraged Rate subsystem	<p>Computes interarrival times for the entities in the queuing system using a variable arrival rate that depends on the following two statistical signals:</p> <ul style="list-style-type: none"> • The #n output from the FIFO Queue block, which reports the number of entities in the queue • The #n output from the Single Server block, which reports the number of entities in the server (either 0 or 1 for a single server) <p>Specifically, when k customers are in the queue-server pair, the arrival rate is $\alpha/(k+1)$. In this model, $\alpha = 1$.</p>
Exponential Service Time Distribution subsystem	Creates a signal representing the service time for entities in the server.

Results and Displays

The model includes these visual ways to understand its performance:

- A Display block that shows the average waiting time in the queue
- A scope comparing empirical and theoretical values for the average waiting time in the queue

Theoretical Results

According to queuing theory, the mean waiting time in the queue equals

$$\frac{1}{\mu^2(1 - e^{-1/\mu})} - \frac{1}{\mu}$$

where μ is the service rate and the arrival rate is 1 when the queuing system is empty.

Experimenting with the Model

Move the Service Rate Gain slider during the simulation and observe the change in the queue waiting time, shown in the Waiting Time Comparison Scope.

References

[1] Kleinrock, Leonard, *Queueing Systems, Volume I: Theory*, New York, Wiley, 1975.

Prioritized Queuing Policy Comparison

This model uses the Priority Queue block to implement queuing policies based on service time. Compare the performance of two prioritized queuing policies with the first-in, first-out (FIFO) queuing policy.

Preemption Policy Comparison

This model compares three policies for working with entities that have been preempted from a server:

- Preemption with discard — A preempted entity cannot reenter the server
- Preemption with restart — A preempted entity reenters the server and requires the same service time that it required the first time
- Preemption with resume — A preempted entity reenters the server and requires only the residual service time that remained at the time of preemption

Application Demos

Application-specific demos include

- “Shared Access Communications Media” on page 13-16
- “Dynamic Voltage Scaling Using Online Gradient Estimation” on page 13-16
- “Comparison of Routing Policies” on page 13-19
- “F-14 Flight Control Over a Network” on page 13-19
- “Selective-Repeat Automatic Repeat Request” on page 13-19
- “Tank Filling Station” on page 13-20
- “Astable Multivibrator Circuit” on page 13-21

Shared Access Communications Media

This model is an abstraction of a communications bus or shared media. Each transmitter that uses the bus has exclusive access until the use is completed. In this model, transmit node 1 is sending to receive node 2, and vice versa. Background traffic from other nodes in the network is modeled in the aggregate. You can adjust the aggregate traffic for the other nodes using a slider in the Average Arrival Rate block.

Using the plots, you can monitor bus/media utilization, queue lengths, average waiting times, and the number of dropped packets.

The bus/media abstraction is enclosed in a subsystem to facilitate copying into another model. For example, you can explore further by modeling the delays incurred over a bus between two processors.

Dynamic Voltage Scaling Using Online Gradient Estimation

This model simulates an AT90S8535 microcontroller that uses a dynamic voltage scaling (DVS) feature to adjust the input voltage depending on the workload. By lowering the input voltage when the workload is low, the microcontroller reduces energy consumption while guaranteeing quality of service. The DVS controller is based on an online gradient estimation

technique called infinitesimal perturbation analysis (IPA). In a *single* simulation of a parameterized system — not the large number of simulations required by a traditional finite-difference approach — IPA can provide sensitivity information that yields a first-order approximation of the system’s performance metrics as a function of the parameters.

Applying IPA to the Controller

The performance metric to minimize is the average cost per job, given by

$$J(\theta) = w P(\theta) + S(\theta) = w c_2 [V_t / (1 - c_1/\theta)]^2 + S(\theta)$$

where

- θ is the average service time of a job, which is a function of the input voltage V . That is, finding the optimal value of θ also yields the optimal value of V .
- w is a weighting constant.
- P is the average energy consumption of a job in Joules.
- S is the average system time for jobs, which measures quality of service. This model uses an M/M/1 queuing system, so a closed-form expression for $S(\theta)$ provides a way to compare the IPA results in the simulation with theoretical results.
- c_1 and c_2 are device-dependent constants.
- V_t is the device’s minimum input voltage.

To find a value of θ for which $dJ/d\theta$ is 0, this model uses a gradient method with constant step size $\Delta = 10^{-5}$. The k th iteration of the optimization, which occurs upon the departure of the k th job, uses the estimate θ_k to produce

$$\theta_{k+1} = \theta_k - \Delta \cdot \left. \frac{dJ}{d\theta} \right|_{\theta_k} = \theta_k - \Delta \cdot \left(\frac{2wc_1c_2V_t^2\theta_k}{(\theta_k - c_1)^3} + \text{IPA estimation of } \left. \frac{dS}{d\theta} \right|_{\theta_k} \right)$$

To learn about the IPA estimation of $dS/d\theta$, see the works listed in “References” on page 13-18.

Structure of the Model

The model includes these components:

- Job Arrivals section — source of jobs that form the workload
- FIFO Queue, Single Server, and other blocks in the blue section — queuing system for jobs
- DVS Optimizer subsystem — uses the queue length, θ_k value, service time for the latest job, and total number of jobs to compute θ_{k+1} and the corresponding updated input voltage

Results and Displays

The model includes these visual ways to understand its performance:

- A dynamic plot showing how the DVS controller varies the voltage during the simulation to reduce the average cost per job
- A Display block that shows the average service time for jobs
- A Display block that shows the corresponding input voltage

To experiment, try changing the value of the Avg Interarrival Time block before running the simulation.

References

- [1] Cassandras, Christos G., and Stéphane Lafortune, *Introduction to Discrete Event Systems*, Boston, Kluwer Academic Publishers, 1999.
- [2] Weiser M., B. Welch, A. Demers, and S. Shenker, “Scheduling for reduced CPU energy,” in *Proceedings of the 1st Symposium on Operating Systems Design and Implementation*, pp. 13-23, 1994.
- [3] Zhuang, S., and C. G. Cassandras, “Optimal Dynamic Voltage Scaling for Wireless Sensor Networks with Real-Time Constraints,” in *Proc. of SPIE Conf. on Intelligent Systems in Design and Manufacturing VI*, Oct. 2005.

Comparison of Routing Policies

This model runs three similar queuing systems in parallel under three routing policies

Policy	Description
Shortest Queue	Each entity advances to the shortest queue
Probabilistic	Each entity is randomly routed to a queue
Round robin	Entities alternate between the two queues

The plot compares the average system time (end-to-end delay) to demonstrate the optimality of the Shortest Queue policy when all service time distributions are exponential. If this assumption is violated, a Shortest Queue policy is not necessarily optimal.

F-14 Flight Control Over a Network

This model modifies the Simulink F-14 flight control demo model by adding a network between the controller and the plant that models the aircraft dynamics. Double-click the Shared Access Network block to view the SimEvents blocks that model the network delays. The network model is similar to the one in the Shared Access Communications Media demonstration.

In the subsystem, double-click the Average Arrival Rate block (red) and use the slider to change the background traffic level. You can examine the impact of background traffic on the accuracy of the flight control loop. At low traffic volume, the delays have little effect. However, at moderate volume, the increased delay causes the control loop to oscillate when the control stick is moved quickly. At higher volume, the delays are sufficient to render the control loop ineffective.

Selective-Repeat Automatic Repeat Request

This model simulates an ideal selective-repeat ARQ system. Packets are created, transmitted over a channel and affected by errors with a probability determined by the Forward Channel. The receiver uses the return channel to send an acknowledgement (ACK) or nonacknowledgement (NAK) message back to the transmitter, depending on whether the packet was received

correctly. The transmitter retransmits only those packets that were not received. The return channel is assumed to be error-free. Scopes display

- The packet IDs of the transmitted and received packets
- The presence of errors in the received packets, where a `crc_check` attribute value of 2 indicates an error and a value of 1 indicates no error
- The reconstructed packet stream
- The theoretical and empirical times required for correct reception

Tank Filling Station

This model contains a tank whose fluid level decreases as fluid flows out of the tank and increases when trucks fill the tank.

A controller aims to keep the tank’s fluid level between a lower level, L, and an upper level, U. The controller responds to varying tank levels by changing its state and the tank outflow rate, as shown below. The number M satisfies $L < M < U$.

Tank Level Threshold	State Change	Outflow Rate
U	Set OVF = 1 (overflow)	20
L	Set URF = 1 (underflow)	0
M	When crossing M from above, set OVF = 0	10
	When crossing M from below, set URF = 0	10


Trucks arrive to fill the tank and wait in a queue. Each truck has attributes that specify the tank inflow rate and the truck’s service time. The tank inflow rate equals

- The truck’s inflow attribute, while the truck is at the filling station.
- Zero, when no truck is present.

Plots display the short-term and long-term views of the system.

Astable Multivibrator Circuit

This model demonstrates how to model a hybrid system containing event-driven and time-driven behavior. The system in this model is the simple astable multivibrator circuit shown below. An astable multivibrator is a two-stage switching circuit that generates a square wave with no external triggering pulse. The circuit switches between its two stable states, remaining in each state for a duration that depends on the discharging of the capacitive circuit.



To focus on the integration of the time-driven and event-driven aspects, the model makes these simplifying assumptions:

- Transistors change their on/off states instantaneously.
- The turn-on voltage for transistors is 0.6 V.

Circuit Operation

- When supply voltage, V_{CC} , is applied, one transistor conducts more than the other due to imbalance in resistance and capacitance values.
- Suppose that Q1 is on and Q2 is off. The voltage at the output of Q1 is approximately V_{cesat} , which is zero.
- Immediately after Q1 turns on, C1 charges with the time constant $R1C1$ toward the supply voltage through R1. That is, V_{be} of Q2 increases asymptotically towards V_{CC} .
- When V_{be} of Q2 crosses the turn-on voltage, Q2 turns on and voltage at the output of Q2 drops to approximately zero.
- Immediately after Q2 turns on, V_{be} of Q1 falls due to capacitive coupling between the collector of Q2 and the base of Q1. As a result, Q1 turns off.
- Immediately after Q2 turns on, C2 charges with the time constant $R2C2$ toward the supply voltage through R2. That is, V_{be} of Q1 increases asymptotically towards V_{CC} .
- When V_{be} of Q1 crosses the turn-on voltage, Q1 turns on again and the process repeats.

Structure of the Demo

The model has a symmetric arrangement of blocks representing the two branches in the circuit. Each branch in the block diagram has a feedback connection to the other branch. Each branch includes these components:

- Enabled subsystem — models the variation in the base voltage when a transistor switches its on/off state. In this subsystem, a Transfer Fcn block models the charging and discharging of the RC circuit.
- Switching Logic subsystem — switches based on the potential barrier. In this subsystem, the Saturation block enables the detection of level crossings, while the Embedded MATLAB Function block contains the logic.
- Signal Latch block from SimEvents — allows the transistor V_{ce} or the input to the RC circuit to have two values at one time. The sequence of these values is controlled by the block's event-based operation, which enables the feedback to occur without delay.

Functions — Alphabetical List

simeventsconfig

Purpose Change settings of Simulink model

Syntax
simeventsconfig
simeventsconfig(sys, 'des')
simeventsconfig(sys, 'hybrid')

Description simeventsconfig changes settings of the current system to values appropriate for discrete-event simulation (DES) modeling.

simeventsconfig(sys, 'des') changes settings of the specified system to values appropriate for discrete-event simulation modeling. In particular, this syntax changes the solver to a variable-step discrete solver.

simeventsconfig(sys, 'hybrid') changes settings of the specified system to values appropriate for modeling systems that combine time-driven and discrete-event-driven behavior. In particular, this syntax changes the solver to a continuous solver, ode45.

This function assigns the following Simulink settings.

Parameter	Setting for DES Systems	Setting for Hybrid Systems
SolverName	VariableStepDiscrete	ode45
SolverType	Variable-step	Variable-step
MaxStep	inf	inf
SaveTime	off	off
SaveOutput	off	off
AlgebraicLoopMsg	error	error
SolverPrmCheckMsg	none	none

For information about the settings in the table, see “Model Parameters” in the Simulink documentation. For information about discrete and continuous solvers, see “Solvers”.

See Also simeventsstartup

Purpose	Open SimEvents library
Syntax	simeventslib
Description	simeventslib opens the main SimEvents library.

simeventsstartup

Purpose Default Simulink model settings for SimEvents

Syntax `simeventsstartup('des')`
`simeventsstartup('hybrid')`

Description `simeventsstartup('des')` changes the default Simulink model settings to values appropriate for discrete-event simulation (DES) modeling.

`simeventsstartup('hybrid')` changes the default Simulink model settings to values appropriate for modeling systems that combine time-driven and discrete-event-driven behavior.

This function assigns the following Simulink settings.

Parameter	Setting for DES Systems	Setting for Hybrid Systems
SolverName	VariableStepDiscrete	ode45
SolverType	Variable-step	Variable-step
MaxStep	inf	inf
SaveTime	off	off
SaveOutput	off	off
AlgebraicLoopMsg	error	error
SolverPrmCheckMsg	none	none

For information about the settings in the table, see “Model Parameters” in the Simulink documentation. For information about discrete and continuous solvers, see “Solvers”.

Changes in default Simulink model settings apply to new models that you create later in the MATLAB session, but not to previously created models. To change the model settings of a previously created model, use `simeventsconfig`.

Tip To install these model settings each time you start MATLAB, invoke `simeventsstartup` from your `startup.m` file.

See Also

`simeventsconfig`, `startup`

Blocks — By Category

Generators (p. 15-2)	Generate entities, events, and signals
SimEvents Sinks (p. 15-3)	View and export data
Attributes (p. 15-3)	Manage data attached to entities
Queues (p. 15-3)	Store entities in queue
Servers (p. 15-4)	Delay entities by service time
Routing (p. 15-4)	Design entity paths
Gates (p. 15-4)	Regulate entity admission
SimEvents Ports and Subsystems (p. 15-5)	Control timing using discrete event subsystems
Timing (p. 15-5)	Compute time that entities spend in region
Probes (p. 15-5)	Report information about the simulation
Event Translation (p. 15-6)	Convert event types

Generators

Entity Generators (p. 15-2)	Create entities
Event Generators (p. 15-2)	Create function-call events
Signal Generators (p. 15-2)	Create numerical signals

Entity Generators

Event-Based Entity Generator	Generate entity upon signal-based event or function call
Time-Based Entity Generator	Generate entities using intergeneration times from signal or statistical distribution

Event Generators

Entity-Based Function-Call Event Generator	Generate function call events corresponding to entities
Signal-Based Function-Call Event Generator	Generate function-call events in response to signal-based events

Signal Generators

Event-Based Random Number	Generate random numbers from specified distribution, parameters, and initial seed
---------------------------	---

SimEvents Sinks

Attribute Scope	Plot data from attribute of arriving entities
Discrete Event Signal to Workspace	Write event-based signal to workspace
Entity Sink	Accept or block entities
Instantaneous Entity Counting Scope	Plot entity count versus time
Instantaneous Event Counting Scope	Plot event count versus time
Signal Scope	Plot data from signal
X-Y Attribute Scope	Plot data from two attributes of arriving entities
X-Y Signal Scope	Plot data from two signals

Attributes

Get Attribute	Output value of entity's attribute
Set Attribute	Assign data to entity

Queues

FIFO Queue	Store entities in sequence for undetermined length of time
LIFO Queue	Store entities in stack for undetermined length of time
Priority Queue	Store entities in sorted sequence for undetermined length of time

Servers

Infinite Server	Delay any number of entities for period of time
N-Server	Serve up to N entities for period of time
Single Server	Serve one entity for period of time

Routing

Input Switch	Accept entities from selected entity input port
Output Switch	Select entity output port for departure
Path Combiner	Merge entity paths
Replicate	Output copies of entity

Gates

Enabled Gate	Permit entity arrivals only when control signal is positive
Release Gate	Permit one pending entity to arrive when event occurs
Signal Latch	Write input signal value to memory and read memory to output signal upon events

SimEvents Ports and Subsystems

Discrete Event Inport	Input port for Discrete Event Subsystem block
Discrete Event Outport	Provide output port for Discrete Event Subsystem block
Discrete Event Subsystem	Subsystem to be executed upon signal-based events
Subsystem Configuration	Configuration for Discrete Event Subsystem block

Timing

Read Timer	Report statistical data about named timer associated with arriving entities
Start Timer	Associate named timer to each arriving entity independently and start timing

Probes

Entity Departure Counter	Count departures and write result to signal port and/or attribute
--------------------------	---

Event Translation

Entity Departure Event to
Function-Call Event

Convert entity departure event into
one or two function calls

Signal-Based Event to Function-Call
Event

Convert signal-based events into
function calls

Blocks — Alphabetical List

Attribute Scope

Purpose Plot data from attribute of arriving entities

Library SimEvents Sinks

Description



This block creates a plot using data from an attribute of arriving entities. Use the **Y attribute name** parameter to specify which attribute to plot along the vertical axis.

Use the **Enable entity OUT port** option to choose whether the entity advances to a subsequent block or whether the block absorbs the arriving entity.

The **Plot type** parameter on the **Plotting** tab determines whether and how the block connects the points that it plots. For details, see “Connections Among Points in Plots” on page 9-4.

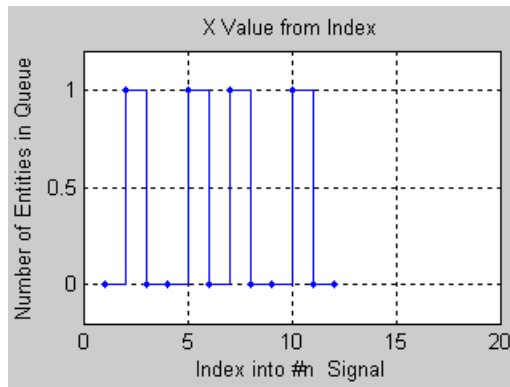

Selecting Data for the Horizontal Axis

Use the **X value from** parameter to select the type of data for the horizontal axis. The table below describes the choices.

Source of X Data	Description of Plot
Event time	Plot of the specified attribute versus simulation time.
Index	Plot of the successive values of the specified attribute against a horizontal axis that represents the index of the values. The first entity's attribute value has an index of 1, the second entity's attribute value has an index of 2, and so on. For example, you might use this option when multiple entities might arrive simultaneously, to help determine the exact sequence among the simultaneous attribute values.

The figures below illustrate the different sources of data for the horizontal axis. The plots look similar, except that the second plot has

uniform horizontal spacing rather than time-based spacing between successive points.



Attribute Scope

Ports

Entity Input Ports

Label	Description
IN	Port for arriving entities, whose attributes contain the data to plot.

Entity Output Ports

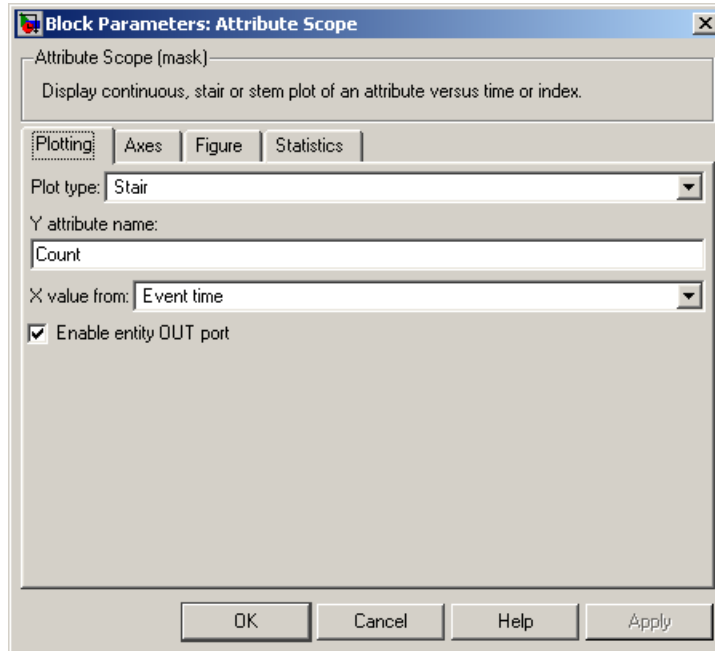
Label	Description
OUT	Port for departing entities. This port appears only if you select Enable entity OUT port .

Signal Output Ports

Label	Description
#a	Number of entities that have arrived at the block since the start of the simulation.

Dialog Box

Plotting Tab



Plot type

The presentation format for the data. See “Connections Among Points in Plots” on page 9-4 for details.

Y attribute name

Name of the attribute to plot along the vertical axis.

X value from


Source of data for the plot’s horizontal axis. See “Selecting Data for the Horizontal Axis” on page 16-2 for details.

Enable entity OUT port

Causes the block to have an entity output port labeled **OUT**, through which the arriving entity departs. If you clear this box, the block absorbs arriving entities.

Attribute Scope

Axes Tab



Initial X axis lower limit, Initial X axis upper limit

The interval shown on the X axis at the beginning of the simulation. The interval might change from this initial setting due to zooming, autoscaling, or the **If X value is beyond limit** setting.

If X value is beyond limit

Determines how the plot changes if one or more X values are not within the limits shown on the X axis. For details, see “Varying Axis Limits Automatically” on page 9-5.

Initial Y axis lower limit, Initial Y axis upper limit

The interval shown on the Y axis at the beginning of the simulation. The interval might change from this initial setting

due to zooming, autoscaling, or the **If Y value is beyond limit** setting.

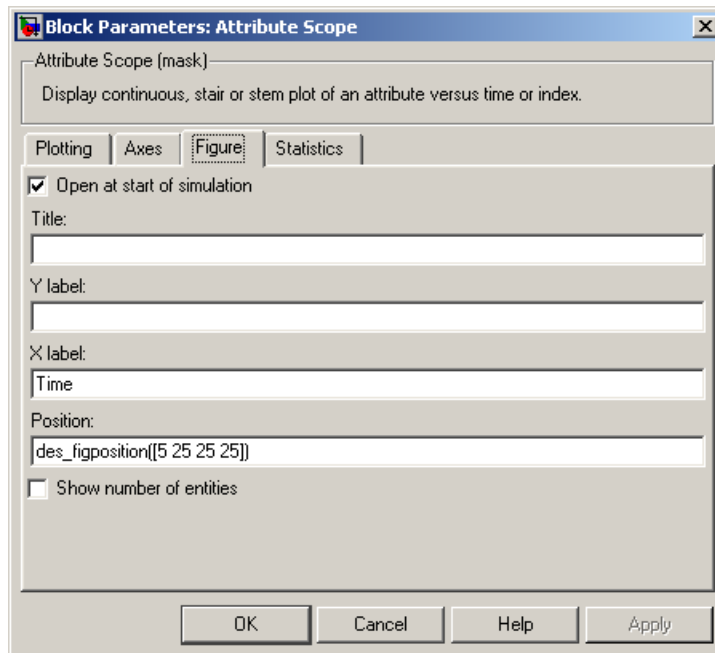
If Y value is beyond limit

Determines how the plot changes if one or more attribute values are not within the limits shown on the Y axis. For details, see “Varying Axis Limits Automatically” on page 9-5.

Show grid

Toggles the grid on and off.

Figure Tab



Open at start of simulation

Selecting this option causes the plot window to open when you start the simulation. If you clear this box, you can open the plot

Attribute Scope

window while the simulation is running by double-clicking the block icon.

Title

Text that appears as the title of the plot, above the axes.

Y label

Text that appears to the left of the vertical axis.

X label

Text that appears below the horizontal axis.

Position


A four-element vector of the form [left bottom width height] specifying the position of the scope window. (0,0) is the lower left corner of the display.

Show number of entities

Displays the number of plotted points using an annotation in the plotting window.

Statistics Tab

These parameters determine whether certain ports output data throughout the simulation, output data only when you stop or pause the simulation, or are omitted from the block. For descriptions of the affected ports, see the “Signal Output Ports” table above.



Number of entities arrived

Controls the presence and behavior of the signal output port labeled **#a**.

Examples

- “Example: Round-Robin Approach to Choosing Inputs” in the Getting Started documentation
- “Example: Setting Attributes” on page 1-15

See Also

X-Y Attribute Scope, Signal Scope, Chapter 9, “Plotting Data”, “Accessing Attributes of Entities” on page 1-20

Discrete Event Inport

Purpose Input port for Discrete Event Subsystem block

Library SimEvents Ports and Subsystems


Description Discrete Event Inport blocks are the links from outside a discrete event subsystem into the subsystem. By default, the names of Discrete Event Inport blocks appear in the subsystem window as **Din**, **Din1**, **Din2**, etc. A discrete event inport represents a real scalar signal of type double.

1

Copying and pasting Discrete Event Inport blocks is supported, but duplicating them is not.

To create a subsystem using the Discrete Event Subsystem block, see “Setting Up Signal-Based Discrete Event Subsystems” on page 8-11.

Dialog Box



Execute subsystem upon signal-based events

If you select this option, Simulink executes the subsystem when a qualifying signal-based event occurs in the signal corresponding to this inport block. If you clear this option, the subsystem reads the signal upon execution but does not respond to its events.

Type of signal-based event

Determines the type of event that is a qualifying event in the signal corresponding to this inport block. This field appears only if you select **Execute subsystem upon signal-based events**.

Type of change in signal value, Trigger type

The type of change in the control signal's value, or the type of trigger, that further restricts the event type specified in **Type of signal-based event**. This field appears only if you set **Type of signal-based event** to Change in signal or Trigger.

Specify event priority for executing subsystem

Select this option to control the sequencing of the subsystem's execution event in response to the signal corresponding to this inport block, relative to other simultaneous events.

Subsystem execution event priority

The priority of the subsystem's execution event, relative to other simultaneous events. This field appears only if you select **Specify event priority for executing subsystem**.

Examples

See "Examples Using Discrete Event Subsystem Blocks" on page 8-16.

See Also

Discrete Event Subsystem, Discrete Event Outport, Chapter 8, "Controlling Timing Using Subsystems"

Discrete Event Output

Purpose	Provide output port for Discrete Event Subsystem block
Library	SimEvents Ports and Subsystems
Description	<p>Discrete Event Output blocks are the links from a discrete event subsystem to a destination outside the subsystem. By default, the names of Discrete Event Output blocks appear in the subsystem window as Dout, Dout1, Dout2, etc.</p> <p>To create a subsystem using the Discrete Event Subsystem block, see “Setting Up Signal-Based Discrete Event Subsystems” on page 8-11.</p>
Examples	See “Examples Using Discrete Event Subsystem Blocks” on page 8-16.
See Also	Discrete Event Subsystem, Discrete Event Inport, Chapter 8, “Controlling Timing Using Subsystems”



Discrete Event Signal to Workspace

Purpose Write event-based signal to workspace

Library SimEvents Sinks

Description This block writes its input to a structure or array in the base MATLAB workspace when the simulation stops or pauses. This block is similar to the To Workspace block in the Simulink Sinks library but is tailored for use with event-based signals.

Output Format

The **Save format** parameter determines the output format. The Structure With Time output format is most appropriate for event-based signals because it indicates when the signal assumes each value. Updates of event-based signals are typically aperiodic.

To convert a structure with time into a two-column matrix containing times in the first column and signal values in the second column, use an assignment like the one below. In place of `simout`, use the name specified in this block's **Variable name** parameter.

```
times_values = [simout.time, simout.signals.values];
```

For descriptions of all output formats, see the reference page for the To Workspace block in the Simulink documentation.

Comparison with To Workspace Block

This block can detect zero-duration values of the input signal, as well as signal updates that do not necessarily correspond to time steps determined by time-based dynamics.

This block does not support data types other than `double`, and has no **Sample time** parameter because event-based signals do not have a true sample time.

The simulation times at which this block records data is typically unrelated to the variable that a model creates if you select **Time** in the **Save to workspace** section of the **Data Import/Export** tab of the Configuration Parameters dialog box. By default, this option is selected

Discrete Event Signal to Workspace

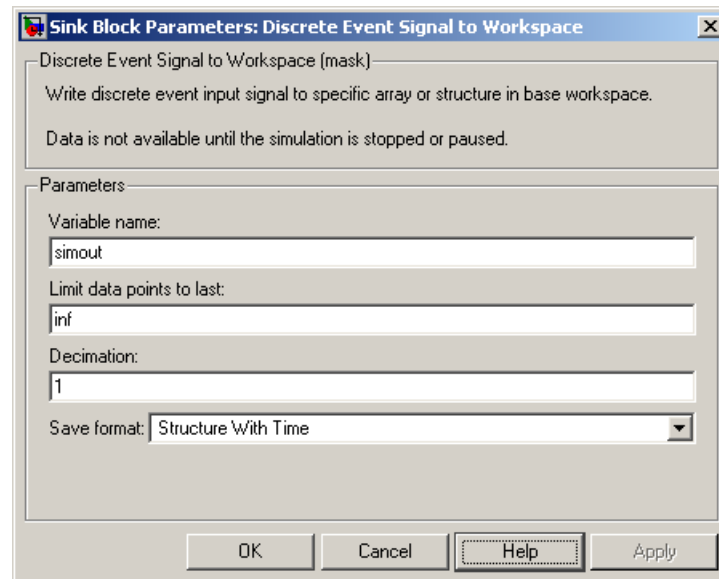
and the variable is called tout. The simeventsconfig function clears the time logging option to avoid confusion between the time steps listed in tout and the update times of event-based signals in the simulation.

Ports

This block has one signal input port for the signal to write to the workspace.

The block has no entity ports, and no signal output port.

Dialog Box



Variable name

The name of the structure or array that holds the data.

Limit data points to last

The maximum number of input samples to be saved.

Decimation

A positive integer, n , that specifies the decimation factor. The block ignores the first $n-1$ out of every n input samples.

Save format

Format in which to save simulation output to the workspace.
The recommended format for event-based signals is Structure
With Time.

Examples

- “Example: Sending Queue Length to the Workspace” on page 3-23
- “Example: Observing Service Completions” on page 2-32

See Also

To Workspace, “Sending Data to the MATLAB Workspace” on page 3-23

Discrete Event Subsystem

Purpose Subsystem to be executed upon signal-based events

Library SimEvents Ports and Subsystems

Description



This block represents a subsystem of the system that contains it, and is configured so that Simulink executes the subsystem upon signal-based events that you specify in one or more Discrete Event Inport blocks within the subsystem. Inputs from the upper level to blocks inside the subsystem are attached to the **Din**, **Din1**, **Din2**, etc., ports on the Discrete Event Subsystem block. Inputs must be real scalar signals of type double. Outputs, if any, from blocks inside the subsystem to the upper level are attached to the **Dout**, **Dout1**, **Dout2**, etc., ports on the Discrete Event Subsystem block.

The number of input ports drawn on the Discrete Event Subsystem block's icon corresponds to the number of nonduplicate Discrete Event Inport blocks inside the subsystem. Similarly, the number of output ports drawn on the block corresponds to the number of Discrete Event Outport blocks inside the subsystem.

Note This block is compatible only with inports and outports from the SimEvents Ports and Subsystems library. To add inports or outports to the Discrete Event Subsystem window, either copy new ports from the SimEvents Ports and Subsystems library (not the Simulink Ports & Subsystems library), or copy and paste the ports that are in the Discrete Event Subsystem window by default.

To create a subsystem using the Discrete Event Subsystem block, see “Setting Up Signal-Based Discrete Event Subsystems” on page 8-11. A discussion of discrete event subsystems and examples using this block are in Chapter 8, “Controlling Timing Using Subsystems”. To view the contents of the subsystem, double-click the Discrete Event Subsystem block; the Model Explorer tool does not show the contents.

“Block execution” in this documentation is shorthand for “block methods execution.” Methods are functions that Simulink uses to solve a model.

Blocks are made up of multiple methods. For details, see “Block Methods” in the Simulink documentation.

Ports

Signal Input Ports

Label	Description
Din , Din1 , Din2 , etc.	Signals that serve as inputs to the blocks in the subsystem. At least one input signal must be present, but others are optional. Each port appears only if the subsystem contains a Discrete Event Inport block of the same name.

Signal Output Ports

Label	Description
Dout , Dout1 , Dout2 , etc.	Optional signals that serve as outputs from the blocks in the subsystem. Each port appears only if the subsystem contains a Discrete Event Outport block of the same name.

Examples

See “Examples Using Discrete Event Subsystem Blocks” on page 8-16.

See Also

Discrete Event Inport, Discrete Event Outport, Chapter 8, “Controlling Timing Using Subsystems”

Enabled Gate

Purpose Permit entity arrivals only when control signal is positive

Library Gates

Description This block represents a gate that is open whenever the control signal at the **en** input port is positive, and closed whenever the signal is zero or negative. By definition, an open gate permits entity arrivals as long as the entities would be able to advance immediately to the next block, while a closed gate forbids entity arrivals. The **en** signal is a numerical signal of type `double`. Because the signal can remain positive for a time interval of arbitrary length, an enabled gate can remain open for a time interval of arbitrary length. The length can be zero or a positive number.

Ports

Entity Input Ports

Label	Description
IN	Port for arriving entities.

Signal Input Ports

Label	Description
en	The gate is open whenever this signal is positive.

Entity Output Ports

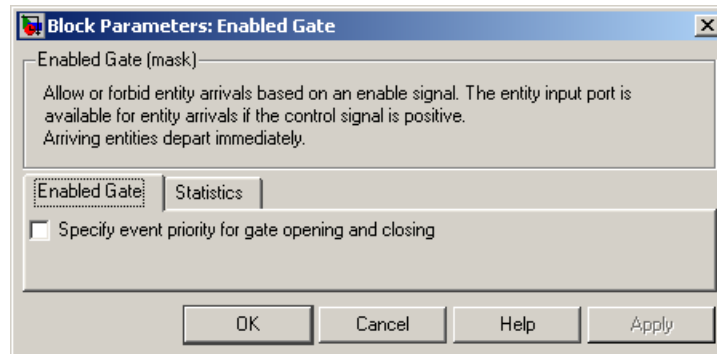
Label	Description
OUT	Port for departing entities.

Signal Output Ports

Label	Description	Time of Update When Statistic Is On
#d	Number of entities that have departed from this block since the start of the simulation.	After entity departure

Dialog Box

Enabled Gate Tab



Specify event priority for gate opening and closing

Select this option to control the processing sequence of the gate-opening and gate-closing events, relative to other simultaneous events.

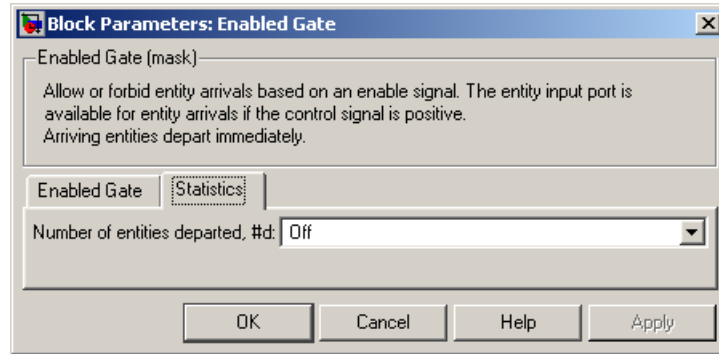
Event priority

The priority of the gate-opening and gate-closing events, relative to other simultaneous events. Gate opening and closing are distinct events that share the same event priority. This field appears only if you select **Specify event priority for gate opening and closing**.

Enabled Gate

Statistics Tab

These parameters determine whether certain ports output data throughout the simulation, output data only when you stop or pause the simulation, or are omitted from the block. For descriptions of the affected ports, see the “Signal Output Ports” table above.



Number of entities departed

Controls the presence and behavior of the signal output port labeled #d.

Examples

- “Example: Controlling Joint Availability of Two Servers” on page 7-4
- “Example: First Entity as a Special Case” on page 7-10

See Also

Release Gate, Chapter 7, “Regulating Arrivals Using Gates”

Purpose Count departures and write result to signal port and/or attribute

Library Probes

Description



This block computes the number of entities that have departed from this block since the start of the simulation or since the last reset, whichever occurred later. The block writes this number to a signal output port and/or an attribute of each departing entity. The departing entity is included in the count.

Resetting the Counter During the Simulation

To reset the entity count to zero based on a signal-based event, use one of the procedures below.

Resetting the Counter upon Trigger Edges

- 1 Set **Reset counter upon** to Trigger from port `tr`.
- 2 Set **Trigger type** to Rising, Falling, or Either to indicate which zero crossings cause the counter to reset.
- 3 If you want to specify an explicit priority for the reset event, then select **Specify event priority for counter reset** and enter the priority using the **Event priority** parameter. Though slightly slower, this option lets you control the processing sequence of simultaneous events.
- 4 Click **OK** or **Apply**. The block now has a signal input port labeled `tr`.
- 5 Connect a signal to the `tr` input port. During the simulation, the block resets its counter whenever the signal at the `tr` port satisfies your specified event criteria.

Resetting the Counter upon Value Changes

- 1 Set **Reset counter upon** to Change in signal from port `vc`.

Entity Departure Counter

- 2 Set **Type of change in signal value** to Rising, Falling, or Either to indicate which value changes cause the counter to reset.
- 3 If you want to specify an explicit priority for the reset event, then select **Specify event priority for counter reset** and enter the priority using the **Event priority** parameter. Though slightly slower, this option lets you control the processing sequence of simultaneous events.
- 4 Click **OK** or **Apply**. The block now has a signal input port labeled **vc**.
- 5 Connect a signal to the **vc** input port. During the simulation, the block resets its counter whenever the signal at the **vc** port satisfies your specified event criteria.

Ports

Entity Input Ports

Label	Description
IN	Port for arriving entities.

Signal Input Ports

Label	Description
tr	When the signal at this port satisfies the specified trigger criteria, the block resets its internal counter and the #d output signal to zero. This port appears only if you set Reset counter upon to Trigger from port tr.
vc	When the signal at this port satisfies the specified value-change criteria, the block resets its internal counter and the #d output signal to zero. This port appears only if you set Reset counter upon to Change in signal from port vc.

Entity Departure Counter

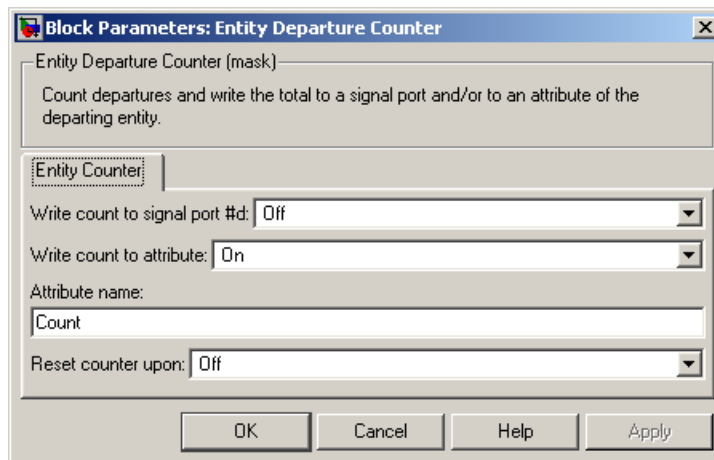
Entity Output Ports

Label	Description
OUT	Port for departing entities.

Signal Output Ports

Label	Description	Time of Update When Statistic Is On
#d	Number of entities that have departed from this block since the start of the simulation or since the last reset.	After entity departure

Dialog Box



Write count to signal port #d

Controls the presence and behavior of the signal output port labeled #d. This parameter determines whether the block outputs the entity count through a signal output port throughout the

Entity Departure Counter

simulation, only when you stop or pause the simulation, or not at all.

Write count to attribute

If you set this to On, the block assigns the entity count to the attribute specified in the **Attribute name** parameter. The entity must already possess an attribute by that name.

Attribute name

The name of the attribute the block uses to record the entity count. This field appears only if you set **Write count to attribute** to On.

Reset counter upon

Determines whether the block resets its internal counter to zero based on trigger edges, based on changes in signal values, or not at all.

Trigger type

Determines whether rising, falling, or either type of trigger edge causes the counter to reset. This field appears only if you set **Reset counter upon** to Trigger from port tr.

Type of change in signal value

Determines whether rising, falling, or either type of value change causes the counter to reset. This field appears only if you set **Reset counter upon** to Change in signal from port vc.

Specify event priority for counter reset

Select this option to control the sequencing of the counter's reset event, relative to other simultaneous events. This field appears only if you set **Reset counter upon** to Trigger from port tr or Change in signal from port vc.

Event priority

The priority of the counter's reset event, relative to other simultaneous events. This field appears only if you set **Reset counter upon** to either Trigger from port tr or Change in signal from port vc, and then select **Specify event priority for counter reset**.

Examples

- “Example: Setting Attributes” on page 1-15
- “Example: Resetting a Counter After a Transient Period” on page 1-23
- “Stopping Based on Entity Count” on page 10-33

See Also


Instantaneous Entity Counting Scope, “Counting Entities” on page 1-21

Entity Departure Event to Function-Call Event

Purpose Convert entity departure event into one or two function calls

Library Event Translation

Description This block converts an entity departure event into one or two function calls that you can use to invoke function-call subsystems, Stateflow blocks, or other blocks that accept function-call inputs. The block can suppress its output under certain conditions.



Criteria for Generating Function Calls

The primary criterion is the departure, or imminent departure, of an entity from the block. You can choose whether the block issues the function call before or after the departure.

To issue up to two function calls per event, select **Generate optional function call f2 after function call f1**. If you configure the block to issue the **f1** function call before the entity departure, then you can independently choose whether the block issues the **f2** function call before or after that departure.

To make the **f1** or **f2** output function call contingent upon a secondary criterion, select **Suppress function call f1 if enable signal e1 is not positive** or **Suppress function call f2 if enable signal e2 is not positive**. The block acquires an additional signal input port, labeled **e1** or **e2**, to which you connect a control signal. If the control signal is zero or negative when the block is about to issue the function call, then the block suppresses the function call. The **e1** and **e2** ports operate independently of each other as secondary criteria for their respective function-call output ports.

Entity Departure Event to Function-Call Event

Ports

Entity Input Ports

Label	Description
IN	Port for arriving entities.

Signal Input Ports

Label	Description
e1	When this signal is 0 or negative, the block does not issue a function call at the f1 output port. This input port appears only if you select Suppress function call f1 if enable signal e1 is not positive .
e2	When this signal is 0 or negative, the block does not issue a function call at the f2 output port. This input port appears only if you select Suppress function call f2 if enable signal e2 is not positive .

Entity Output Ports

Label	Description
OUT	Port for departing entities.

Entity Departure Event to Function-Call Event

Signal Output Ports


Label	Description	Time of Update When Port Is Present	Order of Update
f1	Function-call signal, possibly contingent on e1 input signal.	Before or after entity departure, depending on Timing of function call f1 parameter	1
f2	Function-call signal, possibly contingent on e2 input signal.	Before entity departure if both Timing of function call... parameters are set to Before entity departure; otherwise, after entity departure	2
#d	Number of entities that have departed from this block since the start of the simulation.	After entity departure	4
#f1	Number of function calls the block has generated at the f1 port during the simulation.	After entity departure	3
#f2	Number of function calls the block has generated at the f2 port during the simulation.	After entity departure	3

Output signals having the same number in the Order of Update column above are updated in an arbitrary sequence relative to each other; you should not rely on a specific sequence for your simulation results.

Entity Departure Event to Function-Call Event

Dialog Box

Function Call Tab



Timing of function call f1

Determines whether the **f1** function call occurs before or after the entity departure event.

Suppress function call f1 if enable signal e1 is not positive

Selecting this option causes **f1** function calls to be contingent upon a positive value at the **e1** signal input port.

Generate optional function call f2 after function call f1

Selecting this option causes the block to issue a function call at the optional **f2** output port when appropriate criteria are satisfied.

Timing of function call f2

Determines whether the **f2** function call occurs before or after the entity departure event. This field appears only if you set **Timing of function call f1** to Before entity departure and select **Generate optional function call f2 after function call f1**.


Suppress function call f2 if enable signal e2 is not positive

Selecting this option causes **f2** function calls to be contingent upon a positive value at the **e2** signal input port. This field appears

Entity Departure Event to Function-Call Event

only if you select **Generate optional function call f2 after function call f1**.

Statistics Tab



Number of entities departed

Controls the presence and behavior of the signal output port labeled **#d**.

Number of f1 function calls

Controls the presence and behavior of the signal output port labeled **#f1**.

Number of f2 function calls

Controls the presence and behavior of the signal output port labeled **#f2**. This field is active only if you select **Generate optional function call f2 after function call f1** on the **Function Call** tab of this dialog box.

Examples

- “Sample Use Cases” on page 1-8 — examples related to entity generation based on events
- “Example: Opening a Gate Upon Entity Departures” on page 7-7 — an example involving a gate

Entity Departure Event to Function-Call Event

- “Example: Using Entity-Based Timing for Choosing a Port” on page 8-29 — a computation inside a discrete event subsystem

See Also

Signal-Based Event to Function-Call Event, “Manipulating Events” on page 2-39

Entity Sink

Purpose Accept or block entities

Library SimEvents Sinks

Description This block provides a way to terminate an entity path:



- If you select **Input port available for entity arrivals**, the block always accepts entity arrivals.
- Otherwise, the block never accepts entity arrivals. The simulation issues an error message if an entity attempts to arrive at the block.

Ports

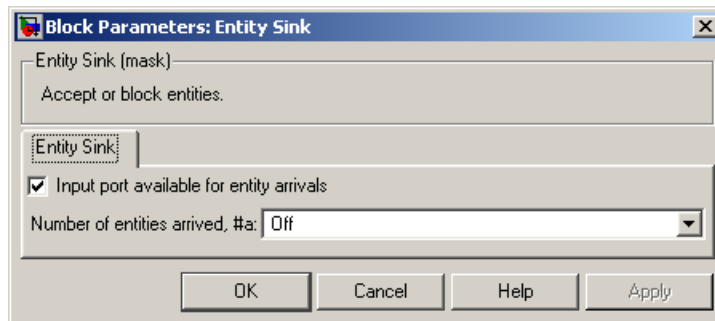
Entity Input Ports

Label	Description
IN	Port for entities that arrive or attempt to arrive.

Signal Output Ports

Label	Description	Time of Update When Statistic Is On
#a	Number of entities that the block has accepted. This port appears only if you select Input port available for entity arrivals and then set Number of entities arrived to either On or Upon stop or pause.	After entity arrival

Dialog Box



Input port available for entity arrivals

Determines whether the block accepts or blocks entities that attempt to arrive.

Number of entities arrived

Controls the presence and behavior of the signal output port labeled **#a**. This field appears only if you select **Input port available for entity arrivals**.

Examples

- “Modeling the Channels” in the Getting Started documentation
- “Example: Using an Attribute to Select an Output Port” in the Getting Started documentation

See Also


Time-Based Entity Generator, Event-Based Entity Generator

Entity-Based Function-Call Event Generator

Purpose Generate function call events corresponding to entities

Library Generators / Event Generators

Description This block generates a function call corresponding to each entity that arrives at the block. You can choose whether the block issues the function call before or after the departure. You can use the function call to invoke function-call subsystems, Stateflow blocks, or other blocks that accept function-call inputs.



This block is similar to the Entity Departure Event to Function-Call Event block, which offers more flexibility.

Ports

Entity Input Ports

Label	Description
IN	Port for arriving entities.

Entity Output Ports

Label	Description
OUT	Port for departing entities.

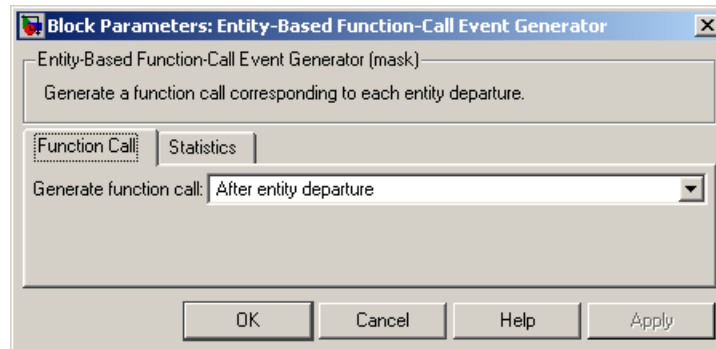
Entity-Based Function-Call Event Generator

Signal Output Ports

Label	Description	Time of Update When Port Is Present	Order of Update
f1	Function-call signal.	Before or after entity departure, depending on Generate function call parameter	1
#d	Number of entities that have departed from this block since the start of the simulation.	After entity departure	3
#f1	Number of function calls the block has generated since the start of the simulation.	After entity departure	2

Dialog Box

Function Call Tab

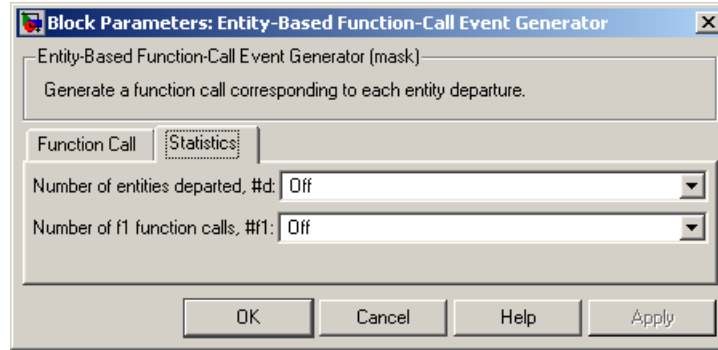


Generate function call

Determines whether the function call occurs before or after the entity departs from this block.

Entity-Based Function-Call Event Generator

Statistics Tab



Number of entities departed

Controls the presence and behavior of the signal output port labeled **#d**.

Number of f1 function calls

Controls the presence and behavior of the signal output port labeled **#f1**.

Examples

See “Example: Performing a Computation on Selected Entity Paths” on page 8-31.

See Also

Entity Departure Event to Function-Call Event, Signal-Based Function-Call Event Generator, “Generating Function-Call Events” on page 2-36

Purpose Generate entity upon signal-based event or function call

Library Generators / Entity Generators

Description This block is designed to generate entities when events of a specified type occur.



When to Generate Entities	Generate entities upon Value
Each time Simulink recomputes and outputs the value of a signal	Sample time hit from port <code>ts</code>
Each time an input signal has a trigger edge	Trigger from port <code>tr</code>
Each time an input signal changes its value	Change in signal from port <code>vc</code>
Each time an input signal issues a function call	Function call from port <code>fcn</code>

For details about these options, see “Generating Entities When Events Occur” on page 1-2.

Note An exceptional case is when the block temporarily suspends its normal entity-generation behavior. See the description of the `Delay first pending entity` option in “Responding to Blockage at Entity Output Port” on page 16-37.

Responding to Blockage at Entity Output Port

Using the **Allow OUT port blocking** and optional **Response during blockage period** parameters, you can choose how this block responds when the subsequent entity input port is not available to accept the newly generated entity. The possible responses are in the table below.

Event-Based Entity Generator

Response to Blockage	Parameter Values
Error message	Clear the Allow OUT port blocking check box.
The block discards all entities generated while the OUT port is blocked. Such entities do not depart from the block.	Select Allow OUT port blocking and set Response during blockage period to Discard generated entities
The block holds one generated entity and temporarily suspends the generation of additional entities. When the subsequent entity input port becomes available, this block outputs the held entity and resumes normal operation.	Select Allow OUT port blocking and set Response during blockage period to Delay first pending entity

Ports

Signal Input Ports

Label	Description
ts	Signal whose updates indicate when to generate entities. This port appears only if you set Generate entities upon to Sample time hit from port ts .
tr	Trigger signal whose edges indicate when to generate entities. This port appears only if you set Generate entities upon to Trigger from port tr .

Event-Based Entity Generator

Label	Description
vc	Signal whose numerical changes in value indicate when to generate entities. This port appears only if you set Generate entities upon to Change in signal from port vc.
fcn	Function-call signal that indicates when to generate entities. This port appears only if you set Generate entities upon to Function call from port fcn.

Entity Output Ports

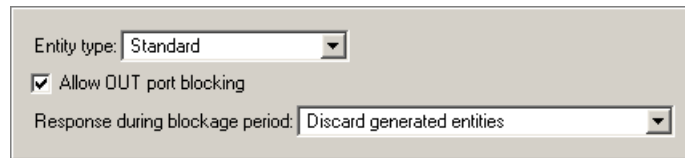
Label	Description
OUT	Port through which generated entities depart.

Signal Output Ports

Label	Description	Time of Update When Statistic Is On	Order of Update
#d	Number of entities that have departed from this block since the start of the simulation.	After entity departure	3
pe	A value of 1 indicates when the block tries and fails to output an entity.	After entity generation if OUT port is blocked, and after entity departure in all cases	2
w	Average intergeneration time, in seconds, for all pairs of successive entities that have departed from this block. The signal value is 0 before the second entity departure.	After entity departure	1

Event-Based Entity Generator

Dialog Box



Entity type: Standard

Allow OUT port blocking

Response during blockage period: Discard generated entities

Entity type

The standard type includes attributes called *Priority* and *Count*, while the blank type includes no attributes.

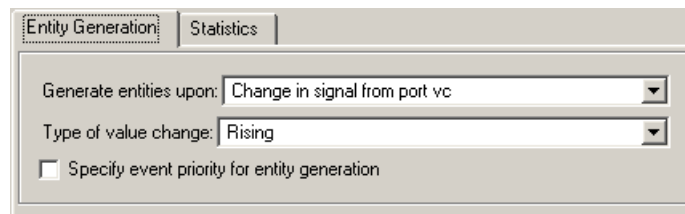
Allow OUT port blocking

If you do not select this option, the simulation halts with an error message if the subsequent entity input port is not available to accept an entity when this block generates it.

Response during blockage period

Determines how the block responds if a generated entity cannot depart immediately because the entity input port of the subsequent block is unavailable; see “Responding to Blockage at the Entity Output Port” on page 16-177. This field appears only if you select **Allow OUT port blocking**.

Entity Generation Tab



Entity Generation | Statistics

Generate entities upon: Change in signal from port vc

Type of value change: Rising

Specify event priority for entity generation

Generate entities upon

The type of event that indicates when the block should generate an entity.

Trigger type

Determines whether rising, falling, or either type of trigger edge causes an entity generation. This field appears only if you set **Generate entities upon** to Trigger from port tr.

Type of value change

Determines whether rising, falling, or either type of value change causes an entity generation. This field appears only if you set **Generate entities upon** to Change in signal from port vc.

Specify event priority for entity generation

Select this option to control the sequencing of the entity-generation event, relative to other simultaneous events.

Generation event priority

The priority of the entity-generation event, relative to other simultaneous events. This field appears only if you select **Specify event priority for entity generation**.

Generate entity at simulation start

If you select this option, the block generates the first entity when the simulation begins. Otherwise, the block generates the first entity upon the first update of the **ts** signal at a nonzero value of time. This field appears only if you set **Generate entities** to Sample time hit from port ts.


Allow entity generation at simulation start

If you select this option, the block responds to function calls at the starting time of the simulation. Otherwise, the block responds only to function calls at subsequent times. This field appears only if you set **Generate entities upon** to Function call from port fcn.

Statistics Tab

These parameters determine whether certain ports output data throughout the simulation, output data only when you stop or pause the simulation, or are omitted from the block. For descriptions of the affected ports, see the “Signal Output Ports” table above.

Event-Based Entity Generator



Number of entities departed

Controls the presence and behavior of the signal output port labeled **#d**.

Status of pending entity departure

Controls the presence and behavior of the signal output port labeled **pe**.

Average intergeneration time

Controls the presence and behavior of the signal output port labeled **w**.

Examples

- “Example: Plotting Event Counts to Find Roundoff Error” on page 9-12
- “Example: Race Conditions at a Switch” on page 2-23
- “Sample Use Cases” on page 1-3 under “Detecting Sample Time Hits” on page 1-2
- “Sample Use Cases” on page 1-4 under “Detecting Changes in Signal Values” on page 1-4
- “Sample Use Cases” on page 1-6 under “Detecting Edges in Trigger Signals” on page 1-5
- “Sample Use Cases” on page 1-8 under “Detecting Function Calls” on page 1-7

See Also

Time-Based Entity Generator, Entity Sink, “Generating Entities When Events Occur” on page 1-2


Purpose

Generate random numbers from specified distribution, parameters, and initial seed

Library

Generators / Signal Generators

Description



This block generates random numbers in an event-based manner, inferring from a subsequent block when to generate a new random number. For example, when connected to the **t** input port of a Single Server block, the Event-Based Random Number block generates a new random number each time an entity arrives at the server.

You specify the distribution from which the block draws random numbers. The seed of the random number generator is reset to the value of the **Initial seed** parameter each time a simulation starts, which makes the random behavior repeatable.

Connecting to Other Blocks

This block has a restricted set of valid connections to other blocks. This is because the Event-Based Random Number block is designed to infer from a subsequent block when to generate a new random number.

Connect the Event-Based Random Number block to

- Exactly one *notifying port*
- Zero or more *monitoring ports*
- No *reactive ports* or other ports

Notifying ports, listed in the table below, notify the preceding block when a certain event has occurred. When the preceding block is the Event-Based Random Number block, it responds to the notification by generating a new random number. The Event-Based Random Number block must be connected to exactly one notifying port.

Event-Based Random Number

Notifying Ports

Signal Input Port	Block	Generate New Random Number Upon
A1, A2, A3, A4	Set Attribute	Entity arrival
in	Signal Latch	Write event
e1, e2	Entity Departure Event to Function-Call Event	Entity arrival
	Signal-Based Event to Function-Call Event	Relevant signal-based event, depending on configuration of block
t	Infinite Server	Entity arrival
	N-Server	Entity arrival
	Single Server	Entity arrival
t	Time-Based Entity Generator	Simulation start and subsequent entity departures
x	X-Y Signal Scope	Sample time hit at in signal input port

Monitoring ports, listed in the next table, help you observe signal values. You can optionally use a branch line to connect the Event-Based Random Number block to one or more monitoring ports. These connections do not cause the Event-Based Random Number block to generate a new random number but merely enable you to observe the signal.

Monitoring Ports

Signal Input Port	Block
Unlabeled	Discrete Event Signal to Workspace
in	Signal Scope
	X-Y Signal Scope
ts, tr, vc	Instantaneous Event Counting Scope

Reactive ports, listed in the next table, listen for updates or changes in the input signal and cause an appropriate reaction in the block possessing the port. For example, the **p** port on a switch listens for changes in the input signal; the block reacts by selecting a new switch port. You cannot connect the Event-Based Random Number block to a reactive port. To create a random signal that can be an input to a reactive port, see the techniques described in “Generating Random Signals Based on Arbitrary Events” on page 3-5 and “Generating Random Time-Based Signals” on page 3-8.

Event-Based Random Number

Reactive Ports

Signal Input Port	Block
p	Input Switch
	Output Switch
	Path Combiner
ts, tr, vc	Entity Departure Counter
	Event-Based Entity Generator
	Release Gate
	Signal-Based Event to Function-Call Event
	Signal-Based Function-Call Event Generator
wts, wtr, wvc, rts, rtr, rvc	Signal Latch

Note Connections from the Event-Based Random Number block to ports other than notifying ports and monitoring ports are not supported.

Distribution Types

The **Distribution** parameter names the type of distribution the block uses to generate random numbers. When you set the **Distribution** parameter, the block changes its dialog box to show additional parameters that determine the probability density function (or probability mass function, for a discrete distribution). The available distributions and the additional parameters for each are described in the sections that follow.

Event-Based Random Number

Distribution	Additional Parameters
Exponential	Mean
Uniform	Minimum, Maximum
Bernoulli	Probability of 1
Binomial	Probability of success in a single trial, Number of trials
Triangular	Minimum, Maximum, Mode
Gamma	Threshold, Scale, Shape
Gaussian (normal)	Mean, Standard deviation
Geometric	Probability of success in a single trial
Poisson	Mean
Lognormal	Threshold, Scale, Shape
Log-logistic	Threshold, Scale
Beta	Minimum, Maximum, Shape parameter a, Shape parameter b
Discrete uniform	Minimum, Maximum, Number of values
Weibull	Threshold, Scale, Shape
Arbitrary continuous	Value vector, Cumulative probability function vector
Arbitrary discrete	Value vector, Probability vector

For information about the definitions and properties of each distribution, see the references listed in “References” on page 16-55 below.

Range of Output Values

Different distributions have different output ranges. Make sure the distribution and parameters you choose are suitable for your

Event-Based Random Number

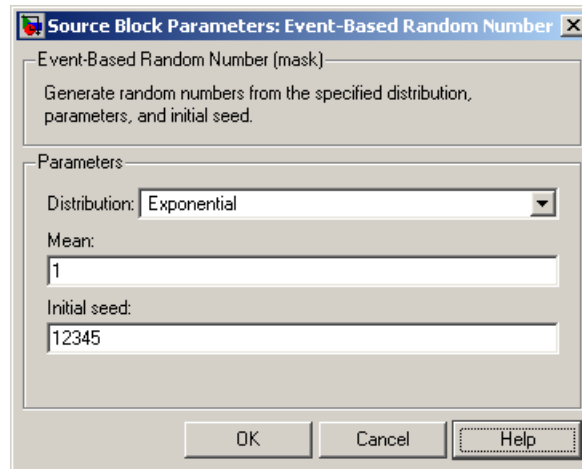
application. For example, when generating random service times, do not use a Gaussian distribution because it can produce negative numbers.

Ports

This block has one signal output port for the random numbers.

The block has no entity ports, and no signal input port.

Dialog Box



Distribution

The distribution from which the block generates random numbers.

Mean

The mean value of an exponential, Gaussian, or Poisson distribution.

Minimum, Maximum

The minimum and maximum values of a uniform, triangular, beta, or discrete uniform distribution.

Probability for output to be 0

The probability of a zero in a Bernoulli distribution.

Probability of success in a single trial

The probability of a successful outcome in each trial used to describe a binomial or geometric distribution.

Number of trials

The number of trials used to describe a binomial distribution.

Mode

The statistical mode of a triangular distribution. The triangular distribution also uses the **Minimum** and **Maximum** parameters to define its density function.

Threshold, Scale, Shape

Parameters that define the density function of a gamma, lognormal, log-logistic, or Weibull distribution. The log-logistic distribution does not use a **Shape** parameter, however.

Standard deviation

The standard deviation of a Gaussian distribution, which also uses the **Mean** parameter to define its density function.

Shape parameter a, Shape parameter b

The first and second shape parameters, respectively, of a beta distribution. The beta distribution also uses the **Minimum** and **Maximum** parameters to define its density function.

Number of values

The number of possible outputs of a discrete uniform distribution, including the values of the **Minimum** and **Maximum** parameters. **Number of values** must exceed 1.

Value vector

A vector of values in ascending order, representing the possible random values in an arbitrary continuous or arbitrary discrete distribution.

Cumulative probability function vector

A vector of values in ascending order representing the cumulative probability function for an arbitrary continuous distribution. The first and last values of the vector must be 0 and 1, respectively.

Event-Based Random Number

This parameter and the **Value vector** parameter must have the same vector length.

Probability vector

A vector of values representing the probability of each value in the **Value vector** function for an arbitrary discrete distribution. This vector must contain nonnegative values that sum to 1. This parameter and the **Value vector** parameter must have the same vector length.

Initial seed

A nonnegative integer that initializes the random number generator.

Examples

See “Examples of Random Event-Based Signals” on page 3-6.

Algorithm

Below are the expressions for f , the probability density functions for the continuous distributions and probability mass functions for the discrete distributions that the block supports.

Exponential Distribution

$$f(x) = \begin{cases} \frac{1}{\mu} \exp\left(-\frac{x}{\mu}\right) & \text{for } x \geq 0 \\ 0 & \text{otherwise} \end{cases}$$

where μ is the **Mean** parameter, a positive number.

A similar function in the Statistics Toolbox is `expnd`.

Uniform Distribution

$$f(x) = \begin{cases} \frac{1}{U-L} & \text{for } L \leq x \leq U \\ 0 & \text{otherwise} \end{cases}$$

where L is the **Minimum** parameter and U is the **Maximum** parameter.

Similar functions are `rand` in MATLAB and `unifrnd` in the Statistics Toolbox.

Bernoulli Distribution

$$f(x) = \begin{cases} p^x(1-p)^{1-x} & \text{for } x = 0, 1 \\ 0 & \text{otherwise} \end{cases}$$

where p is the **Probability of 1** parameter. The value p must be between 0 and 1, inclusive. This is a discrete distribution.

This distribution is a special case of the binomial distribution in which the number of trials is 1.

Binomial Distribution

$$f(x) = \begin{cases} \frac{n!}{x!(n-x)!} p^x q^{(n-x)} & \text{for } x = 0, 1, 2, \dots, n \\ 0 & \text{otherwise} \end{cases}$$

where p is the **Probability of success in a single trial** parameter, $q = 1-p$, and n is the **Number of trials** parameter. The value p must be between 0 and 1, inclusive, while n must be positive. This is a discrete distribution.

A similar function in the Statistics Toolbox is `binornd`.

Triangular Distribution

$$f(x) = \begin{cases} \frac{2(x-L)}{(U-L)(m-L)} & \text{for } L \leq x \leq m \\ \frac{2(U-x)}{(U-L)(U-m)} & \text{for } m < x \leq U \\ 0 & \text{otherwise} \end{cases}$$

Event-Based Random Number

where L is the **Minimum** parameter, U is the **Maximum** parameter, and m is the **Mode** parameter. These parameters must satisfy $L < m < U$.

Gamma Distribution

$$f(x) = \begin{cases} \frac{\left(\frac{x-\theta}{b}\right)^{a-1} \exp\left(-\frac{x-\theta}{b}\right)}{b\Gamma(a)} & \text{for } x \geq \theta \\ 0 & \text{otherwise} \end{cases}$$

where θ is the **Threshold** parameter, b is the **Scale** parameter, and a is the **Shape** parameter. The **Scale** and **Shape** parameters must be positive. Also, Γ is the gamma function (gamma in MATLAB).

A similar function in the Statistics Toolbox is `gamrnd`.

Gaussian (Normal) Distribution

$$f(x) = \frac{\exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right)}{\sigma\sqrt{2\pi}}$$

where μ is the **Mean** parameter and σ is the **Standard deviation** parameter. The standard deviation parameter must be nonnegative.

Similar functions are `randn` in MATLAB and `normrnd` in the Statistics Toolbox.

Geometric Distribution

If the **Probability of success in a single trial** parameter is strictly between 0 and 1, then the probability mass function is defined by

$$f(x) = \begin{cases} pq^x & \text{for } x = 0, 1, 2, \dots \\ 0 & \text{otherwise} \end{cases}$$

where p is the **Probability of success in a single trial** parameter and $q = 1-p$.

In the special case where the **Probability of success in a single trial** parameter is 1, then

$$f(x) = \begin{cases} 1 & \text{for } x = 0 \\ 0 & \text{otherwise} \end{cases}$$

This is a discrete distribution.

A similar function in the Statistics Toolbox is `geornd`.

Poisson Distribution

$$f(x) = \begin{cases} \frac{e^{-\lambda} \lambda^x}{x!} & \text{for } x = 0, 1, 2, \dots \\ 0 & \text{otherwise} \end{cases}$$

where λ is the **Mean** parameter, a positive number. This is a discrete distribution.

A similar function in the Statistics Toolbox is `poissrnd`.

Lognormal Distribution

$$f(x) = \begin{cases} \frac{\exp\left[\frac{-(\ln(x-\theta)-\mu)^2}{2\sigma^2}\right]}{(x-\theta)\sigma\sqrt{2\pi}} & \text{for } x \geq \theta \\ 0 & \text{otherwise} \end{cases}$$

where θ is the **Threshold** parameter, $\exp(\mu)$ is the **Scale** parameter, and σ is the **Shape** parameter. The **Shape** parameter must be positive.

A similar function in the Statistics Toolbox is `lognrnd`.

Log-Logistic Distribution

The log-logistic distribution is derived from the logistic distribution, as follows:

Event-Based Random Number

X = Random variable with logistic distribution

$Y = e^X$ = Random variable with log-logistic distribution

The probability density function for the logistic distribution is

$$f_{\text{logistic}}(x) = \frac{1}{b} \cdot \frac{e^{(x-\theta)/b}}{\left(1 + e^{(x-\theta)/b}\right)^2}$$

where θ is the **Threshold** parameter and b is the **Scale** parameter. The **Scale** parameter must be positive.

Beta Distribution

$$f(x) = \begin{cases} \frac{(x-L)^{a-1}(U-x)^{b-1}}{B(a,b)(U-L)^{a+b-1}} & \text{for } L \leq x \leq U \\ 0 & \text{otherwise} \end{cases}$$

where L is the **Minimum** parameter, M is the **Maximum** parameter, a is the **Shape parameter a** parameter, b is the **Shape parameter b** parameter, and $B(a,b)$ is the beta function defined by

$$B(a,b) = \int_0^1 t^{a-1}(1-t)^{b-1} dt$$

The two shape parameters must be positive.

A similar function in the Statistics Toolbox is `betarnd`.

Discrete Uniform Distribution

$$f(x) = \begin{cases} 1/K & \text{for } x = L + k \frac{(U-L)}{K-1}, k = 0, 1, 2, \dots, K-1 \\ 0 & \text{otherwise} \end{cases}$$

where L is the **Minimum** parameter, U is the **Maximum** parameter, and K is the **Number of values** parameter. This is a discrete

distribution. If $(U-L)/(K-1)$ and L are both integers, then all outputs from this distribution are integers.

A similar function in the Statistics Toolbox is `unidrnd`.

Weibull Distribution

$$f(x) = \begin{cases} \frac{\gamma}{\alpha} \left(\frac{x-\theta}{\alpha} \right)^{\gamma-1} \exp \left[- \left(\frac{x-\theta}{\alpha} \right)^\gamma \right] & \text{for } x \geq \theta \\ 0 & \text{otherwise} \end{cases}$$

where θ is the **Threshold** parameter, α is the **Scale** parameter, and γ is the **Shape** parameter. The **Scale** and **Shape** parameters must be positive.

A similar function in the Statistics Toolbox is `wblrnd`.

See Also

Signal Latch, “Generating Random Signals” on page 3-4

References

- [1] Evans, M., N. Hastings, and B. Peacock, *Statistical Distributions*, Wiley-Interscience, 2000.
- [2] Johnson, N. L., S. Kotz, and N. Balakrishnan, *Continuous Univariate Distributions*, Volume 1, Wiley-Interscience, 1993.
- [3] Johnson, N. L., S. Kotz, and N. Balakrishnan, *Continuous Univariate Distributions*, Volume 2, Wiley-Interscience, 1994.
- [4] Johnson, N. L., S. Kotz, and A. W. Kemp, *Univariate Discrete Distributions*, Wiley-Interscience, 1993.

FIFO Queue

Purpose Store entities in sequence for undetermined length of time

Library Queues

Description



This block stores up to N entities simultaneously, where N is the **Capacity** parameter value. The block attempts to output an entity through the **OUT** port but retains the entity if the **OUT** port is blocked. If the block is storing multiple entities, then entities depart in a first-in, first-out (FIFO) fashion. The length of time that an entity stays in this block cannot be determined in advance.

The **IN** port is unavailable whenever this block stores exactly N entities. In this case, the queue is said to be full.

Ports

Entity Input Ports

Label	Description
IN	Port for arriving entities, which will be stored.

Entity Output Ports

Label	Description
OUT	Port for departing entities.

Signal Output Ports

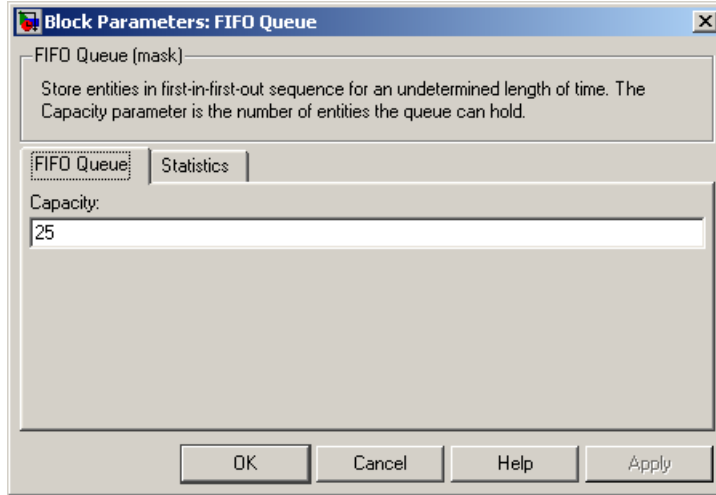
Label	Description	Time of Update When Statistic Is On	Order of Update
#d	Number of entities that have departed from this block since the start of the simulation.	After entity departure	4
#n	Number of entities currently in the queue.	After entity arrival at a nonempty queue and after entity departure	3
pe	A value of 1 indicates when the block tries and fails to output an entity; that is, the queue is nonempty. A value of 0 indicates when the queue is empty.	After entity arrival and after entity departure	2
w	Sample mean of the waiting times in this block for all entities that have departed.	After entity departure	1
len	Average number of entities in the queue over time, that is, the time average of the #n signal.	After entity arrival at a nonempty queue and after entity departure	1

Output signals having the same number in the Order of Update column above are updated in an arbitrary sequence relative to each other; you should not rely on a specific sequence for your simulation results.

FIFO Queue

Dialog Box

FIFO Queue Tab

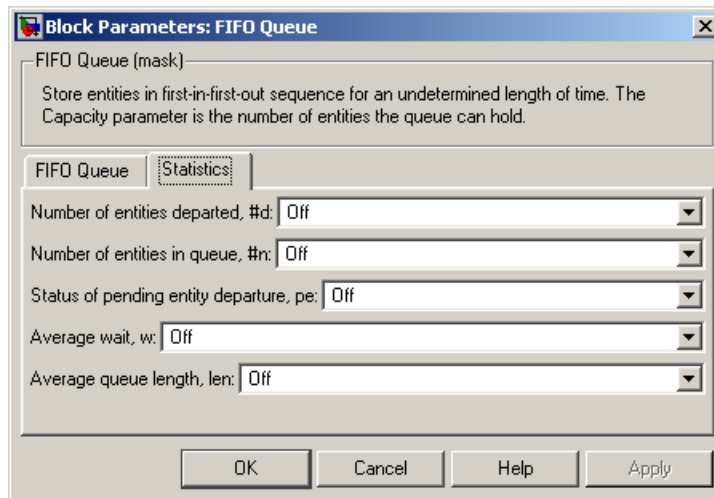


Capacity

Determines how many entities the block can store at a time.

Statistics Tab

These parameters determine whether certain ports output data throughout the simulation, output data only when you stop or pause the simulation, or are omitted from the block. For descriptions of the affected ports, see the “Signal Output Ports” table above.



Number of entities departed

Controls the presence and behavior of the signal output port labeled **#d**.

Number of entities in queue

Controls the presence and behavior of the signal output port labeled **#n**.

Status of pending entity departure

Controls the presence of the signal output port labeled **pe**.

Average wait

Controls the presence and behavior of the signal output port labeled **w**.

Average queue length

Controls the presence and behavior of the signal output port labeled **len**.

Examples

- “Building a Simple Discrete-Event Model” in the Getting Started documentation

FIFO Queue

- “Example: Selecting the First Available Server” in the Getting Started documentation
- “Example: Round-Robin Approach to Choosing Inputs” in the Getting Started documentation
- “Constructs Involving Queues and Servers” in the Getting Started documentation
- “Example of a Logical Queue” in the Getting Started documentation
- “Example: Waiting Time in LIFO Queue” on page 4-2

See Also

LIFO Queue, Priority Queue, “Basic Queues and Servers” in the Getting Started documentation

Purpose Output value of entity's attribute

Library Attributes

Description



This block outputs signals from up to four attributes. For each entity, the block updates the signal at the **A1**, **A2**, **A3**, and/or **A4** signal output port using the value of the attribute named in the corresponding tab of the dialog box. The block also outputs the entity unchanged.

To query fewer than the maximum number of attributes, you can deactivate unused tabs by setting **Send attribute value to signal port A4**, for example, to Off.

To specify the names of attributes you want to query, use the **Attribute name** parameter in the dialog box.

Missing Attributes

If the block is configured to retrieve the value of an attribute that an arriving entity does not possess, then the block can react in one of these ways:

- Issue an error message and halt the simulation.
- Output a default value that you specify. The simulation proceeds.
- Output a default value that you specify and also issue a warning in the MATLAB Command Window. The simulation proceeds.

You determine the reaction using the **Action for missing attribute** parameter on each attribute tab of the dialog. The error and warning options can be useful for debugging a simulation (for example, to locate a mistyped attribute name in a dialog box).

Get Attribute

Ports

Entity Input Ports

Label	Description
IN	Port for arriving entities.

Entity Output Ports

Label	Description
OUT	Port for departing entities.

Signal Output Ports

Label	Description	Time of Update When Statistic Is On	Order of Update
#d	Number of entities that have departed from this block since the start of the simulation.	After entity departure	2
A x , where x = 1, 2, 3, 4	Value of the attribute specified on the A x tab of the dialog. This port appears only if you set Send attribute value to signal port Ax on the A x tab to On.	After entity departure	1

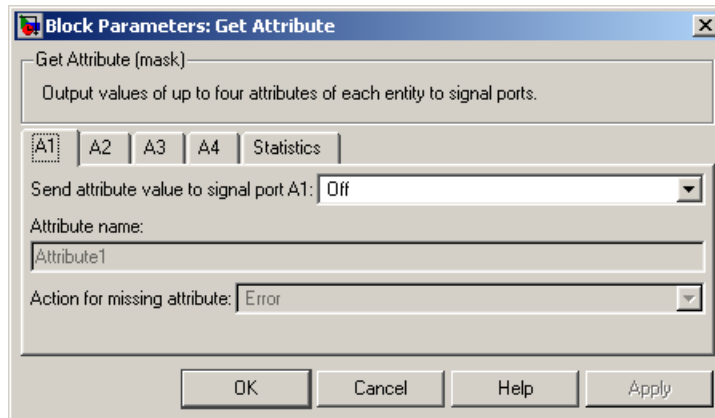
Output signals having the same number in the Order of Update column above are updated in an arbitrary sequence relative to each other; you should not rely on a specific sequence for your simulation results.

Dialog Box

A1, A2, A3, A4 Tabs

The A1, A2, A3, and A4 tabs have similar parameter choices. By assigning different values to the parameters, you can configure this

block to retrieve the values of up to four different attributes for each entity that the block processes.



Send attribute value to signal port Ax, where $x = 1, 2, 3, 4$

Indicates whether the block creates an output signal with the value of an attribute. Choosing *Off* indicates that you are not using this tab of the dialog and makes the parameters below inactive or invisible.

Attribute name

The name of the attribute to query.

Action for missing attribute

The response of the block when the entity does not possess the attribute named above.

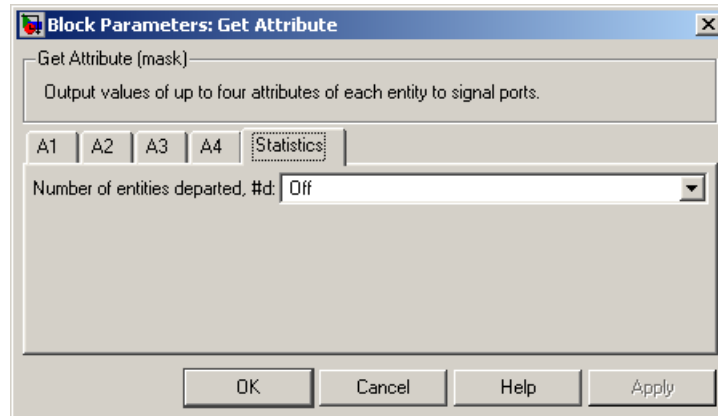
Default value

The value for the corresponding Simulink output signal if the entity does not possess the attribute named above. This field appears only if you set **Action for missing attribute** to *Output default value* or *Output default value and warn*.

Get Attribute

Statistics Tab

These parameters determine whether certain ports output data throughout the simulation, output data only when you stop or pause the simulation, or are omitted from the block. For descriptions of the affected ports, see the “Signal Output Ports” table above.



Number of entities departed

Controls the presence and behavior of the signal output port labeled #d.

Examples

- “Adding Event-Based Behavior” in the Getting Started documentation
- “Example Manipulating Data” on page 1-18

See Also

Set Attribute, “Accessing Attributes of Entities” on page 1-20

Purpose Delay any number of entities for period of time

Library Servers

Description



This block serves any number of entities for a period of time, called the *service time*, and then attempts to output them through the **OUT** port. If the **OUT** port is blocked, then the block holds the entities until the port becomes unblocked. An infinite server is like an infinite set of single servers connected in parallel, followed by a path combiner; the path combiner notifies entities of an unblocked path in the sequence in which the entities completed their service time, until one entity departs.

You specify the service time, which is the duration of service, via a parameter, attribute, or signal, depending on the **Service time from** parameter value. The block determines the service time for an entity upon its arrival.

The **IN** port of an infinite server is always available. You can interpret an infinite server as a mechanism for delaying entities. Some discussions of this block suggest this interpretation by using the word delay instead of serve.

Ports

Entity Input Ports

Label	Description
IN	Port for arriving entities, which will be served.

Signal Input Ports

Label	Description
t	Service time for a newly arrived entity. This port appears only if you set Service time from to Signal port t.

Infinite Server

Entity Output Ports

Label	Description
OUT	Port for departing entities, which have completed their service time.

Signal Output Ports

Label	Description	Time of Update When Statistic Is On	Order of Update
#d	Number of entities that have departed from this block since the start of the simulation.	After entity departure	4
#n	Number of entities in the block.	After entity arrival and after entity departure	3
pe	The value is 1 if the block is holding any pending entities, that is, entities that have completed their service time but cannot depart because the OUT port is blocked. The value of this signal is 0 otherwise.	After service completion and after entity departure	2

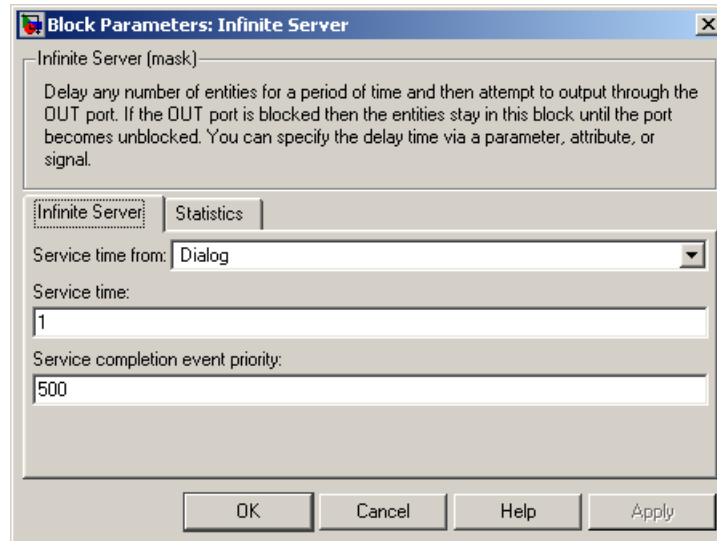
Label	Description	Time of Update When Statistic Is On	Order of Update
#pe	The number of entities that have completed their service but cannot depart because the subsequent entity input port is unavailable.	After service completion and after entity departure	2
w	Sample mean of the waiting times in this block for all entities that have departed. An entity's waiting time might exceed its service time if the OUT port is blocked when the entity completes service.	After entity departure	1

Output signals having the same number in the Order of Update column above are updated in an arbitrary sequence relative to each other; you should not rely on a specific sequence for your simulation results.

Infinite Server

Dialog Box

Infinite Server Tab



Service time from

Determines whether the service time is computed from a parameter in this dialog box, a signal input port, or an attribute of the entity being served.

Service time

The service time for all entities. This field appears only if you set **Service time from** to Dialog.

Attribute name

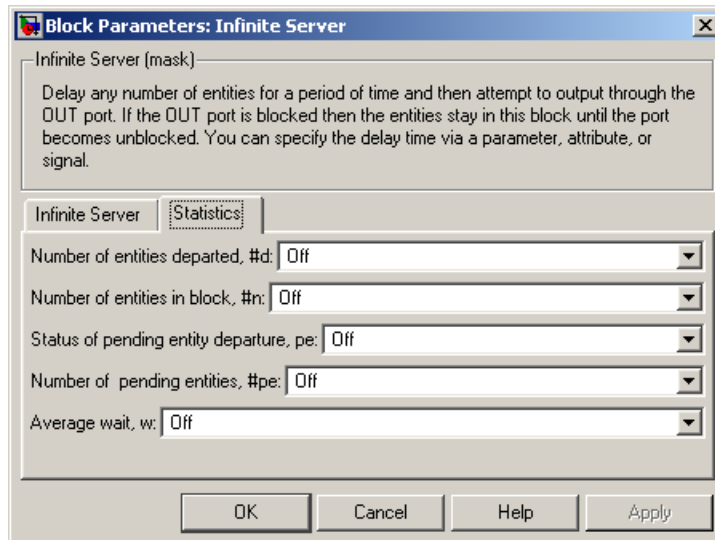
The name of the attribute whose value the block uses as the service time for an entity. This field appears only if you set **Service time from** to Attribute.

Service completion event priority

The priority of the service completion event, relative to other simultaneous events.

Statistics Tab

These parameters determine whether certain ports output data throughout the simulation, output data only when you stop or pause the simulation, or are omitted from the block. For descriptions of the affected ports, see the “Signal Output Ports” table above.



Number of entities departed

Controls the presence and behavior of the signal output port labeled **#d**.

Number of entities in block

Controls the presence and behavior of the signal output port labeled **#n**.

Status of pending entity departure

Controls the presence of the signal output port labeled **pe**.

Number of pending entities

Controls the presence and behavior of the signal output port labeled **#pe**.

Infinite Server

Average wait

Controls the presence and behavior of the signal output port labeled **w**.

Examples

- “Adding Event-Based Behavior” in the Getting Started documentation
- “Restarting a Timer” on page 10-21
- “Example: Counting Simultaneous Departures from a Server” on page 1-22

See Also

Single Server, N-Server, “Basic Queues and Servers” in the Getting Started documentation

Purpose Accept entities from selected entity input port

Library Routing

Description



This block selects exactly one entity input port for potential arrivals. The selected entity input port can change during the simulation. When one entity input port becomes selected, all others become unavailable.

The possible rules the block uses for selecting an entity input port, as well as the corresponding values of the **Switching criterion** parameter in the dialog box, are listed in the table below.

Switching criterion Value	Description
Round robin	At the beginning of the simulation, IN1 is selected. After each departure, the block selects the entity input port next to the last selected port. After exhausting all entity input ports, the block returns to the first one, IN1 .
Equiprobable	At the beginning of the simulation and after each departure, the block randomly chooses which entity input port is selected for the next arrival. All entity input ports are equally likely. The Initial seed parameter initializes the random number generation process.
From signal port p	Selecting this option creates an additional signal input port, labeled p . The signal at this port must have integer values between 1 and the Number of entity input ports parameter value. The block detects changes in this integer value and selects the corresponding entity input port for future arriving entities.

Input Switch

Ports

Entity Input Ports

Label	Description
IN1, IN2, IN3, etc.	Ports for potential entity arrivals. At any given time, one input port is selected and the others are unavailable. The Number of entity input ports parameter determines how many of these entity input ports the block has.

Signal Input Ports

Label	Description
p	Index of the entity input port that is available. Values are 1, 2, 3,..., Number of entity input ports . This port appears only if you set Switching criterion to From signal port p.

Entity Output Ports

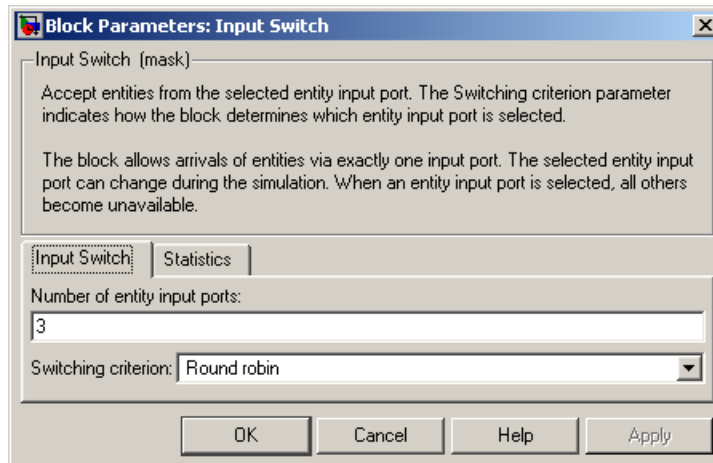
Label	Description
OUT	Port for departing entities.

Signal Output Ports

Label	Description	Time of Update When Statistic Is On	Order of Update
#d	Number of entities that have departed from this block since the start of the simulation.	After entity departure	2
last	Index of the input port that was available the last time an entity departed. The initial value is 0. After an entity has departed, values are 1, 2, 3,..., Number of entity input ports .	After entity departure	1

Dialog Box

Input Switch Tab



Number of entity input ports

Determines how many entity input ports the block has.

Input Switch

Switching criterion

The rule that determines which entity input port is selected for receiving entities.

Initial seed

A nonnegative integer that initializes the random number generator used to select an entity input port. This field appears only if you set **Switching criterion** to Equiprobable.

Specify event priority for port selection

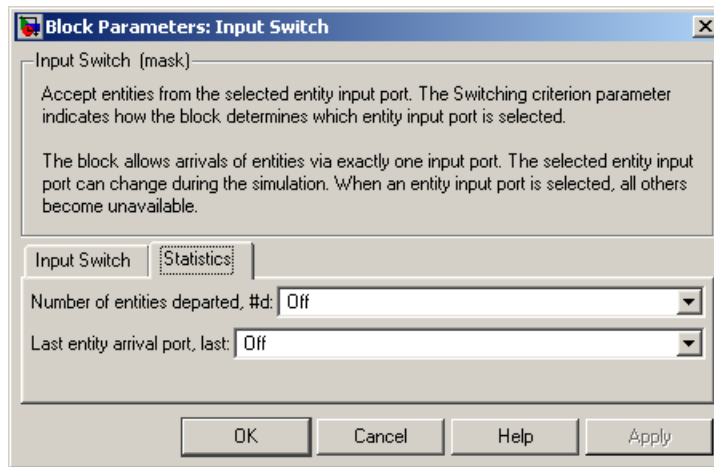
Select this option to control the sequencing of the port-selection event, relative to other simultaneous events. This field appears only if you set **Switching criterion** to From signal port *p*.

Event priority

The priority of the port-selection event, relative to other simultaneous events. This field appears only if you set **Switching criterion** to From signal port *p* and select **Specify event priority for port selection**.

Statistics Tab

These parameters determine whether certain ports output data throughout the simulation, output data only when you stop or pause the simulation, or are omitted from the block. For descriptions of the affected ports, see the “Signal Output Ports” table above.



Number of entities departed

Controls the presence and behavior of the signal output port labeled **#d**.

Last entity arrival port

Controls the presence of the signal output port labeled **last**.

Examples

- “Example: Race Conditions at a Switch” on page 2-23
- “Example: Round-Robin Approach to Choosing Inputs” in the Getting Started documentation
- “Example: Compound Switching Logic” on page 5-5

See Also

Output Switch, “Using the Input Switch” in the Getting Started documentation

Instantaneous Entity Counting Scope

Purpose Plot entity count versus time

Library SimEvents Sinks

Description This block creates a plot by counting arriving entities at each arrival time. The block restarts the count from 1 when the time changes. As a result, the count is cumulative for a given time instant but not cumulative across the entire simulation.



Note If you want to plot the total number of arriving entities across the entire simulation, connect the **#d** signal of the Entity Departure Counter block to the Signal Scope block.

Use the **Enable entity OUT port** option to choose whether the entity advances to a subsequent block or whether the block absorbs the arriving entity.

Ports

Entity Input Ports

Label	Description
IN	Port for arriving entities, which the block counts.

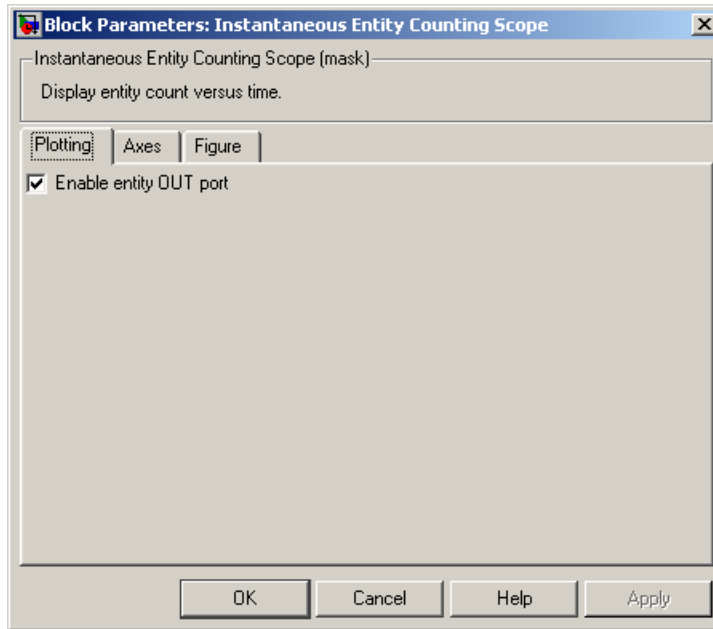
Entity Output Ports

Label	Description
OUT	Port for departing entities. This port appears only if you select Enable entity OUT port .

Instantaneous Entity Counting Scope

Dialog Box

Plotting Tab

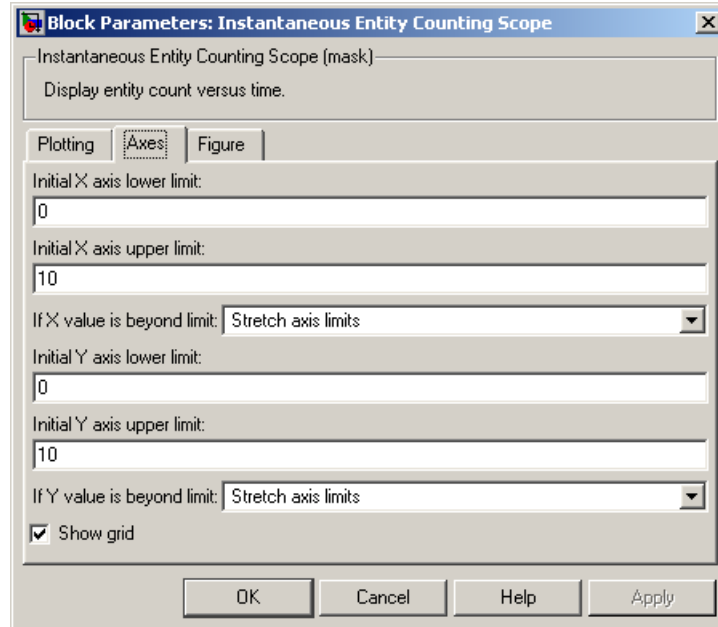


Enable entity OUT port

Causes the block to have an entity output port labeled **OUT**, through which the arriving entity departs. If you clear this box, the block absorbs arriving entities.

Instantaneous Entity Counting Scope

Axes Tab



Initial X axis lower limit, Initial X axis upper limit

The interval shown on the X axis at the beginning of the simulation. The interval might change from this initial setting due to zooming, autoscaling, or the **If X value is beyond limit** setting.

If X value is beyond limit

Determines how the plot changes if one or more X values are not within the limits shown on the X axis. For details, see “Varying Axis Limits Automatically” on page 9-5.

Initial Y axis lower limit, Initial Y axis upper limit

The interval shown on the Y axis at the beginning of the simulation. The interval might change from this initial setting

Instantaneous Entity Counting Scope

due to zooming, autoscaling, or the **If Y value is beyond limit** setting.

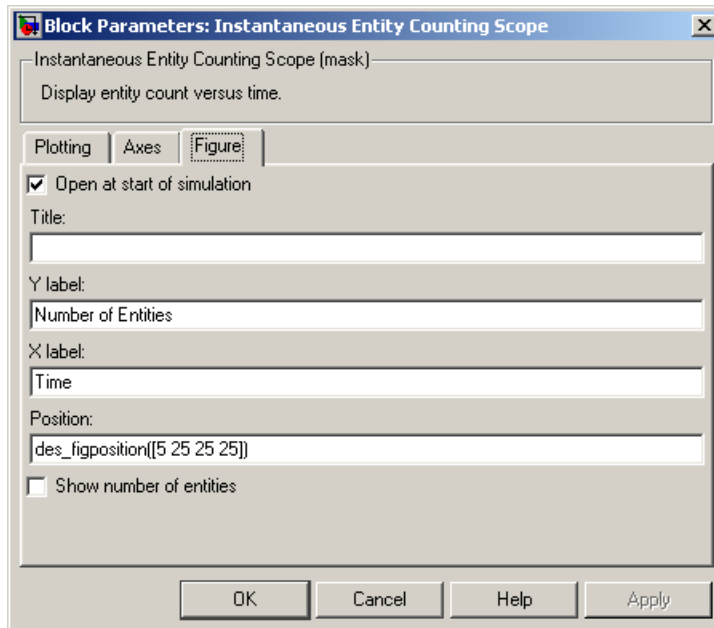
If Y value is beyond limit

Determines how the plot changes if one or more entity counts are not within the limits shown on the Y axis. For details, see “Varying Axis Limits Automatically” on page 9-5.

Show grid

Toggles the grid on and off.

Figure Tab



Open at start of simulation

Selecting this option causes the plot window to open when you start the simulation. If you clear this box, you can open the plot

Instantaneous Entity Counting Scope

window while the simulation is running by double-clicking the block icon.

Title

Text that appears as the title of the plot, above the axes.

Y label

Text that appears to the left of the vertical axis.

X label

Text that appears below the horizontal axis.

Position

A four-element vector of the form [left bottom width height] specifying the position of the scope window. (0,0) is the lower left corner of the display.

Show number of entities

Displays the number of plotted points using an annotation in the plotting window.

Examples

- “Example: Counting Simultaneous Departures from a Server” on page 1-22
- “Example: Synchronizing Service Start Times with the Clock” on page 7-6

See Also

Entity Departure Counter, Instantaneous Event Counting Scope, Chapter 9, “Plotting Data”, “Counting Entities” on page 1-21

Instantaneous Event Counting Scope

Purpose Plot event count versus time

Library SimEvents Sinks

Description



This block creates a plot by counting events. The block restarts the count from 1 when the time changes. As a result, the count is cumulative for a given time instant but not cumulative across the entire simulation.

Ports

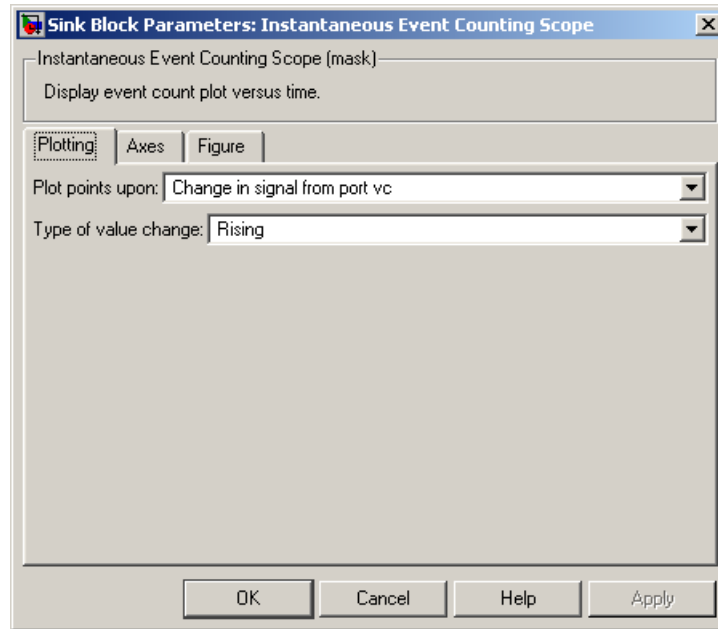
Signal Input Ports

Label	Description
ts	Signal whose updates increment the counter. This port appears only if you set Plot points upon to Sample time hit from port ts.
tr	Trigger signal whose edges increment the counter. This port appears only if you set Plot points upon to Trigger from port tr.
vc	Signal whose numerical changes in value increment the counter. This port appears only if you set Plot points upon to Change in signal from port vc.
fcn	Function-call signal that indicates when to increment the counter. This port appears only if you set Plot points upon to Function call from port fcn.

Instantaneous Event Counting Scope

Dialog Box

Plotting Tab



Plot points upon

The type of event that indicates when the block increments its counter.

Trigger type

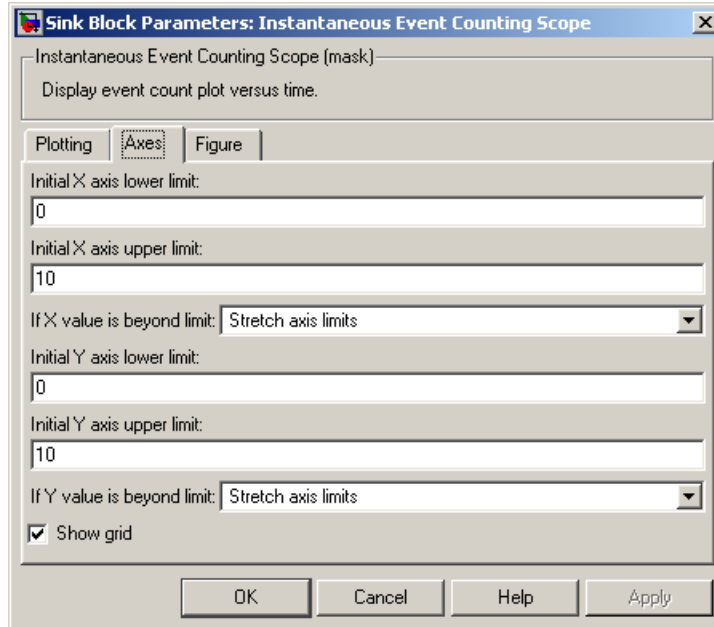
Determines whether rising, falling, or either type of trigger edge causes the block to increment its counter. This field appears only if you set **Plot points upon** to Trigger from port tr.

Type of value change

Determines whether rising, falling, or either type of value change causes the block to increment its counter. This field appears only if you set **Plot points upon** to Change in signal from port vc.

Instantaneous Event Counting Scope

Axes Tab



Initial X axis lower limit, Initial X axis upper limit

The interval shown on the X axis at the beginning of the simulation. The interval might change from this initial setting due to zooming, autoscaling, or the **If X value is beyond limit** setting.

If X value is beyond limit

Determines how the plot changes if one or more X values are not within the limits shown on the X axis. For details, see “Varying Axis Limits Automatically” on page 9-5.

Initial Y axis lower limit, Initial Y axis upper limit

The interval shown on the Y axis at the beginning of the simulation. The interval might change from this initial setting

Instantaneous Event Counting Scope

due to zooming, autoscaling, or the **If Y value is beyond limit** setting.

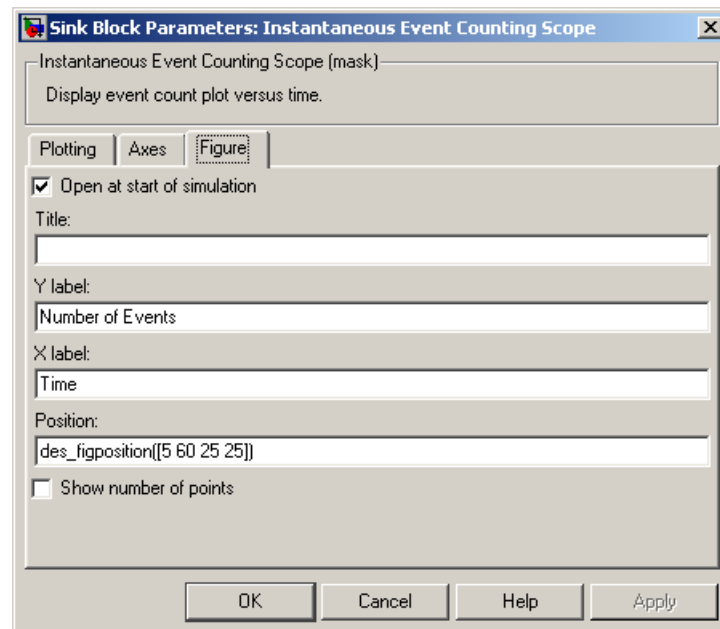
If Y value is beyond limit

Determines how the plot changes if one or more event counts are not within the limits shown on the Y axis. For details, see “Varying Axis Limits Automatically” on page 9-5.

Show grid

Toggles the grid on and off.

Figure Tab



Open at start of simulation

Selecting this option causes the plot window to open when you start the simulation. If you clear this box, you can open the plot

Instantaneous Event Counting Scope

window while the simulation is running by double-clicking the block icon.

Title

Text that appears as the title of the plot, above the axes.

Y label

Text that appears to the left of the vertical axis.

X label

Text that appears below the horizontal axis.

Position

A four-element vector of the form [left bottom width height] specifying the position of the scope window. (0,0) is the lower left corner of the display.

Show number of points

Displays the number of plotted points using an annotation in the plotting window.

Examples

See “Example: Plotting Event Counts to Find Roundoff Error” on page 9-12.

See Also

Signal Scope, Instantaneous Entity Counting Scope, Chapter 9, “Plotting Data”, “Observing Events” on page 2-30

LIFO Queue

Purpose Store entities in stack for undetermined length of time

Library Queues

Description



This block stores up to N entities simultaneously, where N is the **Capacity** parameter value. The block attempts to output an entity through the **OUT** port but retains the entity if the **OUT** port is blocked. If the block is storing multiple entities, then entities depart in a last-in, first-out (LIFO) fashion. The length of time that an entity stays in this block cannot be determined in advance.

The **IN** port is unavailable whenever this block stores exactly N entities. In this case, the queue is said to be full.

Ports

Entity Input Ports

Label	Description
IN	Port for arriving entities, which will be stored.

Entity Output Ports

Label	Description
OUT	Port for departing entities.

Signal Output Ports

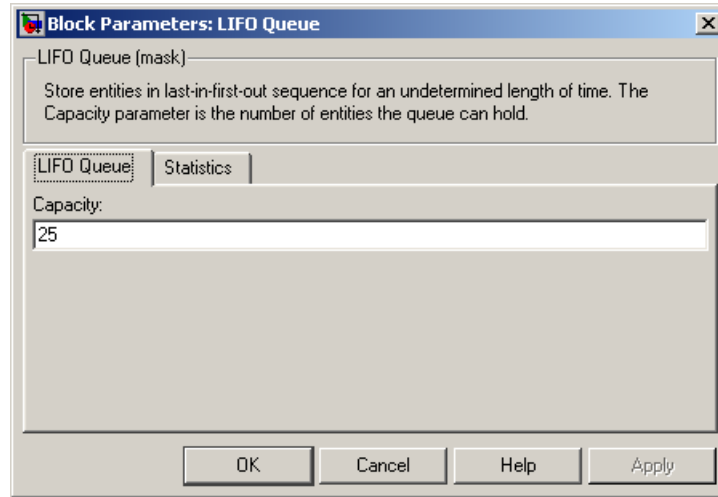
Label	Description	Time of Update When Statistic Is On	Order of Update
#d	Number of entities that have departed from this block since the start of the simulation.	After entity departure	4
#n	Number of entities currently in the queue.	After entity arrival at a nonempty queue and after entity departure	3
pe	A value of 1 indicates when the block tries and fails to output an entity; that is, the queue is nonempty. A value of 0 indicates when the queue is empty.	After entity arrival and before entity departure	2
w	Sample mean of the waiting times in this block for all entities that have departed.	After entity departure	1
len	Average number of entities in the queue over time, that is, the time average of the #n signal.	After entity arrival at a nonempty queue and after entity departure	1

Output signals having the same number in the Order of Update column above are updated in an arbitrary sequence relative to each other; you should not rely on a specific sequence for your simulation results.

LIFO Queue

Dialog Box

LIFO Queue Tab

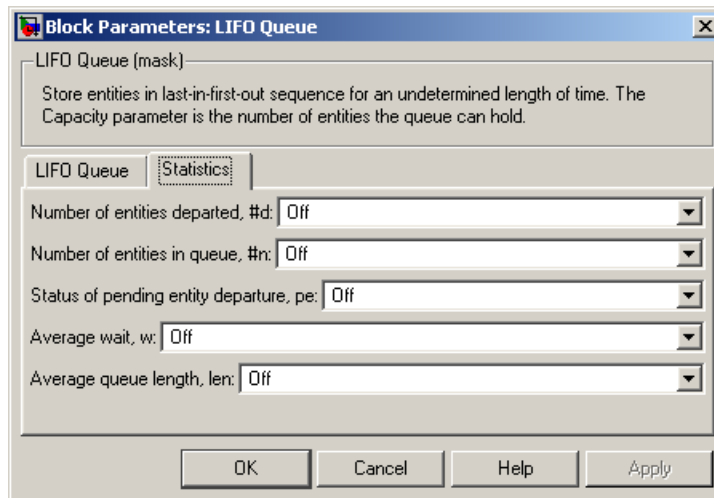


Capacity

Determines how many entities the block can store at a time.

Statistics Tab

These parameters determine whether certain ports output data throughout the simulation, output data only when you stop or pause the simulation, or are omitted from the block. For descriptions of the affected ports, see the “Signal Output Ports” table above.



Number of entities departed

Controls the presence and behavior of the signal output port labeled **#d**.

Number of entities in queue

Controls the presence and behavior of the signal output port labeled **#n**.

Status of pending entity departure

Controls the presence of the signal output port labeled **pe**.

Average wait

Controls the presence and behavior of the signal output port labeled **w**.

Average queue length

Controls the presence and behavior of the signal output port labeled **len**.

Examples

See “Example: Waiting Time in LIFO Queue” on page 4-2.

LIFO Queue

See Also

FIFO Queue, Priority Queue, “Using a LIFO Queuing Discipline” on page 4-2

Purpose Serve up to N entities for period of time

Library Servers

Description



This block stores up to N entities, serving each one independently for a period of time and then attempting to output the entity through the **OUT** port. If the **OUT** port is blocked, then the entity stays in this block until the port becomes unblocked. An N-server is like a set of N single servers connected in parallel, followed by a path combiner; the path combiner notifies entities of an unblocked path in the sequence in which the entities completed their service time, until one entity departs.

You specify the service time, which is the duration of service, via a parameter, attribute, or signal, depending on the **Service time from** parameter value. The block determines the service time for an entity upon its arrival.

All entities that arrive do so via the **IN** port. The **IN** port is unavailable whenever this block contains N entities. In that case, the **IN** port becomes available when at least one of the N entities departs.

Ports

Entity Input Ports

Label	Description
IN	Port for arriving entities, which will be served.

Signal Input Ports

Label	Description
t	Service time for a newly arrived entity. This port appears only if you set Service time from to Signal port t.

N-Server

Entity Output Ports

Label	Description
OUT	Port for departing entities, which have completed their service time.

Signal Output Ports

Label	Description	Time of Update When Statistic Is On	Order of Update
#d	Number of entities that have departed from this block since the start of the simulation.	After entity departure	4
#n	Number of entities currently in the block, between 0 and N.	After entity arrival and after entity departure	3
pe	A value of 1 indicates when the block contains at least one entity that has completed its service but cannot depart because the subsequent entity input port is unavailable.	After service completion and after entity departure	2
#pe	Number of entities that have completed their service but cannot depart because the subsequent entity input port is unavailable.	After service completion and after entity departure	2

Label	Description	Time of Update When Statistic Is On	Order of Update
w	Sample mean of the waiting times in this block for all entities that have departed. An entity's waiting time might exceed its service time if the OUT port is blocked when the entity completes service.	After entity departure	1
util	Utilization of the N-server. If Number of servers is finite, util is the time average of the fraction of servers that are storing an entity. At time values when an entity arrives or departs, util equals 1/N times the time average of the #n signal. If Number of servers is infinite, then util is always zero.	Performance considerations cause the block to suppress signal updates until specific occurrences cause updates. In On mode, updates occur after an entity departure and after an entity arrival. In Upon stop or pause mode, stopping or pausing the simulation causes an update.	1

Output signals having the same number in the Order of Update column above are updated in an arbitrary sequence relative to each other; you should not rely on a specific sequence for your simulation results.

A more precise definition of the utilization signal **util** at an update time $T > 0$ is

$$\frac{1}{T} \sum_k \left(\frac{(\#n)_k}{N} \right) \cdot \text{length}(I_k)$$

where I_k is the k th time interval between successive pairs of times that

util is updated and $(\#n)_k$ is the number of entities the N-Server block is storing during the open interval I_k . If an update of **util** occurs at $T=0$, the value is $\#n/N$.

Dialog Box

N-Server Tab

Block Parameters: N-Server

N-Server (mask):
Serve up to N entities independently for a period of time and then attempt to output each entity through the OUT port. If the OUT port is blocked then the pending entity stays in this block until the port becomes unblocked. You can specify the service time, which is the duration of service, via a parameter, attribute, or signal.
The IN port is unavailable whenever the number of entities stored is N.

N-Server | Statistics

Number of servers:
5

Service time from: Dialog

Service time:
1

Service completion event priority:
500

OK Cancel Help Apply

Number of servers

The number of servers the block represents, N.

Service time from

Determines whether the service time is computed from a parameter in this dialog box, an input signal, or an attribute of the entity being served.

Service time

The service time for all entities. This field appears only if you set **Service time from** to Dialog.

Attribute name

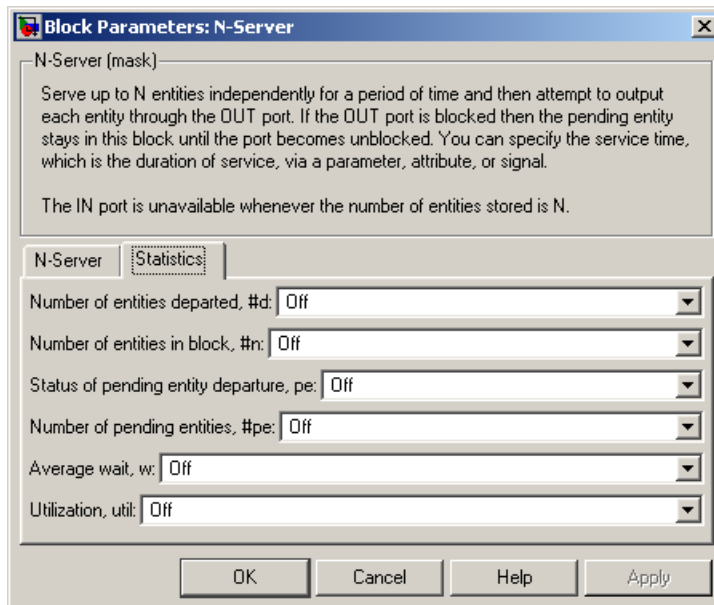
The name of the attribute whose value the block uses as the service time for an entity. This field appears only if you set **Service time from** to Attribute.

Service completion event priority

The priority of the service completion event, relative to other simultaneous events.

Statistics Tab

These parameters determine whether certain ports output data throughout the simulation, output data only when you stop or pause the simulation, or are omitted from the block. For descriptions of the affected ports, see the “Signal Output Ports” table above.



Number of entities departed

Controls the presence and behavior of the signal output port labeled **#d**.

Number of entities in block

Controls the presence and behavior of the signal output port labeled **#n**.

N-Server

Status of pending entity departure

Controls the presence of the signal output port labeled **pe**.

Number of pending entities

Controls the presence and behavior of the signal output port labeled **#pe**.

Average wait

Controls the presence and behavior of the signal output port labeled **w**.

Utilization

Controls the presence and behavior of the signal output port labeled **util**.

Examples

See “Example: M/M/5 Queuing System” on page 4-12.


See Also

Single Server, Infinite Server, “Modeling Multiple Servers” on page 4-12

Purpose Select entity output port for departure

Library Routing

Description



This block receives entities and outputs them through one of multiple entity output ports. The selected port can change during the simulation. The **Switching criterion** parameter indicates how the block determines which entity output port is selected for departure at any given time. When the selected port is not blocked, an arriving entity departs through this port. When the selected port is blocked, this block's entity input port is unavailable and entities cannot arrive.

The values of the **Switching criterion** parameter are described in the table below.

Switching criterion Value	Description
Round robin	The first arriving entity in the simulation departs via the OUT1 port. Upon each subsequent arrival, the block selects the entity output port next to the last selected port. After exhausting all entity output ports, the block returns to the first one, OUT1 .
Equiprobable	At the beginning of the simulation and upon each departure, the block randomly chooses the entity output port through which the next arriving entity departs. All entity output ports are equally likely to be selected. The Initial seed parameter initializes the random number generation process.

Output Switch

Switching criterion Value	Description
First port that is not blocked	When an entity attempts to arrive, the block attempts to output the entity through OUT1 . If that port is blocked, then the block attempts to output the entity through OUT2 , and so on. If all entity output ports are blocked, then this block's IN port is unavailable and the entity cannot arrive.
From signal port p	Selecting this option creates an additional signal input port, labeled p . The signal at this port uses integer values between 1 and the Number of entity output ports parameter value to refer to entity output ports. The block monitors the p signal's value throughout the simulation. When an entity arrives and the entity output port corresponding to the p signal is not blocked, the entity departs through that entity output port. If the indicated entity output port is blocked, then this block does not accept the entity for arrival until either the port becomes unblocked or the p signal changes to indicate an entity output port that is not blocked.
From attribute	The block outputs an arriving entity through the entity output port that corresponds to the value of an attribute of your choice. Name the attribute using the Attribute name parameter. The attribute value must be an integer between 1 and the Number of entity output ports parameter value. If the indicated entity output port is blocked, then this block does not accept the entity for arrival until the entity output port becomes unblocked.

Ports

Entity Input Ports

Label	Description
IN	Port for arriving entities.

Signal Input Ports

Label	Description
p	Index of the entity output port through which an arriving entity departs. Values are 1, 2, 3,..., Number of entity output ports . This port appears only if you set Switching criterion to From signal port p.

Entity Output Ports

Label	Description
OUT1, OUT2, OUT3, etc.	Entity ports through which an arriving entity departs, where the Switching criterion parameter determines which of multiple ports the entity departs through. The Number of entity output ports parameter determines how many of these entity output ports the block has.

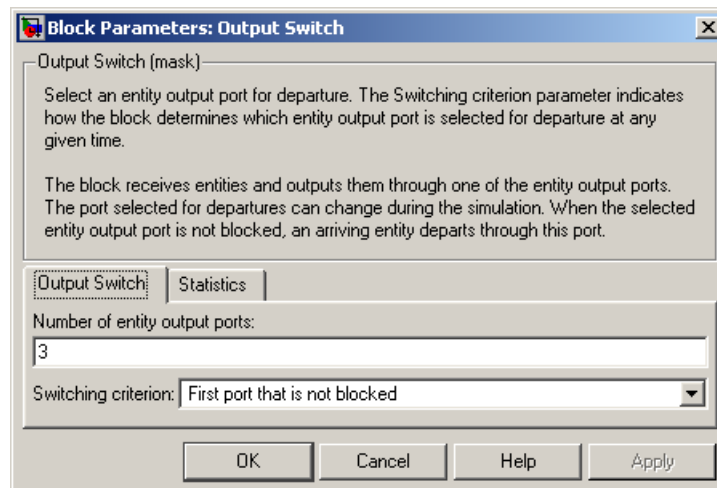
Output Switch

Signal Output Ports

Label	Description	Time of Update When Statistic Is On	Order of Update
#d	Number of entities that have departed from this block since the start of the simulation.	After entity departure	2
last	Index of the output port through which the last entity departed. The initial value is 0. After an entity has departed, values of this signal are 1, 2, 3,..., Number of entity output ports .	After entity departure	1

Dialog Box

Output Switch Tab



Number of entity output ports

Determines how many entity output ports the block has.

Switching criterion

The rule that determines which entity output port an arriving entity departs through.

Initial seed

A nonnegative integer that initializes the random number generator used to select an entity output port. This field appears only if you set **Switching criterion** to Equiprobable.

Attribute name

The name of an attribute used to select an entity output port. This field appears only if you set **Switching criterion** to From attribute.

Specify event priority for port selection

Select this option to control the sequencing of the port-selection event, relative to other simultaneous events. This field appears only if you set **Switching criterion** to From signal port p.

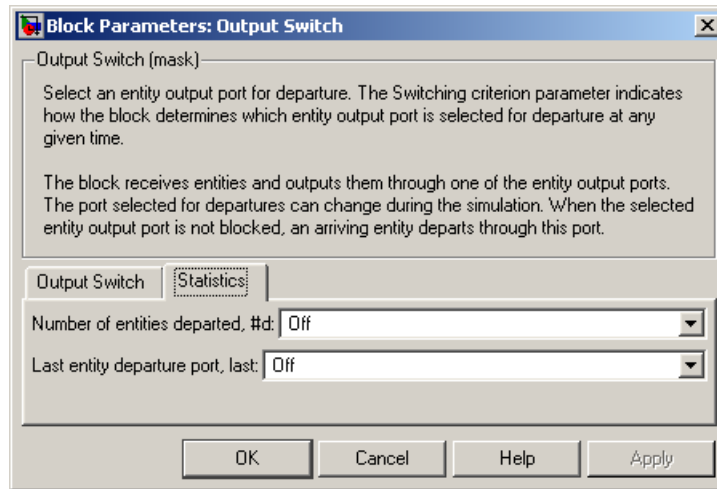
Event priority

The priority of the port-selection event, relative to other simultaneous events. This field appears only if you set **Switching criterion** to From signal port p and select **Specify event priority for port selection**.

Statistics Tab

These parameters determine whether certain ports output data throughout the simulation, output data only when you stop or pause the simulation, or are omitted from the block. For descriptions of the affected ports, see the “Signal Output Ports” table above.

Output Switch



Number of entities departed

Controls the presence and behavior of the signal output port labeled **#d**.

Last entity departure port

Controls the presence of the signal output port labeled **last**.

Examples

- “Example: Selecting the First Available Server” in the Getting Started documentation
- “Example: Using an Attribute to Select an Output Port” in the Getting Started documentation
- “Example: A Packet Switch” in the Getting Started documentation
- “Example: Choosing the Shortest Queue” on page 6-3
- “Example: Using Servers in Shifts” on page 6-11

See Also

Input Switch, “Using the Output Switch” in the Getting Started documentation

Purpose Merge entity paths

Library Routing

Description



This block accepts entities through any entity input port and outputs them through a single entity output port. You specify the number of entity input ports using the **Number of entity input ports** parameter.

If multiple entities arrive at the Path Combiner block simultaneously while the entity output port is not blocked, then the sequence in which the entities depart depends on the sequence of departure events from blocks that precede the Path Combiner block. Although the departure time is the same for all such entities, the sequence might affect the system’s behavior. For example, if the entities advance to a queue, the departure sequence determines their positions in the queue.

Input Port Precedence

The **Input port precedence** parameter indicates how the block determines which entity input port to make available first, whenever the entity output port changes from blocked to unblocked. Choices are in the table below. For an example illustrating when this parameter is significant, see “Combining Entity Paths” in the Getting Started documentation.

Input Port Precedence	Description	Example
IN1 port	Each time the entity output port becomes unblocked, notify entity input ports IN1 , IN2 , IN3 ,... about the change in status until either an entity arrives or all ports are notified.	The sequence of notifications is always IN1 , IN2 , IN3 ,... throughout the simulation.

Path Combiner

Input Port Precedence	Description	Example
Equiprobable	Each time the entity output port becomes unblocked, choose a random entity input port to notify about the change in status. All entity output ports are equally likely to be selected and the Initial seed parameter initializes the random number generation process. If this does not result in an entity arrival, then notify the subsequent ports in turn until either an entity arrives or all ports are notified.	On a block with four entity input ports, if the random number is three, then notify the ports in the sequence IN3, IN4, IN1, IN2 . If the random number is two on the next such occasion, then notify the ports in the sequence IN2, IN3, IN4, IN1 .

Input Port Precedence	Description	Example
Round robin	Each time the entity output port becomes unblocked, notify the port next to the one through which the last departing entity arrived. The IN1 port is considered “next to” the last entity input port on the block. If this does not result in an entity arrival, then notify the subsequent ports in turn until either an entity arrives or all ports are notified.	On a block with four entity input ports, an entity arrives through the IN2 port and advances to a Single Server block, whose entity input port becomes unavailable. Meanwhile, entities attempt to arrive at the Path Combiner block. When the server becomes available, the Path Combiner block notifies the ports in the sequence IN3, IN4, IN1, IN2 . This sequence starts with IN3 because it is next to IN2 , which is the port through which the last departing entity arrived.
From signal port p	Each time the entity output port becomes unblocked, notify the port whose index is the value of the p input signal. If this does not result in an entity arrival, then notify the subsequent ports in turn until either an entity arrives or all ports are notified.	On a block with four entity input ports, if the value of the p signal is three, then notify the ports in the sequence IN3, IN4, IN1, IN2 . If the value of the p signal is two on the next such occasion, then notify the ports in the sequence IN2, IN3, IN4, IN1 .

Path Combiner

Ports

Entity Input Ports

Label	Description
IN1, IN2, IN3, etc.	Port for arriving entities. The Number of entity input ports parameter determines how many of these entity input ports the block has.

Signal Input Ports

Label	Description
p	Index of the entity input port that the block makes available first, upon an event that changes the entity output port from blocked to unblocked. Values are 1, 2, 3,..., Number of entity input ports . This port appears only if you set Input port precedence to From signal port p.

Entity Output Ports

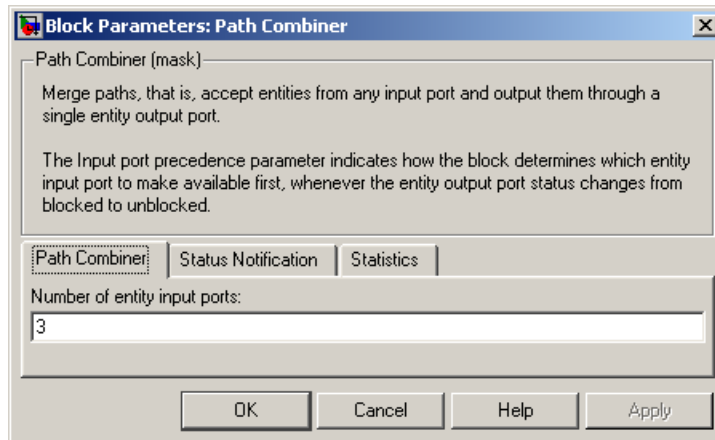
Label	Description
OUT	Port for departing entities.

Signal Output Ports

Label	Description	Time of Update When Statistic Is On	Order of Update
#d	Number of entities that have departed from this block since the start of the simulation.	After entity departure	2
last	Index of the input port through which the last entity arrived. The initial value is 0. After an entity has arrived and departed, values are 1, 2, 3,..., Number of entity input ports .	After entity departure	1

Dialog Box

Path Combiner Tab




Number of entity input ports

Determines how many entity input ports the block has.

Path Combiner

Status Notification Tab



Input port precedence

Determines which entity input port the block makes available first, upon an event that changes the entity output port from blocked to unblocked.

Initial seed

A nonnegative integer that initializes the random number generator used to select an entity input port to notify first about status changes. This field appears only if you set **Input port precedence** to Equiprobable.

Specify event priority for port precedence selection


Select this option to control the sequencing of the event that updates the port precedence, relative to other simultaneous events. This field appears only if you set **Switching criterion** to From signal port p.

Event priority

The priority of the event that updates the port precedence, relative to other simultaneous events. This field appears only if you set **Switching criterion** to From signal port p and select **Specify event priority for port selection**.

Statistics Tab

These parameters determine whether certain ports output data throughout the simulation, output data only when you stop or pause the simulation, or are omitted from the block. For descriptions of the affected ports, see the “Signal Output Ports” table above.



Number of entities departed

Controls the presence and behavior of the signal output port labeled **#d**.

Last entity arrival port

Controls the presence and behavior of the signal output port labeled **last**.

Examples

- “Combining Entity Paths” in the Getting Started documentation
- “Example: A Packet Switch” in the Getting Started documentation
- “Example: First Entity as a Special Case” on page 7-10

See Also


Input Switch, Output Switch, “Combining Entity Paths” in the Getting Started documentation

Priority Queue

Purpose Store entities in sorted sequence for undetermined length of time

Library Queues

Description



This block stores up to N entities simultaneously in a sorted sequence, where N is the **Capacity** parameter value. The queue sorts entities according to the values of an attribute, in either ascending or descending order. Use the **Sorting attribute name** and **Sorting direction** parameters to determine the sorting behavior. The block accepts real numbers, Inf, and -Inf as valid values of the sorting attribute.

The block attempts to output an entity through the **OUT** port but retains the entity if the **OUT** port is blocked. The length of time that an entity stays in this block cannot be determined in advance. The **IN** port is unavailable whenever this block stores exactly N entities. In this case, the queue is said to be full.

While you can view the value of the sorting attribute as an entity priority, this value has nothing to do with event priorities or block priorities.

Ports

Entity Input Ports

Label	Description
IN	Port for arriving entities, which will be stored.

Entity Output Ports

Label	Description
OUT	Port for departing entities.

Signal Output Ports

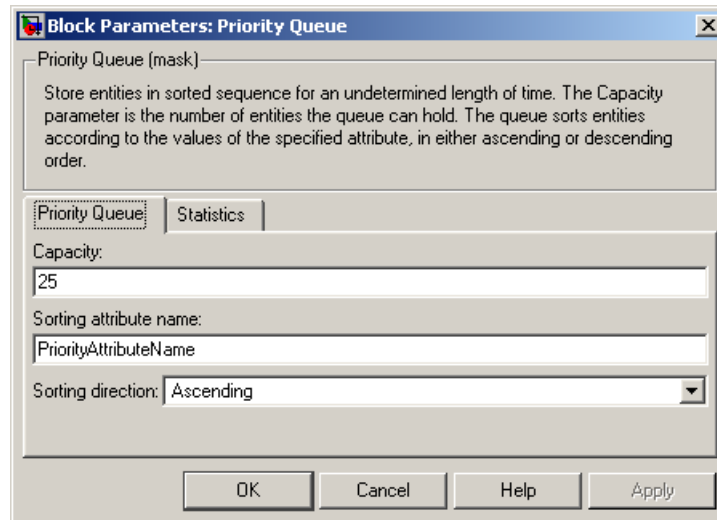
Label	Description	Time of Update When Statistic Is On	Order of Update
#d	Number of entities that have departed from this block since the start of the simulation.	After entity departure	4
#n	Number of entities currently in the queue.	After entity arrival at a nonempty queue and after entity departure	3
pe	A value of 1 indicates when the block tries and fails to output an entity; that is, the queue is nonempty. A value of 0 indicates when the queue is empty.	After entity arrival and after entity departure	2
w	Sample mean of the waiting times in this block for all entities that have departed.	After entity departure	1
len	Average number of entities in the queue over time, that is, the time average of the #n signal.	After entity arrival at a nonempty queue and after entity departure	1

Output signals having the same number in the Order of Update column above are updated in an arbitrary sequence relative to each other; you should not rely on a specific sequence for your simulation results.

Priority Queue

Dialog Box

Priority Queue Tab



Capacity

Determines how many entities the block can store at a time.

Sorting attribute name

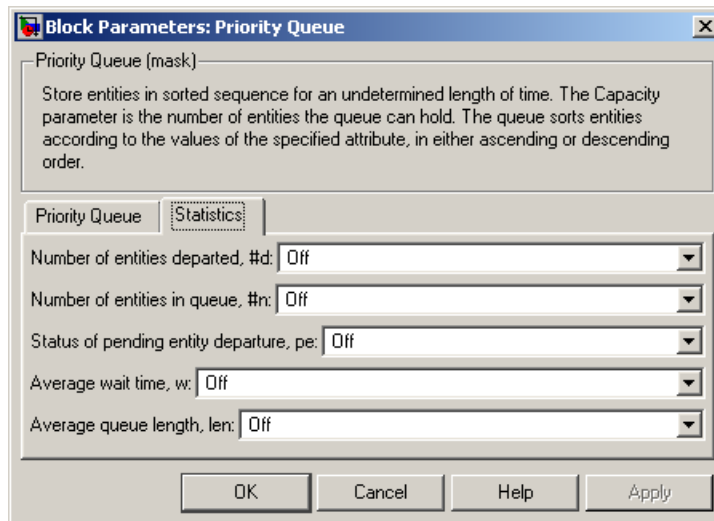
The block uses this attribute to sort entities in the queue.

Sorting direction

Determines whether the entity at the head of the queue is the one with the smallest (Ascending) or largest (Descending) value of the attribute named above. Entities sharing the same attribute value are sorted in FIFO sequence.

Statistics Tab

These parameters determine whether certain ports output data throughout the simulation, output data only when you stop or pause the simulation, or are omitted from the block. For descriptions of the affected ports, see the “Signal Output Ports” table above.



Number of entities departed

Controls the presence and behavior of the signal output port labeled **#d**.

Number of entities in queue

Controls the presence and behavior of the signal output port labeled **#n**.

Status of pending entity departure

Controls the presence of the signal output port labeled **pe**.

Average wait

Controls the presence and behavior of the signal output port labeled **w**.

Average queue length

Controls the presence and behavior of the signal output port labeled **len**.

Examples

- “Example: Serving Preferred Customers First” on page 4-6
- “Example: Preemption by High-Priority Entities” on page 4-10

Priority Queue

See Also

FIFO Queue, LIFO Queue, Single Server, “Sorting by Priority” on page 4-4

Purpose Report statistical data about named timer associated with arriving entities

Library Timing

Description



This block reads the value of a timer that the Start Timer block previously associated with the arriving entity. Using the **Report elapsed time** and **Report average elapsed time** parameters, you can configure the block to report the following statistics via the **et** and **w** signal output ports, respectively:

- The instantaneous value from the named timer associated with the arriving entity
- The average of **et** values among all entities that have arrived at this block during the simulation and possessed a timer of the specified name

Note If the arriving entity does not possess a timer of that name, then you can configure the block to either produce an error or ignore the timer's absence. In the latter case, the output signals maintain their previous values.

The timer continues timing after the entity departs from this block, which is relevant if the same entity arrives at another Read Timer block later in the simulation.

For more information about using this block with the Start Timer block, see “Using Timers” on page 10-18.

Read Timer

Ports

Entity Input Ports

Label	Description
IN	Port for arriving entities.

Entity Output Ports

Label	Description
OUT	Port for departing entities.

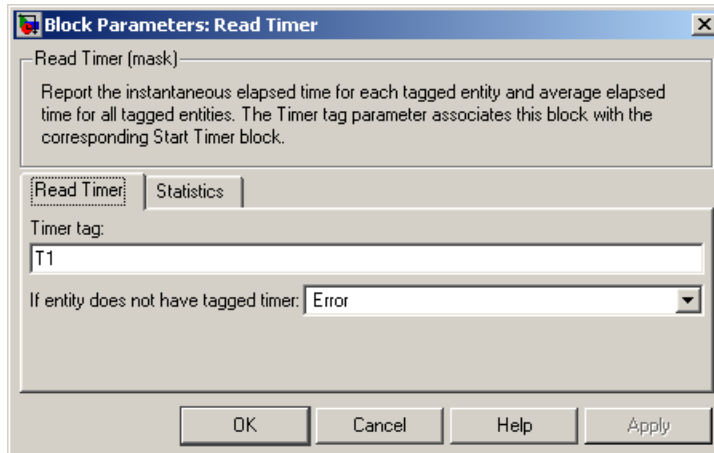
Signal Output Ports

Label	Description	Time of Update When Statistic Is On	Order of Update
et	Instantaneous elapsed time for the arriving entity, if it possesses a timer of the specified name.	After entity departure	2
w	Average among the et values for all entities that have arrived at this block and possessed a timer of the specified name.	After entity departure	1
#d	Number of entities that have departed from this block since the start of the simulation.	After entity departure	3
#t	Total number of entities that have departed from this block and possessed a timer of the specified name.	After entity departure	2

Output signals having the same number in the Order of Update column above are updated in an arbitrary sequence relative to each other; you should not rely on a specific sequence for your simulation results.

Dialog Box

Read Timer Tab



Timer tag

Name of the timer to read. This timer tag corresponds to the **Timer tag** parameter of a Start Timer block in the model.


If entity does not have tagged timer

Behavior of the block if an arriving entity does not possess a timer with the specified timer tag.

Statistics Tab

These parameters determine whether certain ports output data throughout the simulation, output data only when you stop or pause the simulation, or are omitted from the block. For descriptions of the affected ports, see the “Signal Output Ports” table above.

Read Timer



Number of entities departed

Controls the presence and behavior of the signal output port labeled **#d**.

Number of entities departed with specified tag

Controls the presence and behavior of the signal output port labeled **#t**. If you set **If entity does not have tagged timer to Ignore**, then the **#t** value might be less than the **#d** value.

Report elapsed time

Controls the presence of the signal output port labeled **et**.

Report average elapsed time

Controls the presence and behavior of the signal output port labeled **w**.

Examples

- “Basic Procedure for Using Timer Blocks” on page 10-19
- “Timing Multiple Entity Paths with One Timer” on page 10-20
- “Restarting a Timer” on page 10-21
- “Timing Multiple Processes Independently” on page 10-23

See Also

Start Timer, “Using Timers” on page 10-18

Release Gate

Purpose Permit one pending entity to arrive when event occurs

Library Gates

Description



This block permits the arrival of one pending entity when a signal-based event or function call occurs; at all other times, the block's entity input port is unavailable. By definition, the gate's opening permits one pending entity to arrive if able to advance immediately to the next block.

No simulation time passes between the opening and subsequent closing of the gate; that is, the gate opens and then closes in the same time instant. If no entity is already pending when the gate opens, then the gate closes without processing any entities.

The **Open gate upon** parameter determines the type of event that opens the gate. Possible types are

- Edges in a trigger signal
- Changes in the numerical value of a signal
- Function calls

For more details, see "Opening a Gate Instantaneously" on page 7-6.

Ports

Entity Input Ports

Label	Description
IN	Port for arriving entities.

Signal Input Ports

Label	Description
tr	Trigger signal that indicates when to open the gate. This port appears only if you set Open gate upon to Trigger from port tr.
vc	Control signal whose numerical changes indicate when to open the gate. This port appears only if you set Open gate upon to Change in signal from port vc.
fcn	Function-call signal that indicates when to open the gate. This port appears only if you set Open gate upon to Function call from port fcn.

Entity Output Ports

Label	Description
OUT	Port for departing entities.


Signal Output Ports

Label	Description	Time of Update When Statistic Is On
#d	Number of entities that have departed from this block since the start of the simulation.	After entity departure

Release Gate

Dialog Box

Release Gate Tab



Open gate upon

Determines the type of event that causes the gate to open instantaneously.

Trigger type

Determines whether rising, falling, or either type of trigger edge causes the gate to open. This field appears only if you set **Open gate upon** to Trigger from port tr.

Type of change in signal value

Determines whether rising, falling, or either type of value change causes the gate to open. This field appears only if you set **Open gate upon** to Change in signal from port vc.

Specify event priority for gate opening


Select this option to control the sequencing of the gate-opening event, relative to other simultaneous events.

Event priority

The priority of the gate-opening event, relative to other simultaneous events. This field appears only if you select **Specify event priority for gate opening**.

Statistics Tab

These parameters determine whether certain ports output data throughout the simulation, output data only when you stop or pause the simulation, or are omitted from the block. For descriptions of the affected ports, see the “Signal Output Ports” table above.



Number of entities departed

Controls the presence and behavior of the signal output port labeled #d.

Examples

- “Example: Synchronizing Service Start Times with the Clock” on page 7-6
- “Example: First Entity as a Special Case” on page 7-10

See Also

Enabled Gate, Chapter 7, “Regulating Arrivals Using Gates”

Replicate

Purpose Output copies of entity

Library Routing

Description



This block outputs a copy of the arriving entity through each entity output port that is not blocked. You specify the number of copies that the block makes, using the **Number of entity output ports** parameter.

Complete or Partial Replication

The **Replicate entity when** parameter affects the circumstances under which the block accepts an entity to replicate. Choices are in the table below.

Parameter Value	Description
All entity output ports are not blocked	The block accepts an entity to replicate only when all entity output ports are connected to available ports of subsequent blocks.
Any entity output port is not blocked	The block accepts an entity to replicate when at least one entity output port is connected to an available port of a subsequent block.

If you connect multiple copies of this block, you can implement logical combinations of the parameter values in the table.

Departure of Copies

Each time the block replicates an entity, the copies depart in a sequence whose start is determined by the **Departure port precedence** parameter. Choices are in the table below.

Parameter Value	Description	Example
OUT1 port	Each time the block replicates an entity, the copies depart via entity output ports OUT1 , OUT2 , OUT3 ,..., in that sequence.	The sequence of departures is always OUT1 , OUT2 , OUT3 ,... throughout the simulation.
Round robin	Each time the block replicates an entity, the first copy departs via the port after the one that received preference on the last such occasion. The remaining copies depart via the subsequent ports in turn.	On a block with three entity output ports, the first time the block replicates an entity, the copies depart in the sequence OUT1 , OUT2 , OUT3 . The second time, the copies depart in the sequence OUT2 , OUT3 , OUT1 . The third time, the copies depart in the sequence OUT3 , OUT1 , OUT2 . The fourth time is analogous to the first time, and so on.
Equiprobable	Each time the block replicates an entity, the first copy departs via a randomly selected entity output port. All entity output ports are equally likely to be selected and the Initial seed parameter initializes the random number generation process. The remaining copies depart via the subsequent ports in turn.	On a block with four entity output ports, if the random number is three, then the copies depart in the sequence OUT3 , OUT4 , OUT1 , OUT2 . If the random number is two on the next such occasion, then the copies depart in the sequence OUT2 , OUT3 , OUT4 , OUT1 .

An example in which the choice of **Departure port precedence** parameter is relevant is a model that sets an attribute on each replicated entity based on its departure port and then advances all replicated entities along a merged path to a FIFO Queue block. At each replication time during the simulation, the **Departure port**

Replicate

precedence parameter determines the sequence of the replicated entities in the queue.

In some cases, a departure through one entity output port causes another entity output port to become newly blocked. For example, this could happen if two entity output ports connect to a Path Combiner block, which in turn connects to a Single Server block whose service time is nonzero. Use the **If an output port becomes blocked during replication** parameter to determine how the block responds. Choices are in the table below.

Parameter Value	Description
Discard entity	The block discards the entity that is supposed to depart through the newly blocked entity output port.
Warn and discard entity	The block issues a warning message in the MATLAB Command Window, and discards the entity that is supposed to depart through the newly blocked entity output port.
Error	The simulation halts with an error message.

Ports

Entity Input Ports

Label	Description
IN	Port for arriving entities.

Entity Output Ports

Label	Description
OUT1, OUT2, OUT3, etc.	Port for departing entities, which are copies of the arriving entity. The Number of entity output ports parameter determines how many of these entity input ports the block has.

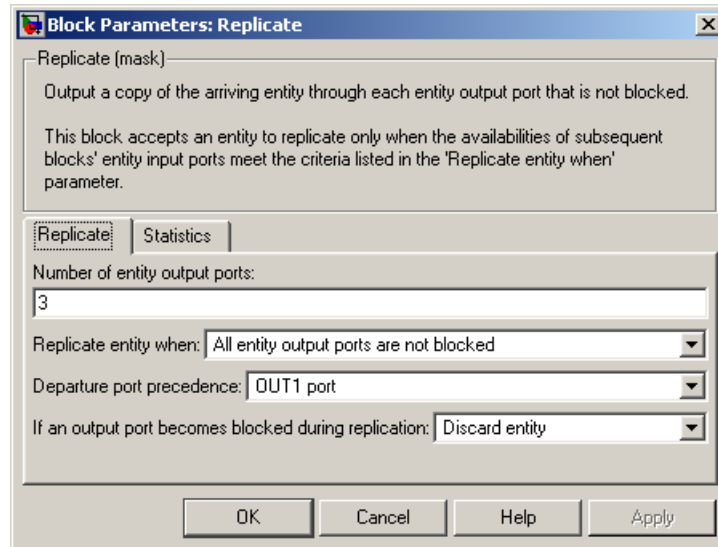
Signal Output Ports

Label	Description	Time of Update When Statistic Is On	Order of Update
#a	Number of entities that have arrived at this block since the start of the simulation.	After entity arrival	1
#d	Number of entities that have departed from this block since the start of the simulation.	After each entity departure	2

Replicate

Dialog Box

Replicate Tab



Number of entity output ports

Determines how many entity output ports the block has; that is, the maximum number of copies the block makes for each arriving entity.

Replicate entity when

Determines whether the block is available to arriving entities whenever at least one entity output port is not blocked, or only when all entity output ports are not blocked.

Departure port precedence

Determines the start of the sequence in which the block outputs the copies, each time the block replicates an entity.

Initial seed

A nonnegative integer that initializes the random number generator used to determine the output sequence. This field

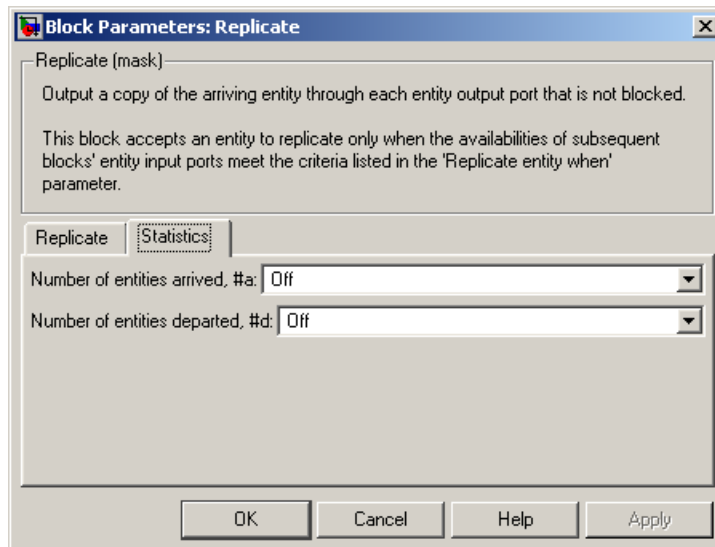
appears only if you set **Departure port precedence** to Equiprobable.

If an output port becomes blocked during replication

Determines whether the block issues a message when a replicated entity is unable to depart because an output port becomes blocked during the replication process. This field appears only if you set **Replicate entity when** to All entity output ports are not blocked.

Statistics Tab

These parameters determine whether certain ports output data throughout the simulation, output data only when you stop or pause the simulation, or are omitted from the block. For descriptions of the affected ports, see the “Signal Output Ports” table above.



Number of entities arrived

Controls the presence and behavior of the signal output port labeled **#a**.

Replicate

Number of entities departed

Controls the presence and behavior of the signal output port labeled **#d**.

Examples

See “Example: Waiting Time in LIFO Queue” on page 4-2.

See Also

Event-Based Entity Generator, Path Combiner, “Replicating Entities on Multiple Paths” on page 1-25

Purpose Assign data to entity

Library Attributes

Description



This block accepts an entity, assigns data to it, and then outputs it. Assigned data is stored in attributes of the entity, where each attribute has a name and a value.

You can assign up to four attributes with a single instance of this block. For each attribute you want to set, configure one of the tabs named **A1**, **A2**, **A3**, and **A4** in the block's dialog box using one of these procedures:

Assigning a Constant Value Using This Dialog Box

- 1 Set **Attribute assignment** to Specify via dialog.
- 2 Specify the name of attribute you want to set, in the **Attribute name** field. All valid MATLAB variable names, except `nan` and `inf` (with any use of case), are valid as attribute names. To determine whether a name is a valid MATLAB variable name, use the `isvarname` function.
- 3 Specify the constant value for the attribute in the **Value** field.

Assigning a Value Using an Input Signal

- 1 Set **Attribute assignment** to From signal port A_x , where x is 1, 2, 3, or 4.
- 2 Specify the name of attribute you want to set, in the **Attribute name** field. All valid MATLAB variable names, except `nan` and `inf` (with any use of case), are valid as attribute names. To determine whether a name is a valid MATLAB variable name, use the `isvarname` function.
- 3 Click **OK** or **Apply**. The block now has a signal input port labeled A_x .
- 4 Connect a signal to the A_x input port. During the simulation, the block assigns the value of this signal to the attribute.

Set Attribute

If you want to set fewer than the maximum number of attributes, then you can deactivate tabs that you are not using by setting those tabs' **Attribute assignment** parameters to Off.

Ports

Entity Input Ports

Label	Description
IN	Port for arriving entities.

Signal Input Ports

Label	Description
Ax	Data to assign for the attribute specified on the Ax tab of the dialog. This port appears only if you set Attribute assignment on the Ax tab to From signal port Ax.

Entity Output Ports

Label	Description
OUT	Port for departing entities, with data assigned to them.

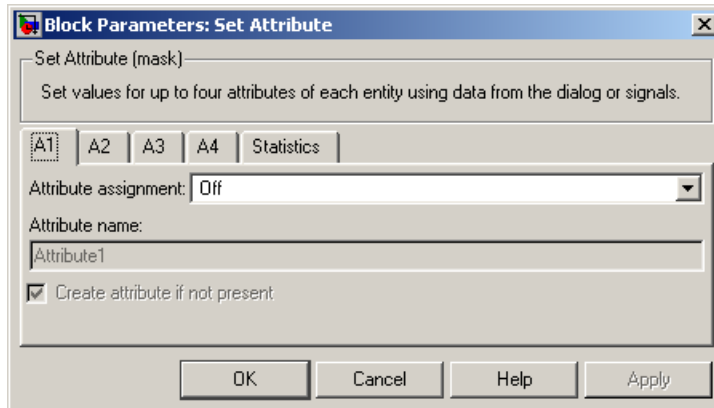
Signal Output Ports

Label	Description	Time of Update When Statistic Is On
#d	Number of entities that have departed from this block since the start of the simulation.	After entity departure

Dialog Box

A1, A2, A3, A4 Tabs

The **A1**, **A2**, **A3**, and **A4** tabs have similar parameter choices. By assigning different values to the parameters, you can configure this block to assign up to four different attributes for each entity that the block processes.



Attribute assignment

Determines whether the data for attribute values comes from the dialog or a signal. Choosing **Off** indicates that you are not using this tab of the dialog and makes the parameters below inactive or invisible.

Attribute name

The name of the attribute to set.

Value

The value to assign to the attribute named above. This parameter appears only if you set **Attribute assignment** to **Specify** via dialog.


Create attribute if not present

Selecting this option enables the block to define new attributes. Otherwise, the block issues an error if the attribute named above does not already exist.

Set Attribute

Statistics Tab

These parameters determine whether certain ports output data throughout the simulation, output data only when you stop or pause the simulation, or are omitted from the block. For descriptions of the affected ports, see the “Signal Output Ports” table above.



Number of entities departed

Controls the presence and behavior of the signal output port labeled **#d**.

Examples

- “Example: Setting Attributes” on page 1-15
- “Example: Using an Attribute to Select an Output Port” in the Getting Started documentation
- “Example: Round-Robin Approach to Choosing Inputs” in the Getting Started documentation
- “Adding Event-Based Behavior” in the Getting Started documentation

See Also

Get Attribute, “Setting Attributes of Entities” on page 1-14

Purpose Write input signal value to memory and read memory to output signal upon events

Library Gates

Description This block stores and outputs the values of the **in** input signal based on events:



- The block writes the value of the **in** signal to an internal memory location when a “write to memory” event occurs. The **Write to memory upon** parameter indicates the type of signal-based event or function call that causes a write event.
- The block reads the memory value and updates the signal at the **out** port, if present, when a “read from memory” event occurs. The **Read from memory upon** parameter indicates the type of internal or external event that causes a read event:
 - If you set **Read from memory upon** to Write to memory event, then every write event causes a read event. The **out** signal is like a resampled version of the **in** signal.
 - Otherwise, the **Read from memory upon** parameter indicates the type of signal-based event or function call that causes a read event. In this case, write and read events occur independently and are not required to alternate. The **out** signal is like a delayed resampled version of the **in** signal.

This block is useful for modeling feedback loops in discrete-event systems in which an output from one component is an input to another component. Because the two components work separately in such a system, the updates of the input and output signals are independent in both causality and timing. This block lets you control the causality and timing associated with storing the output from one component and updating the value received by the other component. For an example that uses this block in a feedback loop, see the Dynamic Voltage Scaling Using Online Gradient Estimation demo.

Signal Latch

Ports

Signal Input Ports

Label	Description
wts	Signal whose updates cause write events. This port appears only if you set Write to memory upon to Sample time hit from port wts.
wtr	Trigger signal whose edges cause write events. This port appears only if you set Write to memory upon to Trigger from port wtr.
wvc	Signal whose numerical changes in value cause write events. This port appears only if you set Write to memory upon to Change in signal from port wvc.
wfcn	Function-call signal that causes write events. This port appears only if you set Write to memory upon to Function call from port wfcn.
rts	Signal whose updates cause read events. This port appears only if you set Read from memory upon to Sample time hit from port rts.
rtr	Trigger signal whose edges cause read events. This port appears only if you set Read from memory upon to Trigger from port rtr.
rvc	Signal whose numerical changes in value cause read events. This port appears only if you set Read from memory upon to Change in signal from port rvc.
rfcn	Function-call signal that causes read events. This port appears only if you set Read from memory upon to Function call from port rfcn.
in	Signal to be resampled and/or delayed.

Signal Output Ports


Label	Description	Time of Update When Statistic Is On	Order of Update
st	0 or 1, depending on whether the block more recently processed a read or write event.	Upon write events and upon read events	1
mem	The value of the block's internal memory when a write event occurs.	Upon write events	1
out	The value of the block's internal memory when a read event occurs.	Upon read events	1

Output signals having the same number in the Order of Update column above are updated in an arbitrary sequence relative to each other; you should not rely on a specific sequence for your simulation results.

Signal Latch

Dialog Box

Write Tab



Initial memory value

The value in the block's internal memory before the first write event occurs.

Write to memory upon

The type of signal-based event or function call that causes a write event.

Trigger type

Determines whether rising, falling, or either type of trigger edge causes a write event. This field appears only if you set **Write to memory upon** to Trigger from port wtr.

Type of change in signal value

Determines whether rising, falling, or either type of value change causes a write event. This field appears only if you set **Write to memory upon** to Change in signal from port wvc.

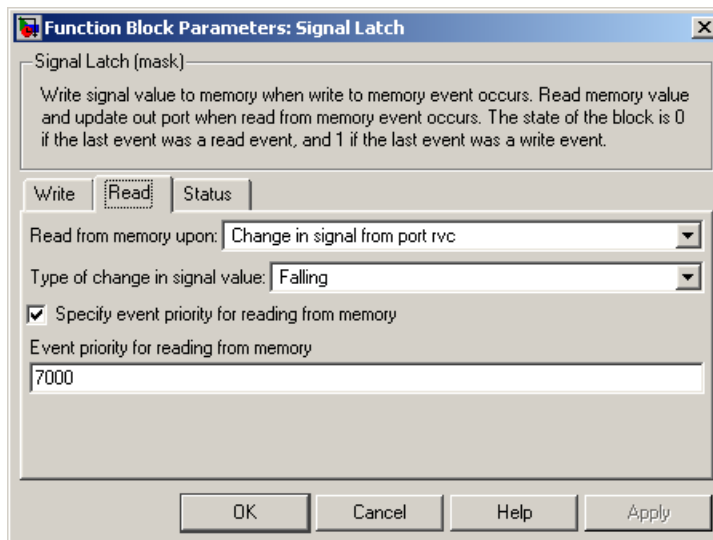
Specify event priority for writing to memory

Select this option to control the sequencing of the write event, relative to other simultaneous events.

Event priority for writing to memory

The priority of the write event, relative to other simultaneous events. This field appears only if you select **Specify event priority for writing to memory**.

Read Tab



Read from memory upon

The type of signal-based event, function call, or internal write event that causes a read event.

Trigger type

Determines whether rising, falling, or either type of trigger edge causes a read event. This field appears only if you set **Read from memory upon** to Trigger from port rtr.

Signal Latch

Type of change in signal value

Determines whether rising, falling, or either type of value change causes a read event. This field appears only if you set **Read from memory upon** to Change in signal from port rvc.

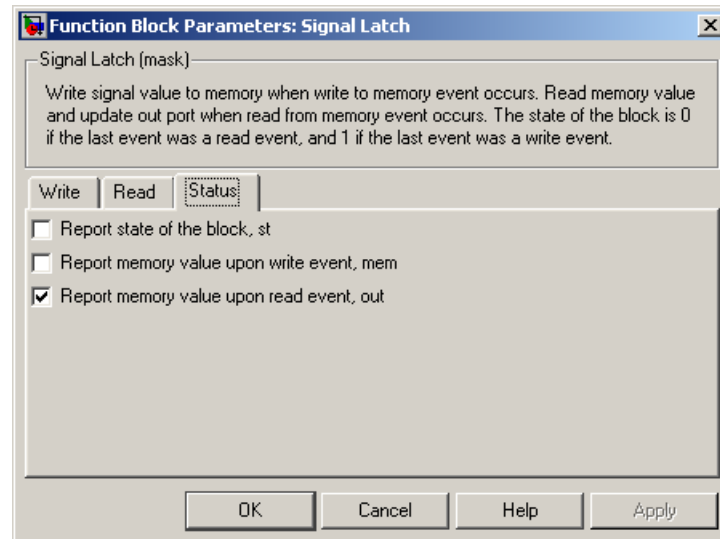
Specify event priority for reading from memory

Select this option to control the sequencing of the write event, relative to other simultaneous events. This field appears only if you set **Read from memory upon** to an option other than Write to memory event.

Event priority for reading from memory

The priority of the read event, relative to other simultaneous events. This field appears only if you select **Specify event priority for writing to memory**.

Status Tab



Report state of the block

Controls the presence of the signal output port labeled **st**.

Report memory value upon write event

Controls the presence of the signal output port labeled **mem**.

Report memory value upon read event


Controls the presence of the signal output port labeled **out**.

Examples

Reading from Memory Upon Each Write Event


In the plot below, the output signal reflects values of the input signal upon each rising or falling value of the **wvc** signal. Between successive write events, the output signal maintains the value from the most recent write event. Before the first write event, the output signal is 0 because of the initial memory value.

Signal Latch



Independent Read and Write Events

In the plot below, the **mem** signal reflects values of the input signal upon each rising or falling value of the **wvc** signal, while the **out** signal reflects values of the **mem** signal upon each rising or falling value of the **rvc** signal.



For examples showing the use of this block in a model, see

- “Example: Creating a Random Signal for Switching” on page 3-7

Signal Latch

- “Generating Random Time-Based Signals” on page 3-8
- “Example: Defining #d Before the First Entity Departure” on page 3-19
- “Example: Resampling a Signal Based on Events” on page 3-20
- “Example: Detecting Collisions by Comparing Events” on page 2-34
- “Example: Compound Switching Logic” on page 5-5

See Also

Data Store Memory, Data Store Read, Data Store Write, “Manipulating Signals” on page 3-18

Purpose Plot data from signal

Library SimEvents Sinks

Description



This block creates a plot using data from a signal. The plot is particularly appropriate for data arising from discrete-event simulations because the plot can include zero-duration values.

The data for the vertical axis comes from the signal connected to the block’s signal input port labeled **in**.

The **Plot type** parameter on the **Plotting** tab determines whether and how the block connects the points that it plots. For details, see “Connections Among Points in Plots” on page 9-4.

Selecting Data for the Horizontal Axis

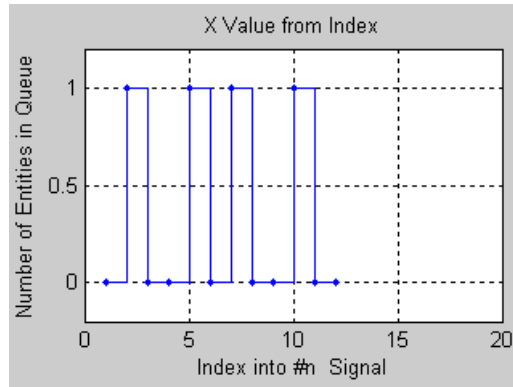
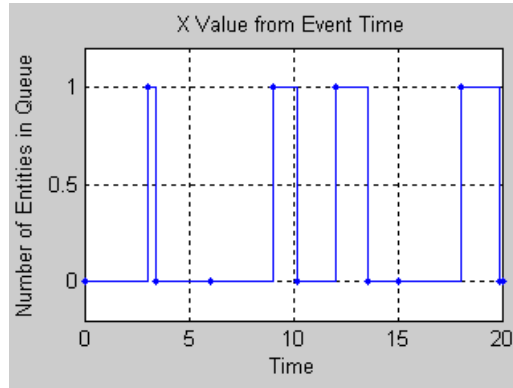
Use the **X value from** parameter to select the type of data for the horizontal axis. The table below describes the choices.

Source of X Data	Description of Plot
Event time	Plot of the in signal versus simulation time. For example, you might use this option to see how the length of a queue changes over time.
Index	Plot of the in signal’s successive values against a horizontal axis that represents the index of the values. The signal’s first value during the simulation has an index of 1, the signal’s second value has an index of 2, and so on. For example, you might use this option for a signal that has zero-duration values, to help determine the exact sequence among values that the signal assumes simultaneously.

The figures below illustrate the different sources of data for the horizontal axis. The plots look similar, except that the second plot has

Signal Scope

uniform horizontal spacing rather than time-based spacing between successive points.



Ports

Signal Input Ports


Label	Description
in	Signal containing data for the Y axis.

Signal Output Ports

Label	Description
#c	Number of points the block has plotted.

Dialog Box

Plotting Tab



Signal Scope

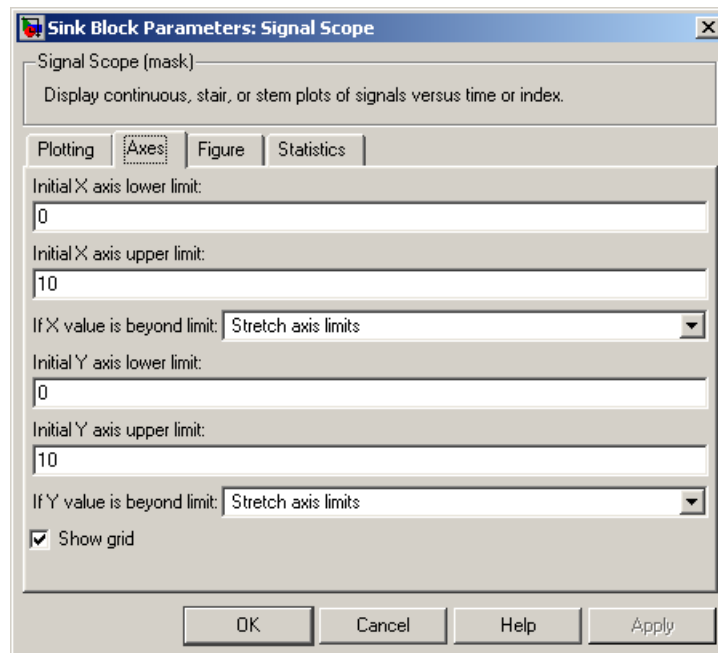
Plot type

The presentation format for the data. See “Connections Among Points in Plots” on page 9-4 for details.

X value from

Source of data for the plot’s horizontal axis. See “Selecting Data for the Horizontal Axis” on page 16-145 for details.

Axes Tab



Initial X axis lower limit, Initial X axis upper limit

The interval shown on the X axis at the beginning of the simulation. The interval might change from this initial setting due to zooming, autoscaling, or the **If X value is beyond limit** setting.

If X value is beyond limit

Determines how the plot changes if one or more X values are not within the limits shown on the X axis. For details, see “Varying Axis Limits Automatically” on page 9-5.

Initial Y axis lower limit, Initial Y axis upper limit

The interval shown on the Y axis at the beginning of the simulation. The interval might change from this initial setting due to zooming, autoscaling, or the **If Y value is beyond limit** setting.

If Y value is beyond limit

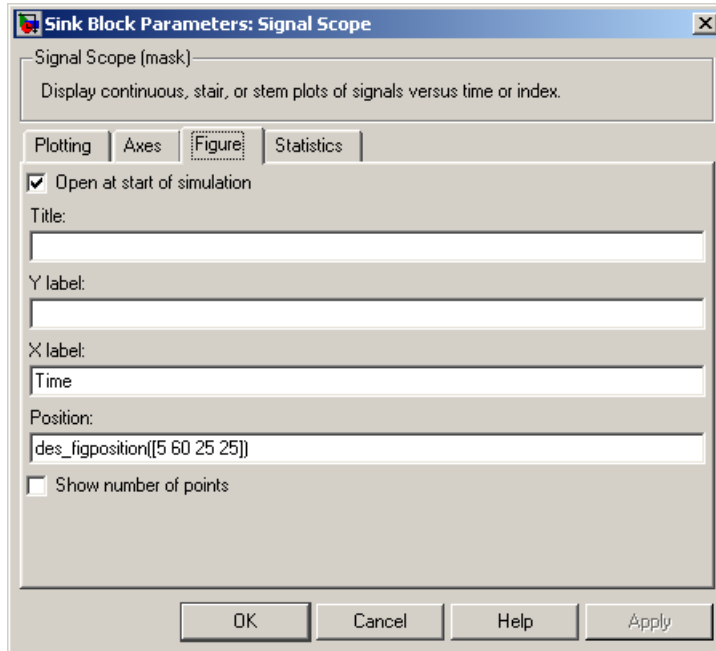
Determines how the plot changes if one or more values of the **in** signal are not within the limits shown on the Y axis. For details, see “Varying Axis Limits Automatically” on page 9-5.

Show grid

Toggles the grid on and off.

Signal Scope

Figure Tab



Open at start of simulation

Selecting this option causes the plot window to open when you start the simulation. If you clear this box, you can open the plot window while the simulation is running by double-clicking the block icon.

Title

Text that appears as the title of the plot, above the axes.

Y label

Text that appears to the left of the vertical axis.

X label

Text that appears below the horizontal axis.

Position


A four-element vector of the form [left bottom width height] specifying the position of the scope window. (0,0) is the lower left corner of the display.

Show number of points

Displays the number of plotted points using an annotation in the plotting window.

Statistics Tab

These parameters determine whether certain ports output data throughout the simulation, output data only when you stop or pause the simulation, or are omitted from the block. For descriptions of the affected ports, see the “Signal Output Ports” table above.



Signal Scope

Number of points plotted

Controls the presence and behavior of the signal output port labeled **#c**.

Examples

- “Building a Simple Discrete-Event Model” and “Observations from Plots” in the Getting Started documentation
- “Example: Using Servers in Shifts” on page 6-11
- “Example: Choosing the Shortest Queue” on page 6-3

See Also

X-Y Signal Scope, Attribute Scope, Chapter 9, “Plotting Data”

Signal-Based Event to Function-Call Event


Purpose

Convert signal-based events into function calls

Library

Event Translation

Description



This block converts a signal-based event or a function-call input into one or two function calls that you can use to invoke function-call subsystems, Stateflow blocks, or other blocks that accept function-call inputs. You specify the type of event the block translates and whether the block suppresses its output under certain conditions.

Criteria for Generating Function Calls

The primary criterion, based on the **Generate function call only upon** parameter, is a signal-based event or a function call. By default, the block issues a function call upon each event of the type you specify.

To issue up to two function calls upon each event, select **Generate optional f2 function call**. If the block issues function calls at both the **f1** and **f2** output ports, then it issues the **f1** call first.

To make the **f1** or **f2** output function call contingent upon a secondary criterion, select **Suppress function call f1 if enable signal e1 is not positive** or **Suppress function call f2 if enable signal e2 is not positive**. The block acquires an additional signal input port, labeled **e1** or **e2**, to which you connect a control signal. If the control signal is zero or negative when the block is about to issue the function call, then the block suppresses the function call. The **e1** and **e2** ports operate independently of each other as secondary criteria for their respective function-call output ports.

Signal-Based Event to Function-Call Event

Ports

Signal Input Ports

Label	Description
ts	Primary criterion is satisfied when Simulink updates the signal at this port. This port appears only if you set Generate function call only upon to Sample time hit from port ts.
tr	Primary criterion is satisfied when the signal has a rising or falling edge, depending on the Trigger type parameter. This port appears only if you set Generate function call only upon to Trigger from port tr.
vc	Primary criterion is satisfied when the signal at this port increases or decreases, depending on the Type of change in signal value parameter. This port appears only if you set Generate function call only upon to Change in signal from port vc.
fcn	Primary criterion is satisfied when this input port detects a function call. This port appears only if you set Generate function call only upon to Function call from port fcn. Do not connect this port to an output port from the same instance of this block.
e1	When this signal is 0 or negative, the block does not issue a function call at the f1 output port. This input port appears only if you select Suppress function call f1 if enable signal e1 is not positive .
e2	When this signal is 0 or negative, the block does not issue a function call at the f2 output port. This input port appears only if you select Suppress function call f2 if enable signal e2 is not positive .

Signal-Based Event to Function-Call Event

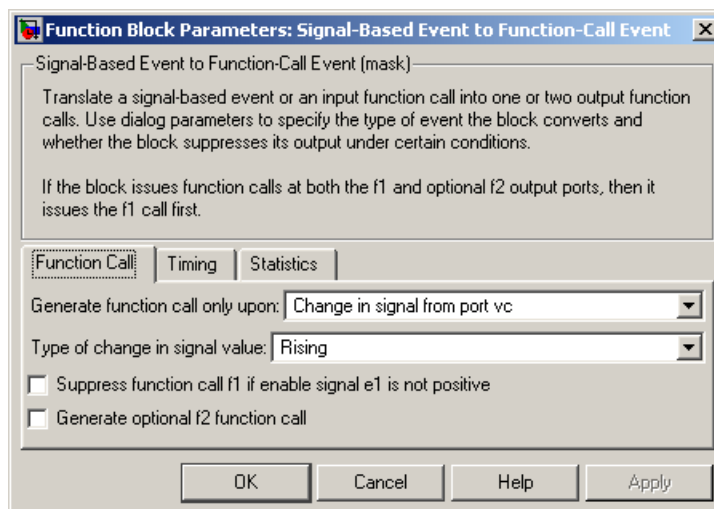
Signal Output Ports

Label	Description	Order of Update
f1	Function call, possibly contingent on e1 input signal	1
f2	Function call, possibly contingent on e2 input signal	2
#f1	Number of function calls the block has generated at the f1 port during the simulation	3
#f2	Number of function calls the block has generated at the f2 port during the simulation	3

Output signals having the same number in the Order of Update column above are updated in an arbitrary sequence relative to each other; you should not rely on a specific sequence for your simulation results.

Dialog Box

Function Call Tab



Signal-Based Event to Function-Call Event

Generate function call only upon

The primary criterion for determining when the block issues a function call. Optional secondary criteria are established by the **Suppress function call...** parameters below.

Trigger type

Determines whether rising, falling, or either type of trigger edge causes the block to generate a function call. This field appears only if you set **Generate function call only upon** to Trigger from port tr.

Type of change in signal value

Determines whether rising, falling, or either type of value change causes the block to generate a function call. This field appears only if you set **Generate function call only upon** to Change in signal from port vc.

Suppress function call f1 if enable signal e1 is not positive

Selecting this option causes **f1** function calls to be contingent upon a positive value at the **e1** signal input port.

Generate optional f2 function call

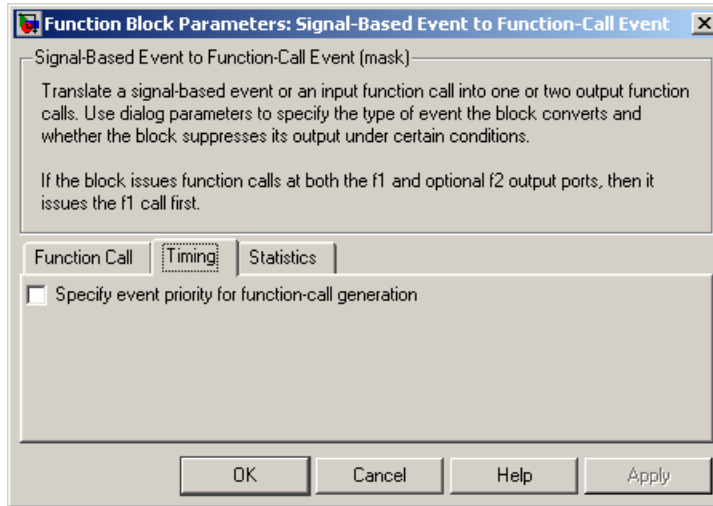
Selecting this option causes the block to issue a function call at the optional **f2** output port when appropriate criteria are satisfied.

Suppress function call f2 if enable signal e2 is not positive

Selecting this option causes **f2** function calls to be contingent upon a positive value at the **e2** signal input port. This field appears only if you select **Generate optional f2 function call**.

Signal-Based Event to Function-Call Event

Timing Tab



Specify event priority for function-call generation

Select this option to control the sequencing of the function-call event, relative to other simultaneous events.

Function-call event priority


The priority of the function-call event, relative to other simultaneous events. This field appears only if you select **Specify event priority for function-call generation**.

Function-call time delay

The delay, in seconds, between the input event and the output function call. A positive value schedules the function call in the future, while a value of zero schedules the function call at the current simulation time.

Signal-Based Event to Function-Call Event

Statistics Tab



Number of f1 function calls

Controls the presence and behavior of the signal output port labeled #f1.

Number of f2 function calls

Controls the presence and behavior of the signal output port labeled #f2. This field is active only if you select **Generate optional f2 function call** on the **Function Call** tab of this dialog box.

Examples

- “Example: Detecting Changes in the Last-Updated Signal” on page 3-10
- “Example: Using a #n Signal as a Trigger” on page 12-11

See Also


Entity Departure Event to Function-Call Event, “Manipulating Events” on page 2-39

Signal-Based Function-Call Event Generator

Purpose Generate function-call events in response to signal-based events

Library Generators / Event Generators

Description



This block generates an output function call corresponding to each signal-based event or input function call. You specify the type of event the block responds to. You can use the function call to invoke function-call subsystems, Stateflow blocks, or other blocks that accept function-call inputs.

This block is similar to the Signal-Based Event to Function-Call Event block, which offers more flexibility.

Ports

Signal Input Ports

Label	Description
ts	The block generates a function call when Simulink updates the signal at this port. This port appears only if you set Generate function call only upon to Sample time hit from port ts.
tr	The block generates a function call when the signal has a rising or falling edge, depending on the Trigger type parameter. This port appears only if you set Generate function call only upon to Trigger from port tr.
vc	The block generates a function call when the signal at this port increases or decreases, depending on the Type of change in signal value parameter. This port appears only if you set Generate function call only upon to Change in signal from port vc.
fcn	The block generates a function call when this input port detects a function call. This port appears only if you set Generate function call only upon to Function call from port fcn. Do not connect this port to an output port from the same instance of this block.

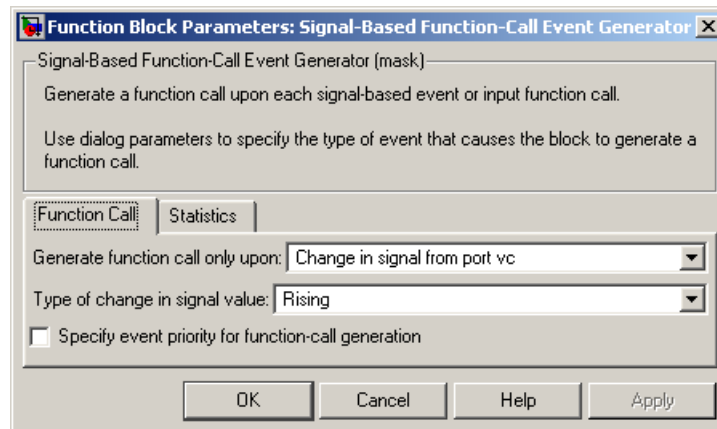
Signal-Based Function-Call Event Generator

Signal Output Ports

Label	Description	Order of Update
f1	Function-call signal.	1
#f1	Number of function calls the block has generated during the simulation.	2

Dialog Box

Function Call Tab



Generate function call only upon

The primary criterion for determining when the block issues a function call. Optional secondary criteria are established by the **Suppress function call...** parameters below.

Trigger type

Determines whether rising, falling, or either type of trigger edge causes the block to generate a function call. This field appears only if you set **Generate function call only upon** to Trigger from port tr.

Signal-Based Function-Call Event Generator

Type of change in signal value

Determines whether rising, falling, or either type of value change causes the block to generate a function call. This field appears only if you set **Generate function call only upon** to **Change** in signal from port **vc**.

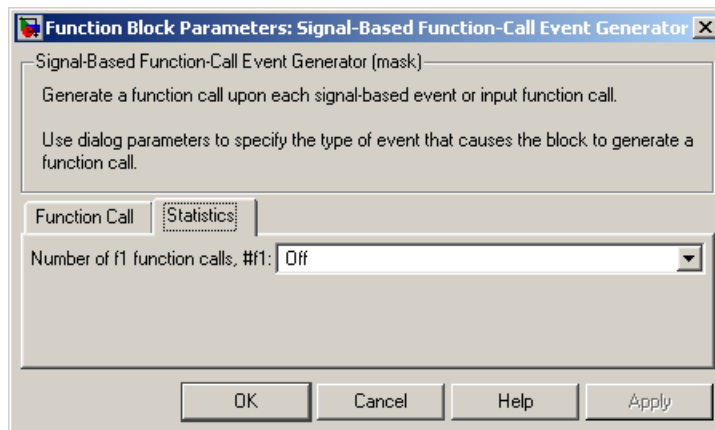
Specify event priority for function-call generation

Select this option to control the sequencing of the function-call event, relative to other simultaneous events.

Function-call event priority

The priority of the function-call event, relative to other simultaneous events. This field appears only if you select **Specify event priority for function-call generation**.

Statistics Tab



Number of f1 function calls

Controls the presence and behavior of the signal output port labeled **#f1**.

Examples

- “Example: Calling a Stateflow Block Upon Changes in Server Contents” on page 2-37

Signal-Based Function-Call Event Generator

- “Example: Counting Events from Multiple Sources” on page 2-41

See Also

Signal-Based Event to Function-Call Event, “Generating Function-Call Events” on page 2-36

Purpose Serve one entity for period of time

Library Servers

Description



This block serves one entity for a period of time and then attempts to output the entity through the **OUT** port. If the **OUT** port is blocked, then the entity stays in this block until the port becomes unblocked. You specify the service time, which is the duration of service, via a parameter, attribute, or signal, depending on the **Service time from** parameter value. The block determines the service time for an entity upon its arrival.

The block permits preemption if you select **Permit preemption based on attribute**. In this case, an entity in the server can depart early via the **P** port. Preemption occurs only if attributes of the current entity and the entity attempting to arrive satisfy specified criteria. For details, see “Preempting an Entity in a Server” on page 4-9.

When the block does not permit preemption, the **IN** port is unavailable whenever this block stores an entity. In this case, the **IN** port becomes available when the entity departs.

Ports

Entity Input Ports

Label	Description
IN	Port for arriving entities, which will be served.

Signal Input Ports

Label	Description
t	Service time for a newly arrived entity. This port appears only if you set Service time from to Signal port t.

Single Server

Entity Output Ports

Label	Description
OUT	Port for departing entities that have completed their service time and have not been preempted.
P	Port for entities that have been preempted by an arriving entity. This port must not be blocked at the time of preemption.

Signal Output Ports

Label	Description	Time of Update When Statistic Is On	Order of Update
#d	Number of entities that have departed from this block's OUT entity output port since the start of the simulation.	After entity departure from the OUT port	4
#n	Number of entities currently in the block, either 0 or 1.	After entity arrival and after entity departure from the OUT port	3
#p	Number of entities that have been preempted from this block since the start of the simulation.	After entity departure from the P port	4
pe	A value of 1 indicates when the block tries and fails to output an entity, that is, an entity has completed its service but cannot depart from the OUT port because the subsequent entity input port is unavailable.	After service completion and after entity departure from the OUT port	2


Label	Description	Time of Update When Statistic Is On	Order of Update
w	Sample mean of the waiting times in this block for all entities that have departed from the OUT port. An entity's waiting time might exceed its service time if the OUT port is blocked when the entity completes service.	After entity departure from the OUT port	1
util	Utilization of the server, which is the fraction of simulation time spent storing an entity. At T=0, the utilization is 0 or 1 depending on whether the server contains an entity.	Performance considerations cause the block to suppress signal updates until specific occurrences cause updates. In On mode, updates occur after an entity departure from the OUT port and after an entity arrival. In Upon stop or pause mode, stopping or pausing the simulation causes an update.	1

Output signals having the same number in the Order of Update column above are updated in an arbitrary sequence relative to each other; you should not rely on a specific sequence for your simulation results.

Single Server

Dialog Box

Single Server Tab



Service time from

Determines whether the service time is computed from a parameter in this dialog box, an input signal, or an attribute of the entity being served.

Service time

The service time for all entities. This field appears only if you set **Service time from** to Dialog.

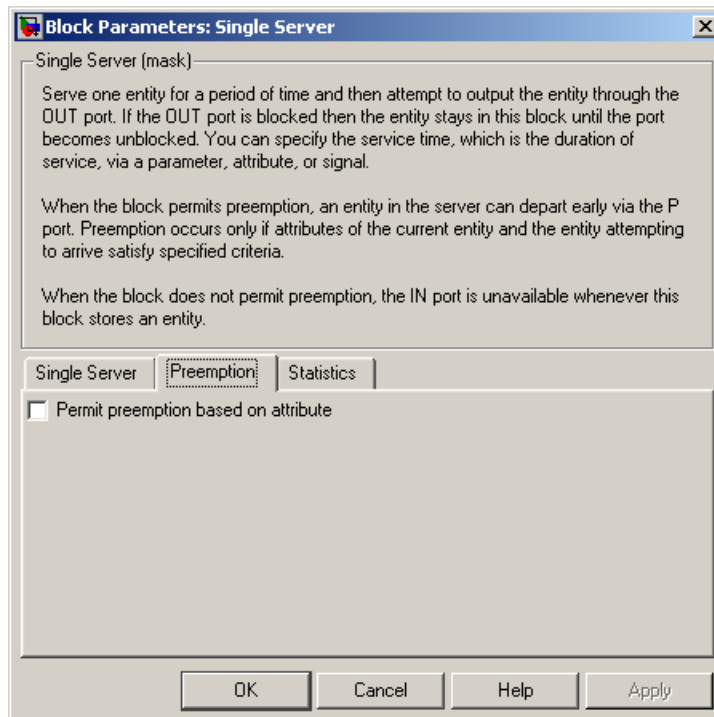
Attribute name

The name of the attribute whose value the block uses as the service time for an entity. This field appears only if you set **Service time from** to Attribute.

Service completion event priority

The priority of the service completion event, relative to other simultaneous events.

Preemption Tab



Permit preemption based on attribute

If you select this option, the block can replace an entity by a higher-priority entity. Otherwise, the block never permits new

Single Server

arrivals when it is storing an entity. Selecting this option also sets **Average wait** on the **Statistics** tab to Off and makes that parameter unavailable.

Sorting attribute name

The block uses this attribute to determine whether a new entity can preempt the one in the server. This field appears only if you select **Permit preemption based on attribute**.

Sorting direction

Preemption occurs when the arriving entity has a strictly smaller (Ascending) or strictly larger (Descending) value of the attribute named above, compared to the attribute value of the entity in the server. This field appears only if you select **Permit preemption based on attribute**.

Write residual service time to attribute

If you select this option, a preemption event causes the block to set an attribute in the preempted entity. The attribute value is the remaining service time the entity would have required if it had not been preempted. This field appears only if you select **Permit preemption based on attribute**.

Residual service time attribute name

The name of the attribute the block uses when recording the residual service time of entities. This field appears only if you select **Write residual service time to attribute**.


Create attribute if not present

Selecting this option enables the block to define a new attribute for the residual service time. Otherwise, the block issues an error if the attribute named above does not already exist. This field appears only if you select **Write residual service time to attribute**.

Statistics Tab

These parameters determine whether certain ports output data throughout the simulation, output data only when you stop or pause

the simulation, or are omitted from the block. For descriptions of the affected ports, see the “Signal Output Ports” table above.



Number of entities departed

Controls the presence and behavior of the signal output port labeled **#d**.

Number of entities in block

Controls the presence and behavior of the signal output port labeled **#n**.

Number of entities preempted

Controls the presence and behavior of the signal output port labeled **#p**.

Single Server

Status of pending entity departure

Controls the presence of the signal output port labeled **pe**.

Average wait

Controls the presence and behavior of the signal output port labeled **w**. This field is available only if you clear the **Permit preemption based on attribute** option on the **Preemption** tab.

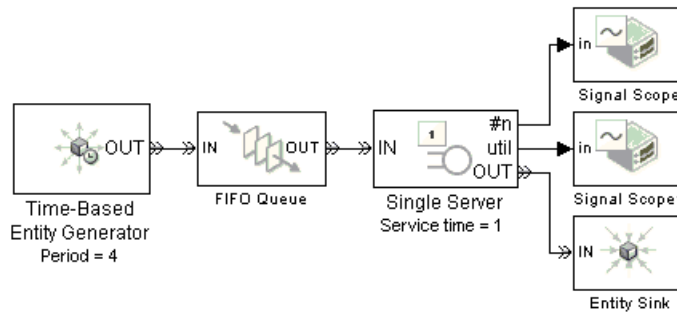
Utilization

Controls the presence and behavior of the signal output port labeled **util**.


Examples

- “Building a Simple Discrete-Event Model” in the Getting Started documentation
- “Example: Selecting the First Available Server” in the Getting Started documentation
- “Constructs Involving Queues and Servers” in the Getting Started documentation
- “Example: Using Servers in Shifts” on page 6-11
- “Example: Preemption by High-Priority Entities” on page 4-10
- “Example: Controlling Joint Availability of Two Servers” on page 7-4
- “Example: Synchronizing Service Start Times with the Clock” on page 7-6


The following example illustrates the timing of updates of the **util** signal, as described in Signal Output Ports on page 16-164.




The server has idle periods that reduce its utilization. However, the server block recomputes the **util** signal only when the number of entities in the server changes. While the definition of utilization says that the utilization is less than 1 at time 3, the **util** signal remains at its previous value of 1 until the next entity arrives at time 4.



Single Server



In a longer simulation, the differences in the value of **util** compared to its theoretical definition become less pronounced.



See Also

N-Server, Infinite Server, “Basic Queues and Servers” in the Getting Started documentation

Purpose Associate named timer to each arriving entity independently and start timing

Library Timing

Description



This block associates a named timer to each arriving entity independently and starts the timer. If the entity was previously associated with a timer of the same name, then the block either continues or restarts that timer, depending on your setting for the **If timer has already started** parameter; the Warn and continue option can be helpful for debugging or preventing modeling errors. Other timers, if any, associated with the arriving entity are unaffected.

This block works with the Read Timer block. To read the value of the timer named in this block, reference the timer name in the Read Timer block. For more information about using this pair of blocks, see “Using Timers” on page 10-18.

Ports

Entity Input Ports

Label	Description
IN	Port for arriving entities.

Entity Output Ports

Label	Description
OUT	Port for departing entities, which have named timers attached to them.

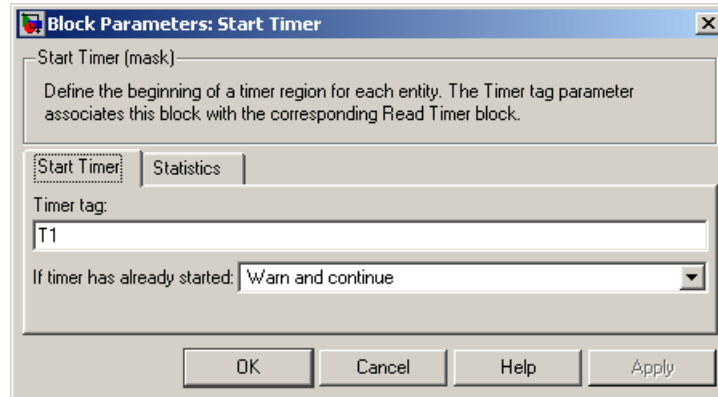
Start Timer

Signal Output Ports

Label	Description	Time of Update When Statistic Is On
#d	Number of entities that have departed from this block since the start of the simulation.	After entity departure

Dialog Box

Start Timer Tab



Timer tag

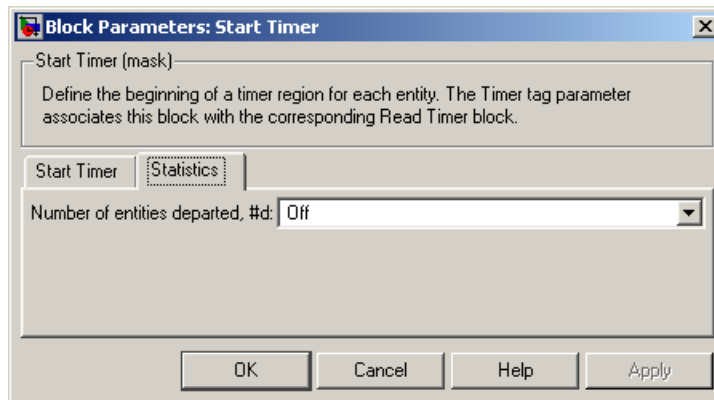
Name of the timer to associate with each entity.

If timer has already started

Behavior of the block if an arriving entity already has a timer with the specified timer tag.

Statistics Tab

These parameters determine whether certain ports output data throughout the simulation, output data only when you stop or pause the simulation, or are omitted from the block. For descriptions of the affected ports, see the “Signal Output Ports” table above.



Number of entities departed

Controls the presence and behavior of the signal output port labeled **#d**.

Examples

- “Basic Procedure for Using Timer Blocks” on page 10-19
- “Timing Multiple Entity Paths with One Timer” on page 10-20
- “Restarting a Timer” on page 10-21
- “Timing Multiple Processes Independently” on page 10-23

See Also

Read Timer, “Using Timers” on page 10-18

Subsystem Configuration

Purpose	Configuration for Discrete Event Subsystem block
Library	SimEvents Ports and Subsystems
Description	Every Discrete Event Subsystem window must contain exactly one copy of this configuration block. Do not delete this block from the subsystem window or else the subsystem will not work properly. To create a subsystem using the Discrete Event Subsystem block, see “Setting Up Signal-Based Discrete Event Subsystems” on page 8-11.
See Also	Discrete Event Subsystem, Discrete Event Inport, Chapter 8, “Controlling Timing Using Subsystems”

Purpose Generate entities using intergeneration times from signal or statistical distribution

Library Generators / Entity Generators

Description This block is designed to generate entities using intergeneration times that satisfy criteria that you specify. The intergeneration time is the time interval between two successive entities' generation events.



Intergeneration Times	Value of Generate entities with Parameter
Distributed according to various parameters in the block's dialog box	Intergeneration time from dialog
Specified using an input signal that the block reads at the start of the simulation and each time it generates an entity	Intergeneration time from port t

For details about these options, see “Introduction to the Time-Based Entity Generator”.

Responding to Blockage at the Entity Output Port

You can choose how this block responds when it generates an entity that the subsequent input port is not available to accept. If applicable, you can also choose how the block responds when that port later becomes available.

- If you set **Response when blocked** to Error, then the simulation halts with an error message if the subsequent entity input port is not available to accept an entity when this block generates it.
- If you set **Response when blocked** to Pause generation, then this block holds the generated entity and temporarily suspends the generation of additional entities. The **Response when unblocked** parameter determines when the held entity departs.

Time-Based Entity Generator

- If you set **Response when unblocked** to Immediate restart, then the held entity departs after this block receives notification that the subsequent entity input port has become available. After the held entity departs, this block schedules the generation of the next entity.
- If you set **Response when unblocked** to Delayed restart, then upon receiving notification that the subsequent entity input port has become available, this block schedules an extra generation event for the held entity. The event time is the current time plus the same intergeneration time used when first generating the held entity. When the extra generation event is processed, the held entity attempts to depart.

For an example, see “Example: Responding to Blockage” on page 16-182.

Ports

Signal Input Ports

Label	Description
t	Time interval between generation events of the current entity and the next entity. The block reads the value after the current entity has departed and the block has updated its output signals, if any. If you do not select Generate entity at simulation start , then the block also reads the value of this signal at the start of the simulation. This port appears only if you set Generate entities with to Intergeneration time from port t.

Entity Output Ports

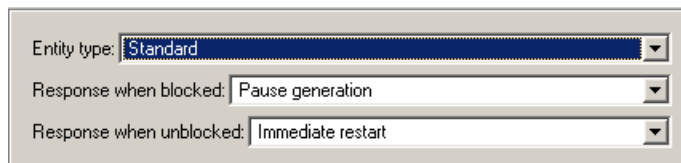
Label	Description
OUT	Port through which generated entities depart.

Time-Based Entity Generator

Signal Output Ports

Label	Description	Time of Update When Statistic Is On	Order of Update
#d	Number of entities that have departed from this block since the start of the simulation.	After entity departure	3
pe	A value of 1 indicates when the block is holding an entity that cannot depart because the subsequent entity input port is unavailable or because of a delayed restart of this block.	After entity generation if OUT port is blocked, and after entity departure in all cases	2
w	Average interdeparture time, in seconds, for all pairs of successive entities that have departed from this block. The signal value is 0 before the second entity departure.	After entity departure	1

Dialog Box



Entity type: Standard

Response when blocked: Pause generation

Response when unblocked: Immediate restart

Entity type

The standard type includes attributes called Priority and Count with default values of 0, while the blank type includes no attributes.

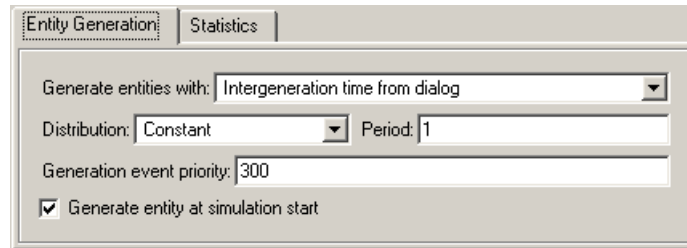
Response when blocked, Response when unblocked

Determines how the block responds if a generated entity cannot depart immediately because the entity input port of the

Time-Based Entity Generator

subsequent block is unavailable; see “Responding to Blockage at the Entity Output Port” on page 16-177.

Entity Generation Tab



Generate entities with

Determines where the block gets instructions about when to generate entities.

Distribution

The statistical distribution of intergeneration times. This field appears only if you set **Generate entities with** to Intergeneration time from dialog.

Period

The time interval between entity generations, in seconds. This field appears only if you set **Generate entities with** to Intergeneration time from dialog and set **Distribution** to Constant.

Initial seed

A nonnegative integer that initializes the random number generator. This field appears only if you set **Generate entities with** to Intergeneration time from dialog and set **Distribution** to Uniform or Exponential.

Minimum, Maximum

The endpoints, in seconds, of the interval over which the distribution is uniform. These fields appear only if you set

Generate entities with to Intergeneration time from dialog and set **Distribution** to Uniform.

Mean

The expected value of the exponential distribution. This field appears only if you set **Generate entities with** to Intergeneration time from dialog and set **Distribution** to Exponential.

Generation event priority

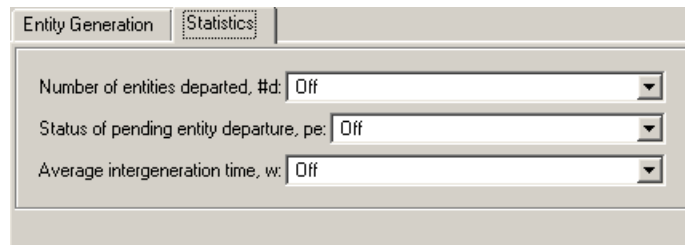
The priority of the entity-generation event, relative to other simultaneous events.

Generate entity at simulation start

If you select this option, the block generates the first entity when the simulation begins and the second entity at the first intergeneration time. Otherwise, the block generates the first entity at the first intergeneration time.

Statistics Tab

These parameters determine whether certain ports output data throughout the simulation, output data only when you stop or pause the simulation, or are omitted from the block. For descriptions of the affected ports, see the “Signal Output Ports” table above.



The screenshot shows a dialog box with two tabs: "Entity Generation" and "Statistics". The "Statistics" tab is active. It contains three dropdown menus, all set to "Off":

- Number of entities departed, #d: Off
- Status of pending entity departure, pe: Off
- Average intergeneration time, w: Off

Number of entities departed

Controls the presence and behavior of the signal output port labeled #d.

Time-Based Entity Generator

Status of pending entity departure

Controls the presence and behavior of the signal output port labeled **pe**.

Average intergeneration time


Controls the presence and behavior of the signal output port labeled **w**.

Examples

- “Building a Simple Discrete-Event Model” in the Getting Started documentation
- “Example: Using Random Intergeneration Times in a Queuing System” in the Getting Started documentation
- “Example: Using a Step Function as Intergeneration Time” in the Getting Started documentation
- “Example: Using an Arbitrary Discrete Distribution as Intergeneration Time” in the Getting Started documentation
- “Using Generation Times from a Vector” on page 1-11

Example: Responding to Blockage

To illustrate the blockage options, consider a Time-Based Entity Generator block followed by a Single Server block, followed by an Entity Sink block.



Suppose the entity generator has **Response when blocked** set to **Pause generation**. Also, suppose that the intergeneration time is 1 with the first entity generated at $T=1$, and that the service times for

Time-Based Entity Generator

the first three entities in the server are 1.5, 2.2, and 1.8. The tables below indicate how the **Response when unblocked** values affect the behavior in the simulation.

Immediate Restart

Time (s)	Behavior
1	Entity generator generates and outputs the first entity to the server. The server's entity input port becomes unavailable. The first entity is scheduled to depart from the server at $T=1+1.5=2.5$.
2	Entity generator generates the second entity and holds it because the OUT port is blocked.
2.5	First entity departs from the server. The server's entity input port becomes available and the second entity advances from the entity generator to the server. The server's entity input port becomes unavailable again. The second entity is scheduled to complete service at $T=2.5+2.2=4.7$. The entity generator schedules the next generation for $T=2.5+1=3.5$.
3.5	Entity generator generates the third entity, and holds it because the OUT port is blocked.
4.7	Second entity departs from the server. The server's entity input port becomes available and the third entity advances from the entity generator to the server. The server's entity input port becomes unavailable again. The third entity is scheduled to complete service at $T=4.7+1.8=6.5$. The entity generator schedules the next generation for $T=4.7+1=5.7$.

Time-Based Entity Generator

Delayed Restart

Time (s)	Behavior
1	Entity generator generates and outputs the first entity to the server. The server's entity input port becomes unavailable. The first entity is scheduled to depart from the server at $T=1+1.5=2.5$.
2	Entity generator generates the second entity and holds it because the OUT port is blocked.
2.5	First entity departs from the server. The server's entity input port becomes available. The entity generator schedules an extra generation event for the second entity at $T=2.5+1=3.5$.
3.5	The second entity advances from the entity generator to the server. The server's entity input port becomes unavailable again. The second entity is scheduled to complete service at $T=3.5+2.2=5.7$. The entity generator schedules the next generation for $T=3.5+1=4.5$.
4.5	Entity generator generates the third entity, and holds it because the OUT port is blocked.
5.7	Second entity departs from the server. The server's entity input port becomes available. The entity generator schedules an extra generation event for the third entity at $T=5.7+1=6.7$.
6.7	The third entity advances from the entity generator to the server. The server's entity input port becomes unavailable again. The third entity is scheduled to complete service at $T=6.7+1.8=8.5$. The entity generator schedules the next generation for $T=6.7+1=7.7$.

See Also

Event-Based Entity Generator, Entity Sink, "Creating Entities Using Intergeneration Times" in the Getting Started documentation

Purpose Plot data from two attributes of arriving entities

Library SimEvents Sinks

Description



This block plots a curve using data from two attributes of arriving entities. Use the **Y attribute name** and **X attribute name** parameters to specify which attributes to plot.

Use the **Enable entity OUT port** option to choose whether the entity advances to a subsequent block or whether the block absorbs the arriving entity.

The **Plot type** parameter on the **Plotting** tab determines whether and how the block connects the points that it plots. For details, see “Connections Among Points in Plots” on page 9-4.

Ports

Entity Input Ports

Label	Description
IN	Port for arriving entities, whose attributes contain the data to plot.

Entity Output Ports

Label	Description
OUT	Port for departing entities. This port appears only if you select Enable entity OUT port .

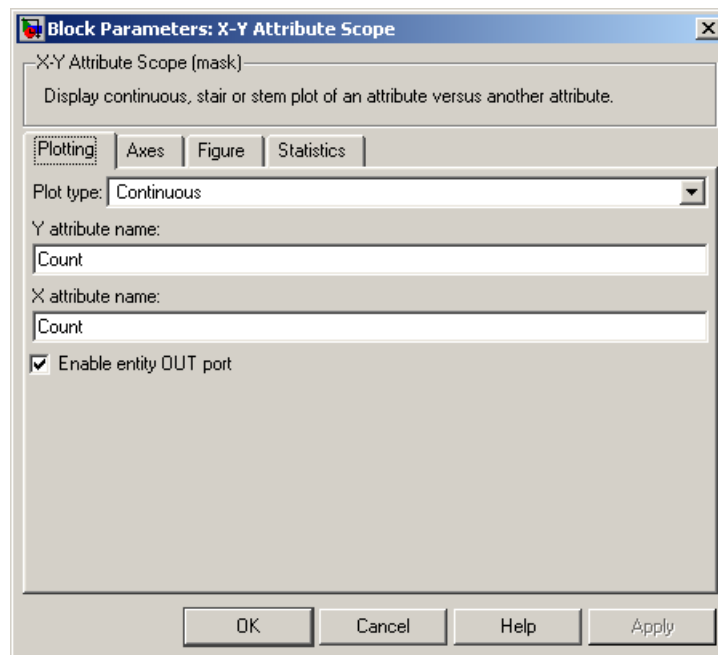
X-Y Attribute Scope

Signal Output Ports

Label	Description
#a	Number of entities that have arrived at the block since the start of the simulation.

Dialog Box

Plotting Tab



Plot type

The presentation format for the data. See “Connections Among Points in Plots” on page 9-4 for details.

Y attribute name

Name of the attribute to plot along the vertical axis.

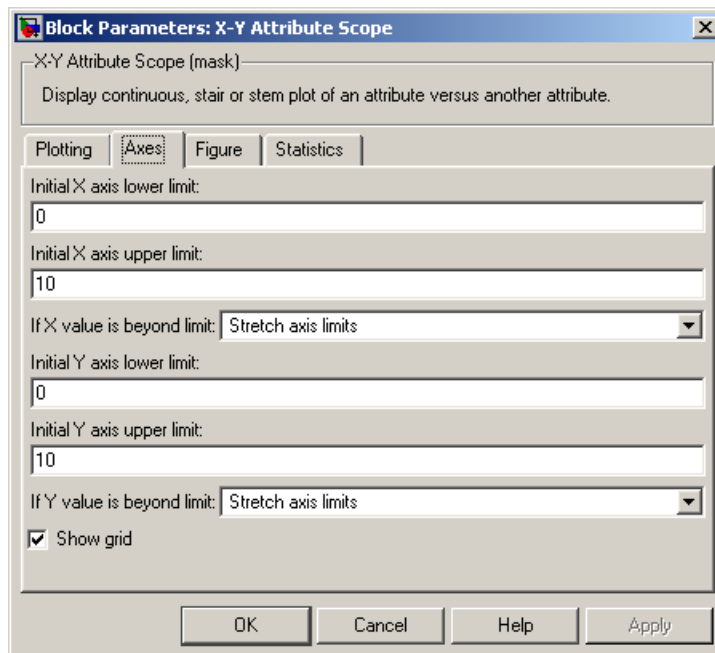
X attribute name

Name of the attribute to plot along the horizontal axis.

Enable entity OUT port

Causes the block to have an entity output port labeled **OUT**, through which the arriving entity departs. If you clear this box, the block absorbs arriving entities.

Axes Tab



The screenshot shows a dialog box titled "Block Parameters: X-Y Attribute Scope" with a close button (X) in the top right corner. The dialog has a tabbed interface with four tabs: "Plotting", "Axes", "Figure", and "Statistics". The "Axes" tab is currently selected. The dialog contains the following fields and controls:

- A text area at the top with the text "X-Y Attribute Scope (mask)" and "Display continuous, stair or stem plot of an attribute versus another attribute."
- Four input fields for axis limits:
 - "Initial X axis lower limit:" with a text box containing "0".
 - "Initial X axis upper limit:" with a text box containing "10".
 - "Initial Y axis lower limit:" with a text box containing "0".
 - "Initial Y axis upper limit:" with a text box containing "10".
- Two dropdown menus for handling values beyond limits:
 - "If X value is beyond limit:" with a dropdown menu set to "Stretch axis limits".
 - "If Y value is beyond limit:" with a dropdown menu set to "Stretch axis limits".
- A checked checkbox labeled "Show grid".
- Four buttons at the bottom: "OK", "Cancel", "Help", and "Apply".

Initial X axis lower limit, Initial X axis upper limit

The interval shown on the X axis at the beginning of the simulation. The interval might change from this initial setting due to zooming, autoscaling, or the **If X value is beyond limit** setting.

X-Y Attribute Scope

If X value is beyond limit

Determines how the plot changes if one or more X values are not within the limits shown on the X axis. For details, see “Varying Axis Limits Automatically” on page 9-5.

Initial Y axis lower limit, Initial Y axis upper limit

The interval shown on the Y axis at the beginning of the simulation. The interval might change from this initial setting due to zooming, autoscaling, or the **If Y value is beyond limit** setting.

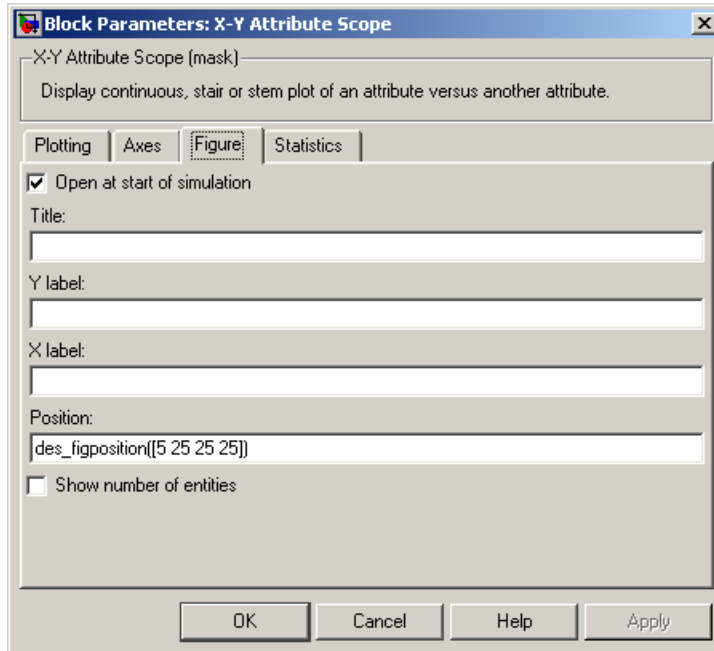
If Y value is beyond limit

Determines how the plot changes if one or more Y values are not within the limits shown on the Y axis. For details, see “Varying Axis Limits Automatically” on page 9-5.

Show grid

Toggles the grid on and off.

Figure Tab



Open at start of simulation

Selecting this option causes the plot window to open when you start the simulation. If you clear this box, you can open the plot window while the simulation is running by double-clicking the block icon.

Title

Text that appears as the title of the plot, above the axes.

Y label

Text that appears to the left of the vertical axis.

X label

Text that appears below the horizontal axis.

X-Y Attribute Scope

Position

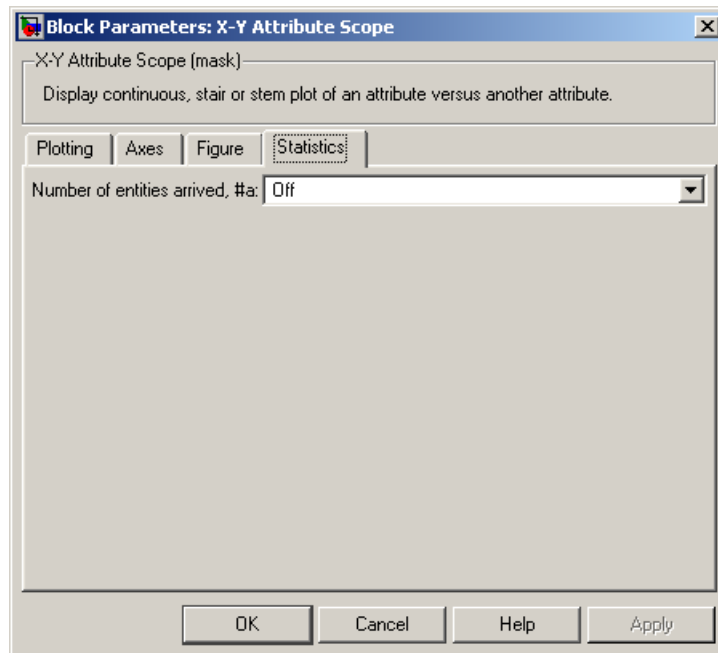
A four-element vector of the form [left bottom width height] specifying the position of the scope window. (0,0) is the lower left corner of the display.

Show number of entities

Displays the number of plotted points using an annotation in the plotting window.

Statistics Tab

These parameters determine whether certain ports output data throughout the simulation, output data only when you stop or pause the simulation, or are omitted from the block. For descriptions of the affected ports, see the “Signal Output Ports” table above.

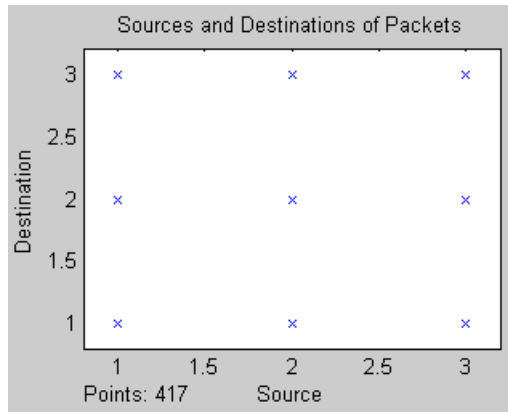


Number of entities arrived

Controls the presence and behavior of the signal output port labeled #a.

Examples

You can modify the example in “Example: A Packet Switch” in the Getting Started documentation to check that all possible combinations of source and destination are used in the simulation. Insert an X-Y Attribute Scope block between the Path Combiner and Output Switch blocks and use it to plot the Destination attribute against the Source attribute.



See Also

Attribute Scope, X-Y Signal Scope, Chapter 9, “Plotting Data”, “Accessing Attributes of Entities” on page 1-20

X-Y Signal Scope

Purpose Plot data from two signals

Library SimEvents Sinks

Description This block plots a curve using data from two input signals. The plot is particularly appropriate for data arising from discrete-event simulations because the plot can include zero-duration values.



The **Plot type** parameter on the **Plotting** tab determines whether and how the block connects the points that it plots. For details, see “Connections Among Points in Plots” on page 9-4.

Ports

Signal Input Ports


Label	Description
in	Signal containing data for Y axis.
x	Signal containing data for X axis.

Signal Output Ports

Label	Description
#c	Number of points the block has plotted.

Dialog Box

Plotting Tab

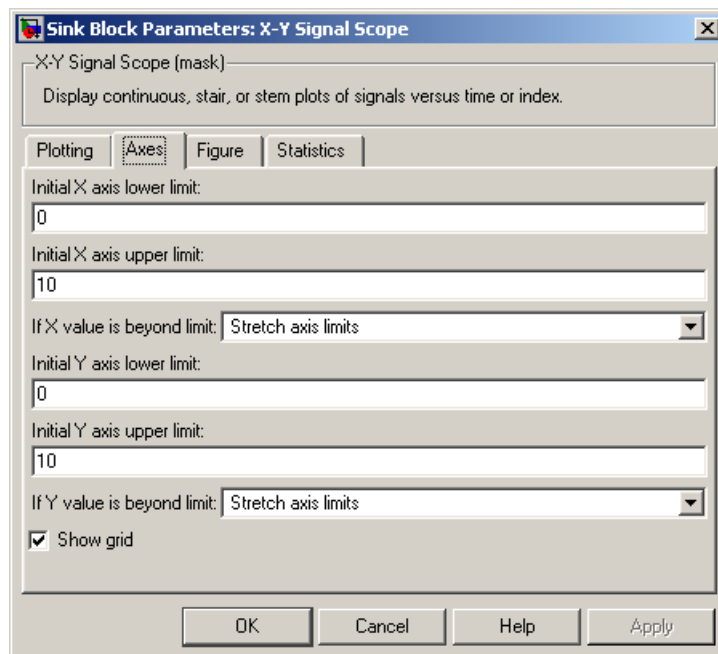


Plot type

The presentation format for the data. See “Connections Among Points in Plots” on page 9-4 for details.

X-Y Signal Scope

Axes Tab



Initial X axis lower limit, Initial X axis upper limit

The interval shown on the X axis at the beginning of the simulation. The interval might change from this initial setting due to zooming, autoscaling, or the **If X value is beyond limit** setting.

If X value is beyond limit

Determines how the plot changes if one or more X values are not within the limits shown on the X axis. For details, see “Varying Axis Limits Automatically” on page 9-5.

Initial Y axis lower limit, Initial Y axis upper limit

The interval shown on the Y axis at the beginning of the simulation. The interval might change from this initial setting

due to zooming, autoscaling, or the **If Y value is beyond limit** setting.

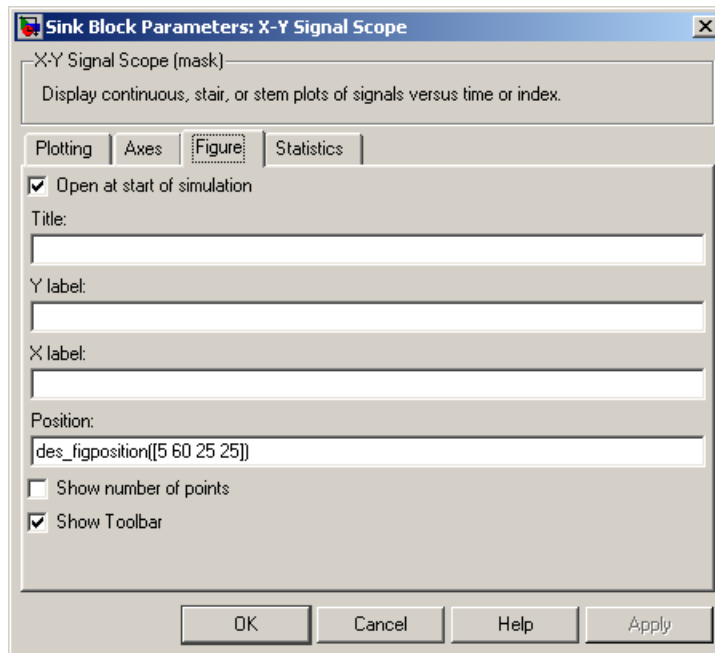
If Y value is beyond limit

Determines how the plot changes if one or more Y values are not within the limits shown on the Y axis. For details, see “Varying Axis Limits Automatically” on page 9-5.

Show grid

Toggles the grid on and off.

Figure Tab



Open at start of simulation

Selecting this option causes the plot window to open when you start the simulation. If you clear this box, you can open the plot

X-Y Signal Scope

window while the simulation is running by double-clicking the block icon.

Title

Text that appears as the title of the plot, above the axes.

Y label

Text that appears to the left of the vertical axis.

X label

Text that appears below the horizontal axis.

Position

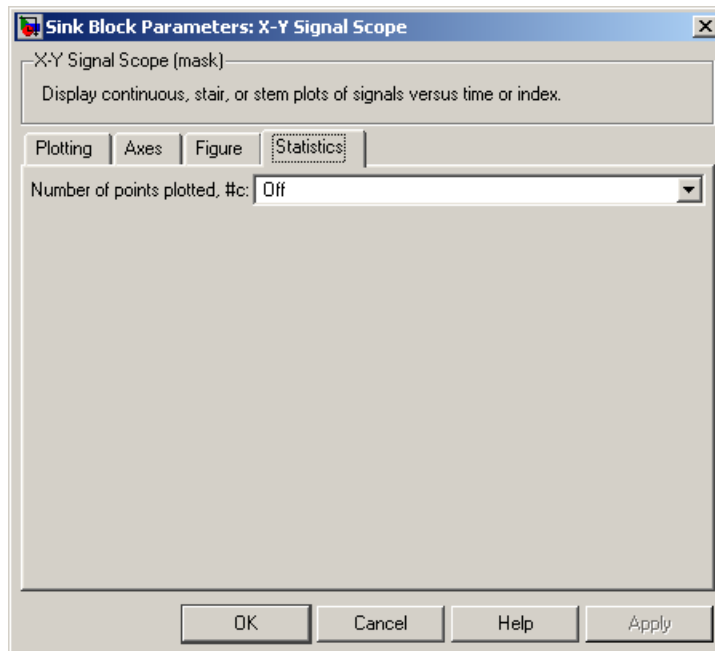
A four-element vector of the form [left bottom width height] specifying the position of the scope window. (0,0) is the lower left corner of the display.

Show number of points

Displays the number of plotted points using an annotation in the plotting window.

Statistics Tab

These parameters determine whether certain ports output data throughout the simulation, output data only when you stop or pause the simulation, or are omitted from the block. For descriptions of the affected ports, see the “Signal Output Ports” table above.



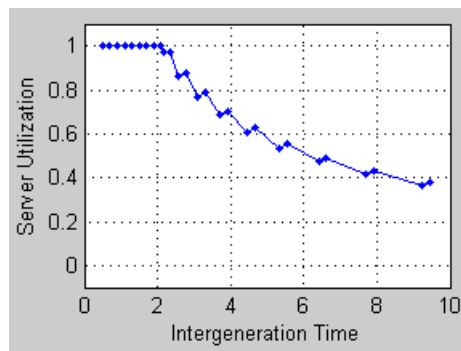
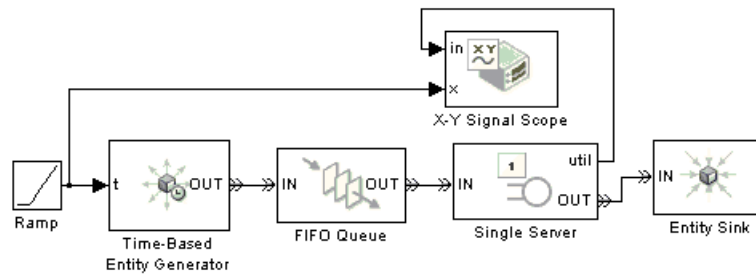
Number of points plotted

Controls the presence and behavior of the signal output port labeled **#c**.

Examples

The model below shows the relationship between the utilization of a server and the interarrival time of entities.

X-Y Signal Scope



See Also

Signal Scope, X-Y Attribute Scope, Chapter 9, “Plotting Data”

advance

To depart from one block and arrive immediately at another block. An entity advances from block to block during a simulation.

arrival

Entrance of an entity to a block via an entity input port. Arrival is the opposite of departure.

attribute

Data associated with an entity.

For example, an entity might be associated with a size, weight, speed, or part number.

available

The state of an entity input port that permits entities to arrive at the block.

For example, when a Single Server block is empty, its entity input port is available. When the block is busy serving, its entity input port is unavailable.

blocked

The state of an entity output port when an entity is trying to depart via the port and the port connects to an unavailable entity input port of another block.

For example, consider a FIFO Queue block whose entity output port is connected to the Single Server block's entity input port. Suppose the queue contains one entity. The queue's entity output port is blocked if the server's entity input port is unavailable, and not blocked if the server's entity input port is available. If the queue is empty, then its entity output port is not blocked because no entity is trying to depart.

departure

Exit of an entity from a block via an entity output port. Departure is the opposite of arrival.

discrete event subsystem

A subsystem containing time-based blocks that is called at the exact time of each qualifying event, rather than at times suggested by the time-based simulation clock. For the Discrete Event Subsystem block, qualifying events are signal-based events; for an appropriately configured Function-Call Subsystem block, qualifying events are function calls.

For example, the subsystem might contain blocks that end the simulation if the length of a queue exceeds 100, and might be configured so that the subsystem executes only at the exact moments when the queue reports an increased length.

entity

An abstract representation of an item of interest in a discrete-event simulation. The specific interpretation of an entity depends on what you are modeling. Entities can carry data, known in SimEvents as attributes.

For example, an entity could represent a packet in a communication network, a person using a bank of elevators, or a part on a conveyor belt.

entity input port

An input port at which an entity can potentially arrive. An entity input port can be available or unavailable; this state, which can change during the simulation, helps determine whether the port actually accepts the arrivals of new entities.

entity output port

An output port from which an entity can potentially depart. An entity output port can have a state of blocked or not blocked; this state, which can change during the simulation, determines whether the port's attempt to output an entity is successful.

entity path

A connection from an entity output port to an entity input port, depicted as a line connecting the entity ports of two blocks. An entity path represents the equivalence between an entity's departure from the first block and arrival at the second block. The connection line depicts a relationship between the two blocks.

An entity path is in active use by an entity only at zero or more discrete times during the simulation. By contrast, a connection line between signal ports represents a signal that has a well-defined value at all times during the simulation.

entity port

An entity input port or an entity output port.

entity priority

A real number associated with an entity, used to determine its sequence in a priority queue.

Contrast with event priority.

entity-departure subsystem

A kind of discrete event subsystem that is called at the exact time of each entity departure from a block or blocks, rather than at times suggested by the time-based simulation clock. See “Creating Entity-Departure Subsystems” on page 8-26 for details.

event

An instantaneous discrete incident that changes a state variable, an output, and/or the occurrence of other events. The prototypical events are arrivals and departures of entities.

Examples of events are the generation of a new data packet in communications, the exit of a person from an elevator, and the placement of a new part on a conveyor belt.

event calendar

The internal list of all events that are scheduled for the current time or future times.

For example, when a server begins its service time on a specific entity, the application inserts an entry into the event calendar for the completion of service on that entity at a future time. In a system representing elevator passengers, this event calendar entry might represent the event whereby a specific person in an elevator reaches the desired floor.

event priority

A positive integer associated with an event, used to sequence the processing of simultaneous events. Simultaneous events having distinct event priorities are processed in ascending order of the event priority values.

Contrast with entity priority.

event translation

Conversion of one event into another. The result of the translation is often a function call, but can be another type of event. The result of the translation can occur at the same time as, or a later time than, the original event.

event-based signal

A signal that can change in response to discrete events. For example, the signal representing the number of entities in a queue changes upon each arrival at or departure from the queue.

event-based simulation

A simulation that permits the system's state transitions to depend on asynchronous discrete incidents called events.

function-call event

A discrete invocation request carried from block to block by a special signal called a function-call signal. A function-call event is also called simply a function call.

intergeneration time

The time interval between successive generations, as in a Time-Based Entity Generator block.

monitoring port

A signal input port that is designed for observing signal values. Contrast with notifying port.

notifying port

A signal input port that notifies the preceding block when a certain event has occurred. When the preceding block is the Event-Based

Random Number block, it responds to the notification by generating a new random number.

For example, the **t** input port of a Single Server block is a notifying port; when connected to this port, the Event-Based Random Number block generates a new random number each time it receives notification that an entity has arrived at the server.

pending entity

An entity that tries and fails to depart from a block. The failure occurs because the entity output port through which the entity would depart is connected to an unavailable entity input port of another block.

preemption

The replacement of an entity in a server block by an entity that satisfies certain criteria.

reactive port

A signal input port that listens for updates or changes in the input signal and causes an appropriate reaction in the block possessing the port. For example, the **p** port on an Input Switch block listens for changes in the input signal; the block reacts by selecting a new entity port for potential arrivals.

sample time hit event

An update in the value of a signal that is connected to a block configured to react to signal updates. The updated value could be the same as or different from the previous value.

signal port

An input or output port that represents a numerical quantity that changes over time and that is defined for all times during the simulation. Unlike an entity port, a signal port has no state and does not have entity arrivals or entity departures.

signal-based event

A sample time hit event, value change event, or trigger event.

simultaneous events

Events that occur at the same value of the simulation clock, even if they are processed sequentially.

For example, in a D/D/1 queuing system where the arrival rate equals the service rate, an entity generation event and a service completion event are simultaneous. Parameters in the model determine which of these events occurs first, though the clock has the same value in both cases.

time-based simulation

A simulation in which state transitions depend on time.

For example, a simulation based solely on differential equations in which time is an independent variable is a time-based simulation.

trigger edge

A rising edge or falling edge of a signal. A rising edge is an increase from a negative or zero value to a positive value (or zero if the initial value is negative). A falling edge is a decrease from a positive or a zero value to a negative value (or zero if the initial value is positive).

trigger event

A trigger edge in a signal that is connected to a block configured to react to trigger edges.

trigger signal

A signal whose trigger edges are used to invoke a behavior during the simulation.

unavailable

The state of an entity input port that prevents entities from arriving at the block.

For example, when a Single Server block is empty, its entity input port is available. When the block is busy serving, its entity input port is unavailable.

value change event

An increase or decrease in the numerical value of a signal that is connected to a block configured to react to relevant changes.

zero-duration value

A value that an event-based signal assumes at an instant in time but that does not persist for a positive duration.

For example, when a full N -server advances one entity to the next block, the statistical signal representing the number of entities in the block assumes the value $N-1$. However, if the departure causes another entity to arrive at the block at the same time instant, then the statistical signal assumes the value N . The value of $N-1$, which does not persist for a positive duration, is a zero-duration value. This phenomenon occurs in many situations; see “Multiple Simultaneous Updates” on page 3-13 for details.

Examples

Use this list to find examples in the documentation.

Attributes of Entities

“Example: Setting Attributes” on page 1-15

“When to Use Attributes” on page 1-17

Counting Entities

“Example: Counting Simultaneous Departures from a Server” on page 1-22

“Example: Resetting a Counter After a Transient Period” on page 1-23

Working with Events

“Example: Comparing Types of Signal-Based Events” on page 2-4

“Example: Race Conditions at a Switch” on page 2-23

“Example: Observing Service Completions” on page 2-32

“Example: Detecting Collisions by Comparing Events” on page 2-34

“Example: Opening a Gate Upon Random Events” on page 2-38

“Example: Counting Events from Multiple Sources” on page 2-41

Queuing Systems

“Example: Event Calendar for a Queue-Server Model” on page 2-11

“Example: Waiting Time in LIFO Queue” on page 4-2

“Example: Serving Preferred Customers First” on page 4-6

“Example: Preemption by High-Priority Entities” on page 4-10

“Example: M/M/5 Queuing System” on page 4-12

“Example: Using Servers in Shifts” on page 6-11

Working with Signals

“Example: Creating a Random Signal for Switching” on page 3-7

“Example: Defining #d Before the First Entity Departure” on page 3-19

“Example: Resampling a Signal Based on Events” on page 3-20

“Example: Sending Queue Length to the Workspace” on page 3-23

Server States

“Example: Failure and Repair of a Server” on page 4-17

“Example: Adding a Warmup Phase” on page 4-19

Routing Entities

“Example: Cascaded Switches with Random Selections” on page 5-2

“Example: Cascaded Switches with Round-Robin Sequence” on page 5-4

“Example: Compound Switching Logic” on page 5-5

“Example: Choosing the Shortest Queue” on page 6-3

Batching

“Example: Varying Fluid Flow Rate Based on Batching Logic” on page 6-6

Gates

“Example: Controlling Joint Availability of Two Servers” on page 7-4

“Example: Synchronizing Service Start Times with the Clock” on page 7-6

“Example: Opening a Gate Upon Entity Departures” on page 7-7

“Example: First Entity as a Special Case” on page 7-10

Discrete Event Subsystems

“Example: Comparing the Lengths of Two Queues” on page 8-16

- “Example: Normalizing a Statistic to Use for Routing” on page 8-17
- “Example: Using Event-Based Timing for a Statistical Computation” on page 8-19
- “Example: Ending the Simulation Upon an Event” on page 8-20
- “Example: Sending Unrepeated Data to the MATLAB Workspace” on page 8-21
- “Example: Focusing on Events, Not Values” on page 8-22
- “Example: Detecting Changes from Empty to Nonempty” on page 8-23
- “Example: Logging Data About the First Entity on a Path” on page 8-24
- “Example: Using Entity-Based Timing for Choosing a Port” on page 8-29
- “Example: Performing a Computation on Selected Entity Paths” on page 8-31

Troubleshooting

- “Example: Plotting Entity Departures to Verify Timing” on page 9-8
- “Example: Plotting Event Counts to Find Roundoff Error” on page 9-12
- “Example: Event Logging” on page 11-6
- “Example: Entity Logging” on page 11-9
- “Example: Discrepancies in Event Times” on page 11-13
- “Example: Time-Based Addition of Event-Based Signals” on page 11-17
- “Example: Incorrect Sequence of Event Priorities” on page 11-20
- “Example: Intergeneration Time of Zero at Simulation Start” on page 11-23
- “Example: Absence of Sample Time Hit at Simulation Start” on page 11-24
- “Example: Faulty Logic in Feedback Loop” on page 11-25
- “Example: Deadlock Resulting from Loop in Entity Path” on page 11-26
- “Example: Invalid Connection of Event-Based Random Number Generator” on page 11-28
- “Example: Sequence of Departures and Statistical Updates” on page 12-5
- “Example: Using a #n Signal as a Trigger” on page 12-11

Statistics

- “Example: Fraction of Dropped Messages” on page 10-8

- “Example: Computing a Time Average of a Signal” on page 10-9
- “Example: Resetting an Average Periodically” on page 10-12
- “Example: Running a Simulation Repeatedly to Gather Results” on page 10-27
- “Example: Running a Simulation and Varying a Parameter” on page 10-29

Timers

- “Basic Example Using Timer Blocks” on page 10-18
- “Timing Multiple Entity Paths with One Timer” on page 10-20
- “Restarting a Timer” on page 10-21
- “Timing Multiple Processes Independently” on page 10-23

A

- arbitrary event sequence 2-20
 - troubleshooting 11-16
- Attribute Scope block 16-2
- attributes of entities
 - plots 9-2
 - reading values 1-20
 - setting values 1-14
 - usage 1-17
- autoscaling 9-7
- averaging signals
 - over samples 10-12
 - over time 10-9
- axis limits 9-5

B

- block-to-block interactions 12-2

C

- cascading switch blocks
 - random 5-2
 - round robin 5-4
- combining events 2-40
- conditional events 2-45
- Configuration Parameters dialog box 11-30
- copying entities 1-25
- counting entities
 - cumulative 1-21
 - instantaneous 1-21
 - reset 1-23
 - storing in attribute 1-24
- counting events 9-2
 - roundoff error 9-12

D

- data types 3-3
- delays

- entities 16-65
 - signal updates 3-17
- discouraged arrivals 13-13
- Discrete Event Inport block 16-10
- Discrete Event Outport block 16-12
- Discrete Event Signal to Workspace block 16-13
- Discrete Event Subsystem block 16-16
 - building subsystems 8-11
- discrete event subsystems 8-7
 - blocks inside 8-10
 - building
 - entity departures 8-27
 - function calls 8-32
 - signal-based events 8-11
 - combinations of events 8-32
 - entity departures 8-26
 - events in terminated signals 8-22
 - function calls 8-32
 - need for 8-2
 - sequence of events 8-8
 - signal-based events 8-14
- discrete state plots 9-4
- discrete-event plots 9-4
 - compared to time-based plots 9-15
 - customizing 9-7
 - saving FIG-file 9-7
 - troubleshooting using 9-8
- dropped messages 10-8

E

- Embedded MATLAB Function blocks 6-3
- Enabled Gate block 16-18
- enabled gates 7-4
- entities
 - counting 1-21
 - event-based generation 1-2
 - logging 11-8
 - replicating 1-25

- entity collisions 2-34
 - entity data
 - plots 9-2
 - reading values 1-20
 - setting values 1-14
 - usage 1-17
 - Entity Departure Counter block 16-21
 - Entity Departure Event to Function-Call
 - Event block 16-26
 - entity generation
 - changes in signal value 1-4
 - event-based 1-2
 - function calls 1-7
 - trigger edges 1-5
 - updates in signal value 1-2
 - vector of times 1-11
 - entity logging 11-8
 - Entity Sink block 16-32
 - Entity-Based Function-Call Event Generator
 - block 16-34
 - entity-departure subsystems 8-26
 - equal event priorities 2-20
 - troubleshooting 11-16
 - event calendar 2-9
 - events on 2-9
 - example 2-11
 - logging 11-2
 - event logging 11-2
 - Event-Based Entity Generator block 16-37
 - Event-Based Random Number block 16-43
 - event-based signals
 - description 3-2
 - feedback loops 11-23
 - initial conditions 3-18
 - integrating 8-3
 - latency 11-19
 - manipulating 3-18
 - MATLAB workspace 3-23
 - random 3-4
 - resampling 3-20
 - troubleshooting 11-17
 - unrepeated values to workspace 8-21
 - update sequence 3-10
 - events
 - conditionalizing 2-45
 - generating 2-36
 - manipulating 2-39
 - observing 2-30
 - priorities 2-19
 - sequence 2-21
 - supported types 2-2
 - translating 2-43
 - troubleshooting 11-13
 - union 2-40
- ## F
- failure modeling
 - conditional events 2-46
 - gates 4-15
 - Stateflow 4-16
 - feedback entity paths
 - troubleshooting 11-26
 - feedback loops
 - troubleshooting 11-23
 - FIFO Queue block 16-56
 - first-order-hold plots 9-4
 - function calls 2-7
 - generating 2-36
 - function-call subsystems
 - discrete event 8-32
- ## G
- G/G/1 queuing systems 13-10
 - gates 7-1
 - combinations 7-9
 - enabled 7-4
 - entity departures 7-7
 - release 7-6

- role in modeling 7-2
- types 7-3

Get Attribute block 16-61

I

independent replications 10-25

Infinite Server block 16-65

initial conditions 3-18

- feedback loops 11-23

initial seeds 10-25

Input Switch block 16-71

Instantaneous Entity Counting Scope block 16-76

Instantaneous Event Counting Scope block 16-81

instantaneous gate openings 7-6

integration

- event-based signals 8-3

intergeneration times

- event generation 2-38

interleaved operations 12-4

L

latency

- interleaved operations 12-4
- signal updates 3-17
- troubleshooting 11-19

LIFO Queue block 16-86

LIFO queues 4-2

Little's law 13-10

logging entities 11-8

logging events 11-2

logic

- block diagrams 6-10
- MATLAB code 6-3
- usage in discrete-event simulations 6-2

loops in entity paths 11-26

M

M/D/1 queuing systems 13-9

M/M/1 queuing systems 13-7

M/M/5 queuing systems 4-12

monitoring ports 16-43

N

N-Server block 16-91

n-servers 4-12

nonstorage blocks 12-4

notifying ports 3-4

O

Output Switch block 16-97

P

Path Combiner block 16-103

plots 9-1

- customizing 9-7
- troubleshooting using 9-8
- zero-duration values 3-14

ports

- event-based random numbers 16-43
- notifying 3-4

preemption in servers 4-9

priorities, entity

- priority queues with preemptive servers 4-10
- queue sequence 4-4
- server preemption 4-9

priorities, event 2-19

- troubleshooting 11-16

Priority Queue block 16-110

Q

queues

- choosing shortest using logic blocks 6-16

- choosing shortest using MATLAB code 6-3
- LIFO vs. FIFO 4-2
- preemptive servers 4-10
- priority 4-4
- queuing systems
 - discouraged arrivals 13-13
 - G/G/1 13-10
 - M/D/1 13-9
 - M/M/1 13-7
 - M/M/5 4-12

R

- race conditions 2-23
- random
 - signals 3-4
- random event sequence 2-20
 - troubleshooting 11-16
- random numbers
 - event-based 3-4
 - switch selection 3-7
 - time-based 3-8
- reactive ports 16-43
- Read Timer block 16-115
- Release Gate block 16-120
- release gates 7-6
- repeating simulations 10-27
- Replicate block 16-124
- replication of entities 1-25
- replications
 - independent 10-25
- resampling signals 3-20
- resetting averages 10-12
- residual service time 4-9
- roundoff error
 - discrete-event plots 9-12
 - event times 11-13
- running simulations
 - repeatedly 10-27
 - varying parameters 10-29

S

- sample means 10-12
- sample time 3-2
- scatter plots 9-4
- scope blocks 9-1
 - zero-duration values 3-14
- seed of random number generator
 - independent replications 10-25
- server states 4-15
- servers
 - failure states 4-15
 - infinite 16-65
 - multiple 4-12
 - preemption 4-9
- Set Attribute block 16-131
- Signal Latch block 16-135
- Signal Scope block 16-145
- Signal-Based Event to Function-Call Event block 16-153
- signal-based events
 - comparison 2-4
 - definition 2-3
- Signal-Based Function-Call Event Generator block 16-159
- signals
 - event-based 3-2
 - random 3-4
- simeventsconfig function 14-2
- simeventslib function 14-3
- simeventsstartup function 14-4
- simulation parameters
 - varying in repeated runs 10-29
- simultaneous events
 - discrete event subsystems 8-8
 - event priorities 2-19
 - interleaved operations 12-4
 - sequential processing 2-21
 - signal updates 3-13
 - troubleshooting 11-13
 - unexpected 11-13

Single Server block 16-163
stack 4-2
stairstep plots 9-4
Start Timer block 16-173
Stateflow and SimEvents 4-16
statistics 10-2
 accessing from blocks 10-4
 custom 10-7
 discrete event subsystems 8-3
 interleaved updates 12-4
 latency 3-17
Statistics tab 10-4
stem plots 9-4
stop time 10-32
 event-based timing 8-2
storage blocks 12-4
 changing processing sequence 12-6
 looped entity paths 11-26
Subsystem Configuration block 16-176
switching entity paths
 random with cascaded blocks 5-2
 repeating sequence 8-4
 round robin with cascaded blocks 5-4
synchronizing entities 7-6

T

time averages 10-9
time-based blocks
 zero-duration values 12-10
Time-Based Entity Generator block 16-177
timer tags 10-19

timers 10-18
 independent 10-23
 multiple entity paths 10-20
 restarting 10-21
triggers
 zero-duration values 12-11

V

visualization 9-1
 zero-duration values 3-14

W

workspace 3-23
 unrepeated signal values 8-21

X

X-Y Attribute Scope block 16-185
X-Y Signal Scope block 16-192

Z

zero-duration values
 definition 3-13
 MATLAB workspace 3-16
 time-based blocks 12-10
 visualization 3-14
zero-order hold
 plots 9-4