


Fixed-Point Toolbox™ 2

User's Guide



MATLAB®

How to Contact The MathWorks



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



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



508-647-7000 (Phone)



508-647-7001 (Fax)



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

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

Fixed-Point Toolbox™ User's Guide

© COPYRIGHT 2004–2009 by The MathWorks, Inc.

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

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

Trademarks

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

Patents

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

Revision History

June 2004	First printing	New for Version 1.0 (Release 14)
October 2004	Online only	Version 1.1 (Release 14SP1)
March 2005	Online only	Version 1.2 (Release 14SP2)
September 2005	Online only	Version 1.3 (Release 14SP3)
October 2005	Second printing	Version 1.3
March 2006	Online only	Version 1.4 (R2006a)
September 2006	Third printing	Version 1.5 (R2006b)
March 2007	Fourth printing	Version 2.0 (R2007a)
September 2007	Online only	Revised for Version 2.1 (R2007b)
March 2008	Online only	Revised for Version 2.2 (R2008a)
October 2008	Online only	Revised for Version 2.3 (R2008b)
March 2009	Online only	Revised for Version 2.4 (R2009a)

Getting Started

1

Product Overview	1-2
Licensing	1-4
Getting Help	1-5
Getting Help in This Document	1-5
Getting Help at the MATLAB Command Line	1-5
Display Settings	1-7
Demos	1-9

Fixed-Point Concepts

2

Fixed-Point Data Types	2-2
Scaling	2-4
Precision and Range	2-5
Range	2-5
Precision	2-6
Arithmetic Operations	2-10
Modulo Arithmetic	2-10
Two's Complement	2-11
Addition and Subtraction	2-12
Multiplication	2-13
Casts	2-19

fi Objects Compared to C Integer Data Types	2-22
Integer Data Types	2-22
Unary Conversions	2-24
Binary Conversions	2-25
Overflow Handling	2-28

Working with fi Objects

3

Constructing fi Objects	3-2
fi Object Syntaxes	3-2
Examples of Constructing fi Objects	3-3
Casting fi Objects	3-13
Overwriting by Assignment	3-13
Ways to Cast with MATLAB Software	3-13
fi Object Properties	3-18
Data Properties	3-18
fimath Properties	3-18
numericType Properties	3-20
Setting fi Object Properties	3-21
fi Object Functions	3-23

Working with fimath Objects

4

Constructing fimath Objects	4-2
fimath Object Syntaxes	4-2
Building fimath Object Constructors in a GUI	4-4
Configuring the MATLAB Default fimath Object	4-5
fimath Object Properties	4-10
Math, Rounding, and Overflow Properties	4-10

Setting fimath Object Properties	4-11
Using fimath Objects to Perform Fixed-Point	
Arithmetic	4-15
Binary Point Arithmetic	4-15
[Slope Bias] Arithmetic	4-19
Using fimath to Specify Rounding and Overflow	
Modes	4-23
Using fimath to Share Arithmetic Rules	4-24
Using fimath ProductMode and SumMode	4-26
Example Setup	4-26
FullPrecision	4-27
KeepLSB	4-28
KeepMSB	4-29
SpecifyPrecision	4-30
fimath Object Functions	4-32

Working with fipref Objects

5

Constructing fipref Objects	5-2
fipref Object Properties	5-3
Display, Data Type Override, and Logging Properties	5-3
Setting fipref Object Properties	5-3
Using fipref Objects to Set Display Preferences	5-5
Using fipref Objects to Set Logging Preferences	5-7
Logging Overflows and Underflows as Warnings	5-7
Accessing Logged Information with Functions	5-9

Using fipref Objects to Set Data Type Override	
Preferences	5-12
Overriding the Data Type of fi Objects	5-12
Using Data Type Override to Help Set Fixed-Point Scaling	5-13
fipref Object Functions	5-15

Working with numerictype Objects

6

Constructing numerictype Objects	6-2
numerictype Object Syntaxes	6-2
Example: Constructing a numerictype Object with Property Name and Property Value Pairs	6-3
Example: Copying a numerictype Object	6-4
Example: Building numerictype Object Constructors in a GUI	6-4
numerictype Object Properties	6-7
Data Type and Scaling Properties	6-7
Setting numerictype Object Properties	6-8
The numerictype Structure	6-11
Valid Values for numerictype Structure Properties	6-11
Properties That Affect the Slope	6-13
Stored Integer Value and Real World Value	6-13
Using numerictype Objects to Share Data Type and	
Scaling Settings of fi objects	6-14
Example 1	6-14
Example 2	6-15
numerictype Object Functions	6-18

Working with quantizer Objects

7

Constructing quantizer Objects	7-2
quantizer Object Properties	7-3
Quantizing Data with quantizer Objects	7-4
Transformations for Quantized Data	7-6
quantizer Object Functions	7-7

Working with the Fixed-Point Embedded MATLAB Subset

8

Supported Functions and Limitations of the Fixed-Point Embedded MATLAB Subset	8-2
Embedded MATLAB MEX	8-9
Using Embedded MATLAB MEX with Fixed-Point Algorithms	8-9
Speeding Up Fixed-Point Execution with the eml mex Function	8-9
Using Data Type Override with Embedded MATLAB MEX	8-12
Embedded MATLAB Function Block	8-13
Composing a MATLAB Language Function in a Simulink Model	8-13
Using Fixed-Point Data Types with the Embedded MATLAB Function Block	8-14
Using the Embedded MATLAB Function Block with Data Type Override	8-14
Using the Model Explorer with a Fixed-Point Embedded MATLAB Function Block	8-15

Embedded MATLAB Coder	8-35
Working with Fixed-Point Compilation Reports	8-36
Generating the Compilation Report	8-36
Opening the Compilation Report	8-37
Viewing Your M-Code	8-37
Viewing Variables in the Variables Tab	8-39
See Also	8-40

Interoperability with Other Products

9

Using fi Objects with Simulink	9-2
Reading Fixed-Point Data from the Workspace	9-2
Writing Fixed-Point Data to the Workspace	9-2
Setting the Value and Data Type of Block Parameters ...	9-6
Logging Fixed-Point Signals	9-6
Accessing Fixed-Point Block Data During Simulation	9-6
Using the Embedded MATLAB Subset	9-8
Using fi Objects with Signal Processing Blockset	
Software	9-9
Reading Fixed-Point Signals from the Workspace	9-9
Writing Fixed-Point Signals to the Workspace	9-9
Using fi Objects with Filter Design Toolbox Software ..	9-14

Index

Getting Started

- “Product Overview” on page 1-2
- “Licensing” on page 1-4
- “Getting Help” on page 1-5
- “Display Settings” on page 1-7
- “Demos” on page 1-9

Product Overview

Fixed-Point Toolbox™ software provides fixed-point data types in MATLAB® technical computing software and enables algorithm development by providing fixed-point arithmetic. The toolbox enables you to create the following types of objects:

- `fi` — Defines a fixed-point numeric object in the MATLAB workspace. Each `fi` object is composed of value data, a `fi` object, and a `numericity` object.
- `fi`math — Governs how overloaded arithmetic operators work with `fi` objects
- `fi`pref — Defines the display, logging, and data type override preferences of `fi` objects
- `numericity` — Defines the data type and scaling attributes of `fi` objects
- `quantizer` — Quantizes data sets

Fixed-Point Toolbox™ software provides you with

- The ability to define fixed-point data types, scaling, and rounding and overflow methods in the MATLAB workspace
- Bit-true real and complex simulation
- Basic fixed-point arithmetic
 - Arithmetic operators `+`, `-`, `*`, `.*` for binary point-only and real [Slope Bias] signals
 - Division for binary point-only signals using the `divide` function or arithmetic operators `/` and `./`
- Arbitrary word length up to `intmax('uint16')` bits
- Logging of minimums, maximums, overflows, and underflows
- Data type override with singles, doubles, or scaled doubles
- Conversions between binary, hex, double, and built-in integers
- Relational, logical, and bitwise operators
- Matrix functions such as `ctranspose` and `horzcat`

- Statistics functions such as `max` and `min`
- Interoperability with Simulink[®], Signal Processing Blockset[™] software, Embedded MATLAB[™] subset, and Filter Design Toolbox[™] software
- Compatibility with the Simulink To Workspace and From Workspace blocks

Licensing

You can use `fi` objects with the `DataType` property set to `double` *without* a Fixed-Point Toolbox™ license when the `fipref LoggingMode` property is set to `off`. A Fixed-Point Toolbox™ license *is* checked out when you

- Use any `fi` object with any `DataType` other than `double`.
- Create any `fi` object when the `fipref LoggingMode` property is set to `on`, including `fi` objects with `DataType double`.
- Load a MAT-file that contains any `fi` object with the `DataType` property set to `single`, `boolean`, `ScaledDouble`, or `Fixed`.

You can prevent the checkout of a Fixed-Point Toolbox™ license when working with Fixed-Point Toolbox™ code by setting the `fipref DataTypeOverride` property to `TrueDoubles`.

Getting Help

In this section...
“Getting Help in This Document” on page 1-5
“Getting Help at the MATLAB Command Line” on page 1-5

Getting Help in This Document

The objects of Fixed-Point Toolbox software are discussed in the following chapters:

- Chapter 3, “Working with fi Objects”
- Chapter 4, “Working with fimath Objects”
- Chapter 5, “Working with fipref Objects”
- Chapter 6, “Working with numericitype Objects”
- Chapter 7, “Working with quantizer Objects”

To get in-depth information about the properties of these objects, refer to the Property Reference.

To get in-depth information about the functions of these objects, refer to the Function Reference.

Getting Help at the MATLAB Command Line

To get command-line help for Fixed-Point Toolbox objects, type

```
help objectname
```

For example,

```
help fi
help fimath
help fipref
help numericitype
help quantizer
```

To get command-line help for Fixed-Point Toolbox functions, type

```
help embedded.fi/functionname
```

For example,

```
help embedded.fi/abs  
help embedded.fi/bitset  
help embedded.fi/sqrt
```

To invoke Help Browser documentation for Fixed-Point Toolbox functions from the MATLAB command line, type

```
doc fixedpoint/functionname
```

For example,

```
doc fixedpoint/int  
doc fixedpoint/add  
doc fixedpoint/savefipref  
doc fixedpoint/quantize
```


Display Settings

In Fixed-Point Toolbox software, the display of `fi` objects is determined by the `fi`pref object. Throughout this User's Guide, code examples of `fi` objects are usually shown as they appear when the `fi`pref properties are set as follows:

- `NumberDisplay` — `'RealWorldValue'`
- `NumericTypeDisplay` — `'full'`
- `FimathDisplay` — `'none'`

For example,

```
p = fi(pref('NumberDisplay', 'RealWorldValue',...  
'NumericTypeDisplay', 'full', 'FimathDisplay', 'none'))
```

```
p =
```

```
      NumberDisplay: 'RealWorldValue'  
NumericTypeDisplay: 'full'  
      FimathDisplay: 'none'  
      LoggingMode: 'Off'  
      DataTypeOverride: 'ForceOff'
```

```
a = fi(pi)
```

```
a =
```

```
3.1416
```

```
      DataTypeMode: Fixed-point: binary point scaling  
      Signedness: Signed  
      WordLength: 16  
      FractionLength: 13
```

In other cases, it makes sense to also show the fimath object display:

- NumberDisplay — 'RealWorldValue'
- NumericTypeDisplay — 'full'
- FimathDisplay — 'full'

For example,

```
p = fipref('NumberDisplay', 'RealWorldValue',...  
         'NumericTypeDisplay', 'full', 'FimathDisplay', 'full')
```

```
p =
```

```
         NumberDisplay: 'RealWorldValue'  
NumericTypeDisplay: 'full'  
         FimathDisplay: 'full'  
         LoggingMode: 'Off'  
         DataTypeOverride: 'ForceOff'
```

```
a = fi(pi)
```

```
a =
```

```
3.1416
```

```
         DataTypeMode: Fixed-point: binary point scaling  
         Signedness: Signed  
         WordLength: 16  
         FractionLength: 13
```

```
         RoundMode: nearest  
         OverflowMode: saturate  
         ProductMode: FullPrecision  
MaxProductWordLength: 128  
         SumMode: FullPrecision  
MaxSumWordLength: 128  
         CastBeforeSum: true
```

For more information, refer to Chapter 5, “Working with fipref Objects”.

Demos

You can access demos in the **Demos** tab of the **Help Navigator** window. Fixed-Point Toolbox software includes the following demos:

- Fixed-Point Basics — Demonstrates the basic use of the fixed-point `fi` object
- Number Circle — Illustrates the definitions of unsigned and signed two's complement integer and fixed-point numbers
- Binary Point Scaling — Explains binary point-only scaling
- Fixed-Point Data Type Override, Min/Max Logging, and Scaling — Steps through the workflow of using doubles override and min/max logging in the toolbox to choose appropriate scaling for a fixed-point algorithm
- Fixed-Point C Development — Shows how to use the parameters from a fixed-point MATLAB program in a fixed-point C program
- Fixed-Point Algorithm Development — Presents the development and verification of a simple fixed-point algorithm
- Fixed-Point Fast Fourier Transform (FFT) — Provides an example of converting a textbook Fast Fourier Transform algorithm into fixed-point MATLAB code and then into fixed-point C code
- Analysis of a Fixed-Point State-Space System with Limit Cycles — Demonstrates a limit cycle detection routine applied to a state-space system
- Quantization Error — Demonstrates the statistics of the error when signals are quantized using various rounding methods
- Fixed-Point Lowpass Filtering Using Embedded MATLAB MEX — Steps through generating a C-MEX function from M code, running the generated C-MEX function, and displaying the results
- Fixed-Point ATAN2 Calculation — Uses the CORDIC algorithm and polynomial approximation to perform a fixed-point calculation of the four quadrant inverse tangent

Fixed-Point Concepts

- “Fixed-Point Data Types” on page 2-2
- “Scaling” on page 2-4
- “Precision and Range” on page 2-5
- “Arithmetic Operations” on page 2-10
- “fi Objects Compared to C Integer Data Types” on page 2-22


Fixed-Point Data Types

In digital hardware, numbers are stored in binary words. A binary word is a fixed-length sequence of bits (1's and 0's). How hardware components or software functions interpret this sequence of 1's and 0's is defined by the data type.

Binary numbers are represented as either fixed-point or floating-point data types. This chapter discusses many terms and concepts relating to fixed-point numbers, data types, and mathematics.

A fixed-point data type is characterized by the word length in bits, the position of the binary point, and whether it is signed or unsigned. The position of the binary point is the means by which fixed-point values are scaled and interpreted.

For example, a binary representation of a generalized fixed-point number (either signed or unsigned) is shown below:



where

- b_i is the i^{th} binary digit.
- wl is the word length in bits.
- b_{wl-1} is the location of the most significant, or highest, bit (MSB).
- b_0 is the location of the least significant, or lowest, bit (LSB).
- The binary point is shown four places to the left of the LSB. In this example, therefore, the number is said to have four fractional bits, or a fraction length of four.

Fixed-point data types can be either signed or unsigned. Signed binary fixed-point numbers are typically represented in one of three ways:

- Sign/magnitude
- One's complement
- Two's complement

Two's complement is the most common representation of signed fixed-point numbers and is the only representation used by Fixed-Point Toolbox documentation. Refer to “Two's Complement” on page 2-11 for more information.

Scaling

Fixed-point numbers can be encoded according to the scheme

$$\text{real-world value} = (\text{slope} \times \text{integer}) + \text{bias}$$

where the slope can be expressed as

$$\text{slope} = \text{fractional slope} \times 2^{\text{fixed exponent}}$$

The integer is sometimes called the *stored integer*. This is the raw binary number, in which the binary point assumed to be at the far right of the word. In Fixed-Point Toolbox documentation, the negative of the fixed exponent is often referred to as the *fraction length*.

The slope and bias together represent the scaling of the fixed-point number. In a number with zero bias, only the slope affects the scaling. A fixed-point number that is only scaled by binary point position is equivalent to a number in [Slope Bias] representation that has a bias equal to zero and a fractional slope equal to one. This is referred to as binary point-only scaling or power-of-two scaling:

$$\text{real-world value} = 2^{\text{fixed exponent}} \times \text{integer}$$

or

$$\text{real-world value} = 2^{-\text{fraction length}} \times \text{integer}$$

Fixed-Point Toolbox software supports both binary point-only scaling and [Slope Bias] scaling.

Note For examples of binary point-only scaling, see the Fixed-Point Toolbox demo “fi Binary Point Scaling.”

Precision and Range

In this section...


“Range” on page 2-5

“Precision” on page 2-6

Note You must pay attention to the precision and range of the fixed-point data types and scalings you choose in order to know whether rounding methods will be invoked or if overflows or underflows will occur.

Range


The range is the span of numbers that a fixed-point data type and scaling can represent. The range of representable numbers for a two’s complement fixed-point number of word length wl , scaling S and bias B is illustrated below:



For both signed and unsigned fixed-point numbers of any data type, the number of different bit patterns is 2^{wl} .

For example, in two’s complement, negative numbers must be represented as well as zero, so the maximum value is $2^{wl-1} - 1$. Because there is only one representation for zero, there are an unequal number of positive and negative numbers. This means there is a representation for -2^{wl-1} but not for 2^{wl-1} :

For slope = 1 and bias = 0:



Overflow Handling

Because a fixed-point data type represents numbers within a finite range, overflows and underflows can occur if the result of an operation is larger or smaller than the numbers in that range.

Fixed-Point Toolbox software allows you to either *saturate* or *wrap* overflows. Saturation represents positive overflows as the largest positive number in the range being used, and negative overflows as the largest negative number in the range being used. Wrapping uses modulo arithmetic to cast an overflow back into the representable range of the data type. Refer to “Modulo Arithmetic” on page 2-10 for more information.

When you create a `fi` object, any overflows are saturated. The `OverflowMode` property of the default `fi` object is `saturate`. You can log overflows and underflows by setting the `LoggingMode` property of the `fi` object to `on`. Refer to “LoggingMode” for more information.

Precision

The precision of a fixed-point number is the difference between successive values representable by its data type and scaling, which is equal to the value of its least significant bit. The value of the least significant bit, and therefore the precision of the number, is determined by the number of fractional bits. A fixed-point value can be represented to within half of the precision of its data type and scaling.

For example, a fixed-point representation with four bits to the right of the binary point has a precision of 2^{-4} or 0.0625, which is the value of its least significant bit. Any number within the range of this data type and scaling can be represented to within $(2^{-4})/2$ or 0.03125, which is half the precision. This is an example of representing a number with finite precision.

Rounding Methods

When you represent numbers with finite precision, not every number in the available range can be represented exactly. If a number cannot be represented exactly by the specified data type and scaling, a rounding method is used to cast the value to a representable number. Although precision is always lost in the rounding operation, the cost of the operation and the amount of bias that is introduced depends on the rounding method itself. To provide you with

greater flexibility in the trade-off between cost and bias, Fixed-Point Toolbox software currently supports the following rounding methods:

- `ceil` rounds to the closest representable number in the direction of positive infinity.
- `convergent` rounds to the closest representable number. In the case of a tie, `convergent` rounds to the nearest even number. This is the least biased rounding method provided by the toolbox.
- `fix` rounds to the closest representable number in the direction of zero.
- `floor`, which is equivalent to two's complement truncation, rounds to the closest representable number in the direction of negative infinity.
- `nearest` rounds to the closest representable number. In the case of a tie, `nearest` rounds to the closest representable number in the direction of positive infinity. This rounding method is the default for `fi` object creation and `fi` arithmetic.
- `round` rounds to the closest representable number. In the case of a tie, the `round` method rounds:
 - Positive numbers to the closest representable number in the direction of positive infinity.
 - Negative numbers to the closest representable number in the direction of negative infinity.

Choosing a Rounding Method. Each rounding method has a set of inherent properties. Depending on the requirements of your design, these properties could make the rounding method more or less desirable to you. By knowing the requirements of your design and understanding the properties of each rounding method, you can determine which is the best fit for your needs. The most important properties to consider are:

- Cost — Independent of the hardware being used, how much processing expense does the rounding method require?
 - Low — The method requires few processing cycles.
 - Moderate — The method requires a moderate number of processing cycles.
 - High — The method requires more processing cycles.

Note The cost estimates provided here are hardware independent. Some processors have rounding modes built-in, so consider carefully the hardware you are using before calculating the true cost of each rounding mode.

- Bias — What is the expected value of the rounded values minus the original values: $E(\hat{\theta} - \theta)$?
 - $E(\hat{\theta} - \theta) < 0$ — The rounding method introduces a negative bias.
 - $E(\hat{\theta} - \theta) = 0$ — The rounding method is unbiased.
 - $E(\hat{\theta} - \theta) > 0$ — The rounding method introduces a positive bias.
- Possibility of Overflow — Does the rounding method introduce the possibility of overflow?
 - Yes — The rounded values may exceed the minimum or maximum representable value.
 - No — The rounded values will never exceed the minimum or maximum representable value.

The following table shows a comparison of the different rounding methods available in both Fixed-Point Toolbox and Simulink® Fixed Point™ products.

Fixed-Point Toolbox Rounding Method	Simulink Fixed Point Rounding Mode	Cost	Bias	Possibility of Overflow
ceil	Ceiling	Low	Large positive	Yes
convergent	Convergent	High	Unbiased	Yes
fix	Zero	Low	<ul style="list-style-type: none"> • Large positive for negative samples • Unbiased for samples with evenly distributed positive and negative values • Large negative for positive samples 	No
floor	Floor	Low	Large negative	No
nearest	Nearest	Moderate	Small positive	Yes
round	Round	High	<ul style="list-style-type: none"> • Small negative for negative samples • Unbiased for samples with evenly distributed positive and negative values • Small positive for positive samples 	Yes
N/A	Simplest (Simulink Fixed Point only)	Low	Depends on the operation	No

Arithmetic Operations


In this section...
“Modulo Arithmetic” on page 2-10
“Two’s Complement” on page 2-11
“Addition and Subtraction” on page 2-12
“Multiplication” on page 2-13
“Casts” on page 2-19

Note These sections will help you understand what data type and scaling choices result in overflows or a loss of precision.

Modulo Arithmetic

Binary math is based on modulo arithmetic. Modulo arithmetic uses only a finite set of numbers, wrapping the results of any calculations that fall outside the given set back into the set.

For example, the common everyday clock uses modulo 12 arithmetic. Numbers in this system can only be 1 through 12. Therefore, in the “clock” system, 9 plus 9 equals 6. This can be more easily visualized as a number circle:



Similarly, binary math can only use the numbers 0 and 1, and any arithmetic results that fall outside this range are wrapped “around the circle” to either 0 or 1.

Two’s Complement

Two’s complement is a way to interpret a binary number. In two’s complement, positive numbers always start with a 0 and negative numbers always start with a 1. If the leading bit of a two’s complement number is 0, the value is obtained by calculating the standard binary value of the number. If the leading bit of a two’s complement number is 1, the value is obtained by assuming that the leftmost bit is negative, and then calculating the binary value of the number. For example,

$$01 = (0 + 2^0) = 1$$

$$11 = \left((-2^1) + (2^0) \right) = (-2 + 1) = -1$$

To compute the negative of a binary number using two’s complement,

- 1 Take the one’s complement, or “flip the bits.”

2 Add a 1 using binary math.

3 Discard any bits carried beyond the original word length.

For example, consider taking the negative of 11010 (-6). First, take the one's complement of the number, or flip the bits:

$$11010 \rightarrow 00101$$

Next, add a 1, wrapping all numbers to 0 or 1:

$$\begin{array}{r} 00101 \\ +1 \\ \hline 00110 \quad (6) \end{array}$$

Addition and Subtraction

The addition of fixed-point numbers requires that the binary points of the addends be aligned. The addition is then performed using binary arithmetic so that no number other than 0 or 1 is used.

For example, consider the addition of 010010.1 (18.5) with 0110.110 (6.75):

$$\begin{array}{r} 010010.1 \quad (18.5) \\ +0110.110 \quad (6.75) \\ \hline 011001.010 \quad (25.25) \end{array}$$

Fixed-point subtraction is equivalent to adding while using the two's complement value for any negative values. In subtraction, the addends must be sign-extended to match each other's length. For example, consider subtracting 0110.110 (6.75) from 010010.1 (18.5):

$$\begin{array}{r} 010010.100 \quad (18.5) \\ -0110.110 \quad (6.75) \\ \hline \end{array} \xrightarrow[\text{and sign extension}]{\text{two's complement}} \begin{array}{r} 010010.100 \quad (18.5) \\ +111001.010 \quad (-6.75) \\ \hline 1001011.110 \quad (11.75) \end{array}$$

Carry bit is discarded

The default `fimath` object has a value of 1 (true) for the `CastBeforeSum` property. This casts addends to the sum data type before addition. Therefore, no further shifting is necessary during the addition to line up the binary points.

If `CastBeforeSum` has a value of 0 (false), the addends are added with full precision maintained. After the addition the sum is then quantized.

Multiplication

The multiplication of two's complement fixed-point numbers is directly analogous to regular decimal multiplication, with the exception that the intermediate results must be sign-extended so that their left sides align before you add them together.

For example, consider the multiplication of 10.11 (-1.25) with 011 (3):

$$\begin{array}{r}
 10.11 \text{ (-1.25)} \\
 \quad 011 \text{ (3)} \\
 \hline
 11011 \\
 \quad 1011 \\
 \hline
 1100.01 \text{ (-3.75)}
 \end{array}$$

The extra 1 is the result of necessary sign extension.

The number of fractional bits of the result is the sum of the number of fractional bits of the factors.

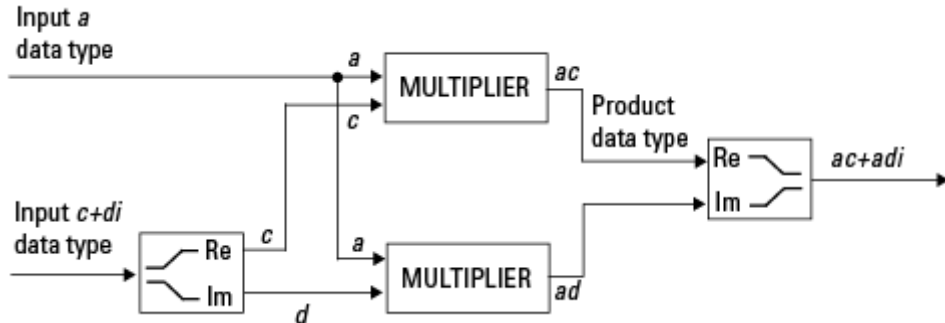
Multiplication Data Types

The following diagrams show the data types used for fixed-point multiplication using Fixed-Point Toolbox software. The diagrams illustrate the differences between the data types used for real-real, complex-real, and complex-complex multiplication.

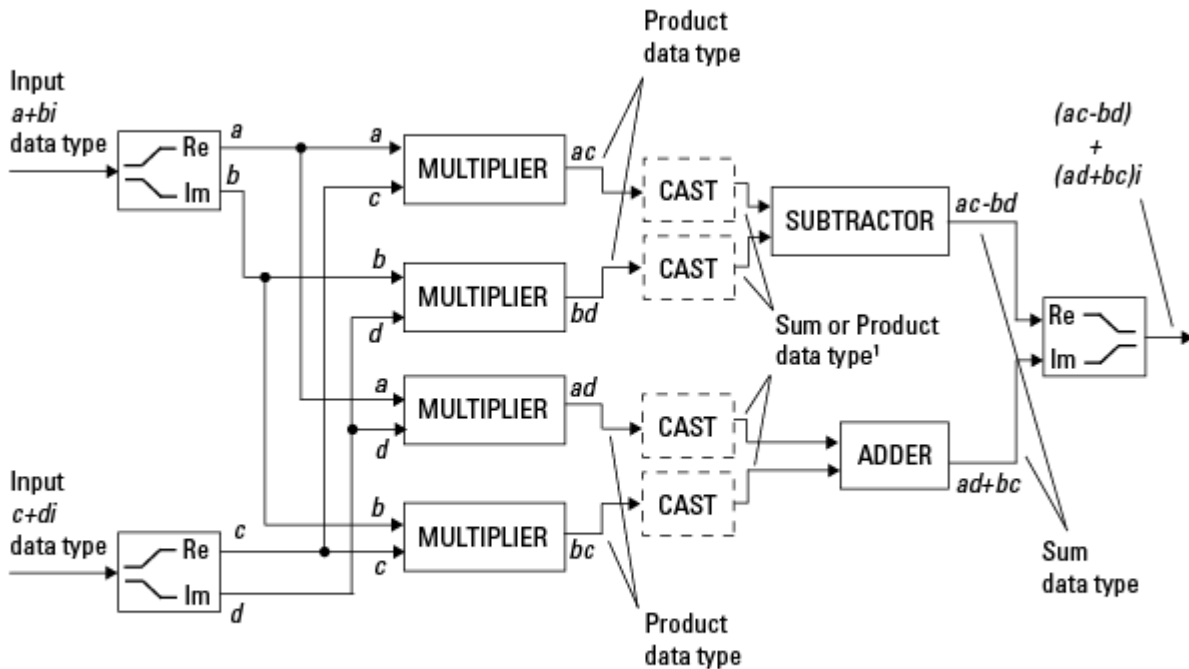
Real-Real Multiplication. The following diagram shows the data types used by the toolbox in the multiplication of two real numbers. The software returns the output of this operation in the product data type, which is governed by the fimath object ProductMode property.



Real-Complex Multiplication. The following diagram shows the data types used by the toolbox in the multiplication of a real and a complex fixed-point number. Real-complex and complex-real multiplication are equivalent. The software returns the output of this operation in the product data type, which is governed by the fimath object ProductMode property:



Complex-Complex Multiplication. The following diagram shows the multiplication of two complex fixed-point numbers. Note that the software returns the output of this operation in the sum data type, which is governed by the fimath object SumMode property. The intermediate product data type is determined by the fimath object ProductMode property.



¹ Sum data type if CastBeforeSum is true,
Product data type if CastBeforeSum is false

When the fimath object CastBeforeSum property is true, the casts to the sum data type are present after the multipliers in the preceding diagram. In C code, this is equivalent to

```
acc=ac;
acc-=bd;
```

for the subtractor, and

```
acc=ad;
acc+=bc;
```

for the adder, where *acc* is the accumulator. When the CastBeforeSum property is false, the casts are not present, and the data remains in the product data type before the subtraction and addition operations.

Multiplication with fimath

In the following examples, let

- `F = fimath('ProductMode','FullPrecision',...
 'SumMode','FullPrecision')`
- `T1 = numerictype('WordLength',24,'FractionLength',20)`
- `T2 = numerictype('WordLength',16,'FractionLength',10)`

Real*Real. Notice that the word length and fraction length of the result `z` are equal to the sum of the word lengths and fraction lengths, respectively, of the multiplicands. This is because the `fimath` `SumMode` and `ProductMode` properties are set to `FullPrecision`:

```
P = fipref;  
P.FimathDisplay = 'none';  
x = fi(5, T1, F)
```

```
x =
```

```
    5
```

```
                  DataTypeMode: Fixed-point: binary point scaling  
                  Signedness: Signed  
                  WordLength: 24  
                  FractionLength: 20
```

```
y = fi(10, T2, F)
```

```
y =
```

```
   10
```

```
                  DataTypeMode: Fixed-point: binary point scaling  
                  Signedness: Signed  
                  WordLength: 16  
                  FractionLength: 10
```

```
z = x*y
```

```
z =
```

```
50
```

```

        DataTypeMode: Fixed-point: binary point scaling
        Signedness: Signed
        WordLength: 40
        FractionLength: 30

```

Real*Complex. Notice that the word length and fraction length of the result z are equal to the sum of the word lengths and fraction lengths, respectively, of the multiplicands. This is because the `fimath` `SumMode` and `ProductMode` properties are set to `FullPrecision`:

```
x = fi(5,T1,F)
```

```
x =
```

```
5
```

```

        DataTypeMode: Fixed-point: binary point scaling
        Signedness: Signed
        WordLength: 24
        FractionLength: 20

```

```
y = fi(10+2i,T2,F)
```

```
y =
```

```
10.0000 + 2.0000i
```

```

        DataTypeMode: Fixed-point: binary point scaling
        Signedness: Signed
        WordLength: 16

```

FractionLength: 10

$z = x*y$

$z =$

$50.0000 + 10.0000i$

DataTypeMode: Fixed-point: binary point scaling
Signedness: Signed
WordLength: 40
FractionLength: 30

Complex*Complex. Complex-complex multiplication involves an addition as well as multiplication, so the word length of the full-precision result has one more bit than the sum of the word lengths of the multiplicands:

$x = \text{fi}(5+6i, T1, F)$

$x =$

$5.0000 + 6.0000i$

DataTypeMode: Fixed-point: binary point scaling
Signedness: Signed
WordLength: 24
FractionLength: 20

$y = \text{fi}(10+2i, T2, F)$

$y =$

$10.0000 + 2.0000i$

DataTypeMode: Fixed-point: binary point scaling
Signedness: Signed
WordLength: 16

```
FractionLength: 10

z = x*y

z =

38.0000 +70.0000i

DataTypeMode: Fixed-point: binary point scaling
Signedness: Signed
WordLength: 41
FractionLength: 30
```

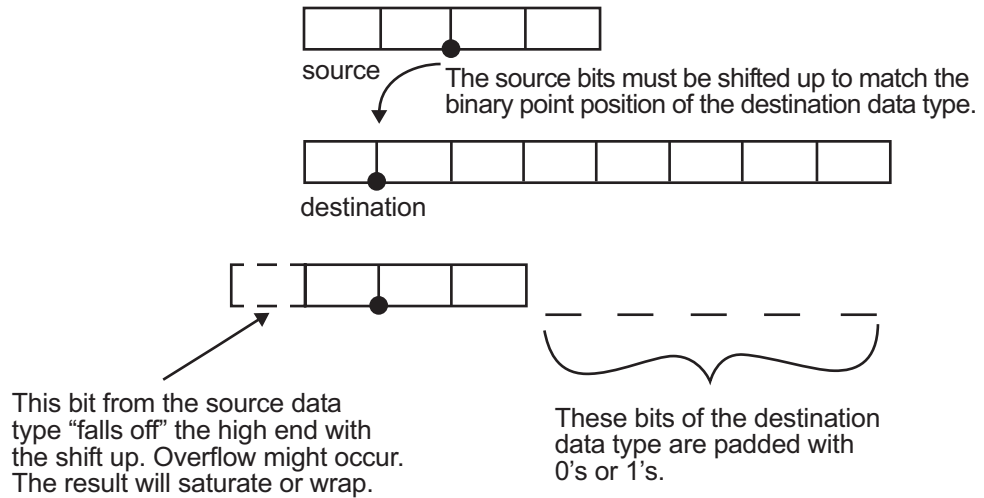
Casts

The `fimath` object allows you to specify the data type and scaling of intermediate sums and products with the `SumMode` and `ProductMode` properties. It is important to keep in mind the ramifications of each cast when you set the `SumMode` and `ProductMode` properties. Depending upon the data types you select, overflow and/or rounding might occur. The following two examples demonstrate cases where overflow and rounding can occur.

Note For more examples of casting, see “Casting fi Objects” on page 3-13.

Casting from a Shorter Data Type to a Longer Data Type

Consider the cast of a nonzero number, represented by a 4-bit data type with two fractional bits, to an 8-bit data type with seven fractional bits:



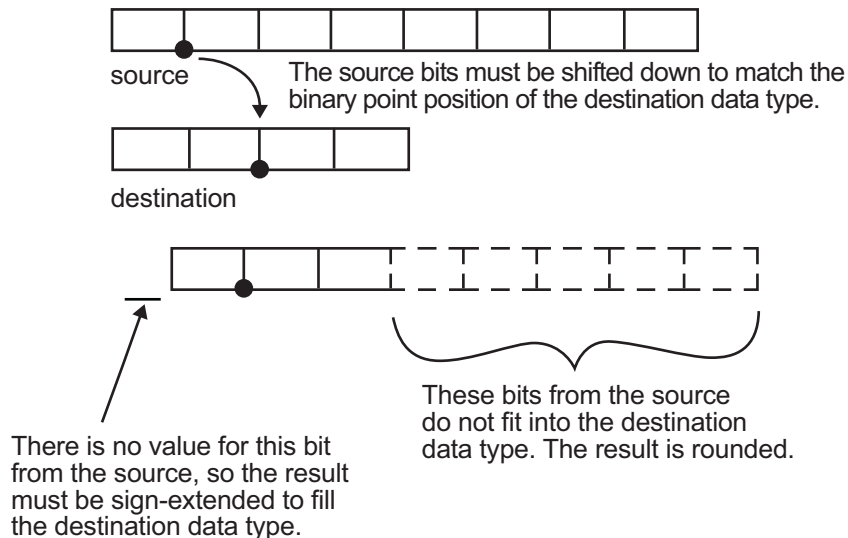
As the diagram shows, the source bits are shifted up so that the binary point matches the destination binary point position. The highest source bit does not fit, so overflow might occur and the result can saturate or wrap. The empty bits at the low end of the destination data type are padded with either 0's or 1's:

- If overflow does not occur, the empty bits are padded with 0's.
- If wrapping occurs, the empty bits are padded with 0's.
- If saturation occurs,
 - The empty bits of a positive number are padded with 1's.
 - The empty bits of a negative number are padded with 0's.

You can see that even with a cast from a shorter data type to a longer data type, overflow can still occur. This can happen when the integer length of the source data type (in this case two) is longer than the integer length of the destination data type (in this case one). Similarly, rounding might be necessary even when casting from a shorter data type to a longer data type, if the destination data type and scaling has fewer fractional bits than the source.

Casting from a Longer Data Type to a Shorter Data Type

Consider the cast of a nonzero number, represented by an 8-bit data type with seven fractional bits, to a 4-bit data type with two fractional bits:



As the diagram shows, the source bits are shifted down so that the binary point matches the destination binary point position. There is no value for the highest bit from the source, so the result is sign-extended to fill the integer portion of the destination data type. The bottom five bits of the source do not fit into the fraction length of the destination. Therefore, precision can be lost as the result is rounded.

In this case, even though the cast is from a longer data type to a shorter data type, all the integer bits are maintained. Conversely, full precision can be maintained even if you cast to a shorter data type, as long as the fraction length of the destination data type is the same length or longer than the fraction length of the source data type. In that case, however, bits are lost from the high end of the result and overflow can occur.

The worst case occurs when both the integer length and the fraction length of the destination data type are shorter than those of the source data type and scaling. In that case, both overflow and a loss of precision can occur.

fi Objects Compared to C Integer Data Types

In this section...
“Integer Data Types” on page 2-22
“Unary Conversions” on page 2-24
“Binary Conversions” on page 2-25
“Overflow Handling” on page 2-28

Note The sections in this topic compare the `fi` object with fixed-point data types and operations in C. In these sections, the information on ANSI C is adapted from Samuel P. Harbison and Guy L. Steele Jr., *C: A Reference Manual*, 3rd ed., Prentice Hall, 1991.

Integer Data Types

This section compares the numerical range of `fi` integer data types to the minimum numerical ranges of ANSI C integer data types.

ANSI C Integer Data Types

The following table shows the minimum ranges of ANSI C integer data types. The integer ranges can be larger than or equal to those shown, but cannot be smaller. The range of a `long` must be larger than or equal to the range of an `int`, which must be larger than or equal to the range of a `short`.

Note that the minimum ANSI C ranges are large enough to accommodate one's complement or sign/magnitude representation, but not two's complement representation. In the one's complement and sign/magnitude representations, a signed integer with n bits has a range from $-2^{n-1} + 1$ to $2^{n-1} - 1$, inclusive. In both of these representations, an equal number of positive and negative numbers are represented, and zero is represented twice.

Integer Type	Minimum	Maximum
signed char	-127	127
unsigned char	0	255
short int	-32,767	32,767
unsigned short	0	65,535
int	-32,767	32,767
unsigned int	0	65,535
long int	-2,147,483,647	2,147,483,647
unsigned long	0	4,294,967,295

fi Integer Data Types

The following table lists the numerical ranges of the integer data types of the `fi` object, in particular those equivalent to the C integer data types. The ranges are large enough to accommodate the two's complement representation, which is the only signed binary encoding technique supported by Fixed-Point Toolbox software. In the two's complement representation, a signed integer with n bits has a range from -2^{n-1} to $2^{n-1} - 1$, inclusive. An unsigned integer with n bits has a range from 0 to $2^n - 1$, inclusive. The negative side of the range has one more value than the positive side, and zero is represented uniquely.

Constructor	Signed	Word Length	Fraction Length	Minimum	Maximum	Closest ANSI C Equivalent
<code>fi(x,1,n,0)</code>	Yes	n (2 to 65,535)	0	-2^{n-1}	$2^{n-1} - 1$	N/A
<code>fi(x,0,n,0)</code>	No	n (2 to 65,535)	0	0	$2^n - 1$	N/A
<code>fi(x,1,8,0)</code>	Yes	8	0	-128	127	signed char
<code>fi(x,0,8,0)</code>	No	8	0	0	255	unsigned char

Constructor	Signed	Word Length	Fraction Length	Minimum	Maximum	Closest ANSI C Equivalent
<code>fi(x,1,16,0)</code>	Yes	16	0	-32,768	32,767	short int
<code>fi(x,0,16,0)</code>	No	16	0	0	65,535	unsigned short
<code>fi(x,1,32,0)</code>	Yes	32	0	-2,147,483,648	2,147,483,647	long int
<code>fi(x,0,32,0)</code>	No	32	0	0	4,294,967,295	unsigned long

Unary Conversions

Unary conversions dictate whether and how a single operand is converted before an operation is performed. This section discusses unary conversions in ANSI C and of `fi` objects.

ANSI C Usual Unary Conversions

Unary conversions in ANSI C are automatically applied to the operands of the unary `!`, `-`, `~`, and `*` operators, and of the binary `<<` and `>>` operators, according to the following table:

Original Operand Type	ANSI C Conversion
char or short	int
unsigned char or unsigned short	int or unsigned int ¹
float	float
Array of T	Pointer to T
Function returning T	Pointer to function returning T

¹If type `int` cannot represent all the values of the original data type without overflow, the converted type is `unsigned int`.

fi Usual Unary Conversions

The following table shows the fi unary conversions:

C Operator	fi Equivalent	fi Conversion
!x	$\sim x = \text{not}(x)$	Result is logical.
~x	bitcmp(x)	Result is same numeric type as operand.
*x	No equivalent	N/A
x<<n	bitshift(x,n) positive n	Result is same numeric type as operand. Round mode is always floor. Overflow mode is obeyed. 0-valued bits are shifted in on the right.
x>>n	bitshift(x,-n)	Result is same numeric type as operand. Round mode is always floor. Overflow mode is obeyed. 0-valued bits are shifted in on the left if the operand is unsigned or signed and positive. 1-valued bits are shifted in on the left if the operand is signed and negative.
+x	+x	Result is same numeric type as operand.
-x	-x	Result is same numeric type as operand. Overflow mode is obeyed. For example, overflow might occur when you negate an unsigned fi or the most negative value of a signed fi.

Binary Conversions

This section describes the conversions that occur when the operands of a binary operator are different data types.

ANSI C Usual Binary Conversions

In ANSI C, operands of a binary operator must be of the same type. If they are different, one is converted to the type of the other according to the first applicable conversion in the following table:

Type of One Operand	Type of Other Operand	ANSI C Conversion
long double	Any	long double

Type of One Operand	Type of Other Operand	ANSI C Conversion
double	Any	double
float	Any	float
unsigned long	Any	unsigned long
long	unsigned	long or unsigned long ¹
long	int	long
unsigned	int or unsigned	unsigned
int	int	int

¹Type long is only used if it can represent all values of type unsigned.

fi Usual Binary Conversions

When one of the operands of a binary operator (+, −, *, .*) is a `fi` object and the other is a MATLAB built-in numeric type, then the non-`fi` operand is converted to a `fi` object before the operation is performed, according to the following table:

Type of One Operand	Type of Other Operand	Properties of Other Operand After Conversion to a fi Object
<code>fi</code>	double or single	<ul style="list-style-type: none"> • Signed = same as the original <code>fi</code> operand • WordLength = same as the original <code>fi</code> operand • FractionLength = set to best precision possible
<code>fi</code>	<code>int8</code>	<ul style="list-style-type: none"> • Signed = 1 • WordLength = 8 • FractionLength = 0

Type of One Operand	Type of Other Operand	Properties of Other Operand After Conversion to a fi Object
fi	uint8	<ul style="list-style-type: none"> • Signed = 0 • WordLength = 8 • FractionLength = 0
fi	int16	<ul style="list-style-type: none"> • Signed = 1 • WordLength = 16 • FractionLength = 0
fi	uint16	<ul style="list-style-type: none"> • Signed = 0 • WordLength = 16 • FractionLength = 0
fi	int32	<ul style="list-style-type: none"> • Signed = 1 • WordLength = 32 • FractionLength = 0
fi	uint32	<ul style="list-style-type: none"> • Signed = 0 • WordLength = 32 • FractionLength = 0
fi	int64	<ul style="list-style-type: none"> • Signed = 1 • WordLength = 64 • FractionLength = 0
fi	uint64	<ul style="list-style-type: none"> • Signed = 0 • WordLength = 64 • FractionLength = 0

Overflow Handling

The following sections compare how ANSI C and Fixed-Point Toolbox software handle overflows.

ANSI C Overflow Handling

In ANSI C, the result of signed integer operations is whatever value is produced by the machine instruction used to implement the operation. Therefore, ANSI C has no rules for handling signed integer overflow.

The results of unsigned integer overflows wrap in ANSI C.

fi Overflow Handling

Addition and multiplication with `fi` objects yield results that can be exactly represented by a `fi` object, up to word lengths of 65,535 bits or the available memory on your machine. This is not true of division, however, because many ratios result in infinite binary expressions. You can perform division with `fi` objects using the `divide` function, which requires you to explicitly specify the numeric type of the result.

The conditions under which a `fi` object overflows and the results then produced are determined by the associated `fimath` object. You can specify certain overflow characteristics separately for sums (including differences) and products. Refer to the following table:

fimath Object Properties Related to Overflow Handling	Property Value	Description
OverflowMode	'saturate'	Overflows are saturated to the maximum or minimum value in the range.
	'wrap'	Overflows wrap using modulo arithmetic if unsigned, two's complement wrap if signed.

fimath Object Properties Related to Overflow Handling	Property Value	Description
ProductMode	'FullPrecision'	<p>Full-precision results are kept. Overflow does not occur. An error is thrown if the resulting word length is greater than MaxProductWordLength.</p> <p>The rules for computing the resulting product word and fraction lengths are given in “ProductMode” in the Property Reference.</p>
	'KeepLSB'	<p>The least significant bits of the product are kept. Full precision is kept, but overflow is possible. This behavior models the C language integer operations.</p> <p>The resulting word length is determined by the ProductWordLength property. If ProductWordLength is greater than is necessary for the full-precision product, then the result is stored in the least significant bits. If ProductWordLength is less than is necessary for the full-precision product, then overflow occurs.</p> <p>The rule for computing the resulting product fraction length is given in “ProductMode” in the Property Reference.</p>

fimath Object Properties Related to Overflow Handling	Property Value	Description
	'KeepMSB'	<p>The most significant bits of the product are kept. Overflow is prevented, but precision may be lost.</p> <p>The resulting word length is determined by the ProductWordLength property. If ProductWordLength is greater than is necessary for the full-precision product, then the result is stored in the most significant bits. If ProductWordLength is less than is necessary for the full-precision product, then rounding occurs.</p> <p>The rule for computing the resulting product fraction length is given in “ProductMode” in the Property Reference.</p>
	'SpecifyPrecision'	<p>You can specify both the word length and the fraction length of the resulting product.</p>
ProductWordLength	Positive integer	<p>The word length of product results when ProductMode is 'KeepLSB', 'KeepMSB', or 'SpecifyPrecision'.</p>
MaxProductWordLength	Positive integer	<p>The maximum product word length allowed when ProductMode is 'FullPrecision'. The default is 128 bits. The maximum is 65,535 bits. This property can help ensure that your simulation does not exceed your hardware requirements.</p>
ProductFractionLength	Integer	<p>The fraction length of product results when ProductMode is 'Specify Precision'.</p>

fimath Object Properties Related to Overflow Handling	Property Value	Description
SumMode	'FullPrecision'	<p>Full-precision results are kept. Overflow does not occur. An error is thrown if the resulting word length is greater than MaxSumWordLength.</p> <p>The rules for computing the resulting sum word and fraction lengths are given in “SumMode” in the Property Reference.</p>
	'KeepLSB'	<p>The least significant bits of the sum are kept. Full precision is kept, but overflow is possible. This behavior models the C language integer operations.</p> <p>The resulting word length is determined by the SumWordLength property. If SumWordLength is greater than is necessary for the full-precision sum, then the result is stored in the least significant bits. If SumWordLength is less than is necessary for the full-precision sum, then overflow occurs.</p> <p>The rule for computing the resulting sum fraction length is given in “SumMode” in the Property Reference.</p>
	'KeepMSB'	<p>The most significant bits of the sum are kept. Overflow is prevented, but precision may be lost.</p> <p>The resulting word length is determined by the SumWordLength property. If SumWordLength is greater than is necessary for the full-precision sum, then the result is stored in the most significant bits. If SumWordLength is less than is necessary for the full-precision sum, then rounding occurs.</p>

fimath Object Properties Related to Overflow Handling	Property Value	Description
		The rule for computing the resulting sum fraction length is given in “SumMode” in the Property Reference.
	'SpecifyPrecision'	You can specify both the word length and the fraction length of the resulting sum.
SumWordLength	Positive integer	The word length of sum results when SumMode is 'KeepLSB', 'KeepMSB', or 'SpecifyPrecision'.
MaxSumWordLength	Positive integer	The maximum sum word length allowed when SumMode is 'FullPrecision'. The default is 128 bits. The maximum is 65,535 bits. This property can help ensure that your simulation does not exceed your hardware requirements.
SumFractionLength	Integer	The fraction length of sum results when SumMode is 'SpecifyPrecision'.

Working with fi Objects

- “Constructing fi Objects” on page 3-2
- “Casting fi Objects” on page 3-13
- “fi Object Properties” on page 3-18
- “fi Object Functions” on page 3-23

Constructing fi Objects

In this section...
“fi Object Syntaxes” on page 3-2
“Examples of Constructing fi Objects” on page 3-3

fi Object Syntaxes

You can create `fi` objects using Fixed-Point Toolbox software in one of two ways:

- You can use the `fi` constructor function to create a new object.
- You can use the `fi` constructor function to copy an existing `fi` object.

To get started, type

```
a = fi(0)
```

to create a `fi` object with the default data type and a value of 0.

```
a =
```

```
0
```

```
      DataTypeMode: Fixed-point: binary point scaling  
      Signedness: Signed  
      WordLength: 16  
      FractionLength: 15
```

A signed `fi` object is created with a value of 0, word length of 16 bits, and fraction length of 15 bits.

To see all of the `fi` object syntaxes, refer to the `fi` constructor function reference page.

Note For information on the display format of fi objects, refer to “Display Settings” on page 1-7.

Examples of Constructing fi Objects

The following examples show you several different ways to construct fi objects. For other, more basic examples of constructing fi objects, see the “Examples” on the fi constructor function reference page.

Constructing a fi Object with Property Name/Property Value Pairs

You can use property name/property value pairs to set fi properties when you create the object:

```
a = fi(pi, 'roundmode', 'floor', 'overflowmode', 'wrap')
```

```
a =
```

```
3.1415
```

```
      DataTypeMode: Fixed-point: binary point scaling  
      Signedness: Signed  
      WordLength: 16  
      FractionLength: 13
```

```
      RoundMode: floor  
      OverflowMode: wrap  
      ProductMode: FullPrecision  
      MaxProductWordLength: 128  
      SumMode: FullPrecision  
      MaxSumWordLength: 128  
      CastBeforeSum: true
```

Constructing a fi Object Using a numerictype Object

You can use a numerictype object to define a fi object:

```
T = numerictype
```

```
T =
```

```
        DataTypeMode: Fixed-point: binary point scaling  
        Signedness: Signed  
        WordLength: 16  
        FractionLength: 15
```

```
a = fi(pi, T)
```

```
a =
```

```
1.0000
```

```
        DataTypeMode: Fixed-point: binary point scaling  
        Signedness: Signed  
        WordLength: 16  
        FractionLength: 15
```

```
        RoundMode: nearest  
        OverflowMode: saturate  
        ProductMode: FullPrecision  
MaxProductWordLength: 128  
        SumMode: FullPrecision  
MaxSumWordLength: 128  
        CastBeforeSum: true
```

You can also use a `fimath` object with a `numericType` object to define a `fi` object:

```
F = fimath
```

```
F =
```

```
        RoundMode: nearest  
        OverflowMode: saturate  
        ProductMode: FullPrecision  
MaxProductWordLength: 128
```



```
                SumMode: FullPrecision
            MaxSumWordLength: 128
            CastBeforeSum: true

a = fi(pi, T, F)

a =

    1.0000

                DataTypeMode: Fixed-point: binary point scaling
                Signedness: Signed
                WordLength: 16
            FractionLength: 15

                RoundMode: nearest
            OverflowMode: saturate
                ProductMode: FullPrecision
    MaxProductWordLength: 128
                SumMode: FullPrecision
    MaxSumWordLength: 128
            CastBeforeSum: true
```

Note The syntax `a = fi(pi,T,F)` is equivalent to `a = fi(pi,F,T)`. You can use both statements to define a `fi` object using a `fimath` object and a `numericType` object.

Constructing a fi Object Using a fimath Object

You can create a `fi` object using a specific `fimath` object. If you do not specify any `numericType` object properties, the word length of the `fi` object defaults to 16 bits and the fraction length is determined by best precision scaling:

```
F = fimath

F =
```

```
        RoundMode: nearest
        OverflowMode: saturate
        ProductMode: FullPrecision
MaxProductWordLength: 128
        SumMode: FullPrecision
MaxSumWordLength: 128
        CastBeforeSum: true
```

```
F.OverflowMode = 'wrap'
```

```
F =
```

```
        RoundMode: nearest
        OverflowMode: wrap
        ProductMode: FullPrecision
MaxProductWordLength: 128
        SumMode: FullPrecision
MaxSumWordLength: 128
        CastBeforeSum: true
```

```
a = fi(pi, F)
```

```
a =
```

```
3.1416
```

```
        DataTypeMode: Fixed-point: binary point scaling
        Signedness: Signed
        WordLength: 16
FractionLength: 13
```

```
        RoundMode: nearest
        OverflowMode: wrap
        ProductMode: FullPrecision
MaxProductWordLength: 128
        SumMode: FullPrecision
MaxSumWordLength: 128
        CastBeforeSum: true
```

You can also create fi objects using a fimath object while specifying various numericity properties at creation time:

```
b = fi(pi, 0, F)
```

```
b =
```

```
3.1416
```

```
      DataTypeMode: Fixed-point: binary point scaling  
      Signedness: Unsigned  
      WordLength: 16  
      FractionLength: 14
```

```
      RoundMode: nearest  
      OverflowMode: wrap  
      ProductMode: FullPrecision  
MaxProductWordLength: 128  
      SumMode: FullPrecision  
MaxSumWordLength: 128  
      CastBeforeSum: true
```

```
c = fi(pi, 0, 8, F)
```

```
c =
```

```
3.1406
```

```
      DataTypeMode: Fixed-point: binary point scaling  
      Signedness: Unsigned  
      WordLength: 8  
      FractionLength: 6
```

```
      RoundMode: nearest  
      OverflowMode: wrap  
      ProductMode: FullPrecision  
MaxProductWordLength: 128  
      SumMode: FullPrecision  
MaxSumWordLength: 128  
      CastBeforeSum: true
```

```
d = fi(pi, 0, 8, 6, F)

d =

    3.1406

        DataTypeMode: Fixed-point: binary point scaling
          Signedness: Unsigned
           WordLength: 8
    FractionLength: 6

           RoundMode: nearest
      OverflowMode: wrap
        ProductMode: FullPrecision
MaxProductWordLength: 128
           SumMode: FullPrecision
    MaxSumWordLength: 128
    CastBeforeSum: true
```

Building fi Object Constructors in a GUI

When you are working with M-files, you can build your `fi` object constructors using the **Insert fi Constructor** dialog box. After specifying the value and properties of the `fi` object in the dialog box, you can insert the prepopulated `fi` object constructor string at a specific location in your M-file.

For example, to create a signed `fi` object with a value of `pi`, a word length of 16 bits and a fraction length of 13 bits, perform the following steps:

- 1** Open the **Insert fi Constructor** dialog box by selecting **Tools > Fixed-Point Toolbox > Insert fi Constructor** from the editor menu.
- 2** Use the edit boxes and drop-down menus to specify the following properties of the `fi` object:
 - **Value** = `pi`
 - **Data type mode** = Fixed-point: binary point scaling
 - **Signedness** = Signed

- **Word length** = 16
- **Fraction length** = 13

Insert fi Constructor

Value:

For example, pi, 2.718, or x/2

numerictype

Data type mode:

Signedness:

Word length:

Fraction length:

fimath

Round mode:

Overflow mode:

Product mode:

Maximum product word length:

Sum mode:

Maximum sum word length:

Cast before sum

OK Cancel Help

- 3** To insert the `fi` object constructor string into your M-file, place your cursor at the desired location in the M-file, and click **OK** on the **Insert fi Constructor** dialog box. Clicking **OK** closes the **Insert fi Constructor** dialog box and automatically populates the `fi` object constructor string in your M-file:

```

8      a = fi(pi, 1, 16, 13, ...
9          'RoundMode', 'Nearest', ...
10         'OverflowMode', 'Saturate', ...
11         'ProductMode', 'FullPrecision', ...
12         'MaxProductWordLength', 128, ...
13         'SumMode', 'FullPrecision', ...
14         'MaxSumWordLength', 128, ...
15         'CastBeforeSum', true)

```

Determining Property Precedence

The value of a property is taken from the last time it is set. For example, create a `numerictype` object with a value of `true` for the `Signed` property and a fraction length of 14:

```
T = numerictype('Signed', true, 'FractionLength', 14)
```

```
T =
```

```

      DataTypeMode: Fixed-point: binary point scaling
      Signedness: Signed
      WordLength: 16
      FractionLength: 14

```

Now, create the following `fi` object in which the `numerictype` property is specified *after* the `Signed` property, so that the resulting `fi` object is signed:

```
a = fi(pi, 'Signed', false, 'numerictype', T)
```

```
a =
```

```
1.9999
```

```

      DataTypeMode: Fixed-point: binary point scaling
      Signedness: Signed
      WordLength: 16
      FractionLength: 14
      RoundMode: nearest

```

```

        OverflowMode: saturate
        ProductMode: FullPrecision
MaxProductWordLength: 128
        SumMode: FullPrecision
MaxSumWordLength: 128
        CastBeforeSum: true

```

Contrast the `fi` object in this code sample with the `fi` object in the following code sample. The `numerictype` property in the following code sample is specified *before* the `Signed` property, so the resulting `fi` object is unsigned:

```
b = fi(pi, 'numerictype', T, 'signed', false)
```

```
b =
```

```
3.1416
```

```

        DataTypeMode: Fixed-point: binary point scaling
        Signedness: Unsigned
        WordLength: 16
        FractionLength: 14

```

```

        RoundMode: nearest
        OverflowMode: saturate
        ProductMode: FullPrecision
MaxProductWordLength: 128
        SumMode: FullPrecision
MaxSumWordLength: 128
        CastBeforeSum: true

```

Copying a fi Object

To copy a `fi` object, simply use assignment, as in the following example:

```
a = fi(pi)
```

```
a =
```

```
3.1416
```

```
        DataTypeMode: Fixed-point: binary point scaling
```

```
Signedness: Signed  
WordLength: 16  
FractionLength: 13
```

```
b = a
```

```
b =
```

```
3.1416
```

```
DataTypeMode: Fixed-point: binary point scaling  
Signedness: Signed  
WordLength: 16  
FractionLength: 13
```


Casting fi Objects

In this section...

“Overwriting by Assignment” on page 3-13

“Ways to Cast with MATLAB Software” on page 3-13

Overwriting by Assignment

Because MATLAB software does not have type declarations, an assignment like `A = B` replaces the type and content of `A` with the type and content of `B`. If `A` does not exist at the time of the assignment, MATLAB creates the variable `A` and assigns it the same type and value as `B`. Such assignment happens with all types in MATLAB—objects and built-in types alike—including `fi`, `double`, `single`, `int8`, `uint8`, `int16`, etc.

For example, the following code overwrites the value and `int8` type of `A` with the value and `int16` type of `B`:

```
A = int8(0);
B = int16(32767);
A = B

A =

    32767

class(A)

ans =

    int16
```

Ways to Cast with MATLAB Software

You may find it useful to cast data into another type—for example, when you are casting data from an accumulator to memory. There are several ways to cast data in MATLAB. The following sections provide examples of three different methods:

- Casting by Subscripted Assignment
- Casting by Conversion Function
- Casting with the Fixed-Point Toolbox `reinterprecast` Function

Casting by Subscripted Assignment

The following subscripted assignment statement retains the type of `A` and saturates the value of `B` to an `int8`:

```
A = int8(0);  
B = int16(32767);  
A(:) = B
```

```
A =
```

```
    127
```

```
class(A)
```

```
ans =
```

```
int8
```

The same is true for `fi` objects:

```
fipref('NumericTypeDisplay', 'short', 'FimathDisplay', 'none');  
A = fi(0, true, 8, 0);  
B = fi(32767, true, 16, 0);  
A(:) = B
```

```
A =
```

```
    127
```

```
    s8,0
```

Note For more information on subscripted assignments, see the `subsasgn` function.

Casting by Conversion Function

You can convert from one data type to another by using a conversion function. In this example, A does not have to be predefined because it is overwritten.

```
B = int16(32767);
A = int8(B)
```

```
A =
```

```
    127
```

```
class(A)
```

```
ans =
```

```
int8
```

The same is true for fi objects:

```
B = fi(32767, true, 16, 0)
A = fi(B, 1, 8, 0)
```

```
B =
```

```
    32767
s16,0
```

```
A =
```

```
    127
s8,0
```

Using a numeric type Object in the fi Conversion Function. Often a specific numeric type is used in many places, and it is convenient to predefine numeric type objects for use in the conversion functions. Predefining these objects is a good practice because it also puts the data type specification in one place.

```
T8 = numeric(1,8,0)
```

```
T8 =
```

```
        DataTypeMode: Fixed-point: binary point scaling
        Signedness: Signed
        WordLength: 8
        FractionLength: 0
```

```
T16 = numericitytype(1,16,0)
```

```
T16 =
```

```
        DataTypeMode: Fixed-point: binary point scaling
        Signedness: Signed
        WordLength: 16
        FractionLength: 0
```

```
B = fi(32767,T16)
```

```
B =
```

```
    32767
    s16,0
```

```
A = fi(B, T8)
```

```
A =
```

```
    127
    s8,0
```

Casting with the reinterpretcast Function

You can convert fixed-point data types without changing the underlying data. The Fixed-Point Toolbox `reinterpretcast` function performs this type of conversion.

In the following example, `B` is an unsigned `fi` object with a word length of 8 bits and a fraction length of 5 bits. The `reinterpretcast` function converts `B` into a signed `fi` object `A` with a word length of 8 bits and a fraction length of 1

bit. The real-world values of A and B differ, but their binary representations are the same.

```
B = fi([pi/4 1 pi/2 4], false, 8, 5)
T = numerictype(true, 8, 1);
A = reinterpretcast(B, T)
```

B =

```
0.7813    1.0000    1.5625    4.0000
```

```
      DataTypeMode: Fixed-point: binary point scaling
      Signedness: Unsigned
      WordLength: 8
      FractionLength: 5
```

A =

```
12.5000   16.0000   25.0000  -64.0000
```

```
      DataTypeMode: Fixed-point: binary point scaling
      Signedness: Signed
      WordLength: 8
      FractionLength: 1
```

To verify that the underlying data has not changed, compare the binary representations of A and B:

```
binary_B = bin(B)
binary_A = bin(A)
```

binary_A =

```
00011001  00100000  00110010  10000000
```

binary_B =

```
00011001  00100000  00110010  10000000
```

fi Object Properties

In this section...
“Data Properties” on page 3-18
“fimath Properties” on page 3-18
“numericity Properties” on page 3-20
“Setting <code>fi</code> Object Properties” on page 3-21

Data Properties

The data properties of a `fi` object are always writable.

- `bin` — Stored integer value of a `fi` object in binary
- `data` — Numerical real-world value of a `fi` object
- `dec` — Stored integer value of a `fi` object in decimal
- `double` — Real-world value of a `fi` object, stored as a MATLAB `double` data type
- `hex` — Stored integer value of a `fi` object in hexadecimal
- `int` — Stored integer value of a `fi` object, stored in a built-in MATLAB integer data type. You can also use `int8`, `int16`, `int32`, `int64`, `uint8`, `uint16`, `uint32`, and `uint64` to get the stored integer value of a `fi` object in these formats
- `oct` — Stored integer value of a `fi` object in octal

To learn more about these properties, see “`fi` Object Properties” in the Fixed-Point Toolbox Reference.

fimath Properties

When you create a `fi` object, a `fimath` object is also automatically created as a property of the `fi` object:

- `fimath` — `fimath` object associated with a `fi` object

The following `fimath` properties are, by transitivity, also properties of a `fi` object. The properties of the `fimath` object listed below are always writable:

- `CastBeforeSum` — Whether both operands are cast to the sum data type before addition
- `MaxProductWordLength` — Maximum allowable word length for the product data type
- `MaxSumWordLength` — Maximum allowable word length for the sum data type
- `OverflowMode` — Overflow mode
- `ProductBias` — Bias of the product data type
- `ProductFixedExponent` — Fixed exponent of the product data type
- `ProductFractionLength` — Fraction length, in bits, of the product data type
- `ProductMode` — Defines how the product data type is determined
- `ProductSlope` — Slope of the product data type
- `ProductSlopeAdjustmentFactor` — Slope adjustment factor of the product data type
- `ProductWordLength` — Word length, in bits, of the product data type
- `RoundMode` — Rounding mode
- `SumBias` — Bias of the sum data type
- `SumFixedExponent` — Fixed exponent of the sum data type
- `SumFractionLength` — Fraction length, in bits, of the sum data type
- `SumMode` — Defines how the sum data type is determined
- `SumSlope` — Slope of the sum data type
- `SumSlopeAdjustmentFactor` — Slope adjustment factor of the sum data type
- `SumWordLength` — The word length, in bits, of the sum data type

To learn more about these properties, see the “`fimath` Object Properties” in the Fixed-Point Toolbox Reference.

numericType Properties

When you create a `fi` object, a `numericType` object is also automatically created as a property of the `fi` object:

`numericType` — Object containing all the data type information of a `fi` object, Simulink signal or model parameter

The following `numericType` properties are, by transitivity, also properties of a `fi` object. The following properties of the `numericType` object become read only after you create the `fi` object. However, you can create a copy of a `fi` object with new values specified for the `numericType` properties:

- `Bias` — Bias of a `fi` object
- `DataType` — Data type category associated with a `fi` object
- `DataTypeMode` — Data type and scaling mode of a `fi` object
- `FixedExponent` — Fixed-point exponent associated with a `fi` object
- `FractionLength` — Fraction length of the stored integer value of a `fi` object in bits
- `Scaling` — Fixed-point scaling mode of a `fi` object
- `Signed` — Whether a `fi` object is signed or unsigned
- `Signedness` — Whether a `fi` object is signed or unsigned

Note `numericType` objects can have a `Signedness` of `Auto`, but all `fi` objects must be `Signed` or `Unsigned`. If a `numericType` object with `Auto Signedness` is used to create a `fi` object, the `Signedness` property of the `fi` object automatically defaults to `Signed`.

- `Slope` — Slope associated with a `fi` object
- `SlopeAdjustmentFactor` — Slope adjustment associated with a `fi` object
- `WordLength` — Word length of the stored integer value of a `fi` object in bits

For further details on these properties, see the [Property Reference](#).

There are two ways to specify properties for `fi` objects in Fixed-Point Toolbox software. Refer to the following sections:

- “Setting Fixed-Point Properties at Object Creation” on page 3-21
- “Using Direct Property Referencing with `fi`” on page 3-22

Setting `fi` Object Properties

You can set `fi` object properties in two ways:

- Setting the properties when you create the object
- Using direct property referencing

Setting Fixed-Point Properties at Object Creation

You can set properties of `fi` objects at the time of object creation by including properties after the arguments of the `fi` constructor function. For example, to set the overflow mode to `wrap` and the rounding mode to `convergent`,

```
a = fi(pi, 'OverflowMode', 'wrap', 'RoundMode', 'convergent')
```

```
a =
```

```
3.1416
```

```
      DataTypeMode: Fixed-point: binary point scaling  
      Signedness: Signed  
      WordLength: 16  
      FractionLength: 13
```

```
      RoundMode: convergent  
      OverflowMode: wrap  
      ProductMode: FullPrecision  
MaxProductWordLength: 128  
      SumMode: FullPrecision  
MaxSumWordLength: 128  
      CastBeforeSum: true
```

Using Direct Property Referencing with fi

You can reference directly into a property for setting or retrieving `fi` object property values using MATLAB structure-like referencing. You do so by using a period to index into a property by name.

For example, to get the `DataTypeMode` of `a`,

```
a.DataTypeMode  
  
ans =  
  
Fixed-point: binary point scaling
```

To set the `OverflowMode` of `a`,

```
a.OverflowMode = 'wrap'  
  
a =  
  
3.1416
```

```
DataTypeMode: Fixed-point: binary point scaling  
Signedness: Signed  
WordLength: 16  
FractionLength: 13
```

```
RoundMode: convergent  
OverflowMode: wrap  
ProductMode: FullPrecision  
MaxProductWordLength: 128  
SumMode: FullPrecision  
MaxSumWordLength: 128  
CastBeforeSum: true
```

fi Object Functions

You can learn about the functions associated with `fi` objects in the [Function Reference](#).

The following data-access functions can be also used to get the data in a `fi` object using dot notation.

- `bin`
- `data`
- `dec`
- `double`
- `hex`
- `int`
- `oct`

For example,

```
a = fi(pi);  
n = int(a)
```

```
n =
```

```
25736
```

```
a.int
```

```
ans =
```

```
25736
```

```
h = hex(a)
```

```
h =
```

```
6488
```

a.hex

ans =

6488

Working with fimath Objects

- “Constructing fimath Objects” on page 4-2
- “fimath Object Properties” on page 4-10
- “Using fimath Objects to Perform Fixed-Point Arithmetic” on page 4-15
- “Using fimath to Specify Rounding and Overflow Modes” on page 4-23
- “Using fimath to Share Arithmetic Rules” on page 4-24
- “Using fimath ProductMode and SumMode” on page 4-26
- “fimath Object Functions” on page 4-32

Constructing fimath Objects

fimath Object Syntaxes

fimath objects define the arithmetic attributes of fi objects. You can create fimath objects in Fixed-Point Toolbox software in one of two ways:

- You can use the `fimath` constructor function to create a new object.
- You can use the `fimath` constructor function to copy an existing fimath object.

To get started, type

```
F = fimath
```

to create a default. If your default fimath object is the MATLAB factory default, you get the following output:

```
F =  
  
      RoundMode: nearest  
      OverflowMode: saturate  
      ProductMode: FullPrecision  
MaxProductWordLength: 128  
      SumMode: FullPrecision  
MaxSumWordLength: 128  
      CastBeforeSum: true
```

To learn how to configure your own MATLAB default fimath object, see “Configuring the MATLAB Default fimath Object” on page 4-5.

To copy a fimath object, simply use assignment as in the following example:

```
F = fimath;  
G = F;  
isequal(F,G)  
  
ans =
```

1

The syntax

```
F = fimath(...'PropertyName',PropertyValue...)
```


allows you to set properties for a `fimath` object at object creation with property name/property value pairs. Refer to “Setting `fimath` Properties at Object Creation” on page 4-11.

Building `fimath` Object Constructors in a GUI

When you are working with M-files, you can build your `fimath` object constructors using the **Insert `fimath` Constructor** dialog box. After specifying the properties of the `fimath` object in the dialog box, you can insert the prepopulated `fimath` object constructor string at a specific location in your M-file.

For example, to create a `fimath` object that uses convergent rounding and wraps on overflow, perform the following steps:

- 1 Open the **Insert `fimath` Constructor** dialog box by selecting **Tools > Fixed-Point Toolbox > Insert `fimath` Constructor** from the editor menu.
- 2 Use the edit boxes and drop-down menus to specify the following properties of the `fimath` object:
 - **Round mode** = Floor
 - **Overflow mode** = Wrap
 - **Product mode** = FullPrecision
 - **Maximum product word length** = 128
 - **Sum mode** = FullPrecision
 - **Maximum sum word length** = 128
 - **Cast before sum** = Checked



- 3** To insert the `fimath` object constructor string into your M-file, place your cursor at the desired location in the M-file, and click **OK** on the **Insert fimath Constructor** dialog box. Clicking **OK** closes the **Insert fimath Constructor** dialog box and automatically populates the `fimath` object constructor string in your M-file:

```

6      F = fimath('RoundMode', 'Floor', ...
7              'OverflowMode', 'Wrap', ...
8              'ProductMode', 'FullPrecision', ...
9              'MaxProductWordLength', 128, ...
10             'SumMode', 'FullPrecision', ...
11             'MaxSumWordLength', 128, ...
12             'CastBeforeSum', true)

```

Configuring the MATLAB Default `fimath` Object

The MATLAB default `fimath` object is the object that `fi` and `fimath` object constructors return when no other `fimath` object is specified in the constructor. When you use property name/property value pairs to partially specify a `fimath` object in a `fi` or `fimath` constructor, the remaining unspecified

properties of the `fimath` object are assigned according to the MATLAB default. The MATLAB factory default `fimath` object has the following properties:

```
RoundMode: nearest
OverflowMode: saturate
ProductMode: FullPrecision
MaxProductWordLength: 128
SumMode: FullPrecision
MaxSumWordLength: 128
CastBeforeSum: true
```

You can configure the default `fimath` object in MATLAB using the Fixed-Point Toolbox `setdefaultfimath`, `savedefaultfimathpref`, and `resetdefaultfimath` functions. The configurable default `fimath` object allows you to define your own `fimath` object to be used as the default in all `fi` and `fimath` constructors. The following sections show you how to:

- Create a `fimath` object
- Set your `fimath` object as the MATLAB default for your current MATLAB session
- Save your `fimath` object as the MATLAB default for all future MATLAB sessions
- Reset the MATLAB default `fimath` object to the MATLAB factory default for your current MATLAB session
- Save the MATLAB factory default `fimath` object as your MATLAB default for all future MATLAB sessions

Setting a User-Configured Default `fimath` Object

To create a `fimath` object that you can later use as the MATLAB default, pick a variable name and use the property name/property value pairs syntax to create a `fimath` object with the properties you desire. For example:

```
Fdefault = fimath('RoundMode','Floor','OverflowMode','Wrap');
```

After you define your `fimath` object, use the `setdefaultfimath` command to set your MATLAB default `fimath` object to be the `fimath` object you just created:

```
setdefaultfimath(Fdefault);
```

From this point forward (in your current MATLAB session), `Fdefault` appears as your MATLAB default `fimath` object in both `fi` and `fimath` constructors:

```
a = fi(pi)
```

```
a =
```

```
3.1415
```

```
      DataTypeMode: Fixed-point: binary point scaling  
      Signedness: Signed  
      WordLength: 16  
      FractionLength: 13
```

```
      RoundMode: floor  
      OverflowMode: wrap  
      ProductMode: FullPrecision  
MaxProductWordLength: 128  
      SumMode: FullPrecision  
MaxSumWordLength: 128  
      CastBeforeSum: true
```

```
F1 = fimath
```

```
F1 =
```

```
      RoundMode: floor  
      OverflowMode: wrap  
      ProductMode: FullPrecision  
MaxProductWordLength: 128  
      SumMode: FullPrecision  
MaxSumWordLength: 128  
      CastBeforeSum: true
```

Saving the Default fimath Object

If you want to use the current default `fimath` object as your MATLAB default `fimath` object in future MATLAB sessions, you must save it as a preference. The `savedefaultfimathpref` function saves the current MATLAB default `fimath` object as a preference to be used in all future MATLAB sessions.

To use the current default `fimath` object as the MATLAB default in all future MATLAB sessions, type

```
savedefaultfimathpref;
```

at the MATLAB command line. MATLAB will now use the current default `fimath` object as the default in all of your future MATLAB sessions.

Resetting the Default fimath Object to the Factory Default

You can reset the MATLAB default `fimath` object to the factory setting at any time by using the `resetdefaultfimath` function. To do so, type

```
resetdefaultfimath;
```

at the MATLAB command line.

After you reset the default `fimath` object, the `fi` and `fimath` constructors use the MATLAB factory default `fimath` object whenever you do not specify a `fimath` object in the constructor:

```
a = fi(pi)
```

```
a =
```

```
3.1416
```

```
      DataTypeMode: Fixed-point: binary point scaling  
      Signedness: Signed  
      WordLength: 16  
      FractionLength: 13  
  
      RoundMode: nearest  
      OverflowMode: saturate  
      ProductMode: FullPrecision
```

```
MaxProductWordLength: 128
    SumMode: FullPrecision
MaxSumWordLength: 128
    CastBeforeSum: true

F1 = fimath

F1 =

    RoundMode: nearest
    OverflowMode: saturate
    ProductMode: FullPrecision
MaxProductWordLength: 128
    SumMode: FullPrecision
MaxSumWordLength: 128
    CastBeforeSum: true
```

To reset the default `fimath` object to the factory default for all future MATLAB sessions, you must again save it as a preference. After setting the MATLAB default `fimath` object to the factory default setting, type

```
savedefaultfimathpref;
```

at the MATLAB command line. The MATLAB factory default `fimath` object becomes your MATLAB default `fimath` object for all future MATLAB sessions.

Using the Default `fimath` Object with the Embedded MATLAB Function Block

For information about using the MATLAB default `fimath` object with the Embedded MATLAB Function block, see “Using `fimath` Objects in Embedded MATLAB Function Blocks” on page 8-18.

fimath Object Properties

In this section...
“Math, Rounding, and Overflow Properties” on page 4-10
“Setting fimath Object Properties” on page 4-11

Math, Rounding, and Overflow Properties

You can always write to the following properties of fimath objects:

Property	Description
CastBeforeSum	Whether both operands are cast to the sum data type before addition
MaxProductWordLength	Maximum allowable word length for the product data type
MaxSumWordLength	Maximum allowable word length for the sum data type
OverflowMode	Overflow-handling mode
ProductBias	Bias of the product data type
ProductFixedExponent	Fixed exponent of the product data type
ProductFractionLength	Fraction length, in bits, of the product data type
ProductMode	Defines how the product data type is determined
ProductSlope	Slope of the product data type
ProductSlopeAdjustmentFactor	Slope adjustment factor of the product data type
ProductWordLength	Word length, in bits, of the product data type
RoundMode	Rounding mode

Property	Description
SumBias	Bias of the sum data type
SumFixedExponent	Fixed exponent of the sum data type
SumFractionLength	Fraction length, in bits, of the sum data type
SumMode	Defines how the sum data type is determined
SumSlope	Slope of the sum data type
SumSlopeAdjustmentFactor	Slope adjustment factor of the sum data type
SumWordLength	Word length, in bits, of the sum data type

For details about these properties, refer to the Property Reference. To learn how to specify properties for `fimath` objects in Fixed-Point Toolbox software, refer to “Setting `fimath` Object Properties” on page 4-11.

Setting `fimath` Object Properties

- “Setting `fimath` Properties at Object Creation” on page 4-11
- “Using Direct Property Referencing with `fimath`” on page 4-12
- “Setting `fimath` Properties in the Model Explorer” on page 4-13

Setting `fimath` Properties at Object Creation

You can set properties of `fimath` objects at the time of object creation by including properties after the arguments of the `fimath` constructor function.

For example, to set the overflow mode to `saturate` and the rounding mode to `convergent`,

```
F = fimath('OverflowMode', 'saturate', 'RoundMode', 'convergent')
```

```
F =
```

```
RoundMode: convergent
OverflowMode: saturate
ProductMode: FullPrecision
MaxProductWordLength: 128
SumMode: FullPrecision
MaxSumWordLength: 128
CastBeforeSum: true
```

Using Direct Property Referencing with fimath

You can reference directly into a property for setting or retrieving `fimath` object property values using MATLAB structure-like referencing. You do so by using a period to index into a property by name.

For example, to get the `RoundMode` of `F`,

```
F.RoundMode
ans =
convergent
```

To set the `OverflowMode` of `F`,

```
F.OverflowMode = 'wrap'
F =
RoundMode: convergent
OverflowMode: wrap
ProductMode: FullPrecision
MaxProductWordLength: 128
SumMode: FullPrecision
MaxSumWordLength: 128
CastBeforeSum: true
```


Setting fimath Properties in the Model Explorer

You can view and change the properties for any `fimath` object defined in the MATLAB workspace in the Model Explorer. Open the Model Explorer by selecting **View > Model Explorer** in any Simulink model, or by typing `daexplr` at the MATLAB command line.

The figure below shows the Model Explorer when you define the following `fimath` objects in the MATLAB workspace:

```
F = fimath
```


```
F =
```

```
      RoundMode: nearest
      OverflowMode: saturate
      ProductMode: FullPrecision
MaxProductWordLength: 128
      SumMode: FullPrecision
      MaxSumWordLength: 128
      CastBeforeSum: true
```

```
G = fimath('OverflowMode','wrap')
```

```
G =
```

```
      RoundMode: nearest
      OverflowMode: wrap
      ProductMode: FullPrecision
MaxProductWordLength: 128
      SumMode: FullPrecision
      MaxSumWordLength: 128
      CastBeforeSum: true
```



Select the **Base Workspace** node in the **Model Hierarchy** pane to view the current objects in the **Contents** pane. When you select a **fimath** object in the **Contents** pane, you can view and change its properties in the **Dialog** pane.

Using fimath Objects to Perform Fixed-Point Arithmetic

In this section...

“Binary Point Arithmetic” on page 4-15

“[Slope Bias] Arithmetic” on page 4-19

Binary Point Arithmetic

The `fimath` object encapsulates the math properties of Fixed-Point Toolbox software, and is itself a property of the `fi` object.

Every `fi` object has a `fimath` object as a property.

```
a = fi(pi)
```

```
a =
```

```
3.1416
```

```

      DataTypeMode: Fixed-point: binary point scaling
      Signedness: Signed
      WordLength: 16
      FractionLength: 13

```

```

      RoundMode: nearest
      OverflowMode: saturate
      ProductMode: FullPrecision
      MaxProductWordLength: 128
      SumMode: FullPrecision
      MaxSumWordLength: 128
      CastBeforeSum: true

```

```
a.fimath
```

```
ans =
```

```
RoundMode: nearest
OverflowMode: saturate
ProductMode: FullPrecision
MaxProductWordLength: 128
SumMode: FullPrecision
MaxSumWordLength: 128
CastBeforeSum: true
```

To perform arithmetic with +, -, .*, or *, two `fi` operands must have the same `fimath` properties.

```
a = fi(pi);
b = fi(8);
isequal(a.fimath, b.fimath)
```

```
ans =
```

```
1
```

```
a + b
```

```
ans =
```

```
11.1416
```

```
DataTypeMode: Fixed-point: binary point scaling
Signedness: Signed
WordLength: 19
FractionLength: 13
```

```
RoundMode: nearest
OverflowMode: saturate
ProductMode: FullPrecision
MaxProductWordLength: 128
SumMode: FullPrecision
MaxSumWordLength: 128
CastBeforeSum: true
```

To perform arithmetic with +, -, .*, or *, two `fi` operands must also have the same data type. For example, you can perform addition on two `fi` objects with data type `double`, but not on an object with data type `double` and one with data type `single`:

```
a = fi(3, 'DataType', 'double')
```

```
a =
```

```
3
```

```
DataTypeMode: double
```

```
b = fi(27, 'DataType', 'double')
```

```
b =
```

```
27
```

```
DataTypeMode: double
```

```
a + b
```

```
ans =
```

```
30
```

```
DataTypeMode: double
```

```
c = fi(12, 'DataType', 'single')
```

```
c =
```

```
12
```

```
DataTypeMode: single
```

```
a + c
```

```
??? Math operations are not allowed on FI objects with  
different data types.
```

Fixed-point `fi` object operands do not have to have the same scaling. Math is permitted between fixed-point and scaled doubles `fi` objects. In this sense, the scaled double data type acts as a fixed-point data type:

```
a = fi(pi)

a =

    3.1416

        DataTypeMode: Fixed-point: binary point scaling
        Signedness: Signed
        WordLength: 16
        FractionLength: 13

        RoundMode: nearest
        OverflowMode: saturate
        ProductMode: FullPrecision
MaxProductWordLength: 128
        SumMode: FullPrecision
MaxSumWordLength: 128
        CastBeforeSum: true

b = fi(magic(2), ...
'DataTypeMode', 'Scaled double: binary point scaling')

b =

     1     3
     4     2

        DataTypeMode: Scaled double: binary point scaling
        Signedness: Signed
        WordLength: 16
        FractionLength: 12

        RoundMode: nearest
        OverflowMode: saturate
        ProductMode: FullPrecision
```

```

MaxProductWordLength: 128
      SumMode: FullPrecision
MaxSumWordLength: 128
      CastBeforeSum: true

a + b

ans =

    4.1416    6.1416
    7.1416    5.1416

      DataTypeMode: Scaled double: binary point scaling
      Signedness: Signed
      WordLength: 18
      FractionLength: 13

      RoundMode: nearest
      OverflowMode: saturate
      ProductMode: FullPrecision
MaxProductWordLength: 128
      SumMode: FullPrecision
MaxSumWordLength: 128
      CastBeforeSum: true

```

Use the `divide` function to perform division with doubles, singles, or binary point-only scaling `fi` objects.

[Slope Bias] Arithmetic

Fixed-point arithmetic using the `fimath` object is supported for all binary point-only signals. Arithmetic is also supported for [Slope Bias] signals with the following restrictions:

- [Slope Bias] signals must be real.
- The `fimath` object `SumMode` and `ProductMode` properties must be set to `'SpecifyPrecision'` for sum and multiply operations, respectively.
- The `fimath` object `CastBeforeSum` property must be set to `'true'`.

- The divide function is not supported for [Slope Bias] signals.

```
f = fimath('SumMode', 'SpecifyPrecision', ...  
          'SumFractionLength', 16)
```

```
f =
```

```
          RoundMode: nearest  
          OverflowMode: saturate  
          ProductMode: FullPrecision  
MaxProductWordLength: 128  
          SumMode: SpecifyPrecision  
          SumWordLength: 32  
          SumFractionLength: 16  
          CastBeforeSum: true
```

```
a = fi(pi, 'fimath', f)
```

```
a =
```

```
3.1416
```

```
          DataTypeMode: Fixed-point: binary point scaling  
          Signedness: Signed  
          WordLength: 16  
          FractionLength: 13
```

```
          RoundMode: nearest  
          OverflowMode: saturate  
          ProductMode: FullPrecision  
MaxProductWordLength: 128  
          SumMode: SpecifyPrecision  
          SumWordLength: 32  
          SumFractionLength: 16  
          CastBeforeSum: true
```

```
b = fi(22, true, 16, 2^-8, 3, 'fimath', f)
```

```
b =
```


22

```

DataTypeMode: Fixed-point: slope and bias scaling
Signedness: Signed
WordLength: 16
Slope: 0.00390625
Bias: 3

RoundMode: nearest
OverflowMode: saturate
ProductMode: FullPrecision
MaxProductWordLength: 128
SumMode: SpecifyPrecision
SumWordLength: 32
SumFractionLength: 16
CastBeforeSum: true

```

a + b

ans =

25.1416

```

DataTypeMode: Fixed-point: binary point scaling
Signedness: Signed
WordLength: 32
FractionLength: 16

RoundMode: nearest
OverflowMode: saturate
ProductMode: FullPrecision
MaxProductWordLength: 128
SumMode: SpecifyPrecision
SumWordLength: 32
SumFractionLength: 16
CastBeforeSum: true

```

Setting the `SumMode` and `ProductMode` properties to `SpecifyPrecision` are mutually exclusive except when performing the `*` operation between matrices.

In this case, both the `SumMode` and `ProductMode` properties must be set to `SpecifyPrecision` for [Slope Bias] signals, because both sum and multiply operations are performed while calculating the result.

Using fimath to Specify Rounding and Overflow Modes

Only rounding and overflow modes set prior to an operation with `fi` objects affect the outcome of those operations. Changing the rounding or overflow mode of a `fi` object after it has been created does not affect its value. For example, consider the `fi` objects `a` and `b`:

```
p = fipref('NumberDisplay', 'RealWorldValue',...
'NumericTypeDisplay', 'none', 'FimathDisplay', 'none');
T = numerictype('WordLength',8,'FractionLength',7);
F = fimath('RoundMode', 'floor', 'OverflowMode', 'wrap');
a = fi(1,T,F)
```

```
a =
```

```
-1
```

```
b = fi(1,T)
```

```
b =
```

```
0.9922
```

Because you create `a` with a `fimath` object `F` that has `OverflowMode` set to `wrap`, the value of `a` is wrapped to `-1`. Conversely, because you create `b` with the default `OverflowMode` value of `saturate`, so its value is saturated to `0.9922`.

Now, assign the `fimath` object `F` to `b`:

```
b.fimath = F
```

```
b =
```

```
0.9922
```

Because the assignment operation and corresponding overflow and saturation already happened when `b` was created, no change happens to `b` when it is assigned the new `fimath` object `F`, even though its rounding mode changed from `saturate` to `wrap`.

Using fimath to Share Arithmetic Rules

You can use a `fimath` object to define common arithmetic rules that you would like to use for many `fi` objects. You can then create multiple `fi` objects, using the same `fimath` object for each. To do so, you also need to create a `numericType` object to define a common data type and scaling. Refer to Chapter 6, “Working with numericType Objects” for more information on `numericType` objects. The following example shows the creation of a `numericType` object and `fimath` object, which are then used to create two `fi` objects with the same `numericType` and `fimath` attributes:

```
T = numericType('WordLength', 32, 'FractionLength', 30)

T =

        DataTypeMode: Fixed-point: binary point scaling
        Signedness: Signed
        WordLength: 32
        FractionLength: 30

F = fimath('RoundMode', 'floor', 'OverflowMode', 'wrap')

F =

        RoundMode: floor
        OverflowMode: wrap
        ProductMode: FullPrecision
        MaxProductWordLength: 128
        SumMode: FullPrecision
        MaxSumWordLength: 128
        CastBeforeSum: true

a = fi(pi, T, F)

a =

    -0.8584
```

```
    DataTypeMode: Fixed-point: binary point scaling
      Signedness: Signed
      WordLength: 32
    FractionLength: 30
```

```
      RoundMode: floor
      OverflowMode: wrap
      ProductMode: FullPrecision
    MaxProductWordLength: 128
      SumMode: FullPrecision
    MaxSumWordLength: 128
    CastBeforeSum: true
```

```
b = fi(pi/2, T, F)
```

```
b =
```

```
    1.5708
```

```
    DataTypeMode: Fixed-point: binary point scaling
      Signedness: Signed
      WordLength: 32
    FractionLength: 30
```

```
      RoundMode: floor
      OverflowMode: wrap
      ProductMode: FullPrecision
    MaxProductWordLength: 128
      SumMode: FullPrecision
    MaxSumWordLength: 128
    CastBeforeSum: true
```

Using fimath ProductMode and SumMode

In this section...

“Example Setup” on page 4-26

“FullPrecision” on page 4-27

“KeepLSB” on page 4-28

“KeepMSB” on page 4-29

“SpecifyPrecision” on page 4-30

Example Setup

The examples in the sections of this topic show the differences among the four settings of the ProductMode and SumMode properties:

- FullPrecision
- KeepLSB
- KeepMSB
- SpecifyPrecision

To follow along, first set the following preferences:

```
p = fipref;  
p.NumericTypeDisplay = 'short';  
p.FimathDisplay = 'none';  
p.LoggingMode = 'on';  
F = fimath('OverflowMode', 'wrap', 'RoundMode', 'floor', ...  
    'CastBeforeSum', false);  
warning off  
format compact
```

Next, define `fi` objects `a` and `b`. Both have signed 8-bit data types. The fraction length is automatically chosen for each `fi` object to yield the best possible precision:

```
a = fi(pi, true, 8)  
a =
```

```

3.1563
s8,5

b = fi(exp(1), true, 8)
b =
2.7188
s8,5

```

FullPrecision

Now, set ProductMode and SumMode for a and b to FullPrecision and look at some results:

```

F.ProductMode = 'FullPrecision';
F.SumMode = 'FullPrecision';
a.fimath = F;
b.fimath = F;
a
a =
3.1563    %011.00101
s8,5

b
b =
2.7188    %010.10111
s8,5

a*b
ans =
8.5811    %001000.1001010011
s16,10

a+b
ans =
5.8750    %0101.11100
s9,5

```

In FullPrecision mode, the product word length grows to the sum of the word lengths of the operands. In this case, each operand has 8 bits, so the product word length is 16 bits. The product fraction length is the sum of the fraction lengths of the operands, in this case $5 + 5 = 10$ bits.

The sum word length grows by one most significant bit to accommodate the possibility of a carry bit. The sum fraction length is aligned with the fraction lengths of the operands, and all fractional bits are kept for full precision. In this case, both operands have 5 fractional bits, so the sum has 5 fractional bits.

KeepLSB

Now, set ProductMode and SumMode for a and b to KeepLSB and look at some results:

```
F.ProductMode = 'KeepLSB';
F.ProductWordLength = 12;
F.SumMode = 'KeepLSB';
F.SumWordLength = 12;
a.fimath = F;
b.fimath = F;
a
a =
    3.1563    %011.00101
         s8,5

b
b =
    2.7188    %010.10111
         s8,5

a*b
ans =
    0.5811    %00.1001010011
         s12,10

a+b
ans =
    5.8750    %0000101.11100
         s12,5
```

In KeepLSB mode, you specify the word lengths and the least significant bits of results are automatically kept. This mode models the behavior of integer operations in the C language.

The product fraction length is the sum of the fraction lengths of the operands. In this case, each operand has 5 fractional bits, so the product fraction length is 10 bits. In this mode, all 10 fractional bits are kept. Overflow occurs because the full-precision result requires 6 integer bits, and only 2 integer bits remain in the product.

The sum fraction length is aligned with the fraction lengths of the operands, and in this model all least significant bits are kept. In this case, both operands had 5 fractional bits, so the sum has 5 fractional bits. The full-precision result requires 4 integer bits, and 7 integer bits remain in the sum, so no overflow occurs in the sum.

KeepMSB

Now, set ProductMode and SumMode for a and b to KeepMSB and look at some results:

```
F.ProductMode = 'KeepMSB';
F.ProductWordLength = 12;
F.SumMode = 'KeepMSB';
F.SumWordLength = 12;
a.fimath = F;
b.fimath = F;
a
a =
    3.1563    %011.00101
         s8,5

b
b =
    2.7188    %010.10111
         s8,5

a*b
ans =
    8.5781    %001000.100101
         s12,6

a+b
ans =
```

```
5.8750    %0101.11100000
          s12,8
```

In `KeepMSB` mode, you specify the word lengths and the most significant bits of sum and product results are automatically kept. This mode models the behavior of many DSP devices where the product and sum are kept in double-wide registers, and the programmer chooses to transfer the most significant bits from the registers to memory after each operation.

The full-precision product requires 6 integer bits, and the fraction length of the product is adjusted to accommodate all 6 integer bits in this mode. No overflow occurs. However, the full-precision product requires 10 fractional bits, and only 6 are available. Therefore, precision is lost.

The full-precision sum requires 4 integer bits, and the fraction length of the sum is adjusted to accommodate all 4 integer bits in this mode. The full-precision sum requires only 5 fractional bits; in this case there are 8, so there is no loss of precision.

SpecifyPrecision

Now set `ProductMode` and `SumMode` for `a` and `b` to `SpecifyPrecision` and look at some results:

```
F.ProductMode = 'SpecifyPrecision';
F.ProductWordLength = 8;
F.ProductFractionLength = 7;
F.SumMode = 'SpecifyPrecision';
F.SumWordLength = 8;
F.SumFractionLength = 7;
a.fimath = F;
b.fimath = F;
a
a =
    3.1563    %011.00101
          s8,5

b
b =
    2.7188    %010.10111
```

```
                s8,5

a*b
ans =
    0.5781    %0.1001010
                s8,7

a+b
ans =
   -0.1250    %1.1110000
                s8,7
```

In `SpecifyPrecision` mode, you must specify both word length and fraction length for sums and products. This example unwisely uses fractional formats for the products and sums, with 8-bit word lengths and 7-bit fraction lengths.

The full-precision product requires 6 integer bits, and the example specifies only 1, so the product overflows. The full-precision product requires 10 fractional bits, and the example only specifies 7, so there is precision loss in the product.

The full-precision sum requires 2 integer bits, and the example specifies only 1, so the sum overflows. The full-precision sum requires 5 fractional bits, and the example specifies 7, so there is no loss of precision in the sum.

fimath Object Functions

You can learn about the functions associated with `fimath` objects in the Function Reference.

Working with fipref Objects

- “Constructing fipref Objects” on page 5-2
- “fipref Object Properties” on page 5-3
- “Using fipref Objects to Set Display Preferences” on page 5-5
- “Using fipref Objects to Set Logging Preferences” on page 5-7
- “Using fipref Objects to Set Data Type Override Preferences” on page 5-12
- “fipref Object Functions” on page 5-15

Constructing fipref Objects

The `fipref` object defines the display and logging attributes for all `fi` objects. You can use the `fipref` constructor function to create a new object.

To get started, type

```
P = fipref
```

to create a default `fipref` object.

```
P =  
    NumberDisplay: 'RealWorldValue'  
    NumericTypeDisplay: 'full'  
    FimathDisplay: 'full'  
    LoggingMode: 'Off'  
    DataTypeOverride: 'ForceOff'
```

The syntax

```
P = fipref(...'PropertyName','PropertyValue'...)
```

allows you to set properties for a `fipref` object at object creation with property name/property value pairs.

Your `fipref` settings persist throughout your MATLAB session. Use `reset(fipref)` to return to the default settings during your session. Use `savefipref` to save your display preferences for subsequent MATLAB sessions.

fipref Object Properties

In this section...

- “Display, Data Type Override, and Logging Properties” on page 5-3
- “Setting fipref Object Properties” on page 5-3

Display, Data Type Override, and Logging Properties

The following properties of `fipref` objects are always writable:

- `FimathDisplay` — Display options for the `fimath` attributes of a `fi` object
- `DataTypeOverride` — Data type override options
- `LoggingMode` — Logging options for operations performed on `fi` objects
- `NumericTypeDisplay` — Display options for the numeric type attributes of a `fi` object
- `NumberDisplay` — Display options for the value of a `fi` object

These properties are described in detail in the Property Reference. To learn how to specify properties for `fipref` objects in Fixed-Point Toolbox software, refer to “Setting fipref Object Properties” on page 5-3.

Setting fipref Object Properties

Setting fipref Properties at Object Creation

You can set properties of `fipref` objects at the time of object creation by including properties after the arguments of the `fipref` constructor function. For example, to set `NumberDisplay` to `bin` and `NumericTypeDisplay` to `short`,

```
P = fipref('NumberDisplay', 'bin', ...
         'NumericTypeDisplay', 'short')
```

```
P =
    NumberDisplay: 'bin'
 NumericTypeDisplay: 'short'
    FimathDisplay: 'full'
```

```
LoggingMode: 'Off'  
DataTypeOverride: 'ForceOff'
```

Using Direct Property Referencing with fipref

You can reference directly into a property for setting or retrieving `fipref` object property values using MATLAB structure-like referencing. You do this by using a period to index into a property by name.

For example, to get the `NumberDisplay` of `P`,

```
P.NumberDisplay  
  
ans =  
  
bin
```

To set the `NumericTypeDisplay` of `P`,

```
P.NumericTypeDisplay = 'full'  
  
P =  
    NumberDisplay: 'bin'  
NumericTypeDisplay: 'full'  
    FimathDisplay: 'full'  
    LoggingMode: 'Off'  
    DataTypeOverride: 'ForceOff'
```


Using fipref Objects to Set Display Preferences

You use the `fipref` object to dictate three aspects of the display of `fi` objects: how the value of a `fi` object is displayed, how the `fimath` properties are displayed, and how the `numericType` properties are displayed.

For example, the following shows the default `fipref` display for a `fi` object:

```
a = fi(pi)

a =

    3.1416

    DataTypeMode: Fixed-point: binary point scaling
    Signedness: Signed
    WordLength: 16
    FractionLength: 13

    RoundMode: nearest
    OverflowMode: saturate
    ProductMode: FullPrecision
    MaxProductWordLength: 128
    SumMode: FullPrecision
    MaxSumWordLength: 128
    CastBeforeSum: true
```

Now, change the `fipref` display properties:

```
P = fipref;
P.NumberDisplay = 'bin';
P.NumericTypeDisplay = 'short';
P.FimathDisplay = 'none'

P =

    NumberDisplay: 'bin'
    NumericTypeDisplay: 'short'
    FimathDisplay: 'none'
    LoggingMode: 'Off'
```

```
        DataTypeOverride: 'ForceOff'  
a  
  
a =  
0110010010001000  
s16,13
```

Using fipref Objects to Set Logging Preferences

In this section...

“Logging Overflows and Underflows as Warnings” on page 5-7

“Accessing Logged Information with Functions” on page 5-9

Logging Overflows and Underflows as Warnings

Overflows and underflows are logged as warnings for all assignment, plus, minus, and multiplication operations when the fipref LoggingMode property is set to on. For example, try the following:

- 1 Create a signed fi object that is a vector of values from 1 to 5, with 8-bit word length and 6-bit fraction length.

```
a = fi(1:5,1,8,6);
```

- 2 Define the fimath object associated with a, and indicate that you will specify the sum and product word and fraction lengths.

```
F = a.fimath;  
F.SumMode = 'SpecifyPrecision';  
F.ProductMode = 'SpecifyPrecision';  
a.fimath = F;
```

- 3 Define the fipref object and turn on overflow and underflow logging.

```
P = fipref;  
P.LoggingMode = 'on';
```

- 4 Suppress the numericType and fimath displays.

```
P.NumericTypeDisplay = 'none';  
P.FimathDisplay = 'none';
```

- 5 Specify the sum and product word and fraction lengths.

```
a.SumWordLength = 16;  
a.SumFractionLength = 15;
```

```
a.ProductWordLength = 16;  
a.ProductFractionLength = 15;
```

- 6** Warnings are displayed for overflows and underflows in assignment operations. For example, try:

```
a(1) = pi  
Warning: 1 overflow occurred in the fi assignment operation.
```

```
a =
```

```
1.9844    1.9844    1.9844    1.9844    1.9844
```

```
a(1) = double(eps(a))/10
```

```
Warning: 1 underflow occurred in the fi assignment operation.
```

```
a =
```

```
0    1.9844    1.9844    1.9844    1.9844
```

- 7** Warnings are displayed for overflows and underflows in addition and subtraction operations. For example, try:

```
a+a
```

```
Warning: 12 overflows occurred in the fi + operation.
```

```
ans =
```

```
0    1.0000    1.0000    1.0000    1.0000
```

```
a-a
```

```
Warning: 8 overflows occurred in the fi - operation.
```

```
ans =
```

```
0    0    0    0    0
```

- 8** Warnings are displayed for overflows and underflows in multiplication operations. For example, try:

```

a.*a
Warning: 4 product overflows occurred in the fi .* operation.

ans =

         0    1.0000    1.0000    1.0000    1.0000

a*a'
Warning: 4 product overflows occurred in the fi * operation.
Warning: 3 sum overflows occurred in the fi * operation.

ans =

    1.0000

```

The final example above is a complex multiplication that requires both multiplication and addition operations. The warnings inform you of overflows and underflows in both.

Because overflows and underflows are logged as warnings, you can use the `dbstop` MATLAB function with the syntax

```
dbstop if warning
```

to find the exact lines in an M-file that are causing overflows or underflows.

Use

```
dbstop if warning fi:underflow
```

to stop only on lines that cause an underflow. Use

```
dbstop if warning fi:overflow
```

to stop only on lines that cause an overflow.

Accessing Logged Information with Functions

When the `fipref` `LoggingMode` property is set to on, you can use the following functions to return logged information about assignment and creation operations to the MATLAB command line:

- `maxlog` — Returns the maximum real-world value
- `minlog` — Returns the minimum value
- `noverflows` — Returns the number of overflows
- `nunderflows` — Returns the number of underflows

`LoggingMode` must be set to on before you perform any operation in order to log information about it. To clear the log, use the function `resetlog`.

For example, consider the following. First turn logging on, then perform operations, and then finally get information about the operations:

```
fipref('LoggingMode','on');  
x = fi([-1.5 eps 0.5], true, 16, 15);  
x(1) = 3.0;  
maxlog(x)
```

```
ans =
```

```
3
```

```
minlog(x)
```

```
ans =
```

```
-1.5000
```

```
noverflows(x)
```

```
ans =
```

```
2
```

```
nunderflows(x)
```

```
ans =
```

```
1
```

Next, reset the log and request the same information again. Note that the functions return empty [], because logging has been reset since the operations were run:

```
resetlog(x)
maxlog(x)

ans =

    []

minlog(x)

ans =

    []

noverflows(x)

ans =

    []

nunderflows(x)

ans =

    []
```

Using fipref Objects to Set Data Type Override Preferences

In this section...

“Overriding the Data Type of fi Objects” on page 5-12

“Using Data Type Override to Help Set Fixed-Point Scaling” on page 5-13

Overriding the Data Type of fi Objects

Use the fipref `DataTypeOverride` property to override `fi` objects with singles, doubles, or scaled doubles. Data type override only occurs when the `fi` constructor function is called. Objects that are created while data type override is on have the overridden data type. They maintain that data type when data type override is later turned off. To obtain an object with a data type that is not the override data type, you must create an object when data type override is off:

```
p = fipref('DataTypeOverride', 'TrueDoubles')

p =

    NumberDisplay: 'RealWorldValue'
 NumericTypeDisplay: 'full'
   FimathDisplay: 'full'
    LoggingMode: 'Off'
  DataTypeOverride: 'TrueDoubles'

a = fi(pi)

a =

    3.1416

    DataTypeMode: double

p = fipref('DataTypeOverride', 'ForceOff')

p =
```



```
        NumberDisplay: 'RealWorldValue'  
NumericTypeDisplay: 'full'  
        FimathDisplay: 'full'  
        LoggingMode: 'Off'  
        DataTypeOverride: 'ForceOff'  
  
a  
  
a =  
  
    3.1416  
  
        DataTypeMode: double  
  
b = fi(pi)  
  
b =  
  
    3.1416  
  
        DataTypeMode: Fixed-point: binary point scaling  
        Signedness: Signed  
        WordLength: 16  
        FractionLength: 13
```

Tip To reset the `fipref` object to its default values use `reset(fipref)` or `reset(p)`, where `p` is a `fipref` object. This is useful to ensure that data type override and logging are off.

Using Data Type Override to Help Set Fixed-Point Scaling

Choosing the scaling for the fixed-point variables in your algorithms can be difficult. In Fixed-Point Toolbox software, you can use a combination of data type override and min/max logging to help you discover the numerical ranges that your fixed-point data types need to cover. These ranges dictate the appropriate scalings for your fixed-point data types. In general, the procedure is

- 1** Implement your algorithm using fixed-point `fi` objects, using initial “best guesses” for word lengths and scalings.
- 2** Set the `fipref` `DataTypeOverride` property to `ScaledDoubles`, `TrueSingles`, or `TrueDoubles`.
- 3** Set the `fipref` `LoggingMode` property to `on`.
- 4** Use the `maxlog` and `minlog` functions to log the maximum and minimum values achieved by the variables in your algorithm in floating-point mode.
- 5** Set the `fipref` `DataTypeOverride` property to `ForceOff`.
- 6** Use the information obtained in step 4 to set the fixed-point scaling for each variable in your algorithm such that the full numerical range of each variable is representable by its data type and scaling.

A detailed example of this process is shown in the Fixed-Point Toolbox “Fixed-Point Data Type Override, Min/Max Logging, and Scaling” demo.

fipref Object Functions

You can learn about the functions associated with `fipref` objects in the Function Reference.

Working with numerictype Objects

- “Constructing numerictype Objects” on page 6-2
- “numerictype Object Properties” on page 6-7
- “The numerictype Structure” on page 6-11
- “Using numerictype Objects to Share Data Type and Scaling Settings of fi objects” on page 6-14
- “numerictype Object Functions” on page 6-18

Constructing numerictype Objects

In this section...

“numerictype Object Syntaxes” on page 6-2

“Example: Constructing a numerictype Object with Property Name and Property Value Pairs” on page 6-3

“Example: Copying a numerictype Object” on page 6-4

“Example: Building numerictype Object Constructors in a GUI” on page 6-4

numerictype Object Syntaxes

numerictype objects define the data type and scaling attributes of `fi` objects, as well as Simulink signals and model parameters. You can create numerictype objects in Fixed-Point Toolbox software in one of two ways:

- You can use the numerictype constructor function to create a new object.
- You can use the numerictype constructor function to copy an existing numerictype object.

To get started, type

```
T = numerictype
```

to create a default numerictype object.

```
T =
```

```
      DataTypeMode: Fixed-point: binary point scaling  
      Signedness: Signed  
      WordLength: 16  
      FractionLength: 15
```

To see all of the numerictype object syntaxes, refer to the numerictype constructor function reference page.

The following examples show different ways of constructing `numerictype` objects. For more examples of constructing `numerictype` objects, see the “Examples” on the `numerictype` constructor function reference page.

Example: Constructing a numerictype Object with Property Name and Property Value Pairs

When you create a `numerictype` object using property name and property value pairs, Fixed-Point Toolbox software first creates a default `numerictype` object, and then, for each property name you specify in the constructor, assigns the corresponding value.

This behavior differs from the behavior that occurs when you use a syntax such as `T = numerictype(s,w)`, where you only specify the property values in the constructor. Using such a syntax results in no default `numerictype` object being created, and the `numerictype` object receives only the assigned property values that are specified in the constructor.

The following example shows how the property name/property value syntax creates a slightly different `numerictype` object than the property values syntax, even when you specify the same property values in both constructors.

To demonstrate this difference, suppose you want to create an unsigned `numerictype` object with a word length of 32 bits.

First, create the `numerictype` object using property name/property value pairs.

```
T1 = numerictype('signed',0,'wordlength',32)
```

```
T1 =
```

```
      DataTypeMode: Fixed-point: binary point scaling
      Signedness: Unsigned
      WordLength: 32
      FractionLength: 15
```

The numerictype object T1 has the same DataTypeMode and FractionLength as a default numerictype object, but the WordLength and Signed properties are overwritten with the values you specified.

Now, create another unsigned 32 bit numerictype object, but this time specify only property values in the constructor.

```
T2 = numerictype(0,32)
```

```
T2 =
```

```
DataTypeMode: Fixed-point: unspecified scaling  
Signedness: Unsigned  
WordLength: 32
```

Unlike T1, T2 only has the property values you specified. The DataTypeMode of T2 is Fixed-Point: unspecified scaling, so no fraction length is assigned.

fi objects cannot have unspecified numerictype properties. Thus, all unspecified numerictype object properties become specified at the time of fi object creation.

Example: Copying a numerictype Object

To copy a numerictype object, simply use assignment as in the following example:

```
T = numerictype;  
U = T;  
isequal(T,U)
```

```
ans =
```

```
1
```

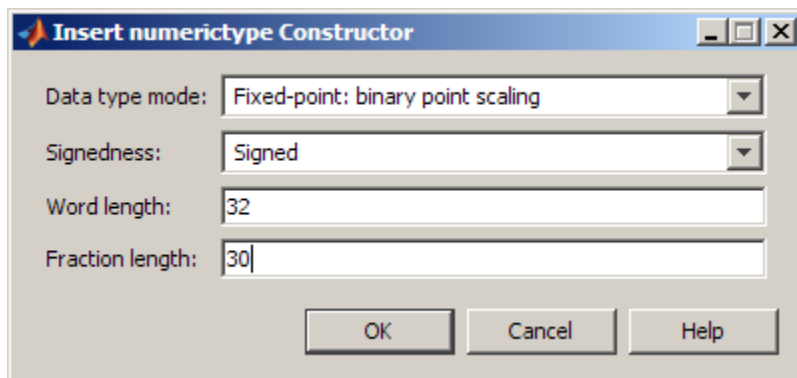
Example: Building numerictype Object Constructors in a GUI

When you are working with M-files, you can build your numerictype object constructors using the **Insert numerictype Constructor** dialog box. After

specifying the properties of the numerictype object in the dialog box, you can insert the prepopulated numerictype object constructor string at a specific location in your M-file.

For example, to create a signed numerictype object with binary-point scaling, a word length of 32 bits and a fraction length of 30 bits, perform the following steps:

- 1 Open the **Insert numerictype Constructor** dialog box by selecting **Tools > Fixed-Point Toolbox > Insert numerictype Constructor** from the editor menu.
- 2 Use the edit boxes and drop-down menus to specify the following properties of the numerictype object:
 - **Data type mode** = Fixed-point: binary point scaling
 - **Signedness** = Signed
 - **Word length** = 32
 - **Fraction length** = 30



- 3 To insert the numerictype object constructor string into your M-file, place your cursor at the desired location in the M-file, and click **OK** on the **Insert numerictype Constructor** dialog box. Clicking **OK** closes the **Insert numerictype Constructor** dialog box and automatically populates the numerictype object constructor string in your M-file:

```
5 T = numerictype(1, 32, 30)
```

numerictype Object Properties

In this section...
“Data Type and Scaling Properties” on page 6-7
“Setting numerictype Object Properties” on page 6-8

Data Type and Scaling Properties

All properties of a `numerictype` object are writable. However, the `numerictype` properties of a `fi` object become read only after the `fi` object has been created. Any `numerictype` properties of a `fi` object that are unspecified at the time of `fi` object creation are automatically set to their default values. The properties of a `numerictype` object are:

- `Bias` — Bias
- `DataType` — Data type category
- `DataTypeMode` — Data type and scaling mode
- `FixedExponent` — Fixed-point exponent
- `SlopeAdjustmentFactor` — Slope adjustment
- `FractionLength` — Fraction length of the stored integer value, in bits
- `Scaling` — Fixed-point scaling mode
- `Signed` — Signed or unsigned
- `Signedness` — Signed, unsigned, or auto
- `Slope` — Slope
- `WordLength` — Word length of the stored integer value, in bits

These properties are described in detail in the Property Reference. To learn how to specify properties for `numerictype` objects in Fixed-Point Toolbox software, refer to “Setting numerictype Object Properties” on page 6-8.

Setting numeric type Object Properties

Setting numeric type Properties at Object Creation

You can set properties of numeric type objects at the time of object creation by including properties after the arguments of the numeric type constructor function.

For example, to set the word length to 32 bits and the fraction length to 30 bits,

```
T = numericType('WordLength', 32, 'FractionLength', 30)
```

```
T =
```

```
      DataTypeMode: Fixed-point: binary point scaling  
      Signedness: Signed  
      WordLength: 32  
      FractionLength: 30
```

Using Direct Property Referencing with numeric type Objects

You can reference directly into a property for setting or retrieving numeric type object property values using MATLAB structure-like referencing. You do this by using a period to index into a property by name.

For example, to get the word length of T,

```
T.WordLength
```

```
ans =
```

```
32
```

To set the fraction length of T,

```
T.FractionLength = 31
```

```
T =
```

```

        DataTypeMode: Fixed-point: binary point scaling
          Signedness: Signed
          WordLength: 32
    FractionLength: 31
    
```

Setting numerictype Properties in the Model Explorer

You can view and change the properties for any numerictype object defined in the MATLAB workspace in the Model Explorer. Open the Model Explorer by selecting **View > Model Explorer** in any Simulink model, or by typing `daexplr` at the MATLAB command line.

The figure below shows the Model Explorer when you define the following numerictype objects in the MATLAB workspace:

```
T = numerictype
```

```
T =
```

```

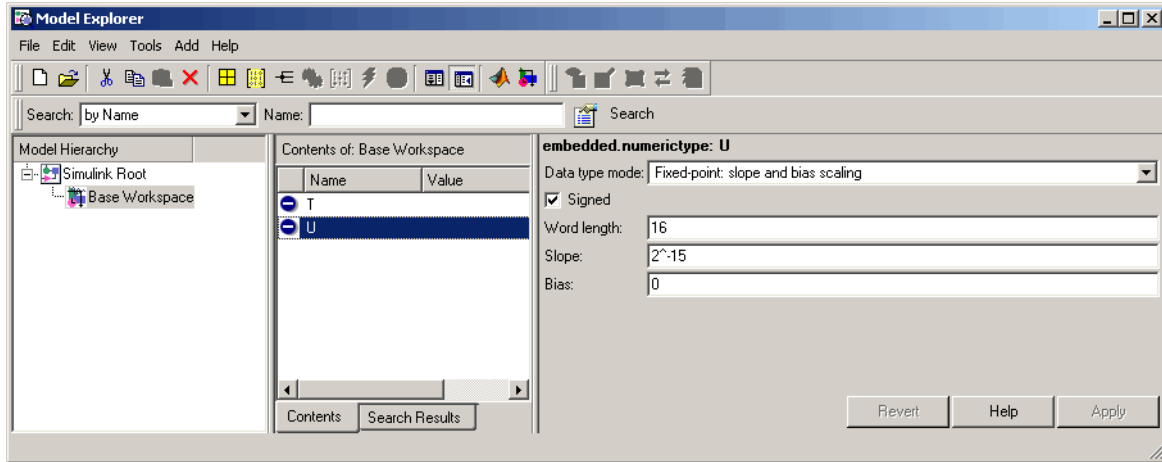
        DataTypeMode: Fixed-point: binary point scaling
          Signedness: Signed
          WordLength: 16
    FractionLength: 15
    
```

```
U = numerictype('DataTypeMode', 'Fixed-point: slope and bias')
```

```
U =
```

```

        DataTypeMode: Fixed-point: slope and bias scaling
          Signedness: Signed
          WordLength: 16
          Slope: 2^-15
          Bias: 0
    
```



Select the **Base Workspace** node in the **Model Hierarchy** pane to view the current objects in the **Contents** pane. When you select a numerictype object in the **Contents** pane, you can view and change its properties in the **Dialog** pane.

The numerictype Structure

In this section...

“Valid Values for numerictype Structure Properties” on page 6-11

“Properties That Affect the Slope” on page 6-13

“Stored Integer Value and Real World Value” on page 6-13

Valid Values for numerictype Structure Properties

The numerictype object contains all the data type and scaling attributes of a fixed-point object. The numerictype object behaves like any MATLAB structure, except that it only lets you set valid values for defined fields. The following table shows the possible settings of each field of the structure.

Note When you create a `fi` object, any unspecified field of the numerictype object reverts to its default value. Thus, if the `DataTypeMode` is set to unspecified scaling, it defaults to binary point scaling when the `fi` object is created. If the `Signedness` property of the numerictype object is set to `Auto`, it defaults to `Signed` when the `fi` object is created.

DataTypeMode	DataType	Scaling	Signedness	Word- Length	Fraction- Length	Slope	Bias
<i>Fixed-point data types</i>							
Fixed-point: binary point scaling	Fixed	BinaryPoint	Signed Unsigned Auto	Positive integer from 1 to 65,536	Positive or negative integer	1	0
Fixed-point: slope and bias scaling	Fixed	SlopeBias	Signed Unsigned Auto	Positive integer from 1 to 65,536	N/A	Any floating- point number	Any floating- point number

DataTypeMode	DataType	Scaling	Signedness	Word- Length	Fraction- Length	Slope	Bias
Fixed-point: unspecified scaling	Fixed	Unspecified	Signed Unsigned Auto	Positive integer from 1 to 65,536	N/A	N/A	N/A
<i>Scaled double data types</i>							
Scaled double: binary point scaling	ScaledDouble	BinaryPoint	Signed Unsigned Auto	Positive integer from 1 to 65,536	Positive or negative integer	1	0
Scaled double: slope and bias scaling	ScaledDouble	SlopeBias	Signed Unsigned Auto	Positive integer from 1 to 65,536	N/A	Any floating- point number	Any floating- point number
Scaled double: unspecified scaling	ScaledDouble	Unspecified	Signed Unsigned Auto	Positive integer from 1 to 65,536	N/A	N/A	N/A
<i>Built-in data types</i>							
double	double	N/A	1 true	64	0	1	0
single	single	N/A	1 true	32	0	1	0
boolean	boolean	N/A	0 false	1	0	1	0

You cannot change the numerictype properties of a fi object after fi object creation.

Properties That Affect the Slope

The **Slope** field of the numerictype structure is related to the SlopeAdjustmentFactor and FixedExponent properties by

$$\text{slope} = \text{slope adjustment factor} \times 2^{\text{fixed exponent}}$$

The FixedExponent and FractionLength properties are related by

$$\text{fixed exponent} = -\text{fraction length}$$

If you set the SlopeAdjustmentFactor, FixedExponent, or FractionLength property, the **Slope** field is modified.

Stored Integer Value and Real World Value

The numerictype StoredIntegerValue and RealWorldValue properties are related according to

$$\text{real-world value} = \text{stored integer value} \times 2^{-\text{fraction length}}$$

which is equivalent to

$$\text{real-world value} = \text{stored integer value} \times (\text{slope adjustment factor} \times 2^{\text{fixed exponent}}) + \text{bias}$$

If any of these properties is updated, the others are modified accordingly.

Using numerictype Objects to Share Data Type and Scaling Settings of fi objects

You can use a numerictype object to define common data type and scaling rules that you would like to use for many fi objects. You can then create multiple fi objects, using the same numerictype object for each.

Example 1

In the following example, you create a numerictype object T with word length 32 and fraction length 28. Next, to ensure that your fi objects have the same numerictype attributes, create fi objects a and b using your numerictype object T.

```
format long g
T = numerictype('WordLength',32,'FractionLength',28)

T =

        DataTypeMode: Fixed-point: binary point scaling
        Signedness: Signed
        WordLength: 32
        FractionLength: 28

a = fi(pi,T)

a =

        3.1415926553309

        DataTypeMode: Fixed-point: binary point scaling
        Signedness: Signed
        WordLength: 32
        FractionLength: 28

        RoundMode: nearest
        OverflowMode: saturate
```

```

        ProductMode: FullPrecision
MaxProductWordLength: 128
        SumMode: FullPrecision
MaxSumWordLength: 128
        CastBeforeSum: true

b = fi(pi/2, T)

b =

        1.5707963258028

        DataTypeMode: Fixed-point: binary point scaling
        Signedness: Signed
        WordLength: 32
        FractionLength: 28

        RoundMode: nearest
        OverflowMode: saturate
        ProductMode: FullPrecision
MaxProductWordLength: 128
        SumMode: FullPrecision
MaxSumWordLength: 128
        CastBeforeSum: true

```

Example 2

In this example, start by creating a numerictype object T with [Slope Bias] scaling. Next, use that object to create two fi objects, c and d with the same numerictype attributes:

```

T = numerictype('scaling','slopebias','slope', 2^2, 'bias', 0)

T =

        DataTypeMode: Fixed-point: slope and bias scaling
        Signedness: Signed
        WordLength: 16
        Slope: 2^2
        Bias: 0

```

```
c = fi(pi, T)
```

```
c =
```

```
4
```

```
    DataTypeMode: Fixed-point: slope and bias scaling  
    Signedness: Signed  
    WordLength: 16  
    Slope: 2^2  
    Bias: 0
```

```
    RoundMode: nearest  
    OverflowMode: saturate  
    ProductMode: FullPrecision  
MaxProductWordLength: 128  
    SumMode: FullPrecision  
MaxSumWordLength: 128  
CastBeforeSum: true
```

```
d = fi(pi/2, T)
```

```
d =
```

```
0
```

```
    DataTypeMode: Fixed-point: slope and bias scaling  
    Signedness: Signed  
    WordLength: 16  
    Slope: 2^2  
    Bias: 0
```

```
    RoundMode: nearest  
    OverflowMode: saturate  
    ProductMode: FullPrecision  
MaxProductWordLength: 128  
    SumMode: FullPrecision  
MaxSumWordLength: 128  
CastBeforeSum: true
```


numerictype Object Functions

You can learn about the functions associated with `numerictype` objects in the Function Reference.

Working with quantizer Objects

- “Constructing quantizer Objects” on page 7-2
- “quantizer Object Properties” on page 7-3
- “Quantizing Data with quantizer Objects” on page 7-4
- “Transformations for Quantized Data” on page 7-6
- “quantizer Object Functions” on page 7-7

Constructing quantizer Objects

You can use quantizer objects to quantize data sets. You can create quantizer objects in Fixed-Point Toolbox software in one of two ways:

- You can use the `quantizer` constructor function to create a new object.
- You can use the `quantizer` constructor function to copy a quantizer object.

To create a quantizer object with default properties, type

```
q = quantizer

q =

    DataMode = fixed
    RoundMode = floor
    OverflowMode = saturate
    Format = [16 15]
```

To copy a quantizer object, simply use assignment as in the following example:

```
q = quantizer;
r = q;
isequal(q,r)

ans =

    1
```

A listing of all the properties of the quantizer object `q` you just created is displayed along with the associated property values. All property values are set to defaults when you construct a quantizer object this way. See “quantizer Object Properties” on page 7-3 for more details.

quantizer Object Properties

The following properties of quantizer objects are always writable:

- `DataMode` — Type of arithmetic used in quantization
- `Format` — Data format of a quantizer object
- `OverflowMode` — Overflow-handling mode
- `RoundMode` — Rounding mode

See the Property Reference for more details about these properties, including their possible values.

For example, to create a fixed-point quantizer object with

- The `Format` property value set to `[16,14]`
- The `OverflowMode` property value set to `'saturate'`
- The `RoundMode` property value set to `'ceil'`

type

```
q = quantizer('datamode','fixed','format',[16,14],...  
             'overflowmode','saturate','roundmode','ceil')
```

You do not have to include quantizer object property names when you set quantizer object property values.

For example, you can create quantizer object `q` from the previous example by typing

```
q = quantizer('fixed',[16,14],'saturate','ceil')
```

Note You do not have to include default property values when you construct a quantizer object. In this example, you could leave out `'fixed'` and `'saturate'`.

Quantizing Data with quantizer Objects

You construct a `quantizer` object to specify the quantization parameters to use when you quantize data sets. You can use the `quantize` function to quantize data according to a `quantizer` object's specifications.

Once you quantize data with a `quantizer` object, its state values might change.

The following example shows

- How you use `quantize` to quantize data
- How quantization affects `quantizer` object states
- How you reset `quantizer` object states to their default values using `reset`

1 Construct an example data set and a `quantizer` object.

```
format long g
randn('state',0);
x = randn(100,4);
q = quantizer([16,14]);
```

2 Retrieve the values of the `maxlog` and `noverflows` states.

```
q.maxlog
ans =
    -1.79769313486232e+308

q.noverflows
ans =
    0
```

Note that `maxlog` is equal to `-realmax`, which indicates that the `quantizer` `q` is in a reset state.

3 Quantize the data set according to the `quantizer` object's specifications.

```
y = quantize(q,x);  
Warning: 15 overflows.
```

4 Check the values of `maxlog` and `noverflows`.

```
q.maxlog  
  
ans =  
  
1.99993896484375  
  
q.noverflows  
  
ans =  
  
15
```

Note that the maximum logged value was taken after quantization, that is, `q.maxlog == max(y)`.

5 Reset the quantizer states and check them.

```
reset(q)  
q.maxlog  
  
ans =  
  
-1.79769313486232e+308  
  
q.noverflows  
  
ans =  
  
0
```

Transformations for Quantized Data

You can convert data values from numeric to hexadecimal or binary according to a quantizer object's specifications.

Use

- `num2bin` to convert data to binary
- `num2hex` to convert data to hexadecimal
- `hex2num` to convert hexadecimal data to numeric
- `bin2num` to convert binary data to numeric

For example,

```
q = quantizer([3 2]);  
x = [0.75  -0.25  
      0.50  -0.50  
      0.25  -0.75  
      0     -1  ];  
b = num2bin(q,x)
```

```
b =  
011  
010  
001  
000  
111  
110  
101  
100
```

produces all two's complement fractional representations of 3-bit fixed-point numbers.

quantizer Object Functions

You can learn about the functions associated with `quantizer` objects in the [Function Reference](#).

Working with the Fixed-Point Embedded MATLAB Subset

- “Supported Functions and Limitations of the Fixed-Point Embedded MATLAB Subset” on page 8-2
- “Embedded MATLAB MEX” on page 8-9
- “Embedded MATLAB Function Block” on page 8-13
- “Embedded MATLAB Coder” on page 8-35
- “Working with Fixed-Point Compilation Reports” on page 8-36

Supported Functions and Limitations of the Fixed-Point Embedded MATLAB Subset

The Embedded MATLAB subset is a restricted subset of the MATLAB language that provides optimizations for:

- Generating efficient, production-quality C code for embedded applications. Embedded MATLAB subset restricts MATLAB semantics to meet the memory and data type requirements of embedded target environments.
- Accelerating fixed-point algorithms.

For more information about the Embedded MATLAB subset, refer to the “Embedded MATLAB” documentation. The Embedded MATLAB subset supports a significant number of Fixed-Point Toolbox functions, which are listed in the table below. The following general limitations always apply to the use of Fixed-Point Toolbox software with the Embedded MATLAB subset:

- `fi`pref and `quantizer` objects are not supported.
- Dot notation is only supported for getting the values of `fi`math and `numericType` properties. Dot notation is not supported for `fi` objects, and it is not supported for setting properties.
- Word lengths greater than 128 bits are not supported.
- It is illegal to change the `fi`math or `numericType` of a given variable once it has been created.
- The `boolean` and `ScaledDouble` values of the `DataTypeMode` and `DataType` properties are not supported.
- The `false` value of the `CastBeforeSum` property is not supported.
- The `numel` function returns the number of elements of `fi` objects in the Embedded MATLAB subset.

To learn about the general limitations of the Embedded MATLAB subset that also apply to use with Fixed-Point Toolbox software, refer to “Overview of the Embedded MATLAB Subset” in the Embedded MATLAB language subset documentation.

Fixed-Point Toolbox Functions Supported for Use with the Embedded MATLAB Language Subset

Function	Remarks/Limitations
abs	—
add	—
all	—
any	—
bitand	• Not supported for slope-bias scaled <code>fi</code> objects.
bitandreduce	—
bitcmp	—
bitconcat	—
bitget	—
bitor	• Not supported for slope-bias scaled <code>fi</code> objects.
bitorreduce	—
bitreplicate	—
bitrol	—
bitror	—
bitset	—
bitshift	—
bitsliceget	—
bitsll	—
bitsra	—
bitsrl	—
bitxor	• Not supported for slope-bias scaled <code>fi</code> objects.
bitxorreduce	—
ceil	—
complex	—
conj	—

Fixed-Point Toolbox Functions Supported for Use with the Embedded MATLAB Language Subset (Continued)

Function	Remarks/Limitations
convergent	—
ctranspose	—
diag	<ul style="list-style-type: none"> • If supplied, the index, k, must be a real and scalar integer value that is not a <code>fi</code> object.
disp	—
divide	<ul style="list-style-type: none"> • Any non-<code>fi</code> input must be constant; that is, its value must be known at compile time so that it can be cast to a <code>fi</code> object. • Complex and imaginary divisors are not supported.
double	—
end	—
eps	<ul style="list-style-type: none"> • Supported for scalar fixed-point signals only. • Supported for scalar, vector, and matrix, <code>fi</code> single and <code>fi</code> double signals.
eq	<ul style="list-style-type: none"> • Not supported for fixed-point signals with different biases.
fi	<ul style="list-style-type: none"> • Use to create a fixed-point constant or variable in the Embedded MATLAB language subset. • The default constructor syntax without any input arguments is not supported. • The syntax <code>fi('PropertyName',PropertyValue...)</code> is not supported. To use property name/property value pairs, you must first specify the value v of the <code>fi</code> object as in <code>fi(v,'PropertyName',PropertyValue...)</code>. • Works for all input values when complete <code>numericType</code> information of the <code>fi</code> object is provided. • Works only for constant input values (value of input must be known at compile time) when complete <code>numericType</code> information of the <code>fi</code> object is not specified.

Fixed-Point Toolbox Functions Supported for Use with the Embedded MATLAB Language Subset (Continued)

Function	Remarks/Limitations
	<ul style="list-style-type: none"> • <code>numericType</code> object information must be available for nonfixed-point Simulink inputs.
<code>fimath</code>	<ul style="list-style-type: none"> • Fixed-point signals coming in to an Embedded MATLAB Function block from Simulink are assigned the <code>fimath</code> object defined in the Embedded MATLAB Function block dialog in the Model Explorer. • Use to create <code>fimath</code> objects in Embedded MATLAB code.
<code>fix</code>	—
<code>floor</code>	—
<code>ge</code>	<ul style="list-style-type: none"> • Not supported for fixed-point signals with different biases.
<code>get</code>	<ul style="list-style-type: none"> • The syntax <code>structure = get(o)</code> is not supported.
<code>getlsb</code>	—
<code>getmsb</code>	—
<code>gt</code>	<ul style="list-style-type: none"> • Not supported for fixed-point signals with different biases.
<code>horzcat</code>	—
<code>imag</code>	—
<code>int, int8, int16, int32</code>	—
<code>iscolumn</code>	—
<code>isempty</code>	—
<code>isfi</code>	—
<code>isfimath</code>	—
<code>isfinite</code>	—
<code>isinf</code>	—
<code>isnan</code>	—
<code>isnumeric</code>	—
<code>isnumericType</code>	—

Fixed-Point Toolbox Functions Supported for Use with the Embedded MATLAB Language Subset (Continued)

Function	Remarks/Limitations
isreal	—
isrow	—
isscalar	—
issigned	—
isvector	—
le	<ul style="list-style-type: none"> • Not supported for fixed-point signals with different biases.
length	—
logical	—
lowerbound	—
lsb	<ul style="list-style-type: none"> • Supported for scalar fixed-point signals only. • Supported for scalar, vector, and matrix, <code>fi</code> single and double signals.
lt	<ul style="list-style-type: none"> • Not supported for fixed-point signals with different biases.
max	—
min	—
minus	<ul style="list-style-type: none"> • Any non-<code>fi</code> input must be constant; that is, its value must be known at compile time so that it can be cast to a <code>fi</code> object.
mpy	—
mrdivide	—
mtimes	<ul style="list-style-type: none"> • Any non-<code>fi</code> input must be constant; that is, its value must be known at compile time so that it can be cast to a <code>fi</code> object.
ndims	—
ne	<ul style="list-style-type: none"> • Not supported for fixed-point signals with different biases.
nearest	—
numberofelements	<ul style="list-style-type: none"> • <code>numberofelements</code> and <code>numel</code> both work the same as MATLAB <code>numel</code> for <code>fi</code> objects in the Embedded MATLAB language subset.

Fixed-Point Toolbox Functions Supported for Use with the Embedded MATLAB Language Subset (Continued)

Function	Remarks/Limitations
<code>numerictype</code>	<ul style="list-style-type: none"> Fixed-point signals coming in to an Embedded MATLAB Function block from Simulink are assigned a <code>numerictype</code> object that is populated with the signal's data type and scaling information. Returns the data type when the input is a nonfixed-point signal. Use to create <code>numerictype</code> objects in Embedded MATLAB code.
<code>permute</code>	—
<code>plus</code>	<ul style="list-style-type: none"> Any non-<code>fi</code> input must be constant; that is, its value must be known at compile time so that it can be cast to a <code>fi</code> object.
<code>pow2</code>	—
<code>range</code>	—
<code>rdivide</code>	—
<code>real</code>	—
<code>realmax</code>	—
<code>realmin</code>	—
<code>reinterprecast</code>	—
<code>repmat</code>	—
<code>rescale</code>	—
<code>reshape</code>	—
<code>round</code>	—
<code>sign</code>	—
<code>single</code>	—
<code>size</code>	—
<code>sort</code>	—
<code>sqrt</code>	<ul style="list-style-type: none"> Complex and [Slope Bias] inputs error out. Negative inputs yield a 0 result.
<code>sub</code>	—

Fixed-Point Toolbox Functions Supported for Use with the Embedded MATLAB Language Subset (Continued)

Function	Remarks/Limitations
subsasgn	—
subsref	—
sum	—
times	<ul style="list-style-type: none"> Any non-<code>fi</code> input must be constant; that is, its value must be known at compile time so that it can be cast to a <code>fi</code> object.
transpose	—
tril	<ul style="list-style-type: none"> If supplied, the index, k, must be a real and scalar integer value that is not a <code>fi</code> object.
triu	<ul style="list-style-type: none"> If supplied, the index, k, must be a real and scalar integer value that is not a <code>fi</code> object.
uint8, uint16, uint32	—
uminus	—
uplus	—
upperbound	—
vertcat	—

Embedded MATLAB MEX

In this section...

“Using Embedded MATLAB MEX with Fixed-Point Algorithms” on page 8-9

“Speeding Up Fixed-Point Execution with the `emlmex` Function” on page 8-9

“Using Data Type Override with Embedded MATLAB MEX” on page 8-12

Using Embedded MATLAB MEX with Fixed-Point Algorithms

Embedded MATLAB MEX converts M-code to C-MEX functions that contain Embedded MATLAB subset optimizations for automatically accelerating fixed-point algorithms to compiled C code speed in MATLAB. For more information, refer to “Working with Embedded MATLAB MEX” in the Embedded MATLAB language subset documentation.

Speeding Up Fixed-Point Execution with the `emlmex` Function

The Embedded MATLAB `emlmex` function can greatly increase the execution speed of your algorithms; however, improper use of the function can also slow execution. In this example, you will use the `emlmex` function to compile different parts of a simple algorithm. By comparing the run times of the two cases, you will see the benefits and best use of the `emlmex` function.

Algorithm

The algorithm used throughout this example replicates the functionality of the MATLAB `sum` function, which sums the columns of a matrix. To see the algorithm, type open `fi_matrix_column_sum.m` at the MATLAB command line.

```
function B = fi_matrix_column_sum(A)
% Sum the columns of matrix A.
%#eml
    [m,n] = size(A);
    w = get(A,'WordLength') + ceil(log2(m));
    f = get(A,'FractionLength');
```

```

B = fi(zeros(1,n),true,w,f,fi(A));
for j = 1:n
    for i = 1:m
        B(j) = B(j) + A(i,j);
    end
end

```

Trial 1: Best Performance

The best way to speed up the execution of the algorithm is to compile the entire algorithm using the `emlmex` function. To evaluate the performance of the `emlmex` function when the entire algorithm is compiled, run the following code. The first portion of m-code executes the algorithm using only MATLAB functions. The second portion of the code compiles the entire algorithm using the Embedded MATLAB `emlmex` function. The MATLAB `tic` and `toc` functions keep track of the run times for each method of execution.

```

% MATLAB
fipref('NumericTypeDisplay','short','FimathDisplay','none');
A = fi(randn(1000,10));
tic
B = fi_matrix_column_sum(A)
t_matrix_column_sum_m = toc

% Embedded MATLAB
emlmex fi_matrix_column_sum -o fi_matrix_column_sum_x -eg {A} ...
-I [matlabroot '/toolbox/fixedpoint/fidemos']
tic
B = fi_matrix_column_sum_x(A);
t_matrix_column_sum_eml = toc

```

Trial 2: Worst Performance

Compiling only the smallest unit of computation using the `emlmex` function leads to much slower execution. In some cases, the overhead that results from calling the `emlmex` function inside a nested loop can cause even slower execution than using MATLAB functions alone. To evaluate the performance of the `emlmex` function when only the smallest unit of computation is compiled, run the following code. The first portion of m-code executes the algorithm using only MATLAB functions. The second portion of the code compiles the

smallest unit of computation with the `emlmex` function, leaving the rest of the computations to MATLAB.

```
% MATLAB
tic
[m,n] = size(A);
w = get(A,'WordLength') + ceil(log2(m));
f = get(A,'FractionLength');
B = fi(zeros(1,n),true,w,f,fimath(A));
for j = 1:n
    for i = 1:m
        B(j) = fi_scalar_sum(B(j),A(i,j));
        % B(j) = B(j) + A(i,j);
    end
end
t_scalar_sum_m = toc

% Embedded MATLAB
emlmex fi_scalar_sum -o fi_scalar_sum_x -eg {B(1),A(1,1)} ...
-I [matlabroot '/toolbox/fixedpoint/fidemos']
tic
[m,n] = size(A);
w = get(A,'WordLength') + ceil(log2(m));
f = get(A,'FractionLength');
B = fi(zeros(1,n),true,w,f,fimath(A));
for j = 1:n
    for i = 1:m
        B(j) = fi_scalar_sum_x(B(j),A(i,j));
        % B(j) = B(j) + A(i,j);
    end
end
t_scalar_sum_eml = toc
```

Ratio of Times

A comparison of Trial 1 and Trial 2 appears in the following table. Your computer may record different times than the ones the table shows, but the ratios should be approximately the same. There is an extreme difference in ratios between the trial where the entire algorithm was compiled using `emlmex` (`t_matrix_column_sum_eml`) and where only the scalar sum was

compiled (t_scalar_sum_eml). Even the M-file with no eml mex compilation (t_matrix_column_sum_m) did better than when only the smallest unit of computation was compiled using eml mex (t_scalar_sum_eml).

X (Overall Performance Rank)	Time	X/Best	X_m/X_eml
Trial 1: Best Performance			
t_matrix_column_sum_m (2)	1.99759	84.4917	84.4917
t_matrix_column_sum_eml (1)	0.0236424	1	
Trial 2: Worst Performance			
t_scalar_sum_m (4)	10.2067	431.71	2.08017
t_scalar_sum_eml (3)	4.90664	207.536	

Using Data Type Override with Embedded MATLAB MEX

Fixed-Point Toolbox software ships with a demonstration of how to generate a C-MEX function from M-code. The M-code takes the weighted average of a signal to create a lowpass filter. To run the demo, click the Fixed-Point Lowpass Filtering Using Embedded MATLAB MEX link and follow the instructions in the right pane of the Help browser.

You can specify data type override in this demo by typing an extra command at the MATLAB prompt in the “Define Fixed-Point Parameters” section of the demo. To turn data type override on, type the following command at the MATLAB prompt after running the reset(fipref) demo command in that section:

```
fipref('DataTypeOverride','TrueDoubles')
```

This command tells Fixed-Point Toolbox software to create all fi objects with type fi double. When you compile the M-file using the eml mex command in the “Compile the M-File into a MEX File” section of the demo, the resulting MEX-function uses floating-point data.

Embedded MATLAB Function Block

In this section...

“Composing a MATLAB Language Function in a Simulink Model” on page 8-13

“Using Fixed-Point Data Types with the Embedded MATLAB Function Block” on page 8-14

“Using the Embedded MATLAB Function Block with Data Type Override” on page 8-14

“Using the Model Explorer with a Fixed-Point Embedded MATLAB Function Block” on page 8-15

Composing a MATLAB Language Function in a Simulink Model

The Embedded MATLAB Function block lets you compose a MATLAB language function in a Simulink model that generates embeddable code using the Embedded MATLAB subset. When you simulate the model or generate code for a target environment, a function in an Embedded MATLAB Function block generates efficient C code. This code meets the strict memory and data type requirements of embedded target environments. In this way, Embedded MATLAB Function blocks bring the power of MATLAB for the embedded environment into Simulink.

For more information about the Embedded MATLAB Function block and the Embedded MATLAB subset, refer to the following documentation:

- Embedded MATLAB Function block reference page in the Simulink documentation
- “Using the Embedded MATLAB Function Block” in the Simulink documentation
- “Working with the Embedded MATLAB Subset” in the Embedded MATLAB documentation

Using Fixed-Point Data Types with the Embedded MATLAB Function Block

A significant number of Fixed-Point Toolbox functions are supported by the Embedded MATLAB subset. Refer to “Supported Functions and Limitations of the Fixed-Point Embedded MATLAB Subset” on page 8-2 for information about which Fixed-Point Toolbox features are supported by the Embedded MATLAB subset.

Note To simulate models using fixed-point data types in Simulink, you must have a Simulink Fixed Point license.

Using the Embedded MATLAB Function Block with Data Type Override

When you use the Embedded MATLAB Function block in a Simulink model that specifies data type override, the block determines the data type override equivalents of the input signal and parameter types and uses these to run the simulation. The following table shows how the Embedded MATLAB Function block determines the data type override equivalent from the data type of the input signal or parameter and the data type override setting in the Simulink model.

Note The Embedded MATLAB Function block does not support the Scaled doubles data type override setting.

Input Signal or Parameter Type	Data Type Override Setting	Data Type Override Equivalent
Inherited single	True doubles	fi double
	True singles	fi single
Specified single	True doubles	Built-in double
	True singles	Built-in single
Inherited double	True doubles	fi double
	True singles	fi single

Input Signal or Parameter Type	Data Type Override Setting	Data Type Override Equivalent
Specified double	True doubles	Built-in double
	True singles	Built-in single
Inherited Fixed	True doubles	fi double
	True singles	fi single
Specified Fixed	True doubles	fi double
	True singles	fi single

For more information about using the Embedded MATLAB Function block with data type override, see “Using Data Type Override with the Embedded MATLAB Function Block” in the Simulink documentation.

Using the Model Explorer with a Fixed-Point Embedded MATLAB Function Block


- “Specifying Fixed-Point Parameters in the Model Explorer” on page 8-15
- “Using fimath Objects in Embedded MATLAB Function Blocks” on page 8-18
- “Sharing Models with Fixed-Point Embedded MATLAB Function Blocks” on page 8-23
- “Example: Implementing a Fixed-Point Direct Form FIR Using the Embedded MATLAB Function Block” on page 8-25

Specifying Fixed-Point Parameters in the Model Explorer

You can specify parameters for an Embedded MATLAB Function block in a fixed-point model using the Model Explorer. Try the following:

- 1** Place an Embedded MATLAB Function block in a new model. The block is located in the Simulink User-Defined Functions library.
- 2** Open the Model Explorer by selecting **View > Model Explorer** from your model.

- Expand the **untitled*** node in the **Model Hierarchy** pane of the Model Explorer and select the **Embedded MATLAB Function** node. The Model Explorer now appears as follows:



The parameters in the **Simulink input signal properties** group box in the **Dialog** pane apply to Embedded MATLAB Function blocks in models that use fixed-point data types.

FIMATH for fixed-point input signals

Defines the `fimath` object to be associated with Simulink fixed-point or integer signals entering the Embedded MATLAB Function block as inputs. You must enter an expression that evaluates to a `fimath` object. You can do so in one of two ways:

- Fully define the `fimath` object in the parameter value box using Fixed-Point Toolbox MATLAB code.
- Enter a variable name of a `fimath` object that is defined in the MATLAB or model workspace.

The default `fimath` object entered for this parameter emulates C-style math.

FIMATH for `fi` and `fimath` constructors

Specifies the `fimath` object properties to be associated with all `fi` and `fimath` objects constructed within the Embedded MATLAB Function block. This setting applies only to unspecified `fimath` object properties.

- Selecting Same as FIMATH for fixed-point input signals assigns all unspecified `fimath` object properties according to the **FIMATH for fixed-point input signals** parameter.
- Selecting MATLAB factory default assigns all unspecified `fimath` object properties as follows:

```
RoundMode: nearest
OverflowMode: saturate
ProductMode: FullPrecision
MaxProductWordLength: 128
SumMode: FullPrecision
MaxSumWordLength: 128
CastBeforeSum: true
```

Note You can also use the MATLAB *user-configured* default `fimath` object as the **FIMATH for fixed-point input signals**. For more information, see “Specifying a User-Defined `fimath` Object in the Embedded MATLAB Function Block” on page 8-20 in the Fixed-Point Toolbox documentation.

Treat these inherited signal types as fi objects

Choose whether to treat inherited fixed-point and integer signals as `fi` objects.

- When you select `Fixed-point`, all fixed-point inputs to the Embedded MATLAB Function block are treated as Fixed-Point Toolbox `fi` objects.
- When you select `Fixed-Point & Integer`, all fixed-point and integer inputs to the Embedded MATLAB Function block are treated as Fixed-Point Toolbox `fi` objects.

Using `fimath` Objects in Embedded MATLAB Function Blocks

The `FIMATH` for fixed-point input signals and `FIMATH` for `fi` and `fimath` constructors parameters allow you to associate one `fimath` object with input signals, and another (possibly different) `fimath` object with any `fi` and `fimath` objects that are constructed within the block.

- **`FIMATH` for fixed-point input signals** — Specifies the default `fimath` properties to be associated with all fixed-point input signals
- **`FIMATH` for `fi` and `fimath` constructors** — Specifies the default `fimath` properties to be associated with all unspecified `fimath` properties of any `fi` or `fimath` object constructed within the Embedded MATLAB Function block

Note To ensure that all fixed-point and integer arithmetic operations are carried out correctly within the Embedded MATLAB Function block, you must set these two `fimath` settings appropriately.

To learn about the possible settings of these parameters, see “Specifying Fixed-Point Parameters in the Model Explorer” on page 8-15.

You can set these parameters on the following dialog box, which is accessible through both “The Model Explorer” and the “Ports and Data Manager”:

Embedded MATLAB Function: Embedded MATLAB Function

Name: [Embedded MATLAB Function](#)

Update method: Sample Time:

Lock Editor

Saturate on integer overflow.

Simulink input signal properties

FIMATH for fixed-point input signals:

```
fimath(...  
'RoundMode', 'floor',...  
'OverflowMode', 'wrap',...  
'ProductMode', 'KeepLSB', 'ProductWordLength', 32,...  
'SumMode', 'KeepLSB', 'SumWordLength', 32,...  
'CastBeforeSum', true)
```

FIMATH for fi and fimath constructors

Treat these inherited signal types as fi objects:

Description:

[Document Link:](#)

- To access this pane through the Model Explorer, select **View > Model Explorer** from your model menu, and then select the Embedded MATLAB Function block from the Model Hierarchy pane on the left side of the Model Explorer.
- To access this pane through the Ports and Data Manager, select **Tools > Edit Data/Ports** from the Embedded MATLAB Editor menu.

For considerations in setting these parameters, see the following sections:

- “Specifying a User-Defined fimath Object in the Embedded MATLAB Function Block” on page 8-20
- “Specifying the MATLAB Factory Default fimath Object in the Embedded MATLAB Function Block” on page 8-22

Specifying a User-Defined fimath Object in the Embedded MATLAB Function Block.

The **FIMATH for fixed-point input signals** edit box allows you to specify a user-configured fimath object to be associated with all fixed-point input signals. You can construct this fimath object in one of two ways:

- Creating the fimath object inside the **FIMATH for fixed-point input signals** edit box
- Creating the fimath object in the MATLAB or model workspace and then entering its variable name in the **FIMATH for fixed-point input signals** edit box.

Note If the fimath object that you want to specify is also the MATLAB default fimath object, you can simply type `fimath` in the edit box. For more information on the user-configurable MATLAB default fimath object, see “Configuring the MATLAB Default fimath Object” on page 4-5.

If you want a user-configured fimath object to be used as the **FIMATH for fi and fimath constructors**, it must be the same fimath object that is specified as the **FIMATH for fixed-point input signals**. To do this, you must perform the following steps:

- 1 Configure the desired MATLAB default `fimath` object in your MATLAB workspace:


```
Fdefault = fimath('RoundMode','Floor','OverflowMode','Wrap')
```

```
Fdefault =
```


```
          RoundMode: floor
          OverflowMode: wrap
          ProductMode: FullPrecision
MaxProductWordLength: 128
          SumMode: FullPrecision
MaxSumWordLength: 128
          CastBeforeSum: true
```

```
setdefaultfimath(Fdefault);
```

- 2 Specify the user-configured default `fimath` object as the **FIMATH for fixed-point input signals** by typing `fimath` in the edit box:



- 3 Select **Same as FIMATH for fixed-point input signals** from the **FIMATH for fi and fimath constructors** menu, and click **Apply**:



Now, all fixed-point input signals to the Embedded MATLAB Function block and all `fi` and `fimath` objects constructed within the block will have the same `fimath` object properties as the user-configured MATLAB default.

Specifying the MATLAB Factory Default `fimath` Object in the Embedded MATLAB Function Block. The `FIMATH for fi and fimath constructors` parameter gives you the option to use the MATLAB factory default `fimath` object at anytime. Selecting this option forces the block to use the MATLAB factory default `fimath` object for all `fi` and `fimath` objects constructed in the Embedded MATLAB Function block, regardless of the current MATLAB default `fimath` object.

If you also want to use the MATLAB factory default `fimath` object as the **FIMATH for fixed-point input signals**, you have the following options:

- If the current MATLAB default `fimath` object is the factory default, you can simply type `fimath` in the **FIMATH for fixed-point input signals** edit box.
- If the MATLAB default `fimath` object is not set to the factory default, you must either:
 - Construct the MATLAB factory default `fimath` object in the **FIMATH for fixed-point input signals** edit box
 - Construct the MATLAB factory default `fimath` object at the MATLAB command line and then specify its variable name in the **FIMATH for fixed-point input signals** edit box

If the arithmetic operations inside the Embedded MATLAB Function block do not require the `fimath` of the fixed-point input signals to match the `fimath` of the `fi` and `fimath` objects constructed within the block, you can specify any `fimath` object in the **FIMATH for fixed-point input signals** edit box.

Sharing Models with Fixed-Point Embedded MATLAB Function Blocks

Sometimes you might need to share a fixed-point model using the Embedded MATLAB Function block with a coworker. When you do, make sure to move any variables you define in the MATLAB workspace, including `fimath` objects, to the model workspace. For example, try the following:


- 1 Place an Embedded MATLAB Function block in a new model. The block is located in the Simulink User-Defined Functions library.
- 2 Define a `fimath` object in the MATLAB workspace that you want to use for any Simulink fixed-point signal entering the Embedded MATLAB Function block as an input:

```
F = fimath('RoundMode','Floor','OverflowMode','Wrap',...
          'ProductMode','KeepLSB','ProductWordLength',32,...
          'SumMode','KeepLSB','SumWordLength',32)
```

```
F =
```

```
RoundMode: floor
OverflowMode: wrap
ProductMode: KeepLSB
ProductWordLength: 32
SumMode: KeepLSB
SumWordLength: 32
CastBeforeSum: true
```

- 3** Open the Model Explorer by selecting **View > Model Explorer** from your model.
- 4** Expand the **untitled*** node in the **Model Hierarchy** pane of the Model Explorer and select the **Embedded MATLAB Function** node.
- 5** Enter the variable **F** into the **FIMATH for fixed-point input signals** parameter on the **Dialog** pane and click **Apply**. You have now defined the **fimath** object for any Simulink fixed-point signal entering the Embedded MATLAB Function block as an input.
- 6** Select the **Base Workspace** node in the **Model Hierarchy** pane. You can see the variable **F** that you have defined in the MATLAB workspace listed in the **Contents** pane. If you were to send this model to a coworker, that coworker would have to define that same variable in the MATLAB workspace to get the same results as you with this model.
- 7** Cut the variable **F** from the base workspace and paste it into the model workspace listed under the node for your model, in this case **untitled***. The Model Explorer now looks like this:



You can now e-mail your model to a coworker, and because the variables needed to run the model are included in the workspace of the model itself, your coworker can run the model and get the correct results without performing any extra steps.

Example: Implementing a Fixed-Point Direct Form FIR Using the Embedded MATLAB Function Block

The following sections lead you through creating a fixed-point, low-pass, direct form FIR filter in Simulink using Fixed-Point Toolbox software and the Embedded MATLAB Function block. You will perform the following tasks in the sequence shown:

- “Program the Embedded MATLAB Function Block” on page 8-26
- “Prepare the Inputs” on page 8-27
- “Create the Model” on page 8-28
- “Define the fimath Object Using the Model Explorer” on page 8-32
- “Run the Simulation” on page 8-34

Program the Embedded MATLAB Function Block.

- 1** Place an Embedded MATLAB Function block in a new model. The block is located in the Simulink User-Defined Functions library.
- 2** Save your model as `eML_fi.mdl`.
- 3** Double-click the Embedded MATLAB Function block in your model to open the Embedded MATLAB Editor. Type or copy and paste the following MATLAB code, including comments, into the Editor:

```
function [yout,zf] = dffirdemo(b, x, zi)
%eML_fi doc model example
%Initialize the output signal yout and the final conditions zf
Ty = numerictype(1,12,8);
yout = fi(zeros(size(x)), 'numerictype', Ty);
zf = zi;

% FIR filter code
for k=1:length(x);
    % Update the states: z = [x(k);z(1:end-1)]
    zf(:) = [x(k);zf(1:end-1)];
    % Form the output: y(k) = b*z
    yout(k) = b*zf;
end

% Plot the outputs only in simulation.
% This does not generate C code.
eml.extrinsic('figure');
eml.extrinsic('subplot');
eml.extrinsic('plot');
eml.extrinsic('title');
eml.extrinsic('grid');
```

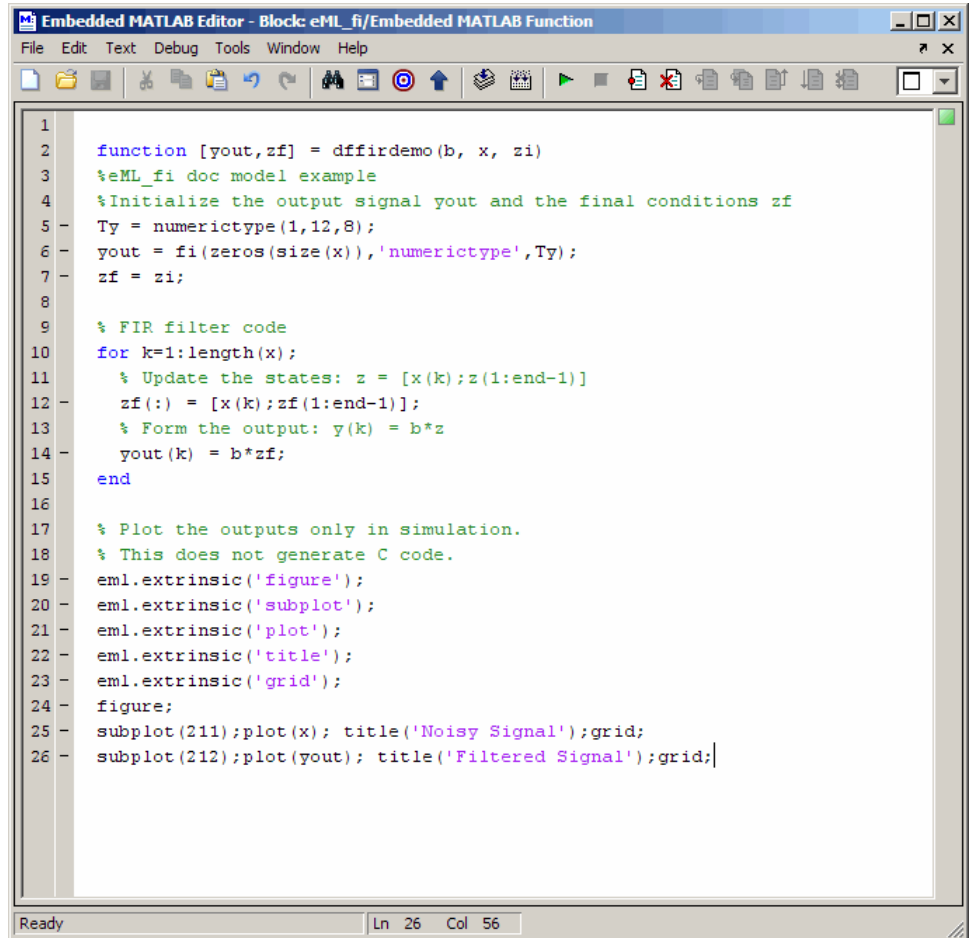


```

figure;
subplot(211);plot(x); title('Noisy Signal');grid;
subplot(212);plot(yout); title('Filtered Signal');grid;

```

The Editor should now appear as follows:



The screenshot shows the Embedded MATLAB Editor window titled "Embedded MATLAB Editor - Block: eML_fi/Embedded MATLAB Function". The editor contains the following code:

```

1
2 function [yout,zf] = dffirdemo(b, x, zi)
3 %eML_fi doc model example
4 %Initialize the output signal yout and the final conditions zf
5 Ty = numerictype(1,12,8);
6 yout = fi(zeros(size(x)), 'numerictype', Ty);
7 zf = zi;
8
9 % FIR filter code
10 for k=1:length(x);
11     % Update the states: z = [x(k);z(1:end-1)]
12     zf(:) = [x(k);zf(1:end-1)];
13     % Form the output: y(k) = b*z
14     yout(k) = b*zf;
15 end
16
17 % Plot the outputs only in simulation.
18 % This does not generate C code.
19 eml.extrinsic('figure');
20 eml.extrinsic('subplot');
21 eml.extrinsic('plot');
22 eml.extrinsic('title');
23 eml.extrinsic('grid');
24 figure;
25 subplot(211);plot(x); title('Noisy Signal');grid;
26 subplot(212);plot(yout); title('Filtered Signal');grid;

```

The status bar at the bottom of the window shows "Ready" and "Ln 26 Col 56".

Prepare the Inputs. Define the filter coefficients b , noise x , and initial conditions z_i by typing the following at the MATLAB command line:

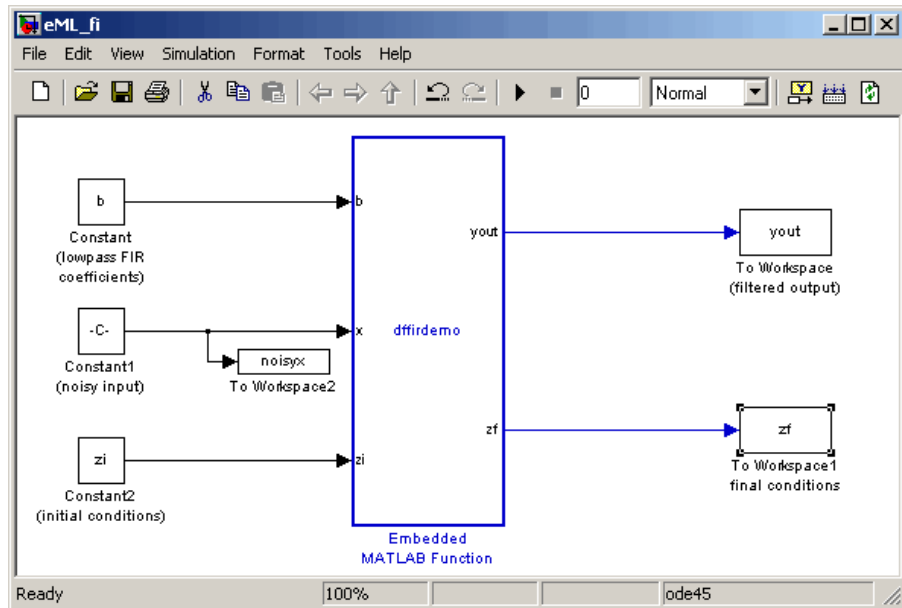
```
b = fi_fir_coefficients;
```

```

load mtlb
x = mtlb;
n = length(x);
noise = sin(2*pi*2140*(0:n-1)'./Fs);
x = x + noise;
zi = zeros(length(b),1);
    
```

Create the Model.

- 1 Add blocks to your model to create the system shown below.



- 2 Set the block parameters in the model to the following values:

Block	Parameter	Value
Constant	Constant value	b
	Interpret vector parameters as 1-D	Unselected
	Sampling mode	Sample based
	Sample time	inf
	Mode	Fixed point
	Sign	Signed
	Scaling	Slope and bias
	Word length	12
	Slope	2^{-12}
	Bias	0
Constant1	Constant value	x+noise
	Interpret vector parameters as 1-D	Unselected
	Sampling mode	Sample based
	Sample time	1
	Mode	Fixed point
	Sign	Signed
	Scaling	Slope and bias
	Word length	12
	Slope	2^{-8}
	Bias	0

Block	Parameter	Value
Constant2	Constant value	zi
	Interpret vector parameters as 1-D	Unselected
	Sampling mode	Sample based
	Sample time	inf
	Mode	Fixed point
	Sign	Signed
	Scaling	Slope and bias
	Word length	12
	Slope	2 ⁻⁸
	Bias	0
To Workspace	Variable name	yout
	Limit data points to last	inf
	Decimation	1
	Sample time	-1
	Save format	Array
	Log fixed-point data as a fi object	Selected

Block	Parameter	Value
To Workspace1	Variable name	zf
	Limit data points to last	inf
	Decimation	1
	Sample time	1
	Save format	Array
	Log fixed-point data as a fi object	Selected
To Workspace2	Variable name	noisyx
	Limit data points to last	inf
	Decimation	1
	Sample time	1
	Save format	Array
	Log fixed-point data as a fi object	Selected

- 3** From the model menu, select **Simulation > Configuration Parameters** and set the following parameters:

Parameter	Value
Stop time	0
Type	Variable-step
Solver	Discrete (no continuous states)

Click **Apply** to save your changes.

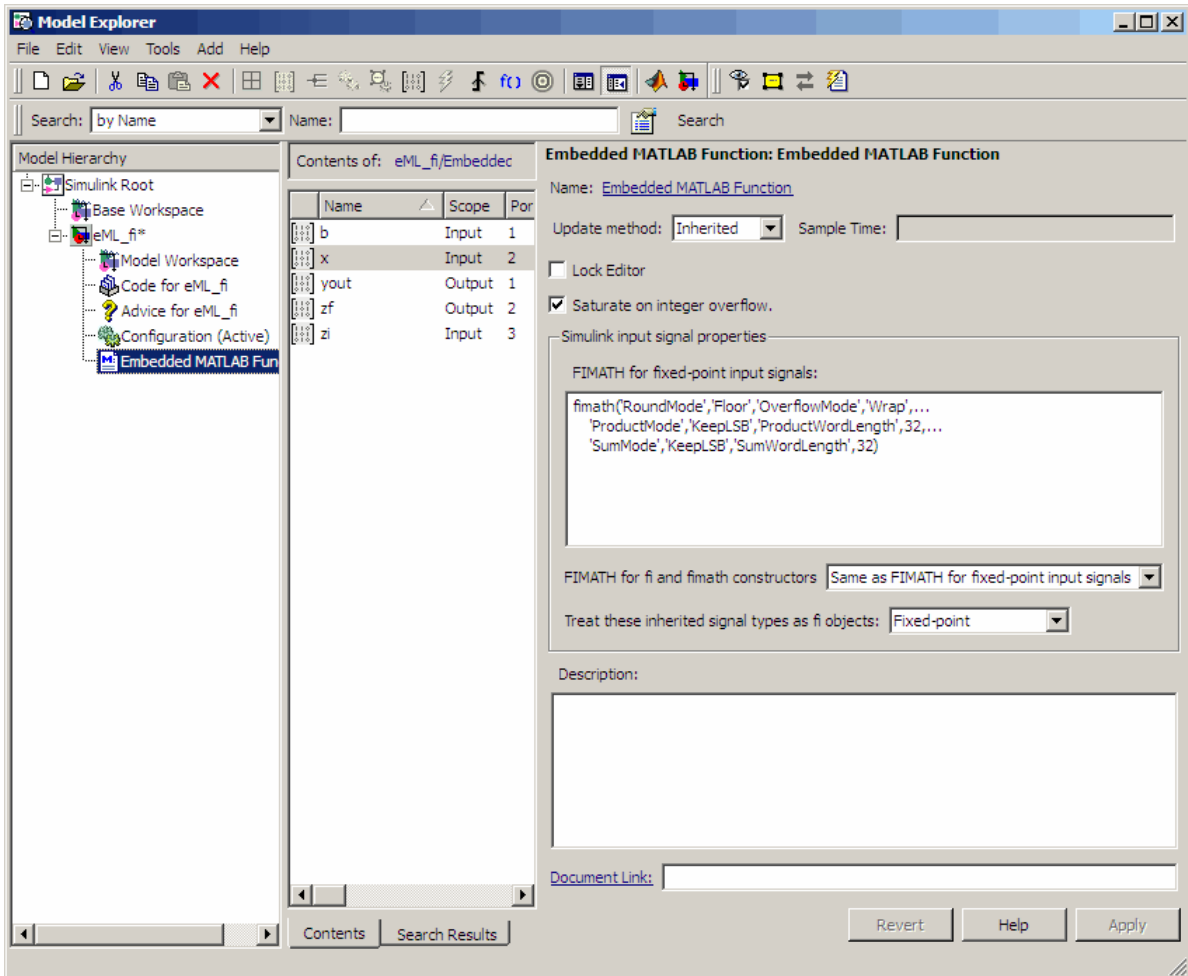
Define the fimath Object Using the Model Explorer.

- 1 Open the Model Explorer for the model by selecting **View > Model Explorer**.
- 2 Click the **eML_fi > Embedded MATLAB Function** node in the **Model Hierarchy** pane. The dialog for the Embedded MATLAB Function block appears in the **Dialog** pane of the Model Explorer.
- 3 Enter the following **fimath** object in the **FIMATH for fixed-point input signals** edit box on the Embedded MATLAB Function block dialog:

```
fimath('RoundMode','Floor','OverflowMode','Wrap',...  
      'ProductMode','KeepLSB','ProductWordLength',32,...  
      'SumMode','KeepLSB','SumWordLength',32)
```

The **fimath** object you define here is assigned to the three inputs to the Embedded MATLAB Function block.

- 4 Set the **FIMATH for fi and fimath constructors** parameter to **Same** as FIMATH for fixed-point input signals and click **Apply**. The Model Explorer now appears as follows:



Now, all `fi` and `fimath` objects that you construct within the Embedded MATLAB Function block are assigned the same `fimath` object properties as the three inputs to the block.

For more information on specifying `fimath` objects in Embedded MATLAB Function blocks, see the “Using `fimath` Objects in Embedded MATLAB Function Blocks” on page 8-18 section of this chapter.

Run the Simulation.

- 1** Run the simulation by selecting your model and typing **Ctrl+T**. While the simulation is running, information outputs to the MATLAB command line. You can look at the plots of the noisy signal and the filtered signal.
- 2** Now build embeddable C code for your model by selecting the model and typing **Ctrl+B**. While the code is building, information outputs to the MATLAB command line. A directory called `eML_fi_grt_rtw` is created in your current working directory.
- 3** Navigate to `eML_fi_grt_rtw > eML_fi.c`. In this file you can see the code that has been generated from your model. Search on the comment in your code

```
%eML_fi doc model example
```

This brings you to the beginning of the section of the code that is generated from your Embedded MATLAB Function block.

Embedded MATLAB Coder

Embedded MATLAB Coder is a Real-Time Workshop® function (`emlc`) that automatically converts M-code directly to C code. It lets you accelerate MATLAB code that uses Fixed-Point Toolbox software. For more information, refer to “Working with Embedded MATLAB Coder” in the Real-Time Workshop product documentation.

Working with Fixed-Point Compilation Reports

In this section...

“Generating the Compilation Report” on page 8-36

“Opening the Compilation Report” on page 8-37

“Viewing Your M-Code” on page 8-37

“Viewing Variables in the Variables Tab” on page 8-39

“See Also” on page 8-40

Generating the Compilation Report

When you compile your M-code with the Embedded MATLAB `emlmex` function or the Real-Time Workshop `emlc` function, you can use the `-report` option to generate a compilation report. This compilation report allows you to examine the data types of the variables and expressions in your code.

To see an example of the compilation report generated by the `emlmex` function, compile `cordic_atan_fixpt.m`. This file ships as a part of the Fixed-Point ATAN2 Calculation demo. You can open the M-file by typing the following at the MATLAB command line:

```
open cordic_atan_fixpt
```

To compile the `cordic_atan_fixpt.m` file, you must provide inputs `x`, `y`, `N`, and `angleLUT`. This example uses the following input values:

```
x = fi(0.23);  
y = x;  
N = 12;  
Tz = numerictype(1,16,13);  
angleLUT = fi(atan(2.^(0:N-1)), 'NumericType', Tz);
```

After you define the input variables in the MATLAB workspace, change your working directory to a local folder and compile the M-file using `emlmex`. Use the `-report` option to generate the compilation report:

```
emlmex cordic_atan_fixpt -eg {x,y,N,angleLUT} -report
```

Opening the Compilation Report

If the compilation is successful, you receive the following message:

C-MEX generation successful: [Open coder report.](#)

Click the `Open coder report` link to view the report.

If the compilation fails, a link to the error report appears:

C-MEX generation failed: [Open error report.](#)

Click the `Open error report` link to view the error report and debug your code. For more information on working with error reports, see “Working with Compilation Reports” in the Embedded MATLAB subset documentation.

Viewing Your M-Code

When the compilation report opens, you can hover your cursor over the variables and expressions in your M-code to see their data type information. The compilation report provides color-coded data type information according to the following legend.

Color	Meaning
Green	Data type information is available for the selected variable at this location in the code.
Orange	There is a warning message associated with the selected variable or expression.
Pink	No data type information is available for the selected variable.
Purple	Data type information is available for the selected expression at this location in the code.
Red	There is an error message associated with the selected variable or expression.

Variables in your M-code that have data type information available appear highlighted in green, as shown in the following figure.

```

Function: cordic_atan_fixpt (h:/Documents/My Models/cordic_atan_fixpt.m)
1 function [z,x,y] = cordic_atan_fixpt(y,x,N,angleLUT) %#eml
2 % Calculate arctangent in range [-pi/2, pi/2] using Vectoring mode CORDIC
3 % algorithm. Both a and y inputs must be real scalar, x must >= 0.
4 % Full precision Fimath is used in all fixed-point operations
5 %
6 % Inputs:
7 % y : y coordinate or imaginary part of the input vector
8 % x : x coordinate or real part of the input vector
9 % N : total number of iterations, must be a non-negative integer
10 % angleLUT : the angle Llook-up table, has same numerictype as the output
11 % angle
12 % Output:
13 % z : angle that equals atan2(y,x), in radius
14 % the output angle range is within [-pi/2, +pi/2]
15 % x : x coordinate of the last vector at the end of the iterations
16 % y : y coordinate of the last vector at the end of the iterations
17 %
18 % Copyright 1984-2008 The MathWorks, Inc.
19 % $Revision: 1.1.4.1 $ $Date: 2008/11/19 14:40:34 $
20
21 % initialization
22 tmp = fi(0, numerictype(x)); % Temporary register
23 z = fi(0, numerictype(angleLUT)); % assum z_{0} is 0
24
25 for i = 0:N-1,
26     if y < 0 % negative y leads to counter clock-wise rotation
27         tmp(:) = x - bitsra(y,i); % x_{i+1} = x_{i} - y_{i}>>i
28         y(:) = y + bitsra(x,i); % y_{i+1} = y_{i} + x_{i}>>i
29         x(:) = tmp;
30         z(:) = z - angleLUT(i);
31     else % positive y leads
32         tmp(:) = x + bitsra(y,i); % x_{i+1} = x_{i} + y_{i}>>i
33         y(:) = y - bitsra(x,i); % y_{i+1} = y_{i} - x_{i}>>i
34         x(:) = tmp;
35         z(:) = z + angleLUT(i);
36     end
37 end
38 end

```

Information for the selected variable: `atan(2*(-1))`

Size	1x1
Complex	Noise rotation
Class	embedded.fi x_{i} + y_{i}>>i
Signed	Yes = y_{i} - x_{i}>>i
WL	16
FL	17+1 = z_{i} - atan(2*(-1))

Expressions in your M-code that have data type information available appear highlighted in purple, as the next figure shows.

```

21  % initialization
22  tmp = fi(0, numerictype(x)); % Temporary register
23  z = fi(0, numerictype(angleLUT)); % assum z_{0} is 0
24
25  for i = 1:10
26      if y < 0 % negative y leads to counter clock-wise rotation
27          Complex No x = bitsra(y,i); % x_{i+1} = x_{i} - y_{i}>>i
28          Class embedded.fi y = x + bitsra(x,i); % y_{i+1} = y_{i} + x_{i}>>i
29          Signed Yes;
30          WL = 16 - angleLUT(i+1); % z_{i+1} = z_{i} + atan(2^{-i})
31          FL = 13;
32      else % positive y leads to clock-wise rotation
33          tmp(:) = x + bitsra(y,i); % x_{i+1} = x_{i} + y_{i}>>i
34          y(:) = y - bitsra(x,i); % y_{i+1} = y_{i} - x_{i}>>i
35          x(:) = tmp;
36          z(:) = z + angleLUT(i+1); % z_{i+1} = z_{i} - atan(2^{-i})
37      end
38  end

```

Information for the selected expression:

Complex	No
Class	embedded.fi
Signed	Yes
WL	16
FL	13

Viewing Variables in the Variables Tab

To see the data type information for all the variables in your M-file, click on the **Variables** tab of the compilation report. You can expand all **fi** and **fi** objects listed in the Variables tab to display the properties of the associated **fi** object, as shown in the following figure.

Summary	All Messages (0)	Variables						
Variable	Type	Size	Complex	Class	Signed	WL	FL	
z	Output variable	1x1	No	embedded.fi	Yes	16	13	
Fimath:								
Round mode: nearest				Sum mode: FullPrecision				
Overflow mode: saturate				Maximum sum word length: 128				
Product mode: FullPrecision				Cast before sum: Yes				
Maximum product word length: 128								
x	Output variable	1x1	No	embedded.fi	Yes	16	17	
y	Output variable	1x1	No	embedded.fi	Yes	16	17	
N	Input variable	1x1	No	double	-	-	-	
angleLUT	Input variable	1x12	No	embedded.fi	Yes	16	13	
tmp	Local variable	1x1	No	embedded.fi	Yes	16	17	
i	Local variable	1x1	No	double	-	-	-	

You can sort the variables by clicking the column headings in the **Variables** tab. To sort the variables by multiple columns, hold down the **Shift** key while clicking on the column headings.

See Also

For more information about using the compilation report with the `emlmex` function, see “Working with Compilation Reports” in the Embedded MATLAB documentation.

For information about using the compilation report with the `emlc` function, see “Working with Compilation Reports” in the Real-Time Workshop documentation.

Interoperability with Other Products

- “Using fi Objects with Simulink” on page 9-2
- “Using the Embedded MATLAB Subset” on page 9-8
- “Using fi Objects with Signal Processing Blockset Software” on page 9-9
- “Using fi Objects with Filter Design Toolbox Software” on page 9-14

Using fi Objects with Simulink

In this section...

“Reading Fixed-Point Data from the Workspace” on page 9-2

“Writing Fixed-Point Data to the Workspace” on page 9-2

“Setting the Value and Data Type of Block Parameters” on page 9-6

“Logging Fixed-Point Signals” on page 9-6

“Accessing Fixed-Point Block Data During Simulation” on page 9-6

Reading Fixed-Point Data from the Workspace

You can read fixed-point data from the MATLAB workspace into a Simulink model via the From Workspace block. To do so, the data must be in a structure format with a **fi** object in the **values** field. In array format, the From Workspace block only accepts real, double-precision data.

To read in **fi** data, the **Interpolate data** parameter of the From Workspace block must not be selected, and the **Form output after final data value by** parameter must be set to anything other than Extrapolation.

Writing Fixed-Point Data to the Workspace

You can write fixed-point output from a model to the MATLAB workspace via the To Workspace block in either array or structure format. Fixed-point data written by a To Workspace block to the workspace in structure format can be read back into a Simulink model in structure format by a From Workspace block.

Note To write fixed-point data to the MATLAB workspace as a **fi** object, select the **Log fixed-point data as a fi object** check box on the To Workspace block dialog. Otherwise, fixed-point data is converted to **double** and written to the workspace as **double**.

For example, you can use the following code to create a structure in the MATLAB workspace with a `fi` object in the `values` field. You can then use the From Workspace block to bring the data into a Simulink model.

```
a = fi([sin(0:10)' sin(10:-1:0)'])
```

```
a =
```

```
      0   -0.5440
  0.8415   0.4121
  0.9093   0.9893
  0.1411   0.6570
 -0.7568  -0.2794
 -0.9589  -0.9589
 -0.2794  -0.7568
  0.6570   0.1411
  0.9893   0.9093
  0.4121   0.8415
 -0.5440      0
```

```
      DataTypeMode: Fixed-point: binary point scaling
      Signedness: Signed
      WordLength: 16
      FractionLength: 15
```

```
      RoundMode: nearest
      OverflowMode: saturate
      ProductMode: FullPrecision
  MaxProductWordLength: 128
      SumMode: FullPrecision
  MaxSumWordLength: 128
      CastBeforeSum: true
```

```
s.signals.values = a
```

```
s =
```

```
signals: [1x1 struct]
```

```
s.signals.dimensions = 2

s =

    signals: [1x1 struct]

s.time = [0:10]'

s =


    signals: [1x1 struct]
    time: [11x1 double]
```

The From Workspace block in the following model has the `fi` structure `s` in the **Data** parameter.

Remember, to write fixed-point data to the MATLAB workspace as a `fi` object, select the **Log fixed-point data as a fi object** check box on the To Workspace block dialog. Otherwise, fixed-point data is converted to `double` and written to the workspace as `double`.

In the model, the following parameters in the **Solver** pane of the **Configuration Parameters** dialog have the indicated settings:

- **Start time** — 0.0
- **Stop time** — 10.0
- **Type** — Fixed-step
- **Solver** — Discrete (no continuous states)
- **Fixed step size (fundamental sample time)** — 1.0



The To Workspace block writes the result of the simulation to the MATLAB workspace as a `fi` structure.

```
simout.signals.values
```

```
ans =
```

```

      0   -8.7041
 13.4634   6.5938
 14.5488  15.8296
   2.2578  10.5117
 -12.1089  -4.4707
 -15.3428 -15.3428
  -4.4707 -12.1089
 10.5117   2.2578
 15.8296  14.5488
   6.5938  13.4634
 -8.7041    0

```

```
DataTypeMode: Fixed-point: binary point scaling
Signedness: Signed
WordLength: 32
FractionLength: 25

RoundMode: nearest
OverflowMode: saturate
ProductMode: FullPrecision
MaxProductWordLength: 128
SumMode: FullPrecision
MaxSumWordLength: 128
CastBeforeSum: true
```

Setting the Value and Data Type of Block Parameters

You can use Fixed-Point Toolbox expressions to specify the value and data type of block parameters in Simulink. Refer to “Block Support for Data and Numeric Signal Types” in the Simulink documentation for more information.

Logging Fixed-Point Signals

When fixed-point signals are logged to the MATLAB workspace via signal logging, they are always logged as `fi` objects. To enable signal logging for a signal, select the **Log signal data** option in the signal’s **Signal Properties** dialog box. For more information, refer to “Logging Signals” in the Simulink documentation.

When you log signals from a referenced model or Stateflow® chart in your model, the word lengths of `fi` objects may be larger than you expect. The word lengths of fixed-point signals in referenced models and Stateflow charts are logged as the next largest data storage container size.

Accessing Fixed-Point Block Data During Simulation

Simulink provides an application program interface (API) that enables programmatic access to block data, such as block inputs and outputs, parameters, states, and work vectors, while a simulation is running. You can use this interface to develop MATLAB programs capable of accessing block data while a simulation is running or to access the data from the MATLAB command line. Fixed-point signal information is returned to you via this API

as `fi` objects. For more information on the API, refer to “Accessing Block Data During Simulation” in the Simulink documentation.

Using the Embedded MATLAB Subset

There are several ways to use Fixed-Point Toolbox software with the Embedded MATLAB subset:

- The Embedded MATLAB `em1mex` function converts your M-code to a C-MEX function and can greatly accelerate the execution speed of your fixed-point algorithms in MATLAB. To use the `em1mex` function, you need a Fixed-Point Toolbox license. For more information, see “Working with Embedded MATLAB MEX” in the Embedded MATLAB documentation.
- The Embedded MATLAB Coder (`em1c`) is a Real-Time Workshop function that automatically converts M-code to C code. Using the Embedded MATLAB Coder allows you to accelerate your MATLAB code that uses Fixed-Point Toolbox software. To use the `em1c` function with Fixed-Point Toolbox software, you also need to have licenses for Simulink and Real-Time Workshop code generation software. For more information, see “Working with Embedded MATLAB Coder” in the Real-Time Workshop documentation.
- The Embedded MATLAB Function block allows you to use MATLAB code in your Simulink models that generate embeddable C code. To use the Embedded MATLAB Function block with Fixed-Point Toolbox software, you also need a Simulink license. For more information on the Embedded MATLAB Function block, see “Using the Embedded MATLAB Function Block” in the Simulink documentation.

For more information on using the Embedded MATLAB subset with Fixed-Point Toolbox software, see Chapter 8, “Working with the Fixed-Point Embedded MATLAB Subset”.

Using fi Objects with Signal Processing Blockset Software

In this section...

“Reading Fixed-Point Signals from the Workspace” on page 9-9

“Writing Fixed-Point Signals to the Workspace” on page 9-9

Reading Fixed-Point Signals from the Workspace

You can read fixed-point data from the MATLAB workspace into a Simulink model using the Signal From Workspace and Triggered Signal From Workspace blocks from Signal Processing Blockset software. Enter the name of the defined `fi` variable in the **Signal** parameter of the Signal From Workspace or Triggered Signal From Workspace block.

Writing Fixed-Point Signals to the Workspace

Fixed-point output from a model can be written to the MATLAB workspace via the Signal To Workspace or Triggered To Workspace block from the blockset. The fixed-point data is always written as a 2-D or 3-D array.

Note To write fixed-point data to the MATLAB workspace as a `fi` object, select the **Log fixed-point data as a fi object** check box on the Signal To Workspace or Triggered To Workspace block dialog. Otherwise, fixed-point data is converted to `double` and written to the workspace as `double`.

For example, you can use the following code to create a `fi` object in the MATLAB workspace. You can then use the Signal From Workspace block to bring the data into a Simulink model.

```
a = fi([sin(0:10)' sin(10:-1:0)'])
```

```
a =
```

```
      0   -0.5440
  0.8415   0.4121
  0.9093   0.9893
  0.1411   0.6570
 -0.7568  -0.2794
 -0.9589  -0.9589
 -0.2794  -0.7568
  0.6570   0.1411
  0.9893   0.9093
  0.4121   0.8415
 -0.5440      0
```

```
      DataTypeMode: Fixed-point: binary point scaling
      Signedness: Signed
      WordLength: 16
      FractionLength: 15
```

```
      RoundMode: nearest
      OverflowMode: saturate
      ProductMode: FullPrecision
      MaxProductWordLength: 128
      SumMode: FullPrecision
      MaxSumWordLength: 128
      CastBeforeSum: true
```

The Signal From Workspace block in the following model has these settings:


- **Signal** — a
- **Sample time** — 1
- **Samples per frame** — 2

- **Form output after final data value by** — Setting to zero

The following parameters in the **Solver** pane of the **Configuration Parameters** dialog have these settings:

- **Start time** — 0.0
- **Stop time** — 10.0
- **Type** — Fixed-step
- **Solver** — Discrete (no continuous states)
- **Fixed step size (fundamental sample time)** — 1.0

Remember, to write fixed-point data to the MATLAB workspace as a `fi` object, select the **Log fixed-point data as a fi object** check box on the Signal To Workspace block dialog. Otherwise, fixed-point data is converted to double and written to the workspace as double.



The Signal To Workspace block writes the result of the simulation to the MATLAB workspace as a `fi` object.

yout =

(:, :, 1) =

0.8415	-0.1319
-0.8415	-0.9561

(:, :, 2) =

1.0504	1.6463
0.7682	0.3324

(:, :, 3) =

-1.7157	-1.2383
0.2021	0.6795

(:, :, 4) =

0.3776	-0.6157
-0.9364	-0.8979

(:, :, 5) =

1.4015	1.7508
0.5772	0.0678

(:, :, 6) =

-0.5440	0
-0.5440	0

```
    DataTypeMode: Fixed-point: binary point scaling
      Signedness: Signed
      WordLength: 17
    FractionLength: 15

      RoundMode: nearest
      OverflowMode: saturate
      ProductMode: FullPrecision
    MaxProductWordLength: 128
      SumMode: FullPrecision
    MaxSumWordLength: 128
    CastBeforeSum: true
```

Using `fi` Objects with Filter Design Toolbox Software

When the `Arithmetic` property is set to `'fixed'`, you can use an existing `fi` object as the input, states, or coefficients of a `dfilt` object in Filter Design Toolbox software. Also, fixed-point filters in the toolbox return `fi` objects as outputs. Refer to the Filter Design Toolbox software documentation for more information.

A

- ANSI C
 - compared with `fi` objects 2-22
- arithmetic
 - fixed-point 4-15
 - with [Slope Bias] signals 4-19
- arithmetic operations
 - fixed-point 2-10

B

- binary conversions 2-25

C

- casts
 - fixed-point 2-19
- complex multiplication
 - fixed-point 2-13

D

- data type override 5-12
- demos 1-9
- display preferences
 - setting 5-5
- display settings 1-7

E

- Embedded MATLAB
 - fixed-point 8-1
- Embedded MATLAB Function block
 - using with Model Explorer and fixed-point models 8-15

F

- `fi` objects
 - constructing 3-2
- `fi`math objects 2-16

- constructing 4-2
- properties
 - setting in the Model Explorer 4-13
 - setting properties in the Model Explorer 4-13

`fi`pref objects

- constructing 5-2
- fixed-point arithmetic 4-15
- fixed-point data
 - reading from workspace 9-2
 - writing to workspace 9-2

fixed-point data types

- addition 2-12
- arithmetic operations 2-10
- casts 2-19
- complex multiplication 2-13
- modular arithmetic 2-10
- multiplication 2-13
- overflow handling 2-5
- precision 2-5
- range 2-5
- rounding 2-6
- saturation 2-5
- scaling 2-4
- subtraction 2-12
- two's complement 2-11
- wrapping 2-5

Fixed-Point Embedded MATLAB 8-1

- fixed-point math 4-15
- fixed-point run-time API 9-6
- fixed-point signal logging 9-6

H

- help
 - getting 1-5

I

- interoperability
 - `fi` objects with Filter Design Toolbox 9-14

- fi objects with Signal Processing Blockset 9-9
- fi objects with Simulink 9-2

L

- licensing 1-4
- logging
 - overflows and underflows 5-7
- logging modes
 - setting 5-7

M

- math
 - with [Slope Bias] signals 4-19
- Model Explorer
 - setting `embedded.fimath` properties 4-13
 - setting `embedded.numericity` properties 6-9
 - using with fixed-point Embedded MATLAB 8-15
- modular arithmetic 2-10
- multiplication
 - fixed-point 2-13

N

- numericity objects
 - constructing 6-2
 - properties
 - setting in the Model Explorer 6-9
 - setting properties in the Model Explorer 6-9

O

- one's complement 2-11
- overflow handling 2-5
 - compared with ANSI C 2-28
- overflows
 - logging 5-7

P

- padding 2-19
- precision
 - fixed-point data types 2-5
- property values
 - quantizer objects 7-3

Q

- quantizer objects
 - constructing 7-2
 - property values 7-3

R

- range
 - fixed-point data types 2-5
- reading fixed-point data from workspace 9-2
- rounding
 - fixed-point data types 2-6
- run-time API
 - fixed-point data 9-6

S

- saturation 2-5
- scaling 2-4
- signal logging
 - fixed-point 9-6
- [Slope Bias] arithmetic 4-19

T

- two's complement 2-11

U

- unary conversions 2-24
- underflows
 - logging 5-7

W

wrapping

fixed-point data types 2-5
writing fixed-point data to workspace 9-2