



# DSP Builder for Intel® FPGAs (Advanced Blockset)

---

## Handbook

Updated for Intel® Quartus® Prime Design Suite: **21.3**



**Online Version**



**Send Feedback**

**HB\_DSPB\_ADV**

ID: **683337**

Version: **2022.03.23**

## Contents

---

<b>1. About DSP Builder for FPGAs.....</b>	<b>12</b>
1.1. DSP Builder for Intel® FPGAs Features.....	12
1.2. DSP Builder for Intel® FPGAs Design Structure.....	12
1.3. DSP Builder for Intel® FPGAs Libraries.....	15
1.4. DSP Builder for Intel® FPGAs Device Support.....	16
<b>2. DSP Builder for Intel FPGAs Advanced Blockset Getting Started.....</b>	<b>17</b>
2.1. Starting DSP Builder in MATLAB*.....	17
2.2. Browsing DSP Builder Libraries and Adding Blocks to a New Model.....	17
2.3. Browsing and Opening DSP Builder Design Examples.....	18
2.4. Creating a New DSP Builder Design with the DSP Builder New Model Wizard.....	19
2.4.1. DSP Builder Menu Options.....	20
2.4.2. DSP Builder New Model Wizard Setup Script Parameters.....	21
2.5. Simulating, Verifying, Generating, and Compiling Your DSP Builder Design.....	21
<b>3. DSP Builder Design Flow.....</b>	<b>23</b>
3.1. Implementing your Design in DSP Builder Advanced Blockset.....	24
3.1.1. Dividing your DSP Builder Design into Subsystems.....	24
3.1.2. Connecting DSP Builder Subsystems.....	24
3.1.3. Creating a New Design by Copying a DSP Builder Design Example.....	33
3.1.4. Vectorized Inputs.....	36
3.2. Verifying your DSP Builder Advanced Blockset Design in Simulink and MATLAB.....	37
3.2.1. Verifying your DSP Builder Advanced Blockset Design with a Testbench.....	37
3.2.2. Running DSP Builder Advanced Blockset Automatic Testbenches.....	38
3.2.3. Using DSP Builder Advanced Blockset References.....	41
3.2.4. Setting Up Stimulus in DSP Builder Advanced Blockset.....	41
3.2.5. Analyzing your DSP Builder Advanced Blockset Design.....	41
3.3. Exploring DSP Builder Advanced Blockset Design Tradeoffs.....	41
3.3.1. Bit Growth.....	42
3.3.2. Managing Bit Growth in DSP Builder Advanced Blockset Designs.....	42
3.3.3. Using Rounding and Saturation in DSP Builder Advanced Blockset Designs.....	42
3.3.4. Scaling with Primitive Blocks.....	43
3.3.5. Changing Data Type with Convert Blocks and Specifying Output Types.....	43
3.4. Verifying your DSP Builder Design with C++ Software Models.....	47
3.4.1. Example CMakefile.....	50
3.5. Verifying your DSP Builder Advanced Blockset Design in the ModelSim Simulator.....	50
3.5.1. Automatic Testbench.....	50
3.5.2. DSP Builder Advanced Blockset ModelSim Simulations.....	51
3.6. Verifying Your DSP Builder Design in Hardware.....	51
3.6.1. Hardware Verification.....	53
3.6.2. Hardware Verification with System-in-the-Loop.....	56
3.7. Integrating Your DSP Builder Advanced Blockset Design into Hardware.....	63
3.7.1. DSP Builder Generated Files.....	63
3.7.2. DSP Builder Designs and the Quartus Prime Project.....	64
3.7.3. Interfaces with a Processor Bus.....	65
<b>4. Primitive Library Blocks Tutorial.....</b>	<b>69</b>
4.1. Creating a Fibonacci Design from the DSP Builder Primitive Library.....	69

4.2. Setting the Parameters on the Testbench Source Blocks.....	71
4.3. Simulating the Fibonacci Design in Simulink.....	72
4.4. Modifying the DSP Builder Fibonacci Design to Generate Vector Signals.....	73
4.5. Simulating the RTL of the Fibonacci Design.....	74
<b>5. IP Tutorial.....</b>	<b>75</b>
5.1. Creating an IP Design.....	75
5.2. Simulating the IP Design in Simulink.....	76
5.3. Viewing Timing Closure and Viewing Resource Utilization for the DSP Builder IP Design.....	77
5.4. Reparameterizing the DSP Builder FIR Filter to Double the Number of Channels.....	77
5.5. Doubling the Target Clock Rate for a DSP Builder IP Design.....	78
<b>6. DSP Builder for Intel FPGAs (Advanced Blockset) Design Examples and Reference Designs.....</b>	<b>79</b>
6.1. DSP Builder Design Configuration Block Design Examples.....	80
6.1.1. Scale .....	80
6.1.2. Local Threshold .....	80
6.2. DSP Builder FFT Design Examples.....	80
6.2.1. FFT .....	81
6.2.2. FFT without BitReverseCoreC Block .....	81
6.2.3. IFFT .....	82
6.2.4. IFFT without BitReverseCoreC Block .....	82
6.2.5. Floating-Point FFT .....	82
6.2.6. Floating-Point FFT without BitReverseCoreC Block .....	82
6.2.7. Floating-Point iFFT .....	83
6.2.8. Floating-Point iFFT without BitReverseCoreC Block .....	83
6.2.9. Multichannel FFT .....	83
6.2.10. Multiwire Transpose .....	83
6.2.11. Parallel FFT.....	83
6.2.12. Parallel Floating-Point FFT .....	83
6.2.13. Single-Wire Transpose .....	83
6.2.14. Switchable FFT/iFFT .....	84
6.2.15. Variable-Size Fixed-Point FFT .....	84
6.2.16. Variable-Size Fixed-Point FFT without BitReverseCoreC Block .....	84
6.2.17. Variable-Size Fixed-Point iFFT .....	84
6.2.18. Variable-Size Fixed-Point iFFT without BitReverseCoreC Block.....	84
6.2.19. Variable-Size Floating-Point FFT.....	84
6.2.20. Variable-Size Floating-Point FFT without BitReverseCoreC Block .....	84
6.2.21. Variable-Size Floating-Point iFFT .....	84
6.2.22. Variable-Size Floating-Point iFFT without BitReverseCoreC Block .....	85
6.2.23. Variable-Size Low-Resource FFT .....	85
6.2.24. Variable-Size Low-Resource Real-Time FFT .....	85
6.2.25. Variable-Size Supersampled FFT.....	85
6.3. DSP Builder DDC Design Example.....	86
6.3.1. DDC Design Example Subsystem.....	93
6.3.2. Building the DDC Design Example.....	96
6.4. DSP Builder Filter Design Examples.....	100
6.4.1. Complex FIR Filter .....	101
6.4.2. Decimating CIC Filter .....	101
6.4.3. Decimating FIR Filter .....	101
6.4.4. Filter Chain with Forward Flow Control .....	102

6.4.5. FIR Filter with Exposed Bus .....	102
6.4.6. Fractional FIR Filter Chain .....	102
6.4.7. Fractional-Rate FIR Filter .....	102
6.4.8. Half-Band FIR Filter .....	103
6.4.9. IIR: Full-rate Fixed-point.....	103
6.4.10. IIR: Full-rate Floating-point.....	103
6.4.11. Interpolating CIC Filter .....	104
6.4.12. Interpolating FIR Filter .....	104
6.4.13. Interpolating FIR Filter with Multiple Coefficient Banks .....	105
6.4.14. Interpolating FIR Filter with Updating Coefficient Banks .....	106
6.4.15. Root-Raised Cosine FIR Filter .....	106
6.4.16. Single-Rate FIR Filter .....	106
6.4.17. Super-Sample Decimating FIR Filter .....	107
6.4.18. Super-Sample Fractional FIR Filter .....	107
6.4.19. Super-Sample Interpolating FIR Filter .....	107
6.4.20. Variable-Rate CIC Filter .....	107
6.5. DSP Builder Finite State Machine Design Example.....	108
6.6. DSP Builder Folding Design Examples.....	108
6.6.1. Position, Speed, and Current Control for AC Motors .....	109
6.6.2. Position, Speed, and Current Control for AC Motors (with ALU Folding) .....	113
6.6.3. About FOC.....	114
6.6.4. Folded FIR Filter .....	114
6.7. DSP Builder Floating Point Design Examples.....	115
6.7.1. Black-Scholes Floating Point .....	115
6.7.2. Double-Precision Real Floating-Point Matrix Multiply .....	115
6.7.3. Fine Doppler Estimator .....	115
6.7.4. Floating-Point Mandelbrot Set .....	116
6.7.5. General Real Matrix Multiply One Cycle Per Output .....	117
6.7.6. Newton Root Finding Tutorial Step 1—Iteration .....	117
6.7.7. Newton Root Finding Tutorial Step 2—Convergence .....	118
6.7.8. Newton Root Finding Tutorial Step 3—Valid .....	118
6.7.9. Newton Root Finding Tutorial Step 4—Control .....	118
6.7.10. Newton Root Finding Tutorial Step 5—Final .....	118
6.7.11. Normalizer .....	118
6.7.12. Single-Precision Complex Floating-Point Matrix Multiply .....	118
6.7.13. Single-Precision Real Floating-Point Matrix Multiply .....	119
6.7.14. Simple Nonadaptive 2D Beamformer .....	119
6.8. DSP Builder Flow Control Design Examples.....	120
6.8.1. Avalon-ST Interface (Input and Output FIFO Buffer) with Backpressure .....	120
6.8.2. Avalon-ST Interface (Output FIFO Buffer) with Backpressure .....	120
6.8.3. Kronecker Tensor Product .....	121
6.8.4. Parallel Loops .....	121
6.8.5. Primitive FIR with Back Pressure .....	121
6.8.6. Primitive FIR with Forward Pressure .....	122
6.8.7. Primitive Systolic FIR with Forward Flow Control .....	123
6.8.8. Rectangular Nested Loop .....	123
6.8.9. Sequential Loops .....	124
6.8.10. Triangular Nested Loop .....	124
6.9. DSP Builder HDL Import Design Example.....	124
6.9.1. Performing a Cosimulation.....	133
6.10. DSP Builder Host Interface Design Examples.....	137

6.10.1.	Memory-Mapped Registers .....	137
6.11.	DSP Builder Platform Design Examples.....	138
6.11.1.	16-Channel DDC .....	138
6.11.2.	16-Channel DUC .....	138
6.11.3.	2-Antenna DUC for WiMAX .....	139
6.11.4.	2-Channel DUC .....	140
6.11.5.	Super-Sample Rate Digital Upconverter .....	140
6.12.	DSP Builder Primitive Block Design Examples.....	140
6.12.1.	8×8 Inverse Discrete Cosine Transform.....	141
6.12.2.	Automatic Gain Control .....	141
6.12.3.	Bit Combine for Boolean Vectors .....	142
6.12.4.	Bit Extract for Boolean Vectors .....	142
6.12.5.	Color Space Converter .....	142
6.12.6.	CORDIC from Primitive Blocks .....	143
6.12.7.	Digital Predistortion Forward Path .....	143
6.12.8.	Fibonacci Series .....	143
6.12.9.	Folded Vector Sort .....	144
6.12.10.	Fractional Square Root Using CORDIC .....	144
6.12.11.	Fixed-point Maths Functions .....	144
6.12.12.	Gaussian Random Number Generator.....	144
6.12.13.	Hello World .....	145
6.12.14.	Hybrid Direct Form and Transpose Form FIR Filter .....	145
6.12.15.	Loadable Counter .....	145
6.12.16.	Matrix Initialization of LUT .....	146
6.12.17.	Matrix Initialization of Vector Memories .....	146
6.12.18.	Multichannel IIR Filter .....	147
6.12.19.	Quadrature Amplitude Modulation .....	147
6.12.20.	Reinterpret Cast for Bit Packing and Unpacking .....	147
6.12.21.	Run-time Configurable Decimating and Interpolating Half-Rate FIR Filter ..	148
6.12.22.	Square Root Using CORDIC .....	148
6.12.23.	Test CORDIC Functions with the CORDIC Block .....	148
6.12.24.	Uniform Random Number Generator.....	148
6.12.25.	Vector Sort—Sequential .....	149
6.12.26.	Vector Sort—Iterative .....	149
6.12.27.	Vector Initialization of Sample Delay .....	149
6.12.28.	Wide Single-Channel Accumulators .....	150
6.13.	DSP Builder Reference Designs.....	150
6.13.1.	1-Antenna WiMAX DDC .....	152
6.13.2.	2-Antenna WiMAX DDC .....	152
6.13.3.	1-Antenna WiMAX DUC .....	153
6.13.4.	2-Antenna WiMAX DUC .....	153
6.13.5.	4-Carrier, 2-Antenna W-CDMA DDC .....	154
6.13.6.	1-Carrier, 2-Antenna W-CDMA DDC .....	155
6.13.7.	4-Carrier, 2-Antenna W-CDMA DUC .....	155
6.13.8.	4-Carrier, 4-Antenna DUC and DDC for LTE.....	156
6.13.9.	1-Carrier, 2-Antenna W-CDMA DDC .....	157
6.13.10.	4-Carrier, 2-Antenna High-Speed W-CDMA DUC at 368.64 MHz with Total Rate Change 32 .....	158
6.13.11.	4-Carrier, 2-Antenna High-Speed W-CDMA DUC at 368.64 MHz with Total Rate Change 48 .....	158

- 6.13.12. 4-Carrier, 2-Antenna High-Speed W-CDMA DUC at 307.2 MHz with Total Rate Change 40 ..... 159
- 6.13.13. Cholesky-based Matrix Inversion..... 160
- 6.13.14. Cholesky Solver Multiple Channels ..... 164
- 6.13.15. Crest Factor Reduction ..... 165
- 6.13.16. Direct RF with Synthesizable Testbench ..... 165
- 6.13.17. Dynamic Decimating FIR Filter ..... 165
- 6.13.18. Multichannel QR Decomposition ..... 166
- 6.13.19. QR Decomposition ..... 166
- 6.13.20. QRD Solver..... 167
- 6.13.21. Reconfigurable Decimation Filter ..... 168
- 6.13.22. Single-Channel 10-MHz LTE Transmitter ..... 168
- 6.13.23. STAP Radar Forward and Backward Substitution ..... 169
- 6.13.24. STAP Radar Steering Generation ..... 169
- 6.13.25. STAP Radar QR Decomposition 192x204 ..... 169
- 6.13.26. Time Delay Beamformer ..... 170
- 6.13.27. Transmit and Receive Modem..... 170
- 6.13.28. Variable Integer Rate Decimation Filter ..... 171
- 6.14. DSP Builder Waveform Synthesis Design Examples..... 171
  - 6.14.1. Complex Mixer ..... 171
  - 6.14.2. Four Channel, Two Banks NCO ..... 171
  - 6.14.3. Four Channel, Four Banks NCO ..... 173
  - 6.14.4. Four Channel, Eight Banks, Two Wires NCO ..... 173
  - 6.14.5. Four Channel, 16 Banks NCO ..... 174
  - 6.14.6. IP ..... 175
  - 6.14.7. NCO ..... 175
  - 6.14.8. NCO with Exposed Bus ..... 175
  - 6.14.9. Real Mixer ..... 175
  - 6.14.10. Super-sample NCO ..... 176
- 7. DSP Builder Design Rules, Design Recommendations, and Troubleshooting..... 177**
  - 7.1. DSP Builder Design Rules and Recommendations..... 177
  - 7.2. Troubleshooting DSP Builder Designs..... 179
    - 7.2.1. About Loops..... 180
    - 7.2.2. DSP Builder Timed Feedback Loops..... 181
    - 7.2.3. DSP Builder Loops, Clock Cycles, and Data Cycles..... 182
- 8. About DSP Builder for Intel FPGAs Optimization..... 183**
  - 8.1. Associating DSP Builder with MATLAB..... 183
  - 8.2. Setting Up Simulink for DSP Builder Designs..... 183
    - 8.2.1. Setting Up Simulink Solver..... 183
    - 8.2.2. Setting Up Simulink Signal Display Option..... 184
  - 8.3. The DSP Builder Windows Shortcut..... 184
  - 8.4. Setting DSP Builder Design Parameters with MATLAB Scripts..... 184
    - 8.4.1. Running Setup Scripts Automatically..... 185
    - 8.4.2. Defining Unique DSP Builder Design Parameters..... 185
    - 8.4.3. Example DSP Builder Custom Scripts..... 185
  - 8.5. Managing your Designs..... 187
    - 8.5.1. Managing Basic Parameters..... 187
    - 8.5.2. Creating User Libraries and Converting a Primitive Subsystem into a Custom Block..... 188
    - 8.5.3. Revision Control..... 188

8.6. How to Manage Latency.....	188
8.6.1. Reading the Added Latency Value for an IP Block.....	189
8.6.2. Zero Latency Example.....	189
8.6.3. Implicit Delays in DSP Builder Designs.....	190
8.6.4. Distributed Delays in DSP Builder Designs.....	191
8.6.5. Latency and $f_{MAX}$ Constraint Conflicts in DSP Builder Designs.....	193
8.6.6. Control Units Delays.....	193
8.7. Flow Control in DSP Builder Designs.....	195
8.8. Reset Minimization.....	197
8.9. About Importing HDL.....	199
<b>9. About Folding.....</b>	<b>201</b>
9.1. ALU Folding.....	201
9.1.1. ALU Folding Limitations.....	202
9.1.2. ALU Folding Parameters.....	202
9.1.3. ALU Folding Simulation Rate.....	202
9.1.4. Using ALU Folding.....	206
9.1.5. Using Automated Verification.....	207
9.1.6. Ready Signal.....	207
9.1.7. Connecting the ALU Folding Ready Signal.....	207
9.1.8. About the ALU Folding Start of Packet Signal.....	208
9.2. Removing Resource Sharing Folding.....	208
<b>10. Floating-Point Data Types.....</b>	<b>210</b>
10.1. DSP Builder Floating-Point Data Type Features.....	211
10.2. DSP Builder Supported Floating-Point Data Types.....	211
10.3. DSP Builder Round-Off Errors.....	212
10.4. Trading Off Logic Utilization and Accuracy in DSP Builder Designs.....	212
10.5. Upgrading Pre v14.0 Designs.....	213
10.6. Floating-Point Sine Wave Generator Tutorial.....	213
10.6.1. Creating a Sine Wave Generator in DSP Builder.....	213
10.6.2. Using Data Type Variables to Parameterize Designs.....	214
10.6.3. Using Data-Type Propagation in DSP Builder Designs.....	215
10.6.4. DSP Builder Testbench Verification.....	215
10.7. Newton-Raphson Root Finding Tutorial.....	218
10.7.1. Implementing the Newton Design.....	219
10.7.2. Improving DSP Builder Floating-Point Designs.....	219
10.8. Forcing Soft Floating-point Data Types with the Advanced Options.....	220
<b>11. Design Configuration Library.....</b>	<b>221</b>
11.1. Avalon Memory-Mapped Agent Settings (AvalonMemoryMappedAgentSettings) .....	221
11.2. Control .....	223
11.2.1. DSP Builder Memory and Multiplier Trade-Off Options.....	226
11.3. Device .....	226
11.4. Edit Params .....	227
11.5. LocalThreshold .....	227
<b>12. IP Library.....</b>	<b>229</b>
12.1. Channel Filter and Waveform Library.....	229
12.1.1. DSP Builder FIR and CIC Filters.....	230
12.1.2. DSP Builder FIR Filters.....	233
12.1.3. Channel Viewer (ChanView) .....	235

- 12.1.4. Complex Mixer (ComplexMixer) ..... 236
- 12.1.5. Decimating CIC ..... 238
- 12.1.6. Decimating FIR ..... 239
- 12.1.7. Fractional Rate FIR ..... 241
- 12.1.8. Interpolating CIC ..... 244
- 12.1.9. Interpolating FIR ..... 245
- 12.1.10. NCO ..... 248
- 12.1.11. Real Mixer (Mixer) ..... 252
- 12.1.12. Scale ..... 254
- 12.1.13. Single-Rate FIR ..... 255
- 12.2. Dependent Delay Library..... 257
- 12.3. FFT IP Library..... 258
  - 12.3.1. Bit Reverse Core C (BitReverseCoreC and VariableBitReverse) ..... 258
  - 12.3.2. FFT (FFT, FFT\_Light, VFFT, VFFT\_Light) ..... 259
- 13. Interfaces Library..... 262**
  - 13.1. Memory-Mapped Library..... 262
    - 13.1.1. Bus Slave (BusSlave) ..... 262
    - 13.1.2. Bus Stimulus (BusStimulus) ..... 262
    - 13.1.3. Bus Stimulus File Reader (Bus StimulusFileReader) ..... 263
    - 13.1.4. External Memory, Memory Read, Memory Write..... 265
    - 13.1.5. Register Bit (RegBit) ..... 269
    - 13.1.6. Register Field (RegField) ..... 270
    - 13.1.7. Register Out (RegOut) ..... 271
    - 13.1.8. Shared Memory (SharedMem) ..... 271
  - 13.2. Streaming Library..... 273
    - 13.2.1. Avalon-ST Input (AStInput) ..... 273
    - 13.2.2. Avalon-ST Input FIFO Buffer (AStInputFIFO) ..... 274
    - 13.2.3. Avalon-ST Output (AStOutput) ..... 274
- 14. Primitives Library..... 276**
  - 14.1. Vector and Complex Type Support..... 276
    - 14.1.1. Vector Type Support..... 276
    - 14.1.2. Complex Support..... 277
  - 14.2. DFT Design Elements Library..... 278
    - 14.2.1. DFT (DFT)..... 278
    - 14.2.2. Reorder (ReorderBlock)..... 280
    - 14.2.3. Reorder and Rescale (ReorderAndRescale)..... 281
  - 14.3. FFT Design Elements Library..... 283
    - 14.3.1. About Pruning and Twiddle for FFT Blocks..... 284
    - 14.3.2. Bit Vector Combine (BitVectorCombine) ..... 285
    - 14.3.3. Butterfly Unit (BFU) ..... 285
    - 14.3.4. Butterfly I C (BFIC) (Deprecated) ..... 286
    - 14.3.5. Butterfly II C (BFIIC) (Deprecated) ..... 287
    - 14.3.6. Choose Bits (ChooseBits) ..... 288
    - 14.3.7. Crossover Switch (XSwitch) ..... 288
    - 14.3.8. Dual Twiddle Memory (DualTwiddleMemoryC) ..... 288
    - 14.3.9. Edge Detect (EdgeDetect) ..... 289
    - 14.3.10. Floating-Point Twiddle Generator (TwiddleGenF) (Deprecated) ..... 289
    - 14.3.11. Fully-Parallel FFTs (FFT2P, FFT4P, FFT8P, FFT16P, FFT32P, and FFT64P) ..... 290
    - 14.3.12. Fully-Parallel FFTs with Flexible Ordering (FFT2X, FFT4X, FFT8X, FFT16X, FFT32X, and FFT64X) ..... 290



14.3.13. General Multitwiddle and General Twiddle (GeneralMultiTwiddle, GeneralMultVTwiddle, GeneralTwiddle, GeneralVTwiddle) .....	291
14.3.14. Hybrid FFT (Hybrid_FFT, HybridVFFT) .....	293
14.3.15. Multiwire Transpose (MultiwireTranspose) .....	294
14.3.16. Parallel Pipelined FFT (PFFT_Pipe) .....	295
14.3.17. Pulse Divider (PulseDivider) .....	295
14.3.18. Pulse Multiplier (PulseMultiplier) .....	296
14.3.19. Single-Wire Transpose (Transpose) .....	296
14.3.20. Split Scalar (SplitScalar) .....	297
14.3.21. Streaming FFTs (FFT2, FFT4, VFFT2, and VFFT4) .....	297
14.3.22. Stretch Pulse (StretchPulse).....	298
14.3.23. Twiddle Angle (TwiddleAngle) .....	298
14.3.24. Twiddle Generator (TwiddleGenC) Deprecated .....	299
14.3.25. Twiddle and Variable Twiddle (Twiddle and VTwiddle) .....	299
14.3.26. Twiddle ROM (TwiddleRom, TwiddleMultRom and TwiddleRomF (deprecated)) .....	300
14.4. Primitive Basic Blocks Library.....	301
14.4.1. Absolute Value (Abs) .....	303
14.4.2. Accumulator (Acc) .....	303
14.4.3. Add .....	305
14.4.4. Add SLoad (AddSLoad) .....	306
14.4.5. AddSub .....	307
14.4.6. AddSubFused .....	307
14.4.7. AND Gate (And) .....	307
14.4.8. Bit Combine (BitCombine) .....	308
14.4.9. Bit Extract (BitExtract) .....	309
14.4.10. Bit Reverse (BitReverse) .....	310
14.4.11. Compare (CmpCtrl) .....	310
14.4.12. Complex Conjugate (ComplexConjugate) .....	310
14.4.13. Compare Equality (CmpEQ) .....	311
14.4.14. Compare Greater Than (CmpGE) .....	312
14.4.15. Compare Less Than (CmpLT) .....	312
14.4.16. Compare Not Equal (CmpNE) .....	312
14.4.17. Constant (Const) .....	313
14.4.18. Constant Multiply (Const Mult) .....	314
14.4.19. Convert .....	314
14.4.20. CORDIC .....	315
14.4.21. Counter .....	325
14.4.22. Count Leading Zeros, Ones, or Sign Bits (CLZ) .....	326
14.4.23. Dual Memory (DualMem) .....	327
14.4.24. Demultiplexer (Demux) .....	329
14.4.25. Divide .....	330
14.4.26. Fanout.....	331
14.4.27. FIFO .....	332
14.4.28. Floating-point Classifier (FloatClass) .....	333
14.4.29. Floating-point Multiply Accumulate (MultAcc) .....	333
14.4.30. ForLoop .....	334
14.4.31. Load Exponent (LdExp) .....	335
14.4.32. Left Shift (LShift) .....	336
14.4.33. Loadable Counter (LoadableCounter) .....	336
14.4.34. Look-Up Table (Lut) .....	337

14.4.35. Loop .....	339
14.4.36. Math .....	340
14.4.37. Minimum and Maximum (MinMax) .....	341
14.4.38. MinMaxCtrl .....	342
14.4.39. Multiply (Mult) .....	343
14.4.40. Multiplexer (Mux) .....	343
14.4.41. NAND Gate (Nand) .....	344
14.4.42. Negate .....	345
14.4.43. NOR Gate (Nor) .....	345
14.4.44. NOT Gate (Not) .....	346
14.4.45. OR Gate (Or) .....	347
14.4.46. Polynomial .....	347
14.4.47. Ready .....	348
14.4.48. Reinterpret Cast (ReinterpretCast) .....	348
14.4.49. Round .....	349
14.4.50. Sample Delay (SampleDelay) .....	349
14.4.51. Scalar Product .....	350
14.4.52. Select .....	351
14.4.53. Sequence .....	352
14.4.54. Shift .....	353
14.4.55. Sqrt .....	354
14.4.56. Subtract (Sub) .....	354
14.4.57. Sum of Elements (SumOfElements) .....	355
14.4.58. Trig .....	356
14.4.59. XNOR Gate (Xnor) .....	358
14.4.60. XOR Gate (Xor) .....	358
14.5. Primitive Configuration Library.....	359
14.5.1. Channel In (ChannelIn) .....	359
14.5.2. Channel Out (ChannelOut) .....	360
14.5.3. General Purpose Input (GPIn) .....	361
14.5.4. General Purpose Output (GPOut) .....	361
14.5.5. Synthesis Information (SynthesisInfo) .....	362
14.6. Primitive Design Elements Library.....	363
14.6.1. Anchored Delay .....	364
14.6.2. Complex to Real-Imag.....	364
14.6.3. Enabled Delay Line.....	364
14.6.4. Enabled Feedback Delay.....	364
14.6.5. Expand Scalar (ExpandScalar).....	364
14.6.6. Finite State Machine.....	364
14.6.7. Nested Loops (NestedLoop1, NestedLoop2, NestedLoop3).....	370
14.6.8. Pause.....	371
14.6.9. Reset-Priority Latch (SRLatch_PS) .....	372
14.6.10. Same Data Type (SameDT).....	372
14.6.11. Set-Priority Latch (SRLatch) .....	372
14.6.12. Single-Cycle Latency Latch (latch_1L) .....	373
14.6.13. Tapped Line Delay (TappedLineDelay).....	373
14.6.14. Variable Super-Sample Delay (VariableDelay).....	373
14.6.15. Vector Fanout (VectorFanout).....	374
14.6.16. Vector Multiplexer (VectorMux).....	374
14.6.17. Zero-Latency Latch (latch_0L) .....	374

<b>15. Utilities Library.....</b>	<b>375</b>
15.1. Analyze and Test Library.....	375
15.1.1. Capture Values.....	375
15.1.2. HDL Import.....	375
15.1.3. HDL Import Config.....	377
15.1.4. Pause.....	378
<b>16. Simulink Supported Blocks.....</b>	<b>379</b>
<b>17. Document Revision History for DSP Builder for Intel FPGAs (Advanced Blockset) Handbook.....</b>	<b>380</b>

## 1. About DSP Builder for FPGAs

---

DSP Builder is a high-level synthesis technology that optimizes the high-level, untimed netlist into low level, pipelined hardware for your target FPGA device and desired clock rate. DSP Builder for FPGAs consists of several Simulink\* libraries that allow you to implement DSP designs quickly and easily. DSP Builder implements the hardware as VHDL or Verilog HDL with scripts that integrate with the software and the simulator.

You can create designs without needing detailed device knowledge and generate designs that run on a variety of FPGA families with different hardware architectures. DSP Builder allows you to manually describe algorithmic functions and apply rule-based methods to generate hardware optimized code. The advanced blockset is particularly suited for streaming algorithms characterized by continuous data streams and occasional control. For example, use DSP Builder to create RF card designs that comprise long filter chains.

After specifying the desired clock frequency, target device family, number of channels, and other top-level design constraints, DSP Builder pipelines the generated RTL to achieve timing closure. By analyzing the system-level constraints, DSP Builder can optimize folding to balance latency versus resources, with no need for manual RTL editing.

DSP Builder advanced blockset includes its own timing-driven IP blocks that can generate high performance FIR, CIC, and NCO models.

1. [DSP Builder for Intel FPGAs Features](#) on page 12
2. [DSP Builder for Intel FPGAs Design Structure](#) on page 12
3. [DSP Builder for Intel FPGAs Libraries](#) on page 15
4. [DSP Builder for Intel FPGAs Device Support](#) on page 16

### 1.1. DSP Builder for Intel® FPGAs Features

- Automatic pipelining to enable timing closure
- Automatic folding
- Easy to compare and target different device families
- High-performance floating-point designs
- Wizard-based interface (system-in-the-loop) to configure, generate, and run hardware verification system.

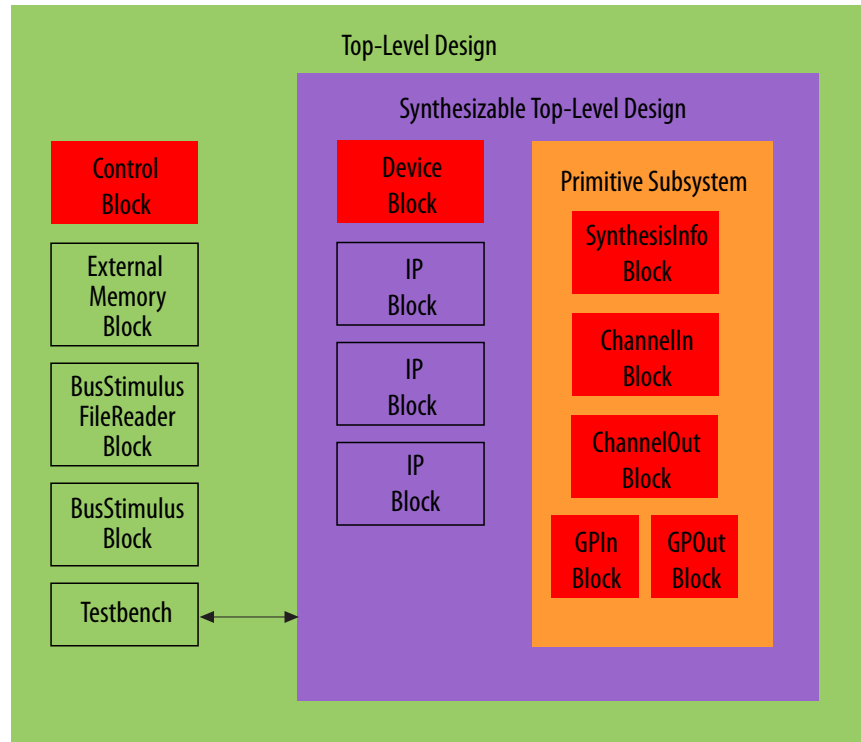
### 1.2. DSP Builder for Intel® FPGAs Design Structure

Organize your DSP Builder designs into hierarchical Simulink subsystems. Every top-level design must contain a **Control** block; the synthesizable top-level design must contain a **Device** block.

**Note:** A DSP Builder design can only have one synthesizable top-level design, which can contain many subsystems (primitive and IP blocks) to help organize your design. Any primitive blocks must be within a primitive subsystem hierarchy and any IP blocks must be outside primitive subsystem hierarchies.

**Figure 1. DSP Builder Design**

Shows the relationship of the synthesizable top-level design and a primitive subsystem. Mandatory blocks are in red.



### The Top-Level Design

A DSP Builder advanced blockset top-level design consists of:

- A Simulink testbench, which provides design inputs and allows you to analyze inputs and outputs
- Top-level configuration blocks
- Optional memory interface specification and stimulus blocks.
  - **External Memory** block to configure an external memory interface
  - **BusStimulus** and **BusStimulusFileReader** blocks to stimulate Avalon-MM interfaces during simulation
  - **Edit Params** block as a shortcut to opening a script `setup_<model name>.m` for editing.

The top-level design must have a **Control** block to specify RTL output directory and top-level threshold parameters

**Note:** Every DSP Builder design must have **Control** block to allow you to simulate or compile your design. Do not place the **Device** block in the top-level design. DSP Builder propagates data types from the testbench to the synthesizable top-level design.

### The Synthesizable Top-Level Design

The synthesizable top-level design is a Simulink subsystem that contains a **Device** block, which sets which family, part, and speed grade to target. The synthesizable top-level design is at the top level of the generated hardware files. The synthesizable top-level design can consist of further level of hierarchies that include primitive subsystems.

**Note:** Only use primitive blocks inside primitive subsystems.

Optionally, you can include more **LocalThreshold** blocks to override threshold settings defined higher up the hierarchy.

### The Primitive Subsystem

Primitive subsystems are scheduled domains for **Primitive** and **IP** library blocks. A primitive subsystem must have:

- A **SynthesisInfo** block, with synthesis style set to **Scheduled**, so that DSP Builder can pipeline and redistribute memories optimally to achieve the desired clock frequency.
- Boundary blocks that delimit the primitive subsystem:
  - **ChannelIn** (channelized input),
  - **ChannelOut** (channelized output),
  - **GPIIn** (general purpose input)
  - **GPOut** (general purpose output).

DSP Builder synchronizes connections that pass through the same boundary block.

Use system interface blocks to delimit the boundaries of scheduled domains within a subsystem. Within these boundary blocks DSP Builder optimizes the implementation you specify by the schematic. DSP Builder inserts pipelining registers to achieve the specified system clock rate. When DSP Builder inserts pipelining registers, it adds equivalent latency to parallel signals that need to be kept synchronous so that DSP Builder schedules them together. DSP Builder schedules signals that go through the same input boundary block (**ChannelIn** or **GPIIn**) to start at the same point in time; signals that go through the same output boundary block (**ChannelOut** or **GPOut**) to finish at the same point in time. DSP Builder adds any pipelining latency that you add to achieve  $f_{MAX}$  in balanced cuts through the signals across the design. DSP Builder applies the correction to the simulation at the boundary blocks to account for this latency in HDL generation. The primitive subsystem as a whole remains cycle accurate. You can specify further levels of hierarchy within primitive subsystems containing primitive blocks, but no further primitive boundary blocks or IP blocks.

Use **SampleDelay** blocks only to specify relative sample offsets of data-streams; do not use for pipelining.

### Related Information

- [Synthesis Information \(SynthesisInfo\)](#) on page 362
- [Primitives Library](#) on page 276

- [DSP Builder Design Rules and Recommendations](#) on page 177
- [Control](#) on page 223
- [Device](#) on page 226
- [Edit Params](#) on page 227
- [LocalThreshold](#) on page 227
- [Primitives Library](#) on page 276
- [Scheduled Synthesis](#) on page 362

### 1.3. DSP Builder for Intel® FPGAs Libraries

**Table 1. Block Types**

This table describes the types of blocks that DSP Builder offers.

Block Type	Description
Configuration blocks	Blocks that configure how DSP Builder synthesizes the design or subsystem
Low-level building blocks (primitives)	Basic operator, logic, and memory primitive blocks for scheduled subsystems delimited by boundary configuration blocks (primitive subsystems).
Common design elements	Common functions for parameterizable subsystems of primitives and within scheduled subsystems delimited by boundary configuration blocks
IP function-level functions (IP)	Stand-alone IP-level blocks comprising functions such as entire FFTs, FIRs and NCOs. Use these blocks only outside of primitive subsystems.
System interface blocks	Blocks that expose Avalon-ST and Avalon-MM interfaces for interaction with other IP (such as external memories) in Platform Designer.
Non-synthesizable blocks	Blocks that play no part in the synthesized design. For example, blocks that provide testbench stimulus, blocks that provide information, or enable design analysis.

**Table 2. Simulink Libraries**

This table lists the Simulink libraries and describes the DSP Builder blocks in those libraries

Library	Description
<b>Design Configuration</b>	Blocks that set the design parameters, such as device family, target $f_{MAX}$ and bus interface signal width.
<b>Primitives</b>	Blocks for primitive subsystems.
<b>Primitives &gt; Primitive Configuration</b>	Blocks that change how DSP Builder synthesizes primitive subsystems, including boundary delimiters.
<b>Primitives &gt; Primitive Basic Blocks</b>	Low-level functions.
<b>Primitives &gt; Primitive Design Elements</b>	Configurable blocks and common design patterns built from primitive blocks.
<b>Primitives &gt; FFT Design Elements</b>	Configurable FFT component blocks built from primitive blocks. Use in primitive subsystems to build custom FFTs.
<b>IP</b>	Full IP functions. Use outside of primitive subsystems.
<b>IP FFT IP</b>	Full FFT IP functions. These blocks are complete primitive subsystems. Click <b>Look under the Mask</b> to see how DSP Builder builds these blocks from the primitive FFT design elements.
<i>continued...</i>	

Library	Description
<b>IP &gt; Channel Filter And Waveform</b>	Functions to construct digital up- and down-conversion chains: FIR, CIC, NCO, mixers, complex mixers, channel view, and scale IP.
<b>Interfaces</b>	Blocks that set and use Avalon interfaces. DSP Builder treats design-level ports that do not route via Avalon interface blocks as individual conduits.
<b>Interfaces &gt; Memory Mapped</b>	Blocks that set and use Avalon-MM interfaces, including memory-mapped blocks, memory-mapped stimulus blocks, and external memory blocks.
<b>Interfaces &gt; Streaming</b>	Avalon-ST blocks.
<b>Utilities</b>	Miscellaneous blocks that support building and refining designs
<b>Utilities &gt; Analyze And Test</b>	Blocks that help with design testing and debugging.
<b>Utilities &gt; Beta Blocks</b>	Blocks that are in development.

**Related Information**

- [Design Configuration Library](#) on page 221
- [IP Library](#) on page 229
- [Interfaces Library](#) on page 262
- [Primitives Library](#) on page 276
- [Utilities Library](#) on page 375
- [Scheduled Synthesis](#) on page 362
- [Avalon Interface Specification](#)  
Avalon interfaces simplify system design by allowing you to easily connect components in an Intel® FPGA

### 1.4. DSP Builder for Intel® FPGAs Device Support

DSP Builder supports devices in Intel® Quartus® Prime Pro Edition.

DSP Builder Advanced blockset supports the following device families:

- Intel Arria® 10
- Intel Agilex™
- Cyclone V devices
- Intel Cyclone® 10 GX
- Intel Stratix® 10

For designs targetting Cyclone V devices, compile the generated RTL with Intel Quartus Prime Standard edition. Before running DSP Builder, set `QUARTUS_ROOTDIR` to refer to an Intel Quartus Prime Standard installation. Then the various Intel Quartus Prime related features in DSP Builder use that Intel Quartus Prime Standard installation.



## 2. DSP Builder for Intel FPGAs Advanced Blockset Getting Started

---

1. [Starting DSP Builder in MATLAB](#) on page 17
2. [Browsing DSP Builder Libraries and Adding Blocks to a New Model](#) on page 17
3. [Browsing and Opening DSP Builder Design Examples](#) on page 18
4. [Creating a New DSP Builder Design with the DSP Builder New Model Wizard](#) on page 19
5. [Simulating, Verifying, Generating, and Compiling Your DSP Builder Design](#) on page 21

### Related Information

- [IP Tutorial](#)
- [Primitives Tutorial](#)

### 2.1. Starting DSP Builder in MATLAB\*

#### STEPS:

1. On Windows\* OS, click **Start > All Programs > Intel FPGA version > DSP Builder > Start in MATLAB version**. On Linux OS, from the command prompt, run "`dsp_builder.sh`".
2. In MATLAB on the **Home** tab, click on the **Simulink Library** icon, to start Simulink.

#### Related Information

- [The DSP Builder Windows Shortcut Menu](#)  
Create the shortcut to set the file paths to DSP Builder and run a batch file with an argument for the MATLAB executable to use.
- [Browsing DSP Builder Libraries and Adding Blocks to a New Model](#)
- [Browsing and Opening DSP Builder Design Examples](#)

### 2.2. Browsing DSP Builder Libraries and Adding Blocks to a New Model

#### BEFORE YOU BEGIN:

Start DSP Builder in MATLAB.

**STEPS:**

1. In the Simulink Library Browser, in the left-hand pane, expand **DSP Builder for Intel FPGAs - Advanced Blockset**.  
Simulink lists the DSP Builder advanced blockset libraries.
2. Click on a library.  
Simulink shows the library in the right-hand pane.
3. To find more information about a block, right click on a block and click **Help for the block**.
4. To add a block to a model, right click on a block and click **Add block to a new model**.

**Related Information**

- [Starting DSP Builder in MATLAB](#)
- [Browsing and Opening DSP Builder Design Examples](#)
- [DSP Builder Advanced Blockset Libraries](#)
- [Creating a DSP Builder Design in Simulink](#)  
Intel recommends you create new designs with the DSP Builder New Model Wizard or copy and rename a design example.

## 2.3. Browsing and Opening DSP Builder Design Examples

**BEFORE YOU BEGIN:**

Start DSP Builder in MATLAB.

**STEPS:**

1. In MATLAB, on the **Home** tab, click the **Help** icon.  
The **Help** window opens.
2. Under **Supplemental software**, click **Examples for DSP Builder for Intel(R) FPGAs - Advanced Blockset**.
3. In the left-hand TOC pane, expand **Examples for DSP Builder for Intel(R) FPGAs - Advanced Blockset Examples**
4. Expand **Floating Point** for example to see the floating-point design examples.
5. Click on a design example to see a description.
6. Click **Open this model**, to open the design example.
7. You can also open a design example by typing a command in the MATLAB window, for example:

```
demo_nco
```

**Related Information**

- [Starting DSP Builder in MATLAB](#) on page 17
- [Starting DSP Builder in MATLAB](#)
- [DSP Builder Advanced Blockset Libraries](#)
- [Browsing DSP Builder Libraries and Adding Blocks to a New Model](#)

- [Creating a DSP Builder Design in Simulink](#)  
Intel recommends you create new designs with the DSP Builder New Model Wizard or copy and rename a design example.

## 2.4. Creating a New DSP Builder Design with the DSP Builder New Model Wizard

Intel recommends you create new designs with the DSP Builder New Model Wizard. Alternatively, you can copy and rename a design example.

### BEFORE YOU BEGIN:

Start DSP Builder in MATLAB.

### STEPS:

1. In the Simulink Library browser, click **New Model**.
2. Click **DSP Builder > New Model Wizard**.  
The **New Model Wizard** opens.
3. Select a fixed- or floating-point model.
4. Select the type (simple or with channelizer).
5. Enter the model name and select where to save the model.
6. Click **Generate**.

DSP Builder creates a new model `<model name>.mdl` and setup script `setup_<model name>.m` that contains everything you need for a DSP Builder model. DSP Builder automatically runs the set-up script when you open the model and before each simulation. To open and edit the script, double click the **Edit Params** block in the model.

*Note:* When you open a model, DSP Builder produces a `model_name_params.xml` file that contains settings for the model. You must keep this file with the model.

1. [DSP Builder Menu Options](#) on page 20
2. [DSP Builder New Model Wizard Setup Script Parameters](#) on page 21

### Related Information

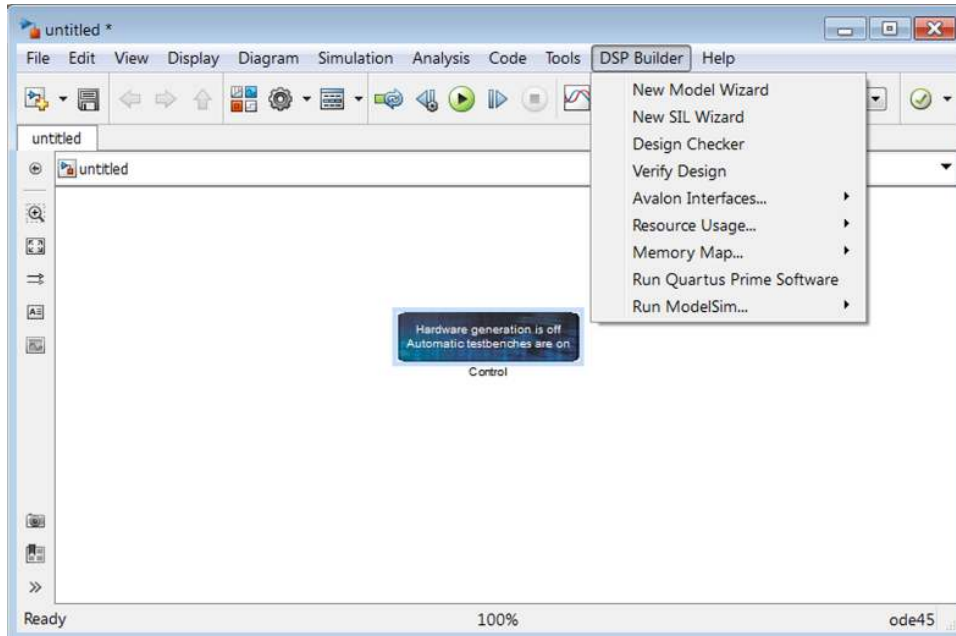
- [Starting DSP Builder in MATLAB](#) on page 17
- [Starting DSP Builder in MATLAB](#)
- [DSP Builder Advanced Blockset Libraries](#)
- [Simulating, Generating, and Compiling Your Design](#)
- [DSP Builder Menu Options](#)  
Simulink includes a **DSP Builder** menu on any Simulink model window. Use this menu to easily start all the common tasks you need to perform on your DSP Builder model.
- [DSP Builder New Model Wizard Setup Script Parameters](#)  
Use the setup script to set name-spaced workspace variables that DSP Builder uses to configure the design.

- [DSP Builder Design Rules and Recommendations](#)  
Use the design rules and recommendations to ensure your design performs correctly.

### 2.4.1. DSP Builder Menu Options

Simulink includes a **DSP Builder** menu on any Simulink model window. Use this menu to start all the common tasks you need to perform on your DSP Builder model.

**Figure 2. DSP Builder Menu**



**Table 3. DSP Builder Menu Options**

Action	Menu Option	Description
Create new design	New Model Wizard	Create a new model from a simple template.
	New SIL Wizard	Create a version of the existing design setup for hardware cosimulation.
Verification	Design Checker	Verify your design against basic design rules.
	Verify Design	Verify the Simulink simulation matches ModelSim simulations of the generated hardware by batch running the automatically generated testbenches.
Parameterization	Avalon Interfaces ...	Configure the memory mapped interface.
Generated hardware details	Resource Usage ...	View resource estimates of the generated hardware.
	Memory Map...	View the generated memory map interface.
Run other software tools	Run Quartus Prime Software	Run a Quartus Prime project for the generated hardware.
	Run ModelSim	Verify the Simulink simulation matches ModelSim simulation of the generated hardware by running an automatically generated testbench in an open ModelSim window.

## 2.4.2. DSP Builder New Model Wizard Setup Script Parameters

The setup script sets name-spaced workspace variables that DSP Builder uses to configure the design

The setup script offers the following options:

- Fixed-point IP (simple testbench)
- Fixed-point IP (with **Channelizer**)
- Fixed-point Primitive subsystem (simple testbench)
- Fixed-point Primitive subsystem (with **Channelizer**)
- Floating-point Primitive subsystem (simple testbench)
- Floating-point Primitive subsystem (with **Channelizer**)

**Table 4. Setup Script Parameters**

Option	Description
Floating	The testbench propagates single precision floating-point data into the synthesizable system.
Fixed	The testbench propagates signed fixed-point data into the synthesizable system.
'(simple testbench)'	The testbench consists of simple Simulink source blocks.
Channelizer	The testbench consists of a <b>Channelizer</b> block, which outputs data from a MATLAB array in the DSP Builder valid-channel-data protocol
'IP'	The synthesizable system has two IP function-level subsystems ( <b>IP</b> library blocks): a <b>FIR</b> block and a <b>Scale</b> block
'Primitive'	The synthesizable system is a scheduled primitive subsystem with <b>ChannelIn</b> and <b>ChannelOut</b> boundary blocks. Use this start point to create your own function using low-level (primitive) building blocks .

### Related Information

- [Creating a New DSP Builder Design with the DSP Builder New Model Wizard](#)  
Intel recommends you create new designs with the DSP Builder New Model Wizard. Alternatively, you can copy and rename a design example.
- [DSP Builder Menu Options](#)  
Simulink includes a **DSP Builder** menu on any Simulink model window. Use this menu to easily start all the common tasks you need to perform on your DSP Builder model.

## 2.5. Simulating, Verifying, Generating, and Compiling Your DSP Builder Design

### BEFORE YOU BEGIN:

- Create a design
- Check your design for errors

### STEPS:

1. In Simulink, click **Simulation** ► **Run**.

*Note:* Simulink generates the HDL then starts the simulation

2. Analyze the simulation results.
3. Verify generated hardware (optional).
  - a. Click **DSP Builder Verify Design**.
  - b. Turn on **Verify at subsystem level**, turn off **Run Quartus Prime Software**, and click **Run Verification**.

*Note:* If you turn on **Run Quartus Prime Software**, the verification script also compiles the design in the Quartus Prime software. MATLAB reports the postcompilation resource usage details in the verification window.

MATLAB verifies that the Simulink simulation results match a simulation of the generated HDL in the ModelSim simulator.

- c. Close both verification windows when MATLAB completes the verification.
4. Examine the generated resource summaries:
  - a. Click **Simulation > Start**.
  - b. Click **Resource Usage > Design** for a top-level design summary.
5. View the Avalon-MM register memory map:
  - a. Click **Simulation > Start**.
  - b. Click **Memory Map > Design**. DSP Builder highlights in red any memory conflicts.

*Note:* DSP Builder also generates the memory map in the `<design name>_mmap.h` file.


6. Compile your design in the Quartus Prime software by clicking **Run Quartus Prime**. When the Quartus Prime software opens, click **Processing > Start Compilation**.

### Related Information

- [DSP Builder Generated Files](#) on page 63
- [Creating a New DSP Builder Design with the DSP Builder New Model Wizard](#)  
Intel recommends you create new designs with the DSP Builder New Model Wizard. Alternatively, you can copy and rename a design example.
- [Creating a New Design by Copying a DSP Builder Design Example](#)
- [DSP Builder Advanced Blockset Generated Files](#)  
DSP Builder generates the files in a directory structure at the location you specify in the **Control** block, which defaults to `..\rtl` (relative to the working directory that contains the `.mdl` file)
- [Control](#)  
The **Control** block specifies information about the hardware generation environment and the top-level memory-mapped bus interface widths.

### 3. DSP Builder Design Flow

Figure 3. Design Flow



1. [Implementing your Design in DSP Builder Advanced Blockset](#) on page 24
2. [Verifying your DSP Builder Advanced Blockset Design in Simulink and MATLAB](#) on page 37
3. [Exploring DSP Builder Advanced Blockset Design Tradeoffs](#) on page 41
4. [Verifying your DSP Builder Design with C++ Software Models](#) on page 47
5. [Verifying your DSP Builder Advanced Blockset Design in the ModelSim Simulator](#) on page 50
6. [Verifying Your DSP Builder Design in Hardware](#) on page 51

7. [Integrating Your DSP Builder Advanced Blockset Design into Hardware](#) on page 63

## 3.1. Implementing your Design in DSP Builder Advanced Blockset

1. [Dividing your DSP Builder Design into Subsystems](#) on page 24
2. [Connecting DSP Builder Subsystems](#) on page 24
3. [Creating a New Design by Copying a DSP Builder Design Example](#) on page 33
4. [Vectorized Inputs](#) on page 36

### 3.1.1. Dividing your DSP Builder Design into Subsystems

1. Consider how to divide your design into subsystems. A hierarchical approach makes a design easier to manage, more portable thus easier to update, and easier to debug. If it is a large design, it also makes design partition more manageable.
2. Decide on your timing constraints. DSP Builder advanced blockset achieves timing closure based on your timing constraints, namely sample rate and clock rate. A modular design with well-defined subsystem boundaries, allows you to precisely manage latency and speed of different modules thus achieving timing closure effortlessly.
3. Consider the following factors when dividing your design into subsystems:
  - Identify the functionality of each submodule of your algorithm, and if you can partition your design into different functional subsystems.
  - In multirate designs consider the sample rate variation at different stages of a datapath. Try not to involve too many different sample rates within a subsystem.
  - If your design has a tight latency requirement, use latency management to define the boundary of a subsystem. DSP Builder advanced blockset applies latency constraints on a subsystem basis.
4. To simplify synchronization, implement modules, which DSP Builder can compute in parallel, in the same subsystem. DSP Builder can apply the same rules more easily to each of the parallel paths. Do not worry about constraining the two paths that may otherwise have different latencies.

### 3.1.2. Connecting DSP Builder Subsystems

To connect DSP Builder **IP** library blocks or **Primitive** subsystems, connect `<valid, channel, data>` sets.

1. [DSP Builder Block Interface Signals](#) on page 25
2. [Periods](#) on page 28
3. [Sample Rate](#) on page 28
4. [Building Multichannel Systems](#) on page 29
5. [Channelization for Two Channels with a Folding Factor of 3](#) on page 29
6. [Channelization for Four Channels with a Folding Factor of 3](#) on page 30
7. [Synchronization and Scheduling of Data with the Channel Signal](#) on page 31



8. [Simulink vs Hardware Design Representations](#) on page 32

### 3.1.2.1. DSP Builder Block Interface Signals

DSP Builder designs have three basic interface signals: `valid`, `channel`, and `data`.

The `channel` (`uint8`) signal is a synchronization counter for multiple channel data on the data signals. Typically, it increments from 0 with the changing channels across the data signals within a frame of data

The data signals can be any number of synchronized signals carrying single or multichannel data.

The `valid` (`ufix(1)` or `bool`) signal indicates whether the concurrent data and channel signals have valid information (1), are unknown (0), or do not care (0).

DSP Builder uses these three synchronized signals, to internally connect **IP** or synthesized subsystems and externally connect upstream and downstream blocks. Thus these three signals connect most of the blocks in a DSP Builder advanced blockset design.

Only one set of `valid`, `channel`, and `data` signals can exist in a **IP** and synthesized subsystem. But multiple data signals can exist in a customized synthesizable subsystem.

Data on the `data` wire is only valid when DSP Builder asserts `valid` high. During this clock cycle, `channel` carries an 8-bit integer channel identifier. DSP Builder preserves this channel identifier through the datapath, so that you can easily track and decode data.

This simple protocol is easy to interface with external circuitry. It avoids balancing delays, and counting cycles, because you can simply decode the `valid` and `channel` signals to determine when to capture the data in any downstream blocks. DSP Builder distributes the control structures in each block of your design.

In Primitive subsystems, DSP Builder guarantees all signals that connect to **ChannelOut** blocks line up in the same clock cycle. That is, the delays balance on all paths from and to these blocks. However, you must ensure all the signals arrive at a **ChannelIn** block in the same clock cycle.

The **IP** library blocks follow the same rules. Therefore, it is easy to connect **IP** blocks and Primitive subsystems.

The **IP** library filters all use the same protocol with an additional simplification—DSP Builder produces all the channels for a frame in a multichannel filter in adjacent cycles, which is also a requirement on the filter inputs. If a FIR filter needs to use flow control, pull down the `valid` signal between frames of data—just before you transmit channel 0 data.

The same `<data, valid, channel>` protocol connects all **CIC** and **FIR** filter blocks and all subsystems with **Primitive** library blocks. The blocks in the **Channel Filter** and **Waveform** library support separate real and imaginary (or sine and cosine) signals. The design may require some splitting or combining logic when using the mixer blocks. Use a Primitive subsystem to implement this logic.

1. [Multichannel Systems with IP Library Blocks](#) on page 26

- 2. Valid, Channel, and Data Examples on page 26

**Related Information**

- Channel In (ChannelIn) on page 359
- Channel Out (ChannelOut) on page 360

**3.1.2.1.1. Multichannel Systems with IP Library Blocks**

**IP** library blocks are vectorizable, if data going into a block is a vector requiring multiple instances. For example, for a FIR filter, DSP Builder creates multiple FIR blocks in parallel behind a single **IP** block. If a decimating filter requires a smaller vector on the output, DSP Builder multiplexes data from individual subfilters onto the output vector automatically, to avoid custom glue logic.


**IP** library blocks typically take a channel count as a parameter, which is simple to conceptualize. DSP Builder numbers the channels 0 to  $(N - 1)$ , and you can use the channel indicator at any point to filter out some channels. To merge two streams, DSP Builder creates some logic to multiplex the data. Sequence and counter blocks regenerate `valid` and `channel` signals.

**3.1.2.1.2. Valid, Channel, and Data Examples**

In your design you have a clock rate  $N$  (MHz) and a per-channel sample rate  $M$  (Mps). If  $N = M$ , DSP Builder receives one new data sample per channel every clock cycle.

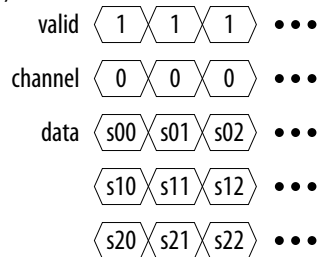
**Figure 4. Single Channel Design**

The frame length, which is the number of clock cycles between data updates for a particular channel, is 1. The out channel count starts (from zero) every clock cycle.  $sPQ$  = the  $Q$ th data sample for channel  $P$ .




**Figure 5. Multichannel Design**

If the data is spread across multiple wires, even for multiple channels, the frame length is 1. The channel signal number, which is a channel synchronization counter, rather than an explicit number expressing the actual channels, is again zero on each clock cycle.




**Figure 6. Single Channel  $n > M$**

DSP Builder receives new data samples only every  $N/M$  clocks. If  $N = 300$  MHz and  $M = 100$  Msps, DSP Builder gives new data every 3 clock cycles. DSP Builder does not know what the data is on the intervening clocks, and sets the valid to low (0). X is unknown or do not care. The frame length is 3 because of a repeating pattern of channel data every 3 clock cycles




**Figure 7. Single Channel  $n > M$  and Two Data Channels**

If  $N = 300$  MHz and  $M = 100$  Msps, with two data channels, the data wire carries the sample for the first channel, the data for the second channel, then a cycle of unknown: The channel signal now increments as DSP Builder receives the different channel data through the frame.




**Figure 8. Three Channels**

If  $N = 300$  MHz and  $M = 100$  Msps, the frame is full along the single data wire.




**Figure 9. Four Channels**

The data now spreads across multiple data signals as one wire is not enough to transmit four channels of data in three clock cycles. DSP Builder attempts to distribute the channels evenly on the wires that it has to use:




**Figure 10. Five Channels**

The data spreads across two data signals that transmit five channels of data in three clock cycles. DSP Builder packs the five channels of data as three on the first wire and two on the second. The channel signal still counts up from zero at the start of each frame and that it specifies a channel synchronization count, rather than expressing all the channels received on a particular clock (which requires as many channel signals as data signals). The valid signal also remains one-dimensional, which can under-specify the validity of the concurrent data if, in a particular frame, channel 0 is valid but channel 3 (received on the same clock) is not. In the five-channel example, DSP Builder receives data for channel 2 on the first data signal at the same time as the invalid data on the second data signal. You require some knowledge of the number of channels transmitted.



**Figure 11. Single Channel  $n < M$**

DSP Builder receives multiple ( $M/N$ ) data samples for a particular channel every clock cycle—super-sample data. If  $N = 200$  MHz and  $M = 800$  Msps, you see a single channel with four new data samples every clock



### 3.1.2.2. Periods

For any data signal in a DSP Builder design, the FPGA clock rate to sample rate ratio determines the *period* value of this data signal. In a multirate design, the signal sample rate can change as the data travels through a decimation or interpolation filter. Therefore *period* at different stages of your design may be different.

In a multichannel design, *period* also decides how many channels you can process on a wire, or on one signal. Where you have more channels than you can process on one path, or wire, in a conventional design, you need to duplicate the datapath and hardware to accommodate the channels that do not fit in a single wire. If the processing for each channel or path is not exactly the same, DSP Builder advanced blockset supports vector or array data and performs the hardware and datapath duplication for you. You can use a wire with a one dimensional data type to represent multiple parallel datapaths. DSP Builder **IP** and **Primitive** library blocks, such as adder, delay and multiplier blocks, all support vector inputs, or fat wires, so that you can easily connect models using a single bus as if it is a single wire.

### 3.1.2.3. Sample Rate

The DSP Builder sample rate may exceed the FPGA clock rate, such as in a super sample rate system, for example in high-speed wireless front-end designs. In a radar or direct RF system with GHz digital-to-analog converters (DAC), the signal driving the DAC can have a sample rate in the GHz range. These high-speed systems require innovative architectural solutions and support for high-speed parallel processing. DSP Builder advanced blockset interpolation filter IP has built in support for super-sample rate signals, and the vector support of its **Primitive** library makes it easy for you to design your super-sample rate module. However, for a super-sample rate design, you must understand how channels are distributed across multiple wires as arrays, and how they are allocated among time slots available on each wire.

Use the following variables to determine the number of wires and the number of channels each wire carries by parameterization:

- *ClockRate* is the system clock frequency.
- *SampleRate* is the data sample rate per channel (MSPS).
- *ChanCount* is the number of channels.

*Note:* Channels are enumerated from 0 to  $ChanCount - 1$ .

- The *Period* (or folding factor) is the ratio of the clock rate to the sample rate and determines the number of available time slots:  
 $Period = \max(1, \text{floor}(ClockRate/SampleRate))$
- The *WiresPerChannel* is the number of wires per channel:  
 $WiresPerChannel = \text{ceil}(SampleRate/ClockRate)$
- The *WireGroups* is the number of wire groups to carry all the channels regardless of channel rate:  
 $WireGroups = \text{ceil}(ChanCount / Period);$
- The number of channel wires the design requires to carry all the channels is the number of channels divided by the folding factor (except for supersampled filters):  
 $ChanWireCount = WiresPerChannel \times WireGroups$
- The number of channels carried per wire is the number of channels divided by the number of channels per wire:  
 $ChanCycleCount = \text{ceil}(ChanCount/WireGroups)$

*Note:* The channel signal counts through 0 to  $ChanCycleCount - 1$ .

#### 3.1.2.4. Building Multichannel Systems

To build multichannel systems, use the required channel count, rather than a single channel system and scaling it up. **Primitive** subsystems contain **ChannelIn** and **ChannelOut** blocks, but do not have explicit support for multiple channels.

1. To create multichannel logic, draw out the logic required for your design to create a single channel version.
2. To transform to a multichannel system, increase all the delays by the channel count required.
3. Use a mask variable to create a parameterizable component.

##### 3.1.2.4.1. Multichannel Systems with IP Library Blocks

**IP** library blocks are vectorizable, if data going into a block is a vector requiring multiple instances. For example, for a FIR filter, DSP Builder creates multiple FIR blocks in parallel behind a single **IP** block. If a decimating filter requires a smaller vector on the output, DSP Builder multiplexes data from individual subfilters onto the output vector automatically, to avoid custom glue logic.

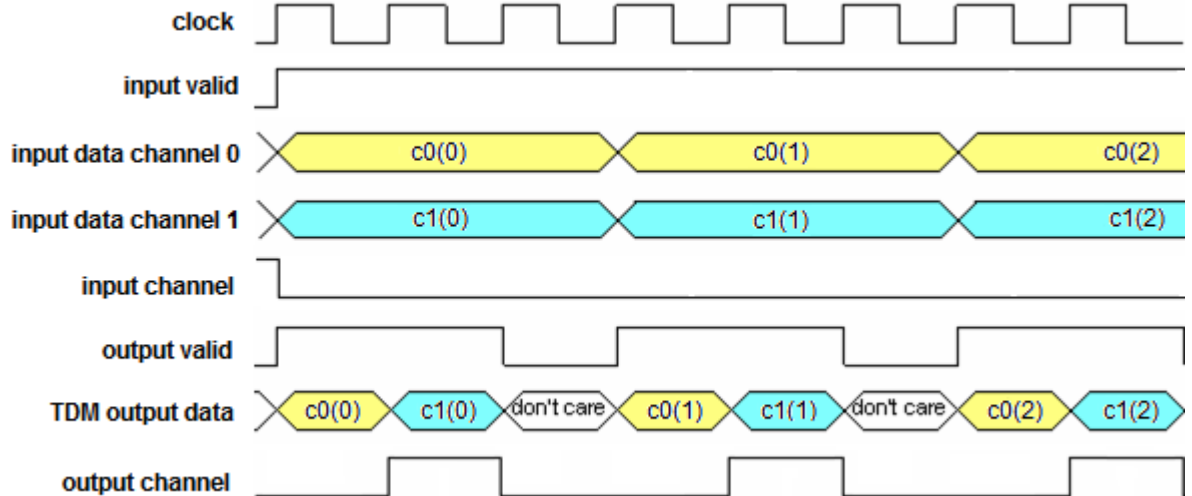
**IP** library blocks typically take a channel count as a parameter, which is simple to conceptualize. DSP Builder numbers the channels 0 to  $(N - 1)$ , and you can use the channel indicator at any point to filter out some channels. To merge two streams, DSP Builder creates some logic to multiplex the data. Sequence and counter blocks regenerate valid and channel signals.

##### 3.1.2.5. Channelization for Two Channels with a Folding Factor of 3

If the number of channels is greater than the period, multiple wires are required. Each **IP** block in your design is internally vectorized to build multiple blocks in parallel.

**Figure 12. Channelization for Two Channels with a Folding Factor of 3**

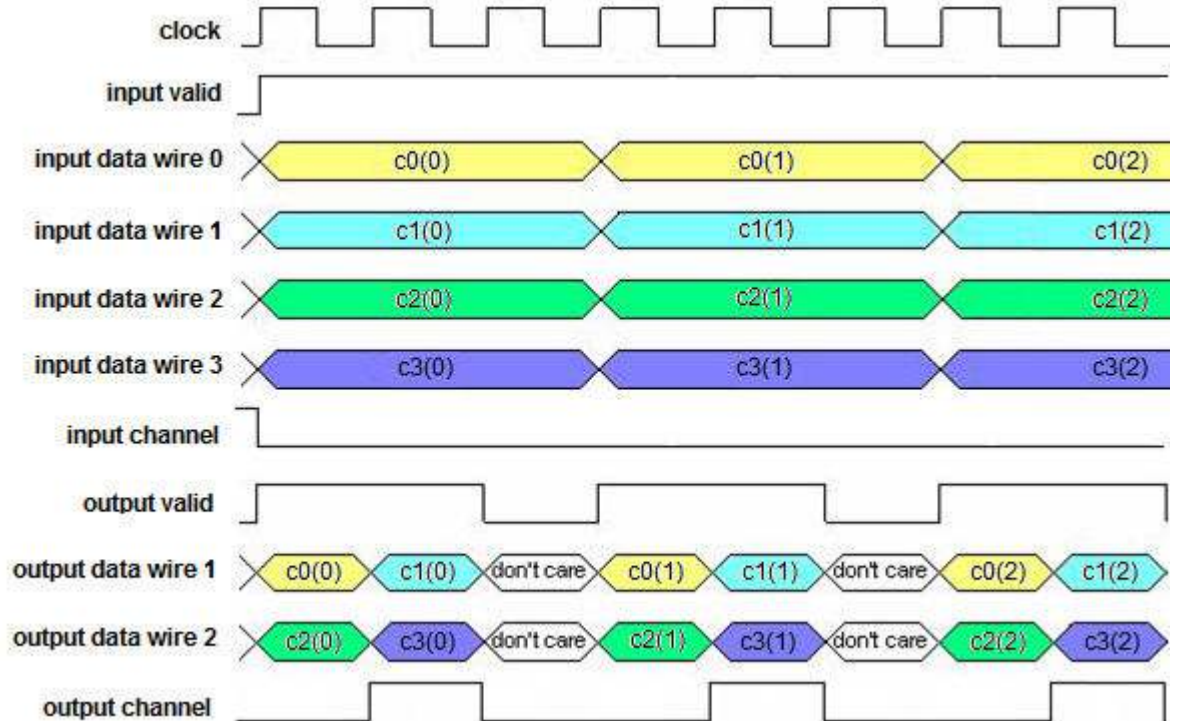
Combines two input channels into a single output wire (ChanCount = 2, ChanWireCount = 1, ChanCycleCount = 2). Three available time slots exist in the output channel and every third time slot has a "don't care" value when the valid signal is low. The value of the channel signal while the valid signal is low does not matter.



### 3.1.2.6. Channelization for Four Channels with a Folding Factor of 3

**Figure 13. Channelization for Four Channels with a Folding Factor of 3**

Combines four input channels into two wires (ChanCount = 4, ChanWireCount = 2, ChanCycleCount = 2). In two wires are required to carry the four channels and the cycle count is two on each wire. DSP Builder distributes the channels evenly on each wire leaving the third time slot as do not care on each wire




*Note:* The generated Help page for the block shows the input and output data channel format that the FIR or CIC filter use after you have run a Simulink simulation.

### 3.1.2.7. Synchronization and Scheduling of Data with the Channel Signal

DSP Builder specifies the channel data separation per wire. The channel signal counts from 0 to `ChanCycleCount - 1` in synchronization with the data. Thus, for `ChanCycleCount = 1`, the channel signal is the same as the channel count, enumerated 0 to `ChanCount - 1`.


For more than a single data wire, it is not equal to the channel count on data wires, but specifies the synchronous channel data alignment across all the data wires. For example,

**Figure 14. Four Channels on One Wire with no invalid cycles.**



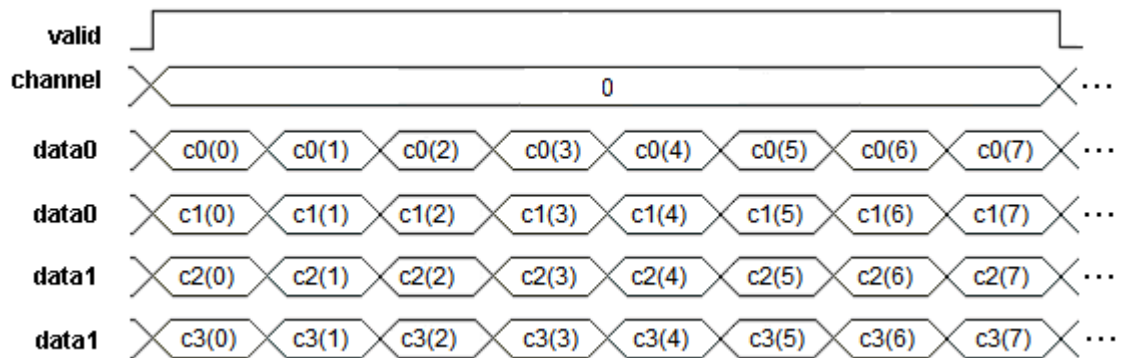
For a single wire, the channel signal is the same as a channel count. However, for `ChanWireCount > 1`, the channel signal specifies the channel data separation per wire, rather than the actual channel number: it counts from 0 to `ChanCycleCount - 1` rather than 0 to `ChanCount - 1`.

**Figure 15. Four Channels on Two Wires with no invalid cycles.**



The channel signal remains a single wire, not a wire for each data wire. It counts over 0 to `ChanCycleCount - 1`.


**Figure 16. Four Channels on Four Wires**



### 3.1.2.8. Simulink vs Hardware Design Representations

Simulink shows the **IP** block as a single block, but the data input and output wires from the **IP** blocks show as a vector with multiple dimension. Multiple wires accommodate all the channels and the Simulink model uses a vector of width 2.

**Figure 17. Simulink and Hardware Representations of a Single Rate FIR Filter**



*Note:* To display the ChanWireCount in Simulink, point to **Port/Signal Displays** in the Format menu and click **Signal Dimensions**.

In a typical wideband CDMA macro-cell system, the DUC module in the RF card needs to process eight inphase (I) and quadrature (Q) data pairs, resulting in 16 independent channels on the datapath. The input sample rate to a DUC is at sample rate 3.84 MHz as defined in the 3GPP specification. A high-performance FPGA running at 245.76 MHz typically maximizes parallel processing power.

**Figure 18. 16-channel WCDMA DUC Design Shows how channel's distribution on wires change in a multirate system.**





**Table 5. 16-channel WCDMA DUC Design**

Signals	Clock Rate (MHz)	ChanCount	Data Sample Rate (MSPS)	Period	Data Signal Pattern	Interpolation Factor
Input to FIR1	245.76	16	3.48	64	I1, I2, ...I8, Q1, ... Q8, zeros(1, 64-16)	2
Input to FIR2	245.76	16	7.68	32	I1, I2, ...I8, Q1, ... Q8, zeros(1, 32-16)	2
Input to CIC	245.75	16	15.36	16	I1, I2, ...I8, Q1, ... Q8	8
Output of CIC	245.75	16	122.88	2	I1, I2, I3, I4, I5, I6, I7, I8, Q1, Q2, Q3, Q4, Q5, Q6, Q7, Q8	8

In this example, the input data at low sample rate 3.84 MSPS can accommodate all channels on a single wire. So the *ChanWireCount* is 1. In fact more time slots are available for processing, since *period* is 64 and only 16 channels are present to occupy the 64 time slots. Therefore the *ChanCycleCount* is 16, which is the number of cycles occupied on a wire. As the data travels down the up conversion chain, its sample rate increases and in turn *period* reduces to a smaller number. At the output of CIC filter, the data sample rate increases to 122.88 Msps, which means only two time slots are available on a wire. As there are 16 channels, spread them out on 8 wires, where each wire supports two channels. At this point, the *ChanWireCount* becomes 8, and *ChanCycleCount* becomes 2. The *ChanCycleCount* does not always equal *period*, as the input data to FIR1 shows.

For most systems, sample rate is less than clock rate, which gives *WirePerChannel*=1. In this case, *ChanWireCount* is the same as *WireGroups*, and it is the number of wires to accommodate all channels. In a super-sample rate system, a single channel's data needs to be split onto multiple wires. Use parallel signals at a clock rate to give an equivalent sample rate that exceeds the clock rate. In this case, *WiresPerChannel* is greater than one, and *ChanWireCount* = *WireGroups* × *WiresPerChannel* because one channel requires multiple wires.

When connecting two modules in DSP Builder, the output interface of the upstream module must have the same *ChanWireCount* and *ChanCycleCount* parameters as the input interface of the downstream module.

### Related Information

[AN 544: Digital Modem Design with the DSP Builder Advanced Blockset.](#)

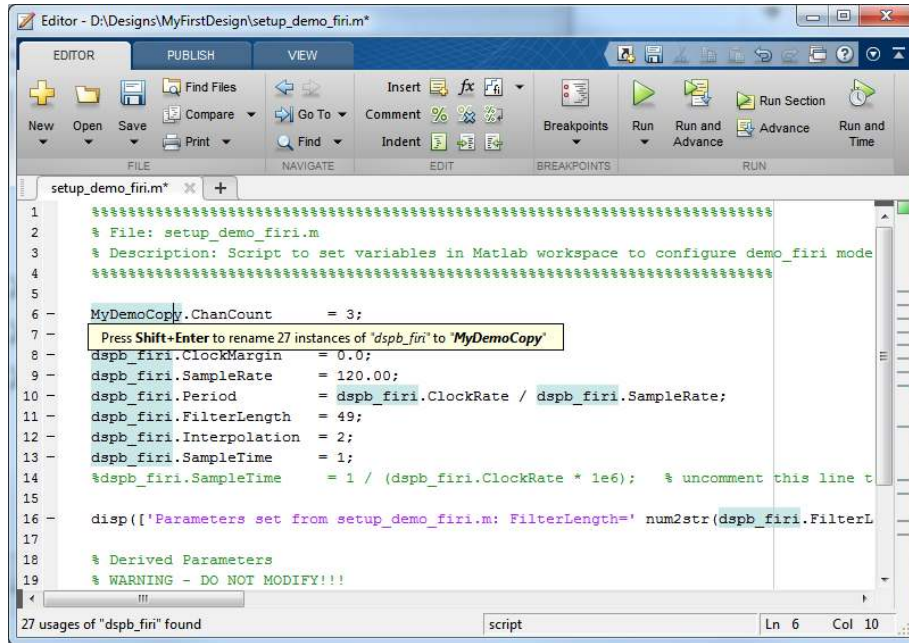
For more information about channelization in a real design

### 3.1.3. Creating a New Design by Copying a DSP Builder Design Example

Start DSP Builder in MATLAB.

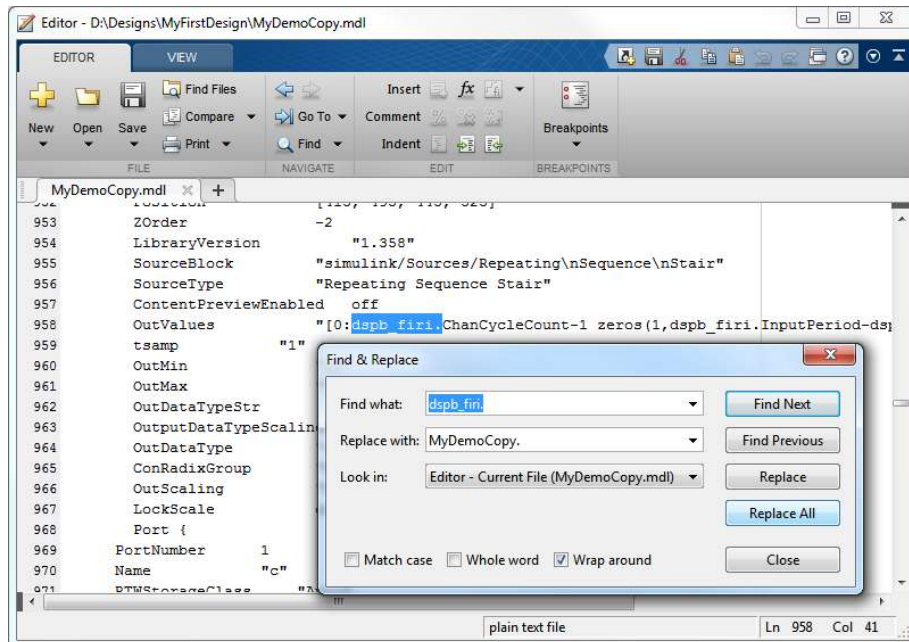
1. Copy and rename the model file to `<model_name>.mdl` (MDL format, not SLX) and the set-up script to `setup_<model_name>.m`.
2. Open the set-up script in the MATLAB Editor.
3. Change the name of the parameter structure so that it does not conflict with the original design example.

Figure 19. Renaming the Parameter



4. Open the new model file as text and globally replace the parameter structure to match.

Figure 20. Replace Parameter Structure



5. Open the model.
6. Click **File > Model Properties > Model Properties > Callbacks** to call the new set-up script.
7. Save the model in .mdl format  
Intel recommends that you create a Simulink project for your new design.

#### Related Information

- [Starting DSP Builder in MATLAB](#)
- [DSP Builder Advanced Blockset Libraries](#)
- [Simulating, Generating, and Compiling Your Design](#)
- [DSP Builder Menu Options](#)  
Simulink includes a **DSP Builder** menu on any Simulink model window. Use this menu to easily start all the common tasks you need to perform on your DSP Builder model.

### 3.1.3.1. Creating a New Design From the DSP Builder FIR Design Example and Changing the Namespaces

1. Open the FIR design example (**demo\_firi**) from the Filters directory, by typing the following command at the MATLAB command prompt:  

```
demo_firi
```
2. In the **demo\_firi** window (the schematic), double-click on the **EditParams** block to open the setup script `setup_demo_firi.m` in the MATLAB Editor.
3. In the Editor, click **File > Save As** and save as `setup_mytutorial.m` in a different directory, for example `\myexamples`.
4. In the **demo\_firi** window, click **File > Save As** and save as `mytutorial.mdl` in the `\myexamples` directory.
5. In the main MATLAB window, navigate to the `\myexamples` directory.
6. In the Editor, click **Edit > Find And Replace**, enter `dspb_firi` in **Find what:** and `my_tutorial` in **Replace with:**. Click **Replace All**. Click **Close**. This step ensures all the setup variables do not interfere with any other workspace variables.
7. Save **setup\_mytutorial.m**.
8. On the Debug menu click **Run setup\_mytutorial.m** to run the script, which creates the workspace variables to use the schematic design.
9. To ensure MATLAB runs the setup script on opening (so that the design displays correctly) and just before simulation (so that the parameters are up-to-date and reflect any edits made since opening), perform the following steps:
  - a. In the mytutorial window (schematic), on the File menu click **Model Properties**.
  - b. On the Callbacks tab click on **PreLoadFcn** and replace `setup_demo_firi;` with `setup_mytutorial;`
  - c. Repeat for the **InitFcn**.
  - d. Click **OK**.

10. In the mytutorial window, double-click on the **FilterSystem** subsystem, then double-click on **InterpolatingFIR** block. Replace all instances of **dspb\_firi** with **mytutorial**. Click **OK**. These parameters set up the FIR filter.
11. Double-click on **InChanView**, block replace all instances of **dspb\_firi** with **mytutorial**, click **OK**.
12. Repeat for the **OutChanView** block and the following blocks:
  - The input stimulus generation blocks:
    - **Sine Wave**
    - **Const**
    - **Impulse**
    - **Random**
    - **Channel Counter**
    - **Valid Sequence**
  - The downsample blocks:
    - **InDownsample**
    - **OutDownsample**.
  - Spectrum analyzers
    - **InSpectrum** (on the axis properties tab)
    - **Outspectrum** (on the axis properties tab)
13. Change the simulation stop time from **20000\*dspb\_firi.SampleTime** to **20000\*mytutorial.SampleTime**.
14. Change the title and save your new design.

### 3.1.4. Vectorized Inputs

Use vector data inputs and outputs for DSP Builder **IP** and **Primitive** library blocks when the clock rate is insufficiently high to carry the total aggregate data. For example, 10 channels at 20 MSPS require  $10 \times 20 = 200$  MSPS aggregate data rate. If the system clock rate is set to 100 MHz, two wires must carry this data, and so the Simulink model uses a vector of width 2.

Unlike traditional methods, you do not need to manually instantiate two **IP** blocks and pass a single wire to each in parallel. Each **IP** block internally vectorizes. DSP Builder uses the same paradigm on outputs, where it represents high data rates on multiple wires as vectors.

Each **IP** block determines the input and output wire counts, based on the clock rate, sample rate, and number of channels.

Any rate changes in the **IP** block affect the output wire count. If a rate change exists, such as interpolating by two, the output aggregate sample rate doubles. DSP Builder packs the output channels into the fewest number of wires (vector width) that supports that rate. For example, an interpolate by two FIR filter may have two wires at the input, but three wires at the output.

The **IP** block performs any necessary multiplexing and packing. The blocks connected to the inputs and outputs must have the same vector widths, which Simulink enforces. Resolve vector width errors by carefully changing the sample rates.

*Note:* Most **Primitive** library blocks also accept vector inputs.

## 3.2. Verifying your DSP Builder Advanced Blockset Design in Simulink and MATLAB

Use this early verification to focus on the functionality of your algorithm, then iterate the design implementation if needed. DSP Builder generates synthesizable VHDL for the design at the start of every Simulink simulation. DSP Builder generates an automatic testbench for the whole design and each subsystem. You can use these testbenches to play data that the Simulink simulation captures through the generated VHDL in ModelSim and confirm the results are identical.

1. [Verifying your DSP Builder Advanced Blockset Design with a Testbench](#) on page 37
2. [Running DSP Builder Advanced Blockset Automatic Testbenches](#) on page 38
3. [Using DSP Builder Advanced Blockset References](#) on page 41
4. [Setting Up Stimulus in DSP Builder Advanced Blockset](#) on page 41
5. [Analyzing your DSP Builder Advanced Blockset Design](#) on page 41

### 3.2.1. Verifying your DSP Builder Advanced Blockset Design with a Testbench

A DSP Builder design testbench is all the subsystems above the subsystem with the **Device** block. Many of the features of DSP Builder are more accessible if you develop the testbench flexibly.

1. Before you start implementing your algorithm, consider the modules that connect to and from your design. Understanding the interface to neighboring modules helps you to use the correct stimulus.
2. Consider the sequence of events that you want to test.
3. If multiple channels of data enter your design, align them properly to follow the DSP Builder advanced blockset data format.
4. Plan your testbench, before you start your design, to allow you to verify and debug your implementation during the design phase.
5. DSP Builder advanced blockset uses a standard interface protocol. Ensure every **IP** or customized block follows this protocol. The input and output signals of your hierarchical design have a common interface.
6. Bring the output signals of subsystems to the top-level design.
7. When you have the top-level testbench in place, debug your subsystems at all levels with the visualization features in Simulink and MATLAB.

### 3.2.1.1. Visualization Features

When designing with DSP Builder advanced blockset, use the following visualization features of MATLAB and Simulink:

- **OutScope** block. In addition to exporting data to work space for analysis, you can use the **OutScope** block to visualize a signal or multiple signals. The **OutScope** block probes and displays data on a wire or a bus relative to the time samples, which is useful when debugging your design.
- **OutputSpectrum** block. You can also use the **OutputSpectrum** block, which displays the signal spectrum in real time, when your design has filtering or FFT.
- Fixed-point toolbox. When dealing with bit growth and quantization, the fixed-point toolbox can be a valuable tool. You can even visualize the dynamic range of a signal by looking at the histogram of the signal.

### 3.2.2. Running DSP Builder Advanced Blockset Automatic Testbenches

Generally, for testbenches, click **DSP Builder > Verify Design**. To run a single subsystem (or the whole design) in an open ModelSim window, click via **DSP Builder > Run ModelSim**. You can use the command line if you want to script testing flows.

- To get a list of the blocks in a design that have automatic testbenches, run the following command in MATLAB:

```
getBlocksWithATBs('model')
```

- To load an automatic testbench from the ModelSim simulator, use the following command:

```
source <subsystem>_atb.do
```

Alternatively, in ModelSim click **Tools > Execute Macro** and select the required .do file.

- You can run an automatic testbench targeting a subsystem or a **IP** block in your design, or you can run an automatic testbench on all of your design.
- To run an automatic testbench from the MATLAB command line on a single entity, use the command `dspba.runModelsimATB`.

- To run testbenches for all subsystems and the device level and set testbench options: in the simulink window, click **DSP Builder > Verify Design** or type:

```
run_all_atbs(<model name>, Run simulation? (0:1), run Quartus (0:1))
```

- To run the device level testbench in the ModelSim simulator, click **DSP Builder > Run ModelSim**.

1. [The dspba.runModelsimATB Command Syntax](#) on page 38
2. [Running All Automatic Testbenches](#) on page 39
3. [The command run\\_all\\_atbs Command Syntax](#) on page 39
4. [Testbench Error Messages](#) on page 40

#### 3.2.2.1. The dspba.runModelsimATB Command Syntax

Use this command to run ModelSim tests.

The `dspba.runModelsimATB` command has the following syntax:

```
dspba.runModelsimATB('model', 'entity', ['rtl_path']);
```

where:

- **model** = design name (without extension, in single quotes)
- **entity** = entity to test (the name of a Primitive subsystem or a ModelIP block, in single quotes)
- **rtl\_path** = optional path to the generated RTL (in single quotes, if not specified the path is read from the **Control** block in your model)

For example:

```
dspba.runModelsimATB('demo_fft16_radix2', 'FFTChip');
```

The return values are in the format [**pass**, **status**, **result**] where:

- **pass** = 1 for success, or 0 for failure
- **status** = should be 0
- **result** = should be a string such as:  
"# \*\* Note: Arrived at end of stimulus data on clk <clock name>"

DSP Builder writes an output file with the full path to the component under test in the working directory. DSP Builder creates a new file with an automatically incremented suffix each time the testbench is run. For example:

```
demo_fft_radix2_DUT_FFTChip_atb.6.out
```

This output file includes the ModelSim transcript and is useful for debugging if you encounter any errors.

### 3.2.2.2. Running All Automatic Testbenches

To automatically run all the individual automatic testbenches in a design use the **command `run_all_atbs`**. Run this command from the same directory that contains the `.mdl` file.

### 3.2.2.3. The command `run_all_atbs` Command Syntax

This command has the syntax:

```
run_all_atbs('model', [runSimulation], [runFit]);
```

where:

- **model** = design name (without extension, in single quotes)
- **runSimulation** = optional flag that runs a simulation when specified (if not specified, a simulation must run previously to generate the required files)
- **runFit** = optional flag which runs the Quartus Prime Fitter when specified

For example:

```
run_all_atbs('demo_agc');
```



```
run_all_atbs('demo_agc', true);
run_all_atbs('demo_agc', false, true);
run_all_atbs('demo_agc', true, true);
```

The return value is 1 if all tests are successful or 0 if any tests fail. The output is written to the MATLAB command window.

### 3.2.2.4. Testbench Error Messages

Typical error messages have the following form:

```
# ** Error (vcom-13) Recompile <path>altera_mf.altera_mf_components because
<path>iee.std_logic_1164 has changed.
```

...

```
# ** Error: <path>mdl_name_system_subsystem_component.vhd(30): (vcom-1195)
Cannot find expanded name: 'altera_mf.altera_mf_components'.
```

...

```
# ** Error: <path>vcom failed.
```

...

```
# At least one module failed to compile, not starting simulation.
```

These errors may occur when a ModelSim precompiled model is out of date, but not automatically recompiled. A similar problem may occur after making design changes when ModelSim has cached a previously compiled model for a component and does not detect when it changes. In either of these cases, delete the **rtl** directory, resimulate your design and run the `dspba.runModelsimATB` or **run\_all\_atbs** command again.

If you run the Quartus Prime Fitter, the command also reports whether the design achieves the target  $f_{MAX}$ . For example:

```
Met FMax Requirement (FMax(291.04) >= Required(200))
```

A summary also writes to a file `results.txt` in the current working directory. For example:

```
Starting demo_agc Tests at 2009-01-23 14:58:48
```

```
demo_agc: demo_agc/AGC_Chip/AGC hardware matches simulation (atb#1):
```

```
PASSED
```

```
demo_agc: Quartus Prime compilation was successful.
```

```
(Directory=../quartus_demo_agc_AGC_Chip_2): PASSED
```

```
demo_agc: Met FMax Requirement (FMax(291.04) >= Required(200)):
```

```
PASSED
```



Finished demo\_agc Tests at 2009-01-23 15:01:59 (3 Tests, 3 Passes, 0 Skipped, 0 Failed (fmax), 0 Failed (non-fmax))

### 3.2.3. Using DSP Builder Advanced Blockset References

All the signals in a DSP Builder advanced blockset design use the built-in Simulink fixed-point types. Be careful if you compare your design with a floating-point reference.

1. Compare your implementation against a reference—a C/C++ bit accurate model, a MATLAB model, or a Simulink design.
2. For a C/C++ model, save your output into a text file and write your C/C++ comparison script.
3. For a MATLAB model, output the DSP Builder advanced blockset testbench data into a workspace or save it into data files.
4. For a Simulink design, put the Simulink model in parallel with your synthesizable design.
5. Use the Simulink scope to compare the two designs.

### 3.2.4. Setting Up Stimulus in DSP Builder Advanced Blockset

1. In your top-level testbench, generate stimulus at real time for both data and control signals. Commonly used test data signals include sine waves, random noise, step functions and constants.
2. Generate channel signals and valid signals as repeated sequences.
3. For simulations and tests, format your `data`, `valid`, or `channel` pair according to the DSP Builder advanced blockset interface protocol in MATLAB or Simulink.

### 3.2.5. Analyzing your DSP Builder Advanced Blockset Design

1. Use Simulink scope blocks.
2. Use the **SpectrumScope** block to check signal spectrum properties, especially in evaluating filter performance.

## 3.3. Exploring DSP Builder Advanced Blockset Design Tradeoffs

Get early estimates of resource utilization before you go to hardware verification, which allows you to experiment with various implementation optimizations early. Access memory-logic tradeoff, or logic-multiplier tradeoff by modifying threshold parameters. You may not need to physically modify the design. DSP Builder can automate design space exploration based on your tradeoff options

1. [Bit Growth](#) on page 42
2. [Managing Bit Growth in DSP Builder Advanced Blockset Designs](#) on page 42
3. [Using Rounding and Saturation in DSP Builder Advanced Blockset Designs](#) on page 42

4. [Scaling with Primitive Blocks](#) on page 43
5. [Changing Data Type with Convert Blocks and Specifying Output Types](#) on page 43

### 3.3.1. Bit Growth

DSP Builder uses the built-in Simulink fixed-point types to specify all fixed-point data. You can display the signals as familiar floating-point types.

Using fixed-point types preserves the extra information of binary point position through hardware blocks, so that it is easy to perform rounding and shifting operations without having to manually track the interpretation of an integer value. A fixed-point type change propagates through your design, with all downstream calculations automatically adjusted.

In a typical mathematical algorithm involving multiplication and addition, data width grows as signals travel through the arithmetic blocks. A large data width implies better accuracy generally, but more hardware resources and potentially lower  $f_{MAX}$  (such as in large adders).

### 3.3.2. Managing Bit Growth in DSP Builder Advanced Blockset Designs

Manage bit growth after you update your design or run a simulation.

1. To display the signal type and width turn on Simulink display of signal types.
2. Manage and control bit width at various stages of your design, either because of hardware resource limitation or  $f_{MAX}$  speed concerns.
3. Track bit growth by studying the algorithm and determining bit width at various stage of the design from the mathematical model of the design.
4. Use Simulink Fixed-Point Toolbox to visualize the bit width distribution at various places of the design. The fixed-point toolbox displays the histogram of datapath signals you log.
5. To log a data signal in your Simulink design, right-click on the wire and select **Signal Properties**.
6. With the histogram decide how many MSBs are unused in the current fixed-point representation, which helps you decide how many MSBs to discard, thus maximizing the dynamic range of your scaled data.

### 3.3.3. Using Rounding and Saturation in DSP Builder Advanced Blockset Designs

**IP** library blocks such as FIR filters produce output data that use full resolution. DSP Builder performs no rounding or saturation on the output data.

1. Use a **Scale** block to provide scaling and control your bit growth before data enters the next stage of your IP or primitive subsystems.  
*Note:* For primitive subsystems, use a **Convert** block to apply rounding and saturation. The **Convert** block does not perform scaling.
2. To reduce bit width of a wide word, use a **Convert** block instead of just forcing output data type in an arithmetic block.

Whether you choose the **Scale** block or **Convert** block to perform rounding and saturation, depends on your algorithm and resource requirement. The **Convert** block does not support scaling, although you can combine a few **Primitive** library blocks to implementing scaling. The **Scale** block allows you to use a different scaling factor on a cycle basis. It supports both amplification and attenuation of data.

### 3.3.4. Scaling with Primitive Blocks

Use **Primitive** library blocks to build your own run-time reconfigurable scaling.

1. Use the left shift operation to remove redundant MSBs; use bit extract to remove LSBs and preserve the MSBs.
2. Choose the number of MSBs to discard with the run-time reconfigurable parameter that comes from an input port.
3. Use a control register to connect to this port, and update the shift value by a processor such as a Nios II processor.
4. If the FPGA clock is low, use this implementation to realize different scaling for different channels. If it is a high speed application and your processor bus updates much slower than logic clock rate, you cannot use this circuit to apply different scaling for different channels.

### 3.3.5. Changing Data Type with Convert Blocks and Specifying Output Types

1. Preserve the real-world value using a **Convert** block.
2. Preserve bit pattern by setting the output data type mode on any other **Primitive** library block or use a **Reinterpretcast** block.

#### Related Information

[Convert](#) on page 314

#### 3.3.5.1. The Convert Block and Real-world Values

The **Convert** block converts a data type to preserve the real-word value and optionally rounds and saturates the data type when not possible. **Convert** blocks can sign extend or discard bits as necessary. Similarly you can convert the same number of bits while preserving the real world value (as far as possible, subject to rounding and saturation).

Figure 21. Convert Block Changing Data Type while preserving real-world value

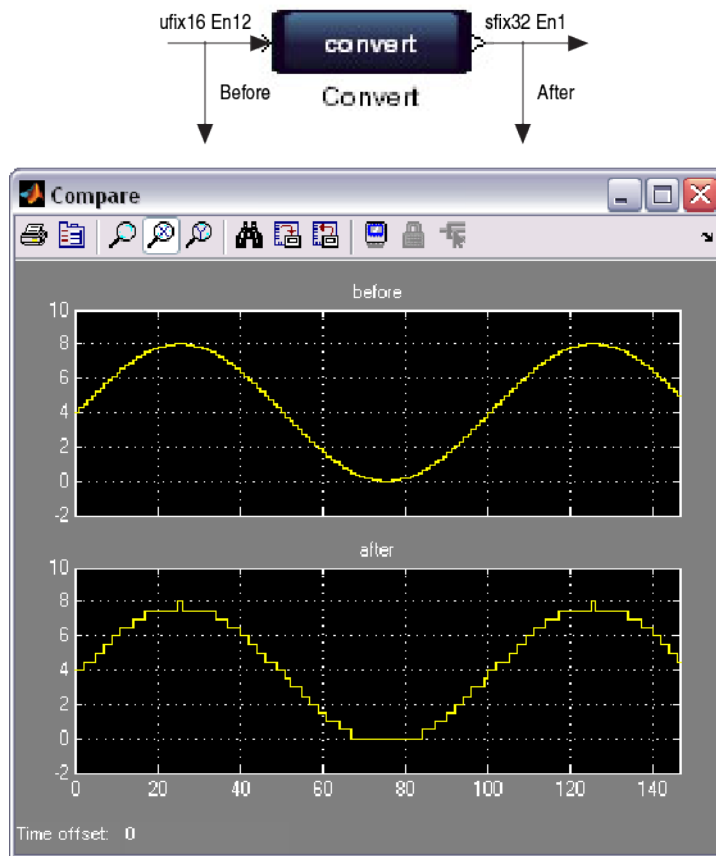
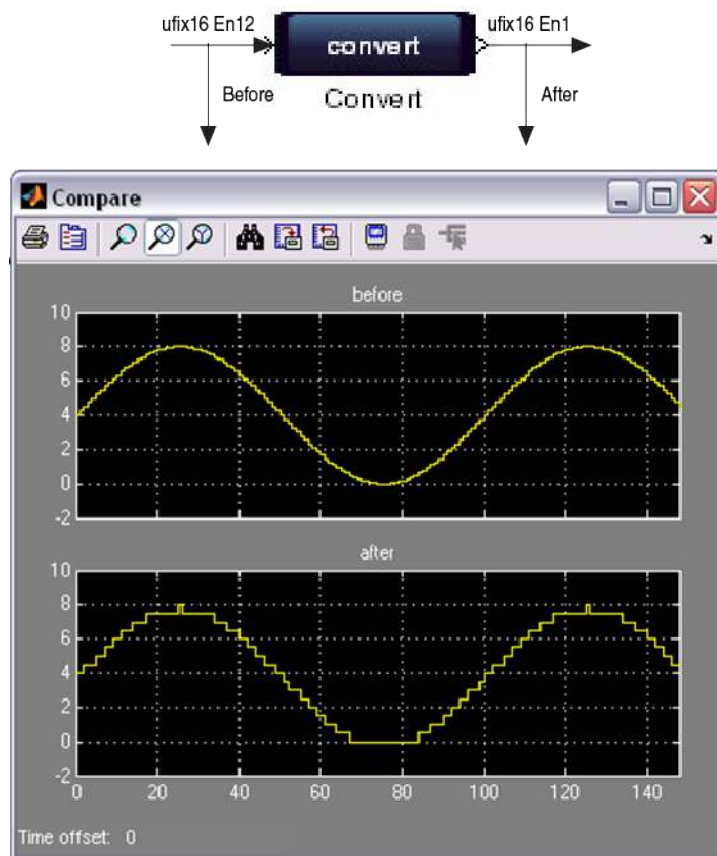


Figure 22. Convert Block Using Same Number of Bits while preserving real-world value



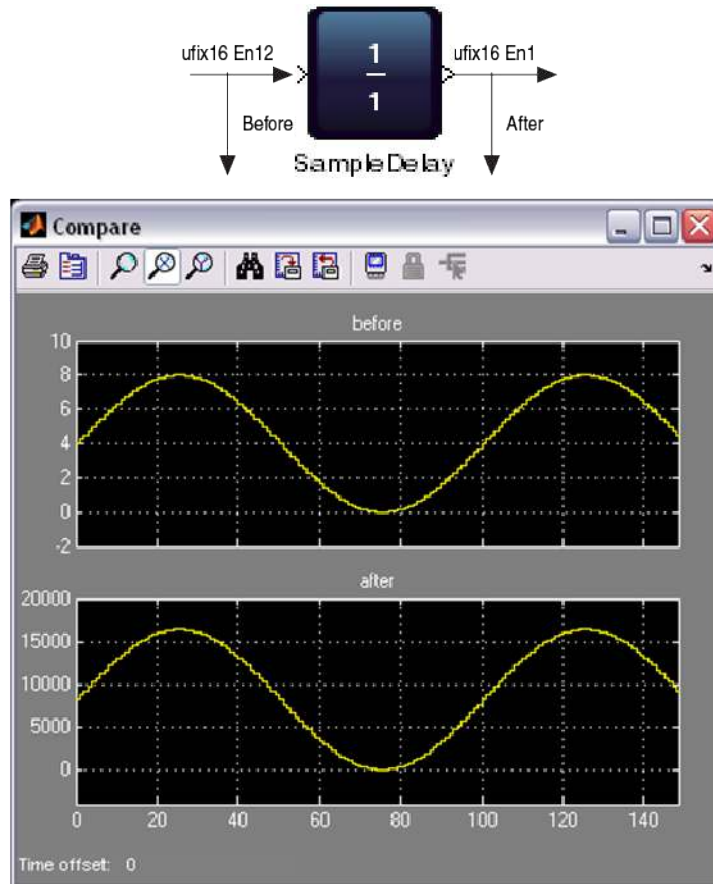
#### Related Information

[Convert](#) on page 314

### 3.3.5.2. Output Data Types on Primitive Blocks

Set the output data type with **Specify via dialog** on any(except **Convert**) DSP Builder **Primitive** library block. For example you can use a zero-length sample delay. Specifying the output type with the dialog box is a casting operation. This operation does not preserve the numerical value, it just preserves the underlying bits. This operation never adds hardware to a block—it just changes the interpretation of the output bits. DSP Builder implements this reinterpretation by aligning the LSBs of the old and new data types. For example, if the new data type has fewer fractional bits than the old data type, the new numerical value is larger than the old numerical value.

Figure 23. SampleDelay Block with number of delays set to 0



For example, a **Mult** block with both input data types specified as **sfix16\_En15** naturally has an output type of **sfix32\_En30**. The specified output data type has two fewer fractional bits than the natural input data type. Therefore, if you specify the output data type as **sfix32\_En28**, the output numerical value is effectively multiplied by four, and a  $1*1$  input gives an output value of 4.

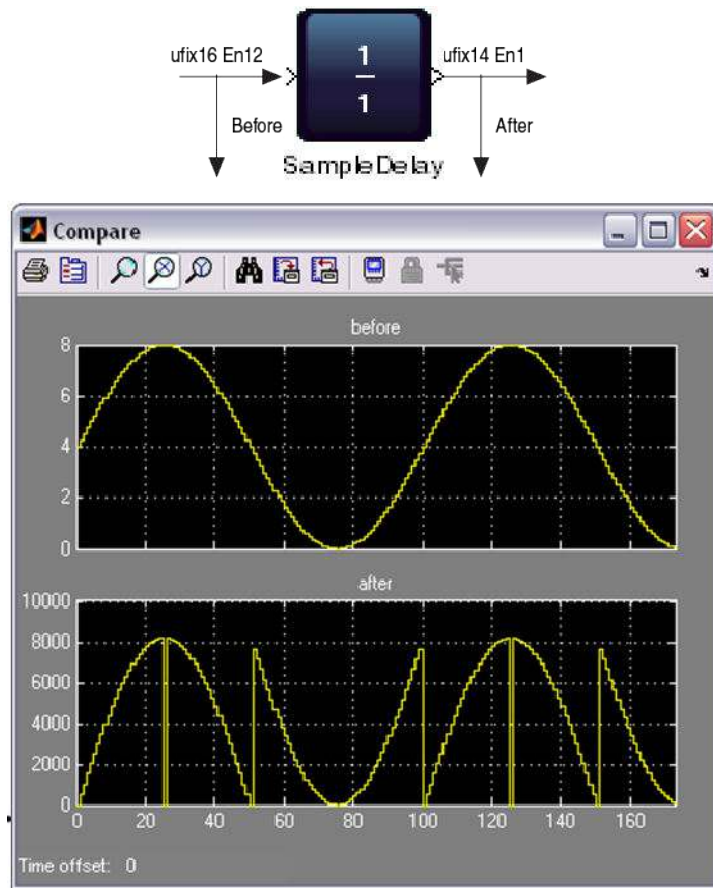
If you specify output data type of **sfix32\_En31**, the output numerical value is effectively divided by two and a  $1*1$  input gives an output value of 0.5.

If you want to change the data type format in a way that preserves the numerical value, use a **Convert** block, which adds the corresponding hardware. Adding a **Convert** block directly after a **Primitive** library block allows you to specify the data type in a way that preserves the numerical value. For example, a **Mult** block followed by a **Convert** block, with input values  $1*1$  always gives output value 1.

To reinterpret the bit pattern and also discard bits, if the type you specify with the **Output data type** is smaller than the natural (inherited) output type, DSP Builder discards the MSBs (most significant bits).

Never set **Specify via dialog** to be bigger than the natural (inherited) bit pattern—DSP Builder performs no zero-padding or sign extension, and the result may generate hardware errors due to signal width mismatches. Use the **Convert** block for any sign extension or zero padding.

Figure 24. SampleDelay Block and Reducing Bit Width



If you want to use sign extends and zero pads to reinterpret the bit pattern, you can combine these methods.

To set a specific format so that DSP Builder can resolve types, for example, in feedback loops, set **Specify via dialog** on an existing **Primitive** library block or insert a zero-cycle sample delay (which generates no hardware and just casts the type interpretation).

To ensure the data type is equal to some other signal data type, force the data type propagation with a Simulink data type propagation block.

#### Related Information

[Primitives Library](#) on page 276

### 3.4. Verifying your DSP Builder Design with C++ Software Models

DSP Builder supports C++ software models for designs that support bit-accurate simulation.

The software model includes a testbench, which is an executable program to check the output of the software models matches the output of Simulink simulation. The generated CMake script creates projects and makefiles (depending on parameters)

that you can use to compile the software model and testbench. The testbench and the CMake script allow you to verify the model functionality. Also, you can use the testbench as a starting point for integration of generated models into a larger, system-level, simulation.

1. In the design's **Control** block turn on **Generate software model**.

The default language is `cpp03` (C++ 2003 standard conformant) and **Generate an ATB (automatic testbench) and CMake build script** is turned on (by default).

2. Turn on **Bit Accurate Simulation** on the **SynthesisInfo** blocks in all subsystems.

You must enable bit-accurate simulation for all subsystems otherwise DSP Builder generates incomplete software models.

3. Compile the design.

DSP Builder creates a directory, `cmodel`, which contains the following files:

- A `cs1.*` header file containing utility functions and implementation details for the generated models.

A `[model/subsystem name]_CModel(.h/.cpp)` pair for each subsystem and the device level system.

A `[model/subsystem name]_atb.cpp` file containing the device level test bench for the model.

A `CMakeFiles.txt/CMakeLists.txt` file containing CMake build scripts for building the ATB executable and model files.

4. Generate the project or makefiles using `CMakeLists.txt`.

For example, to generate Visual Studio 2017 projects, run:

```
cmake -G "Visual Studio 15 2017 Win64"
```

Or to generate a makefile for the release build with symbols on Linux:

```
cmake -G "Unix Makefiles" -DCMAKE_BUILD_TYPE=RelWithDebInfo
```

Refer to the CMake documentation for more options.

The minimum supported compilers are `gcc6.3` and `MSVC14.0`. Other compilers may work but are not supported.

5. Set the `MPFR_INC_PATH`, `MPFR_LIB_PATH`, `MPFR_INC_PATH`, `MPFR_LIB_PATH` options to the include and library directories of builds of the `mpfr` or `mpir` libraries if they are required by the build scripts.

You must set these options otherwise the generation may fail. Your design requires these libraries if the bit widths in the software model are larger than 64-bits or when modeling certain floating-point configurations. These bit-widths may appear on internal signals, not just the inputs and outputs whose widths you specify.

Build instructions and prebuilt binaries are on the `mpfr` or `mpir` websites: <https://www.mpfr.org/> and <http://mpir.org/>



The following example commands specifying both mpir and mpfr paths:

- Windows:

```
cmake -G "Visual Studio 15 2017 Win64" -DMPPIR_INC_PATH="C:\path\to\mpir\3.0.0\win64-vc14\include" -DMPPIR_LIB_PATH="C:\path\to\mpir\3.0.0\win64-vc14\lib\release" -DMPFR_INC_PATH="C:\path\to\mpfr\4.0.1\win64-vc14\include" -DMPFR_LIB_PATH="C:\path\to\mpfr\4.0.1\win64-vc14\lib\release"
```

- Linux:

```
export CC=/path/to/gcc/6.3.0/1/linux64/bin/gcc
export CXX=/path/to/gcc/6.3.0/1/linux64/bin/g++
cmake -G "Unix Makefiles" -DMPPIR_INC_PATH="path/to/mpir/3.0.0/linux64-gc63/include" -
DMPPIR_LIB_PATH="path/to/mpir/3.0.0/linux64-gc63/lib/release" -
DMPFR_INC_PATH="path/to/mpfr/4.0.1/linux64-gc63/include" -
DMPFR_LIB_PATH="path/to/mpfr/4.0.1/linux64-gc63/lib/release"
```

6. On Windows, open the generated solution file and run the compilation. On Linux, run make.

After compilation, DSP Builder creates an executable of the same name as the generated testbench, *<design name>\_CModel\_atb.exe*.

7. Run the .exe with the cmodel directory as the working directory so that the generated stimulus file paths are correct.

If simulation was successful, the executable produces the following output to stdout:

```
Opening stimulus files...
Simulating...
Success! Stimulation matches output stimulus results.
```

8. Refer to the testbench to see how you can integrate the generated models into an existing system.

Subsystems contain structs representing their inputs and outputs. These structs have a generated constructor that reads values from a stimulus file for the testbench.

```
struct IO_xIn
{
    int64_t v; int64_t c; int64_t x; int64_t y;

    IO_xIn()
    : v(0)
    , c(0)
    , x(0)
    , y(0)
    {
    }

    IO_xIn(csl::StimulusFile& stm)
    {
        stm.Get<1>(v); stm.Get<8>(c); stm.Get<27>(x); stm.Get<27>(y);
    }
};
```

When integrating the model, replace the stimulus file constructor by manually setting the input or output values on the struct before using them to drive the model using `read()`, `write()`, or `execute()` functions.

### 3.4.1. Example CMakeList File

Generate project or makefiles using CMakeLists.txt

```
cmake_minimum_required (VERSION 2.11) project (simple_dut_CModel_atb)
set (simple_dut_CModel_atb 1)
set (simple_dut_CModel_atb 0)

# set by user as a hint
set (MPIR_INC_PATH "" CACHE PATH "MPIR include path (hint)") set (MPIR_LIB_PATH
" CACHE PATH "MPIR library path (hint)") set (MPFR_INC_PATH "" CACHE PATH
"MPFR include path (hint)") set (MPFR_LIB_PATH "" CACHE PATH "MPFR library path
(hint)")

option(USE_MPIR "Include and link against the MPIR library for models that
require arbitrary precision" OFF)
option(USE_MPFR "Include and link against the MPFR library for models that
require arbitrary precision floating point" OFF)

include("CMakeFiles.txt")

add_executable(simple_dut_CModel_atb ${cmodeL_SRC}) add_definitions(-
D_CRT_SECURE_NO_WARNINGS)
if (MSVC)
else()
set(CMAKE_CXX_FLAGS_RELEASE "-O1 -DNDEBUG")
endif()

if(USE_MPIR)
add_definitions(-DCSL_USE_MPIR) find_path(MPIR_INC
NAMES mpir.h
HINTS ${MPIR_INC_PATH}
)
find_library(MPIR_LIB NAMES mpir altera_mpir HINTS ${MPIR_LIB_PATH}
)
```

## 3.5. Verifying your DSP Builder Advanced Blockset Design in the ModelSim Simulator

Verify your design in Simulink or the ModelSim simulator with the automatic testbench flow. Also, compare Simulink results with the generated RTL, on all synthesizable IP and primitive subsystems. This final verification before you port the design to system-level integration ensures you should not need to iterate your design.

*Note:* Intel recommends the automatic testbench flow.

1. [Automatic Testbench](#) on page 50
2. [DSP Builder Advanced Blockset ModelSim Simulations](#) on page 51

### 3.5.1. Automatic Testbench

Each **IP** library block, and each synthesized **Primitive** library block writes out test vectors to a stimulus file (\*.stm) during a Simulink simulation run. DSP Builder creates an RTL testbench for each separate entity in your design (that is, for each **IP** block and **Primitive** subsystem). These testbenches replay the test vectors through the generated RTL, and compare the output from the RTL to the output from the

Simulink model. If a mismatch at any cycle exists, the simulation stops and DSP Builder indicates an error. Use these DSP Builder automatic testbenches, to verify the correct behavior of the synthesis engine.

The automatic testbench flow uses a stimulate-and-capture method and is therefore not restricted to a limited set of source blocks. The Simulink simulation stores data at the inputs and outputs of each entity during simulation. Then the testbench for each entity uses this data as a stimulus and compares the ModelSim output to the Simulink captured output. The result indicates whether the outputs match when the `valid` signal is high.

### 3.5.1.1. DSP Builder Advanced Blockset Automatic Testbench Files

**Table 6. Files for an Automatic Testbench**

File Name	Description
<name>.vhd	The HDL that is generated as part of the design (regardless of automatic testbenches).
<name>_stm.vhd	An HDL file that reads in data files of captured Simulink simulation inputs and outputs on <name>
<name>_atb.vhd	A wrapper HDL file that performs the following actions: <ul style="list-style-type: none"> <li>• Declares &lt;name&gt;_stm and &lt;name&gt; as components</li> <li>• Wires the input stimuli read by &lt;name&gt;_atb to the inputs of &lt;name&gt; and the output stimuli and the outputs of &lt;name&gt; to a validation process that checks the captured Simulink data</li> <li>• Channel matches the VHDL simulation of &lt;name&gt; for all cycles where <code>valid</code> is high</li> <li>• Checks that the <code>valid</code> signals match</li> </ul>
<input>/<output>.stm	The captured Simulink data that the <b>ChannelIn</b> , <b>ChannelOut</b> , <b>GPIn</b> , <b>GPout</b> and <b>IP</b> blocks write.

Each block writes a single stimulus file capturing all the signals through it writing them in columns as doubles with one row for each timestep.

The device-level testbenches use these same stimulus files, following connections from device-level ports to where the signals are captured. Device-level testbenches are therefore restricted to cases where the device-level ports are connected to stimulus capturing blocks.

### 3.5.2. DSP Builder Advanced Blockset ModelSim Simulations

ModelSim simulations compare the complete Simulink model with hardware. This comparison uses the same stimulus capture and comparison method as the automatic testbenches.

DSP Builder captures stimulus files on the device level inputs and records Simulink output data on the device level outputs. It creates a ModelSim testbench that contains the HDL generated for the device that the captured inputs feed. It compares the Simulink outputs to the ModelSim simulation outputs in an HDL testbench process, reports any mismatches, and stops the ModelSim simulation.

## 3.6. Verifying Your DSP Builder Design in Hardware

Alternatively, verify the hardware with the system in the loop.

1. Set up verification structures around the DUT using on-chip RAMs. If the design interfaces to off-chip RAM for reading and storing data, the design requires no additional verification structures.
  - a. Add buffers to load with test vectors for DUT inputs and logic to drive DUT inputs with this data.
  - b. Add buffers to store the DUT results.
    - Use a **SharedMem** block from the **Interface** library to implement buffers. DSP Builder automatically generates processor interface to these blocks that it requires to load and read the buffers from MATLAB (with MATLAB API).
    - Use **Counter** blocks from the Primitive library or custom logic to implement a connection between the test buffers and DUT inputs and outputs.
    - Consider using **RegField**, **RegBit**, and **RegOut** blocks from the **Interface** library to control the system and poll the results from MATLAB. DSP Builder automatically generates a processor interface for these blocks.
2. Assemble the high-level system in Platform Designer.
3. Use appropriate Platform Designer library blocks to add debugging interfaces and data storage.
  - a. Add PLLs to generate clocks with the required frequency. You can use separate clocks for the processor interface clock and system clock of the DSP Builder design, if you generate the DSP Builder design with **Use separate bus clock** option.
  - b. Add debug host (JTAG/USB). All memory-mapped read and write requests go through this IP core. Connect it to DSPBA processor interface (Avalon memory-mapped agent) and any other IP that needs to be accessed from host.
  - c. Add the DSP Builder top-level design with the source and sink buffers.
  - d. If you assemble a system with a DSP Builder design that connects to off-chip memory, add an appropriate block to the Platform Designer system and connect it to the DSP Builder block interfaces (Avalon memory-mapped host). Also, connect the debug host to off-chip RAM so the host can access it.
4. Create a Quartus Prime project.
5. Add your high-level Platform Designer system into a top-level module and connect up all external ports.
6. Provide port placement constraints.

If you are using on-chip RAMs for testing and JTAG-based debugging interface, you mainly need to place clock and reset ports. If you use off-chip RAM for data storage, provide more complex port assignments. Other assignments may be required based on the specific design and external interfaces it uses.
7. Provide timing constraints.

Compile the design and load it into the FPGA.

### 3.6.1. Hardware Verification

DSP Builder provides an interface for accessing the FPGA directly from MATLAB. This interface allows you to use MATLAB data structures to provide stimuli for the FPGA and read the results from the FPGA.

This interface provides memory-mapped read and write accesses to your design running on an FPGA using the System Console system debugging tool.

**Table 7. Methods for SystemConsole Class**

Call these methods as `SystemConsole.<method_name>[arguments]`. Use these methods to scan and establish the master connection to the FPGA.

Method	Description
<code>executeTcl(script)</code>	Executes a Tcl script specified through <code>&lt;script&gt;</code> string in SystemConsole.
<code>designLoad(path)</code>	Loads the design (.sof) file specified through <code>&lt;path&gt;</code> parameter to FPGA.
<code>refreshMasters</code>	Detects and lists all available master connections.
<code>openMaster(index)</code>	Creates and returns a master connection to a specified master link. The <code>&lt;index&gt;</code> specifies the index (starting 1) of the connection from the list returned by <code>refreshMasters</code> function. For example, <code>M=SystemConsole.openMaster(1);</code>

**Table 8. Methods for Master Class**

Read and write through a master connection. Call these methods on a master object returned by `SystemConsole.openMaster(index)` method.

Method	Description
<code>close()</code>	Closes the connection associated with the master object. <i>Note:</i> Always call this method when you finish working with current master connection.
<code>setTimeoutValue(timeout)</code>	Use this method to override the default timeout value of 60 seconds for the master connection object. The specified <code>&lt;timeout&gt;</code> value in seconds.
<code>read(type, address, size [, timeout])</code>	Returns a list of <code>&lt;size&gt;</code> number of values of type <code>&lt;type&gt;</code> read from memory on FPGA starting at address <code>&lt;address&gt;</code> . For example, <code>data = masterObj.read('single', 1024, 10)</code> Reads consequent 10 4-byte values (40 bytes overall) with a starting address of 1,024 and returns the results as list of 10 'single' typed values.
<code>write(type, address, data [, timeout])</code>	Writes <code>&lt;data&gt;</code> (a list of values of type <code>&lt;type&gt;</code> ) to memory starting at address <code>&lt;address&gt;</code> . For example: <code>masterObj.write('uint16', 1024, 1:10);</code> Writes values 1 to 10 to memory starting address 1,024, where each value occupies 2 bytes in memory (overall 20 bytes are written).

**Table 9. Parameters for read(type, address, size [, timeout])**

Parameter	Description
<code>&lt;type&gt;</code>	The type of each element in returned array. <ul style="list-style-type: none"> <li>1 byte: 'char', 'uint8', 'int8'</li> <li>2 bytes: 'uint16', 'int16'</li> <li>4 bytes: 'uint32', 'int32', 'single'</li> <li>8 bytes: 'uint64', 'int64', 'double'</li> </ul>
<code>&lt;address&gt;</code>	The start address for the read operation. You can specify as a hexadecimal string.

*continued...*

Parameter	Description
	<i>Note:</i> The address should specify a byte address
<size>	The number of <type> (type specifies 1/2/4/8 bytes based on value) values to read.
<timeout>	An optional parameter to override the default timeout value for this operation only.

**Table 10. Parameters for write(type, address, data [, timeout])**

Parameter	Description
<type>	The type each element in specified <data>. Each type specifies 1/2/4/8 bytes: <ul style="list-style-type: none"> <li>• 1 byte : 'char', 'uint8', 'int8'</li> <li>• 2 bytes: 'uint16', 'int16'</li> <li>• 4 bytes: 'uint32', 'int32', 'single'</li> <li>• 8 bytes: 'uint64', 'int64', 'double'</li> </ul>
<address>	The start address for the write operation. You can specify as a hexadecimal string. <i>Note:</i> The address should be specified as a byte address
<data>	An array or single element data to be written to memory.
<timeout>	An optional parameter to override the default timeout value for this operation only.

1. [Hardware Verification Design Example](#) on page 54

### 3.6.1.1. Hardware Verification Design Example

DSP Builder design example for off chip source and sink buffers.

**Figure 25. Top-Level System**

Shows source and sink buffers and DUT.

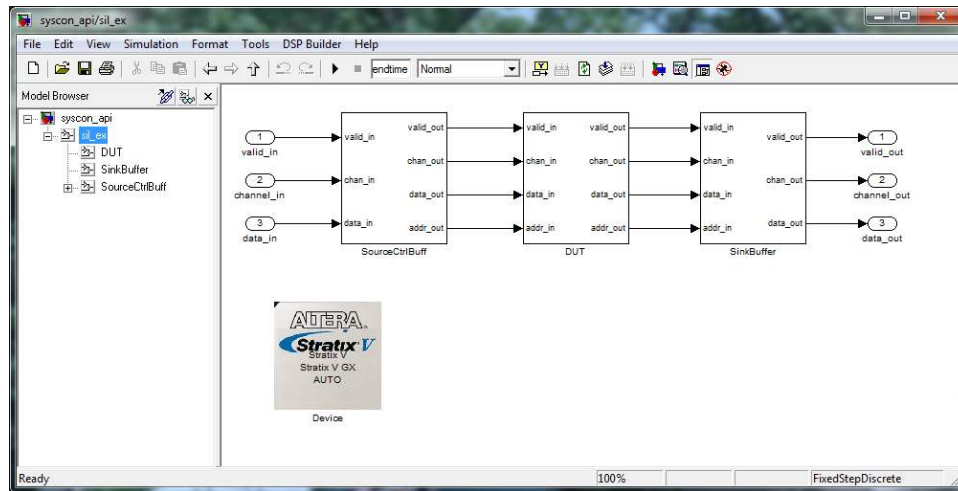


Figure 26. Source Buffer

**AddressGen** block triggers reads from **SharedMem** to drive DUT input. **RegField** initiates execution of the **AddressGen** block from host.

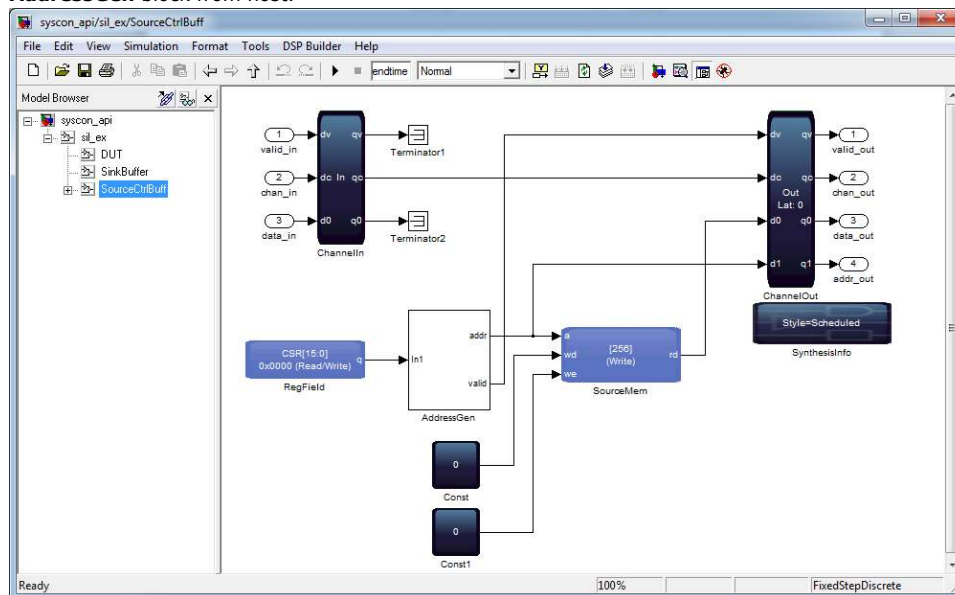


Figure 27. Sink Buffer

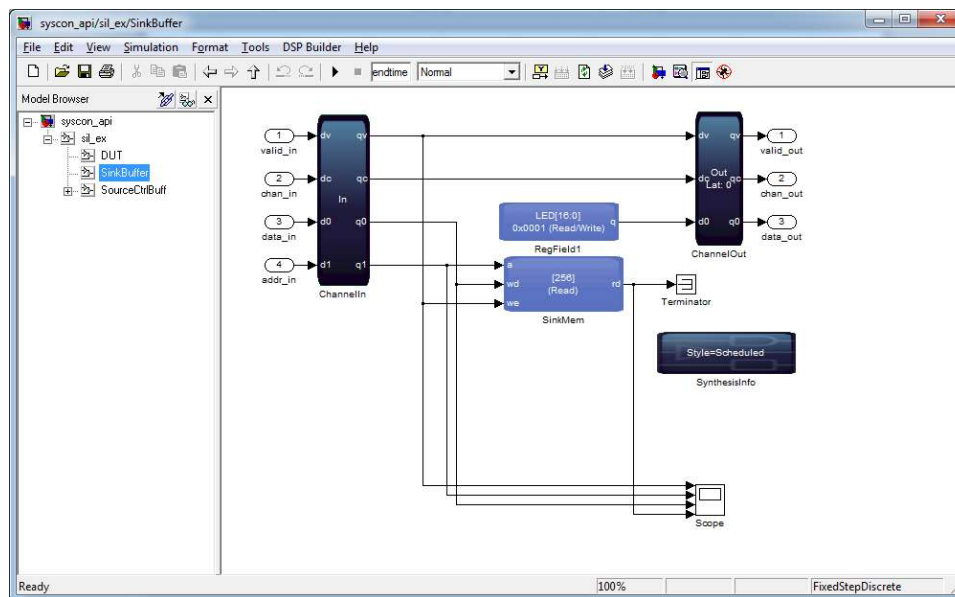


Figure 28. Platform Designer System

**master\_0** is the instance of JTAG debug master. **syscon\_api\_sil\_ex\_0** is the instance of the top-level DSP Builder system, which contains test buffers.

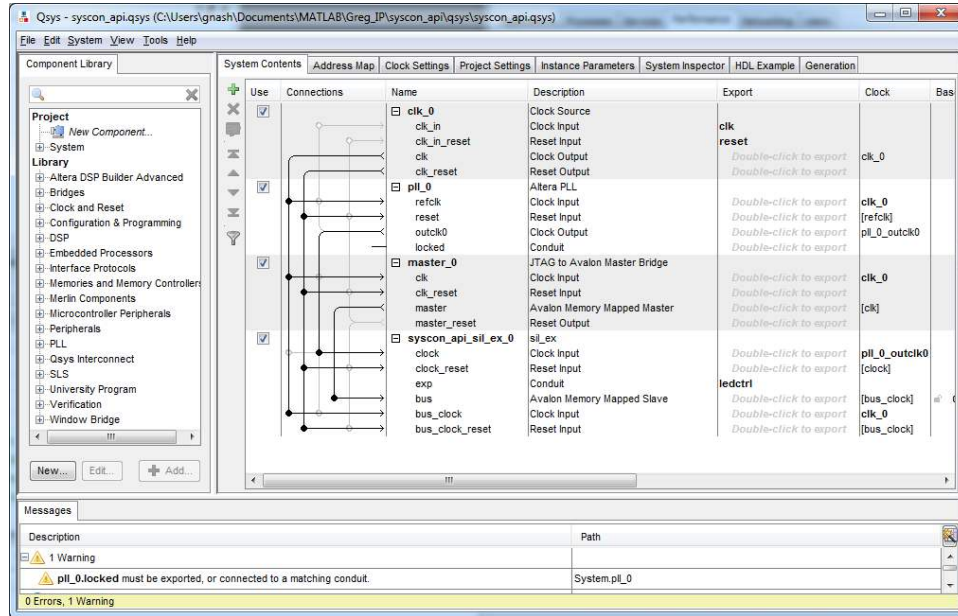
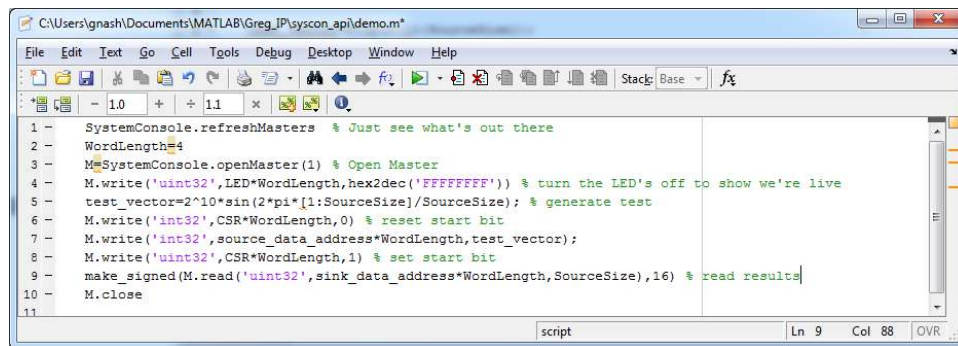


Figure 29. MATLAB API



### 3.6.2. Hardware Verification with System-in-the-Loop

Intel provides the system-in-the-loop flow for hardware verification.



System-in-the-loop:

- Automatically generates HW verification system for DSP Builder designs based on your configuration.
- Provides a wizard-based interface to configure, generate, and run HW verification system.
- Provides two separate modes:
  - **Run Test Vectors** loads and runs test vectors with large chunks (based on test memory size on target verification platform)
  - **Data Sample Stepping** loads one set sample at a time while stepping through Simulink simulation

**Data Sample Stepping** generates a copy of the original model and replaces the DSP Builder block with a special block providing connection to the FPGA to process data.

1. [Preparing for DSP Builder System-In-The-Loop](#) on page 57
2. [System-In-The-Loop Supported Blocks](#) on page 57
3. [Building Custom JTAG-Based Board Support Packages](#) on page 58
4. [Running System-In-the-Loop](#) on page 60
5. [System-In-The-Loop Parameters](#) on page 61

### 3.6.2.1. Preparing for DSP Builder System-In-The-Loop

1. Ensure you have a full installation of the Intel FPGA for OpenCL SDK.
2. Ensure ALTERAOCLSDKROOT variable points to the installation root directory.
3. For Windows, if you intend to use **Data Sample Stepping**, add the following suffix to PATH environment variable:

```
<YOUR_OPEN_CL_INSTALLATION_ROOT>/host/windows64/  
bin:<YOUR_DSPBA_INSTALLATION_ROOT>/backend/windows64
```

### 3.6.2.2. System-In-The-Loop Supported Blocks

System-in-the-loop only supports DSP Builder device-level blocks. The block interface may have complex and vector type ports.

All block input and output ports should pass through a single DSP Builder **ChannelIn** or **ChannelOut** interface, or be connected to a single **IP** block. The block may contain memory-mapped registers and memory blocks (accessible through the autogenerated Avalon memory-mapped agent interface). Observe the following limitations:

- The design should use the same clock for system and bus interfaces. The design does not support separate clocks.
- For autogenerated Avalon memory-mapped agent interfaces, use the name `bus`.
- The design does not support any other combination of DSP Builder block interface, including Avalon memory-mapped host interfaces.

The overall bitwidth of block input and output ports should not exceed 512 bits (excluding the valid signal).

Running hardware verification with **Data Sample Stepping** loads a new set of test data to FPGA every simulation step (if the data set is valid), which gives big timing gaps between two subsequent cycles for DSP Builder blocks running on hardware. If your DSP Builder block implementation cannot handle such gaps, system-in-the-loop simulation results may be incorrect.

### 3.6.2.3. Building Custom JTAG-Based Board Support Packages

1. [Setting up Board Support Package for 28 nm Device Families](#) on page 58
2. [Setting up Board Support Packages for Other Device Families](#) on page 58
3. [Publishing the Package in the System-In-The-Loop Wizard](#) on page 59
4. [System-in-the-Loop Third-Party Board Support Packages](#) on page 60
5. [Template Values in the System-in-the-Loop boardinfos.xml File](#) on page 60

#### 3.6.2.3.1. Setting up Board Support Package for 28 nm Device Families

1. Copy the `<ALTERAOCLSDKROOT>/board/dspb_sil_jtag` directory to a location where you have write access (for example `CUSTOM_BOARDS`).
2. Change `<CUSTOM_BOARDS>/dspba_sil_jtag/hardware` directory
3. Rename `jtag_c5soc` directory to desired name (for example `jtag_myboard`), remove the second directory.
4. Change to the `jtag_myboard` directory
5. In the `top.qsf` file, in board specific section:
  - a. Change the device family and name setting according to device in your board
  - b. Change `clk` and `resetn` port location and IO standard assignments according to your board specification
6. In `top.sdc` file, in the call for `create_clock` command. Change the clock period according to a clock specification you have set in `top.qsf` file
7. Open `board.qsys` file in Platform Designer
  - a. Update `REF_CLK_RATE` parameter value for `kernel_clk_generator` instance according to clock specification you have set in `top.qsf/sdc` files
  - b. Update the device family setting according to your board specification
8. Open `system.qsys` file in Platform Designer. Update device family setting according to your board specification.
9. In the `<CUSTOM_BOARDS>/dspba_sil_jtag/board_env.xml` file change the default hardware name to your directory (for example `jtag_myboard`).

#### 3.6.2.3.2. Setting up Board Support Packages for Other Device Families

1. In the `scripts/post_flow.tcl` file remove the following line:

```
source $::env(ALTERAOCLSDKROOT)/ip/board/bsp/adjust_pll.tcl
```
2. Open `board.qsys` file in Platform Designer

- a. Remove the `kernel_clk_generator` instance.
- b. Add instance of Intel PLL with one output clock. Set the reference clock frequency.
- c. Export the `refclk` clock input interface with `kernel_pll_refclk` name
- d. Connect `outclk0` to the initial source of `kernel_clk_generator.kernel_clk` output.
- e. Connect the `global_rest_in.out_reset` output to reset input of PLL instance.
- f. Set the generated clock frequency

*Note:* The design must meet timing on this clock domain. Intel advises that you use a low target frequency.

### 3.6.2.3.3. Publishing the Package in the System-In-The-Loop Wizard

1. Create a file named **boardinfos.xml** file in the directory where you copy the **dspba\_sil\_jtag** directory.
2. Before you start MATLAB, set the `DSPBA_SIL_BSP_ROOT` variable to point to the directory where this file and your custom board is located.

**Figure 30. Boardinfo.xml File Content**

```

boardinfos.xml [s:\data\akarspet] - Gvim
File Edit Tools Syntax Buffers Window Help
<?xml version="1.0" encoding="utf-8" ?>
<bsp xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:noNamespaceSchemaLocation="boardinfos.xsd">

<!-- Example explained
NOTE: In order for the custom board to be picked by SIL Wizard, please assign
DSPBA_SIL_BOARD_ROOT=<directory_where_this_file_is_located>

name - Display name for the board package
directory - Directory for the BSP relative to DSPBA_SIL_BOARD_ROOT
variant - The name of the actual board inside <directory>/hardware
type - Always set to 'custom'
template - Specify which board package has been used as template for the custom board
device.family - Device family for the specified concrete board
device.name - The name of the actual device against which the package should be compiled
device.id - Id of the device in case this is not a single device in the board
debugInterface.type - The type of the communication media.
Choices: Jtag, PCIe
dataStorage.type - The type of test data storage.
Choices: On-Chip RAM, DDR3
dataStorage.size - Specify the size of test data storage
dataStorage.unit - Specify the size of test data storage

NOTE: For 14.1, always set dataStorage.modifyable=0
-->

<bsp name="Display Name for the Board"
directory="<CUSTOM_BOARD>/dspba_sil_jtag"
variant="jtag_myboard"
type="custom"
template="dspba_sil_jtag">

<device family="Family Name" name="Device Name" id="1"/>

<debugInterface type="Jtag"/>

<dataStorage type="On-Chip RAM" modifyable="0" size="256" sizeUnit="KB"/>
</bsp>
</bsp>
22,17 All

```

### 3.6.2.3.4. System-in-the-Loop Third-Party Board Support Packages

System-in-the-loop supports the following third-party board support packages that are available for OpenCL:

- Bittware
- Nallatech
- ProcV

These packages are not available in the system-in-the-loop wizard by default. After you install these boards, publish the packages to the system-in-the-loop wizard.

### 3.6.2.3.5. Template Values in the System-in-the-Loop boardinfos.xml File

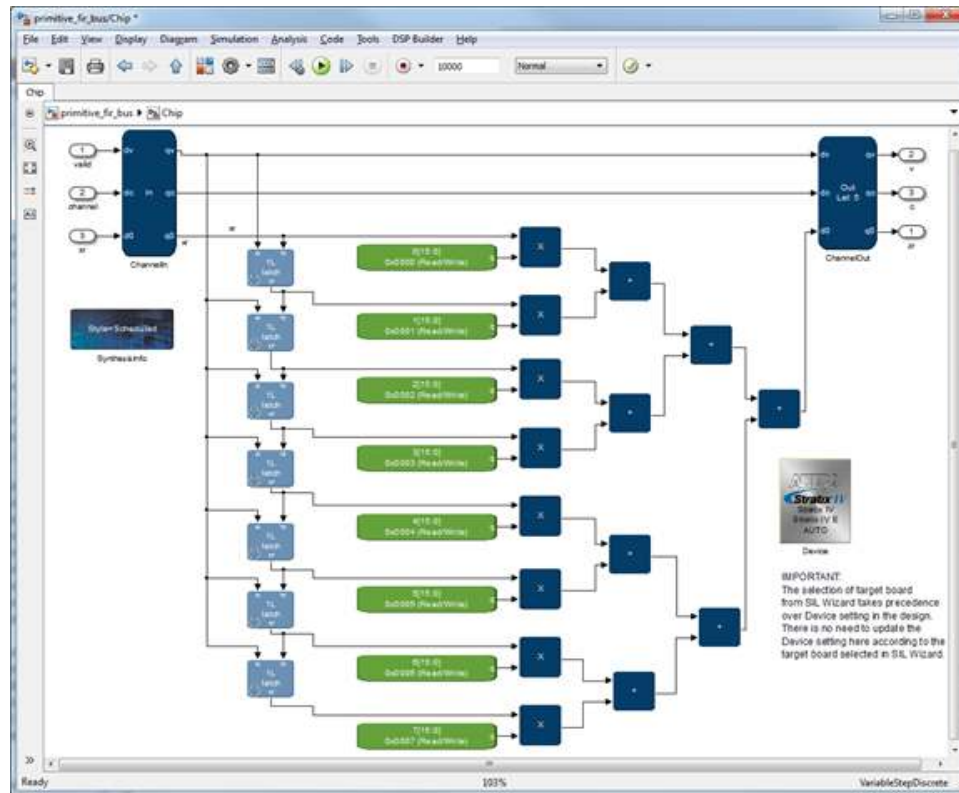
**Table 11. Template Values in boardinfos.xml File**

Board Support Package	Value of 'Template' field
Bittware	bittware_s5phg
Nallatech	nallatech_pcie3x5
Gidel ProcV	gidel_procev
Custom packages based on dspa_sil_jtag	dspba_sil_jtag
Custom packages based on dspba_sil_pci	dspba_sil_pcie

### 3.6.2.4. Running System-In-the-Loop

This walkthrough uses a DSP Builder design that implements a primitive FIR filter with memory-mapped registers for storing coefficients

Figure 31. FIR Filter with Memory-Mapped Registers



1. In the design’s Control block ensure you turn on **Generate Hardware**.
2. Simulate the model to generate RTL.
3. Select a device-level sub-system in your design and click **DSP Builder** ► **New SIL Wizard**.
4. On the **Parameters** tab, specify the parameters.
5. Click the **Run** tab and specify the run parameters

### 3.6.2.5. System-In-The-Loop Parameters

The design interface settings only generate appropriate adapters between the DSP Builder **ChannelIn** and **ChannelOut** interfaces and test Avalon-ST interface. The hardware platform always runs at fixed clock rate.

Table 12. System-In-The-Loop Parameters

Section	Option	Description
BSP Settings	BSP	Select the target BSP you want to run the hardware test on.
	Device	Device on the selected board.
BSP Memory Allocation	Total Memory Size	Specify the total size for test memory to use.
	Input Memory Size	Specify the amount of memory (from total memory size) for storing input test data. The remaining memory is for storing output data.

*continued...*

Section	Option	Description
		You might require several iterations to load and process all input test vectors because of memory limitations.
Design Interface	Clock Rate	Specify the same value as in the DSP Builder block setup file.
	Sample Rate	Specify the same value as in the DSP Builder block setup file.
	Number of Channels	The number of channels for the DSP Builder block. Specify the same value as in the DSP Builder block setup file.
	Frame Size	This value represents a number of valid data samples that you should supply to the DSP Builder block without timing gaps in between. If this value is more than 1, the wizard inserts a specific block in between test data provider and the DSP Builder block. This block enables data transmission to the DSP Builder block only when the specified amount of data is already available. An example of such a design is a folded multichannel design.
-	Destination Directory	Specify the directory where DSP Builder should generate the system-in-the-loop related files. You should change to this directory to simulate the system-in-the-loop generated model with an FPGA proxy.

**Table 13. System-In-The-Loop Run Settings**

Setting	Description
Select SIL Flow	Select the system- in-the-loop flow to use. The options are: <b>Run Test Vectors</b> runs all test vectors through the hardware verification system. The test vectors are based on simulation data recorded in DSP Builder .stm format files during Simulink simulation. <b>Step Through Simulation</b> allows processing every different set of valid input data on hardware separately, while simulating a design from Simulink. The wizard generates a separate model <model_name>_SIL in the SIL destination directory, which you should use for hardware verification. The original DSP Builder device level block is replaced with a specific block providing communication with the FPGA. You should change to SIL destination directory before you can simulate this model. If you change the flow, regenerate and recompile the system into a new destination directory.
Generate	Generates the infrastructure, files, and blocks for the hardware verification platform.
Compile	Compiles the entire hardware verification system in the Quartus Prime software to the generation configuration file. Allow at least 10-15 minutes for this step to run (more time for large DSP Builder designs). During this time the MATLAB input interface is unavailable.
Select JTAG Cable	Press Scan to scan available JTAG connections for programming the board. Choose the required JTAG cable from the discovered list.
Program	Program the board through selected JTAG cable. Go directly to this step if you have a pregenerated design with no changes with the flow parameters and in DSP Builder design under test.
Run	Run the test on hardware. <b>Run Test Vectors</b> only. The hardware test automatically detects and executes write requests over the DSP Builder autogenerated Avalon memory-mapped agent interface. The wizard cannot keep the sequence of transfers for write requests over Avalon memory-mapped agent interface and the DSP Builder data interface on hardware exactly the same as during simulation. Therefore, you may see data mismatches for a few sets of output samples at points where write requests are issued.
Compare	Compare the hardware verification results with simulation outputs. <b>Run Test Vectors</b> only.

*continued...*

Setting	Description
	The wizard compares only valid output samples.
Simulate <original_model>_SIL system	<p>During simulation, the FPGA proxy block that replaces the original DSP Builder design in the system-in-the-loop:</p> <ul style="list-style-type: none"> <li>• Every time you update inputs, it loads data to DSP Builder if valid input is high</li> <li>• Every time you request outputs, it populates outputs with data read from hardware if output memory contains valid sample.</li> </ul> <p><b>Step through simulation</b> only.</p> <p>Because the FPGA proxy updates its output only with valid samples, you see the same results repeated on the outputs until hardware has a new valid set of data. This behavior may differ from simulation results, where outputs are populated at every simulation cycles with available values.</p>

### 3.7. Integrating Your DSP Builder Advanced Blockset Design into Hardware

Integrate your DSP Builder advanced blockset design as a black-box design in your top-level design. Integrate into Platform Designer to create a complete project that integrates a processor, memory, datapath, and control.

1. [DSP Builder Generated Files](#) on page 63
2. [DSP Builder Designs and the Quartus Prime Project](#) on page 64
3. [Interfaces with a Processor Bus](#) on page 65

#### 3.7.1. DSP Builder Generated Files

DSP Builder generates the files in a directory structure at the location you specify in the **Control** block, which defaults to `../rtl` (relative to the working directory that contains the `.mdl` file). If you turn on the **Generate Hardware** option in the parameters for the DSP Builder **Control** block, every time the simulation runs, the underlying hardware synthesizes, and VHDL writes out into the specified directory.

**Table 14. Generated Files**

DSP Builder creates a directory structure that mirrors the structure of your design. The root to this directory can be an absolute path name or a relative path name. For a relative path name (such as `../rtl`), DSP Builder creates the directory structure relative to the MATLAB current directory.

File	Description
rtl directory	
<model name>.xml	An XML file that describes the attributes of your model.
<model name>_entity.xml	An XML file that describes the boundaries of the system.
<model name>_params.xml	When you open a model, DSP Builder produces a <code>model_name_params.xml</code> file that contains settings for the model. You must keep this file with the model.
rtl/<model name> subdirectory	
<block name>.xml	An XML file containing information about each block in the advanced blockset, which translates into HTML on demand for display in the MATLAB Help viewer and for use by the DSP Builder menu options.
<i>continued...</i>	

File	Description
<code>&lt;model name&gt;.vhd</code>	This is the top-level testbench file. It may contain non-synthesizable blocks, and may also contain empty black boxes for Simulink blocks that are not fully supported.
<code>&lt;model name&gt;.add.tcl</code>	This script loads the VHDL files in this subdirectory and in the subsystem hierarchy below it into the Quartus Prime project.
<code>&lt;model name&gt;.qip</code>	This file contains information about all the files DSP Builder requires to process your design in the Quartus Prime software. The file includes a reference to any <code>.qip</code> file in the next level of the subsystem hierarchy.
<code>&lt;model name&gt;_&lt;block name&gt;.vhd</code>	DSP Builder generates a VHDL file for each component in your model.
<code>&lt;model name&gt;_&lt;subsystem&gt;_entity.xml</code>	An XML file that describes the boundaries of a subsystem as a black-box design.
<code>&lt;subsystem&gt;.xml</code>	An XML file that describes the attributes of a subsystem.
<code>*.stm</code>	Stimulus files.
<code>safe_path.vhd</code>	Helper function that the <code>.qip</code> and <code>.add.tcl</code> files reference to ensure that pathnames read correctly in the Quartus Prime software.
<code>safe_path_msim.vhd</code>	Helper function that ensures a path name reads correctly in ModelSim.
<code>&lt;subsystem&gt;_atb.do</code>	Script that loads the subsystem automatic testbench into ModelSim.
<code>&lt;subsystem&gt;_atb.wav.do</code>	Script that loads signals for the subsystem automatic testbench into ModelSim.
<code>&lt;subsystem&gt;/&lt;block&gt;/*.hex</code>	Files that initialize the RAM in your design for either simulation or synthesis.
<code>&lt;subsystem&gt;.sdc</code>	Design constraint file for TimeQuest support.
<code>&lt;subsystem&gt;.tcl</code>	This Tcl script exists only in the subsystem that contains a <b>Device</b> block. You can use this script to setup the Quartus Prime project.
<code>&lt;subsystem&gt;_hw.tcl</code>	A Tcl script that loads the generated hardware into Platform Designer.

### Related Information

- [Simulating the Fibonacci Design in Simulink](#) on page 72
- [Simulating the IP Design in Simulink](#) on page 76
- [Simulating, Verifying, Generating, and Compiling Your DSP Builder Design](#) on page 21
- [Simulating, Verifying, Generating, and Compiling Your DSP Builder Design](#) on page 21
- [Control](#) on page 223

### 3.7.2. DSP Builder Designs and the Quartus Prime Project

DSP Builder creates a Quartus Prime project in the design directory that contains the `.mdl` file when you add your design to the Quartus Prime project.

The Quartus Prime project file (`.qpf`), Quartus Prime settings file (`.qsf`), and `.qip` files have the same name as the subsystem in your design that contains the **Device** block. For example, DSP Builder creates the files `DDCChip.qpf`, `DDCChip.qsf`, and `DDCChip.qip` for the **demo\_ddc** design.



These files contain all references to the files in the hardware destination directory that the **Control** block specifies. DSP Builder generates these files when you run a Simulink simulation. The project automatically loads into the Quartus Prime software.

When you compile your design the project compiles with the `.tcl` scripts in the hardware destination directory.

The `.qip` file references all the files that the project requires. Use the **Archive Project** command in the Quartus Prime software to use this file to archive the project.

For information about archiving projects, refer to the Quartus Prime Help.

### 3.7.2.1. Adding a DSP Builder Advanced Blockset Design to an Existing Quartus Prime Project

1. Find the `.add.tcl` file in the subsystem that contains the **Device** block.
2. Alternatively, you can add a reference to the `.qip` file in this subsystem from the `.qip` file for the top-level Quartus Prime project.

For information about using Tcl files in the Quartus Prime software, refer to the Quartus Prime Help.

### 3.7.3. Interfaces with a Processor Bus

DSP Builder designs can interface with a processor bus. You can drag and drop any register in your design without manually creating address decoding logic and memory-mapped switch fabric generation.

1. [Assigning Base Addresses in DSP Builder Designs](#) on page 65
2. [Adding a DSP Builder Design to a Platform Designer System](#) on page 66
3. [Updating Registers with the Nios II Processor](#) on page 68

#### 3.7.3.1. Assigning Base Addresses in DSP Builder Designs

You can add or drop **IP** or **Primitive** library control register fields and memory into your design.

1. Record the base address of the modules that connect to the Avalon-MM interface.
2. Start from address 0 in your design or any other arbitrary integer. The base address is a relative address and is expressed as an integer.

The base address depends on the data width of the bus and the width of the parameterizable variables. For example, FIR coefficients, where a register for a 32-bit FIR coefficient requires two words on a 16-bit bus.

3. Note the relative base addresses of modules within your design. When you integrate your model into a Platform Designer system, Platform Designer generates a base address for the entire DSP Builder model. Platform Designer references individual modules within the `.mdl` design based on the model base address (autogenerated) and relative base address you assign in the `.mdl` file or its setup script.
4. Manage base addresses, by specifying the bus data width in the **Control** block.

5. For **IP** designs consider the number of registers each IP core needs and the number of words each register requires
6. For **Primitive** subsystems, treat registers independently.
7. Ensure each **IP** library block and register or memory in a **Primitive** subsystem has a unique base address.

### 3.7.3.2. Adding a DSP Builder Design to a Platform Designer System

You can use the DSP Builder design with other Platform Designer components that have Avalon Streaming (Avalon-ST) interfaces. You should design the system in Platform Designer, where you can connect the Avalon-ST interfaces. Hardware Tcl (`_hw.tcl`) files describe the interfaces.

The output of a DSP Builder design is a source of Avalon-ST data for downstream components. It supplies data (and corresponding valid, channel, and start and end of packet information) and accepts a Boolean flag input from the downstream components, which indicates the downstream block is ready to accept data.

The input of the DSP Builder design is a sink of Avalon-ST data for upstream components. It accepts data (and corresponding valid, channel, and start and end of packet information) and provides a Boolean flag output to the upstream component, which indicates the DSP Builder component is ready to accept data.

1. Simulate your design with **Hardware Generation** turned on in **Control** block. DSP Builder generates a `<model>_hw.tcl` file for the subsystem containing the **Device** block. This file marks the boundary of the synthesizable part of your design and ignores the testbench blocks.

2. Add the synthesizable model to Platform Designer by including `<model>_hw.tcl` at the IP search path.

Platform Designer native streaming data interface is the Avalon Streaming (Avalon-ST) interface, which DSP Builder advanced blockset does not support. The DSP Builder advanced blockset native interface `<valid, channel, data>` ports are exported to the top-level as conduit signals.

3. Add DSP Builder components to Platform Designer by adding a directory that contains generated hardware to the **IP Search Path** in the Platform Designer **Options** dialog box.
4. Define Avalon-ST interfaces to build system components that Platform Designer can join together.

Upstream and downstream components are part of the system outside of the DSP Builder design.

5. Register all paths across the DSP builder design to avoid algebraic loops.

A design may have multiple Avalon-ST input and output blocks.

6. Generate the Platform Designer system.

In the `hw.tcl` file, the name of the Avalon-ST masked subsystem block is the name of the interface.

7. Add FIFO buffers on the output (and if required on the input) to build designs that supporting backpressure, and declare the collected signals as an Avalon-ST interface in the `hw.tcl` file generated for the device level.

These blocks do not enforce Avalon-ST behavior. They encapsulate the common Avalon-ST signals into an interface.

1. [Modifying Avalon-ST Blocks](#) on page 67
2. [Restrictions for DSP Builder Designs with Avalon Streaming Interface Blocks](#) on page 67

#### Related Information

[Interfaces Library](#) on page 262

#### 3.7.3.2.1. Modifying Avalon-ST Blocks

Modify Avalon-ST blocks to add more ports, to add custom text or to extend the blocks

1. Look under the mask to see the implementation of the DSP Builder Avalon-ST blocks masked subsystems.
2. Extend the definition further, by breaking the link and adding further ports that the `hw.tcl` file declares, or add text that DSP Builder writes unevaluated directly into the interface declaration in the `hw.tcl` file.

*Note:* When you edit the mask do not edit the mask type, as DSP Builder uses it to identify the subsystems defining the interfaces.

3. Add more ports to Avalon ST blocks by connecting these ports internally in the same way as the existing signals. For example, with FIFO buffers.
4. If you add inputs or output ports that you connect to the device level ports, tag these ports with the role the port takes in the Avalon-ST interface. For example, you may want to add **error** and **empty** ports.
5. Add custom text to the description field.  
Any text you write to the description field of the DSP Builder masked subsystem writes with no evaluation into the `hw.tcl` file immediately after the standard parameters for the interface and before the port declarations. Ensure you correctly add the text of any additions.

#### Related Information

[Streaming Library](#) on page 273

#### 3.7.3.2.2. Restrictions for DSP Builder Designs with Avalon Streaming Interface Blocks

You can place the Avalon streaming interface blocks in different levels of hierarchy. However, never place Simulink, **IP** or **Primitive** library blocks between the interface and the device level ports.

The Avalon streaming interface specification only allows a single data port per interface. Thus you may not add further data ports, or even using a vector through the interface and device-level port (which creates multiple data ports).

To handle multiple data ports through a single Avalon streaming interface, pack them together into a single (not vector or bus) signal, then unpack on the other side of the interface. For the maximum width for a data signal, refer to the *Avalon Interface Specifications*.

Use the **BitCombine** and **BitExtract** blocks to pack and unpack.

### Related Information

- [Streaming Library](#) on page 273
- [Avalon Interface Specification](#)

### 3.7.3.3. Updating Registers with the Nios II Processor

You can use a processor such as a Nios II processor to read or modify a control register, or to update memory contents in your DSP Builder design.

1. Identify the Platform Designer base address assigned to your DSP Builder design. You can also find the base address information in **system.h** file in your Nios II project, after you have loaded the SOPC library information into your Nios II IDE.
2. Identify the base address for the **IP** block of interest in your DSP Builder advanced blockset design. It is the base address assigned to your DSP Builder advanced blockset model for step 1, plus the address offset you specified in your **IP** block or in your setup script. You can also identify the address offset by right clicking on the **IP** block and selecting **Help**.
3. Identify the base address for the register of interest in your DSP Builder advanced blockset design. It is the base address assigned to your DSP Builder advanced blockset model which you identify in step 1, plus the address offset you specified in your register or in your setup script.
  - a. Identify the address offset in the `<design name>_mmap.h` file, which DSP Builder generates with each design.
  - b. Alternatively, identify the address offset by right clicking on the register and select **Help**.
4. When you identify the base address, use **IOWR** and **IORD** commands to write and read registers and memory. For example:

```
IOWR(base_addr_SOPC + base_addr_FIR, coef_x_offset, data)
```

```
IORD(base_addr_SOPC + base_addr_FIR, coef_x_offset)
```

## 4. Primitive Library Blocks Tutorial

---

This tutorial shows how to build a simple design example that uses blocks from the **Primitive** library to generate a Fibonacci sequence.

The Fibonacci sequence is the sequence of numbers that you can create when you add 1 to 0 then successively add the last two numbers to get the next number: 0, 1, 1, 2, 3, 5, 8, 13, 21, 34, 55, 89, 144, 233, 377, 610, ...

Each **Primitive** library block in the design example is parameterizable. When you double-click a block in the model, a dialog box appears where you can enter the parameters for the block. Click the **Help** button in these dialog boxes to view help for a specific block.

You can use the `demo_fibonacci.mdl` model in the `<DSP Builder Advanced install path>/Examples/Primitive` directory or you can create your own Fibonacci model.

1. [Creating a Fibonacci Design from the DSP Builder Primitive Library](#) on page 69
2. [Setting the Parameters on the Testbench Source Blocks](#) on page 71
3. [Simulating the Fibonacci Design in Simulink](#) on page 72
4. [Modifying the DSP Builder Fibonacci Design to Generate Vector Signals](#) on page 73
5. [Simulating the RTL of the Fibonacci Design](#) on page 74

### 4.1. Creating a Fibonacci Design from the DSP Builder Primitive Library

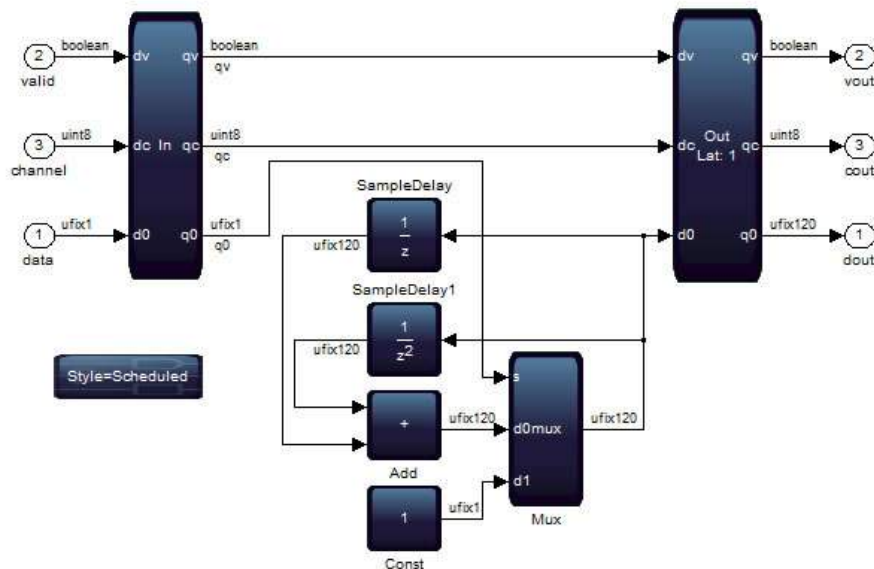
Start DSP Builder in MATLAB

1. From an open Simulink model click **DSP Builder > New Model Wizard**.  
*Note:* When you open a model, DSP Builder produces a `model_name_params.xml` file that contains settings for the model. You must keep this file with the model.
2. Specify the following **New Model Settings**:
  - **Fixed**
  - **Fixed-point Primitive (simple)**
  - **my\_fibonacci**
3. Browse to an appropriate output directory.
4. Click **Generate**.

5. In the Simulink Library Browser, click **DSP Builder Advanced Blockset** > **Primitive** > **Primitive Basic Blocks** .
6. Open the **my\_fibonacci** generated model.
7. Open the **dut** > **prim** subsystem, which has a **ChannelIn** and **ChannelOut** block.
8. Drag and drop two **SampleDelay** blocks into your model.

*Note:* Specify the data type because this design contains loops and DSP Builder cannot determine the type if one of the inherit data options is set.

**Figure 32. Updated Fibonacci Subsystem**



9. Select both of the **SampleDelay** blocks and point to **Rotate and Flip** on the popup menu and click **Flip Block** to reverse the direction of the blocks.
10. Drag and drop **Add** and **Mux** blocks into your model.
11. Drag and drop a **Const** block. Double-click the block and:
  - a. Select **Specify via Dialog** for **Output data type mode**.
  - b. For **Output type** enter **ufix(1)**.
  - c. For **Scaling** enter **1**
  - d. For **Value** enter **1**.
  - e. Click **OK**.
12. Connect the blocks.
13. Double-click on the second **SampleDelay** block (**SampleDelay1**) to display the **Function Block Parameters** dialog box and change the **Number of delays** parameter to **2**.
14. Double-click on the **Add** block to display the **Function Block Parameters** dialog box and set the parameters.
  - a. For **Output data type mode**, select **Specify via Dialog** .

- b. For **Output type** enter **ufix(120)**.
- c. For **Output scaling value** enter **2^-0**
- d. For **Number of inputs** enter **2**.
- e. Click **OK**.

#### Related Information

- [Starting DSP Builder in MATLAB](#) on page 17
- [Starting DSP Builder in MATLAB](#) on page 17
- [DSP Builder Menu Options](#) on page 20

## 4.2. Setting the Parameters on the Testbench Source Blocks

Set the testbench parameters to finish the DSP Builder Fibonacci design.

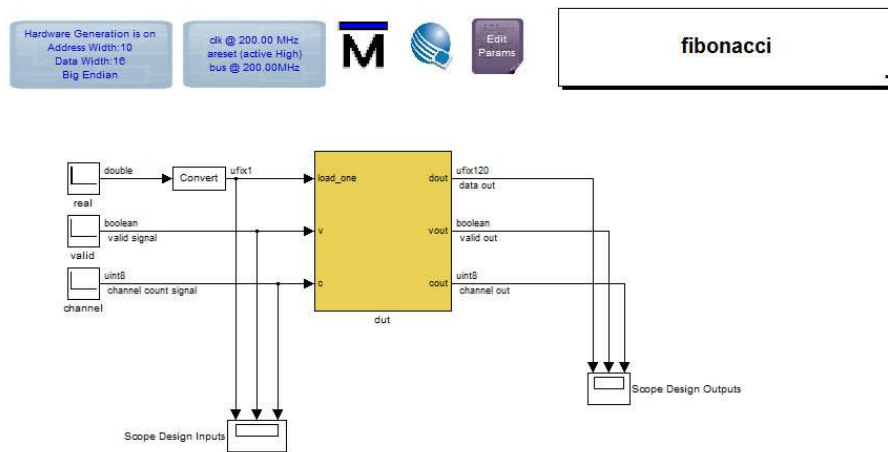
1. Double-click on the **Real** block to display the **Source Block Parameters** dialog box.
2. Set the **Vector of output values** to **[0 1 1 1 zeros(1,171)]'** in the **Main** tab.
3. Switch to the Editor window for `setup_my_fibonacci.m`.
4. Change the parameters to:
  - **my\_fibonacci\_param.ChanCount = 1;**
  - **my\_fibonacci\_param.SampleRate = fibonacci\_param.ClockRate;**
  - **my\_fibonacci\_param.input\_word\_length = 1;**
  - **my\_fibonacci\_param.input\_fraction\_length = 0;**
5. In the top-level design, delete the **ChannelView**, the **Scope Deserialized Outputs** scope and any dangling connections.
6. Double-click the **Convert** block and make the input unsigned by changing:

```
fixdt(1, fibonacci_param.input_word_length, fibonacci_param.input_fraction_length)
```

to:

```
fixdt(0, fibonacci_param.input_word_length, fibonacci_param.input_fraction_length)
```
7. Save the Fibonacci model.

Figure 33. Completed Fibonacci Model

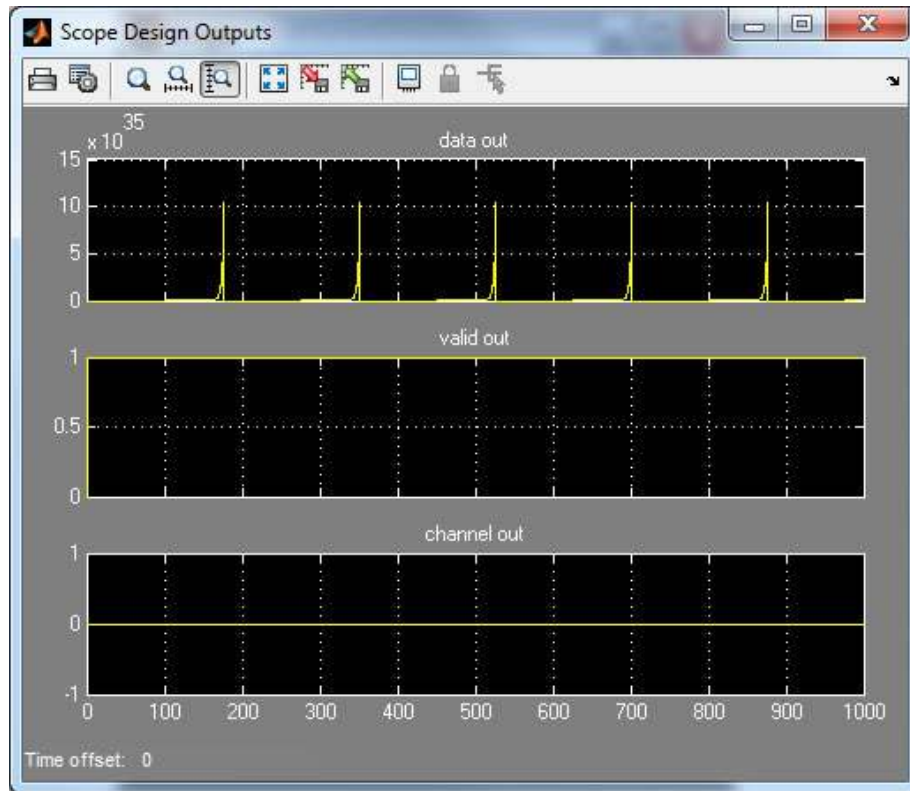


### 4.3. Simulating the Fibonacci Design in Simulink

1. Click **Simulation** ► **Run**.
2. Double-click on the **Scope** block and click **Autoscale** in the scope to display the simulation results .



Figure 34. Fibonacci Sequence in the Simulink Scope



*Note:* You can verify that the **fib** output continues to increment according to the Fibonacci sequence by simulating for longer time periods.

The sequence on the **fib** output starts at 0, and increments to 1 when **q\_v** and **q\_c** are both high at time 21.0. It then follows the expected Fibonacci sequence incrementing through 0, 1, 1, 2, 3, 5, 8, 13 and 21 to 34 at time 30.0.

#### Related Information

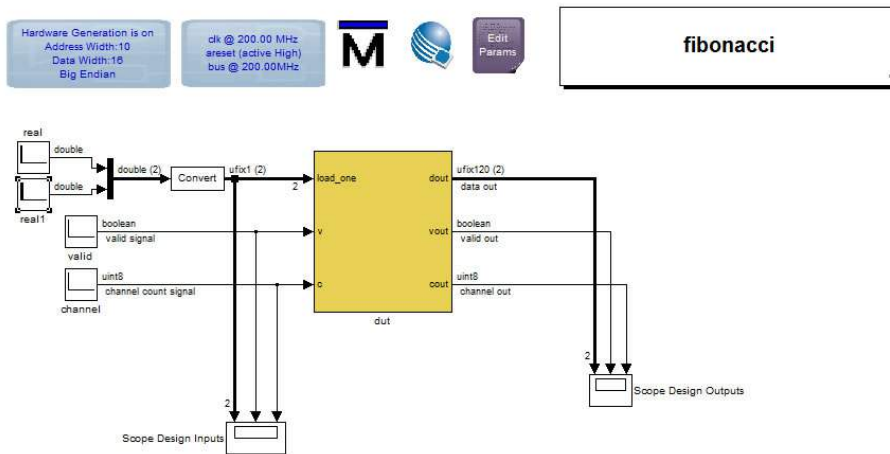
- [DSP Builder Generated Files](#) on page 63
- [Simulating, Verifying, Generating, and Compiling Your DSP Builder Design](#) on page 21

## 4.4. Modifying the DSP Builder Fibonacci Design to Generate Vector Signals

Simulate the design.

1. In the Scope Design Outputs scope zoom in on the y axis to see the short Fibonacci cycle.

Figure 35. Fibonacci Scope Design Outputs

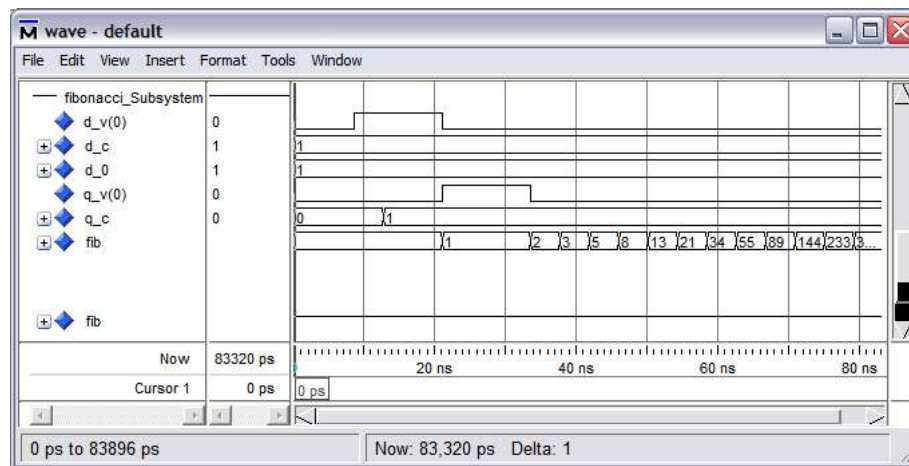


2. Copy the real input block, add a Simulink mux and connect to the Convert block.
3. Edit the timing of the real1 block, for example [0 1 1 1 zeros(1,50)].

## 4.5. Simulating the RTL of the Fibonacci Design

1. To verify that DSP Builder gives the same results when you simulate the generated RTL, click on the **Run ModelSim** block.

Figure 36. Fibonacci Sequence in the ModelSim Wave Window



Compile the design in the Quartus Prime software.

### Related Information

[Simulating, Verifying, Generating, and Compiling Your DSP Builder Design](#) on page 21

## 5. IP Tutorial

---

This tutorial demonstrates how to use blocks from the DSP Builder **IP** library. It shows how you can double the number of channels through a filter, increase the  $f_{MAX}$ , and target a different device family by editing top-level parameters in Simulink.

1. [Creating an IP Design](#) on page 75
2. [Simulating the IP Design in Simulink](#) on page 76
3. [Viewing Timing Closure and Viewing Resource Utilization for the DSP Builder IP Design](#) on page 77
4. [Reparameterizing the DSP Builder FIR Filter to Double the Number of Channels](#) on page 77
5. [Doubling the Target Clock Rate for a DSP Builder IP Design](#) on page 78

### 5.1. Creating an IP Design

Start DSP Builder in MATLAB.

1. Type `demo_firi` at the MATLAB command prompt, which opens the FIR design example.

*Note:* When you open a model, DSP Builder produces a `model_name_params.xml` file that contains settings for the model. You must keep this file with the model.

2. From an open Simulink model click **DSP Builder menu** ► **New Model Wizard**.

*Note:* You must have a model open, before you use the DSP Builder menu.

3. Specify the following **New Model Settings**:

- **Fixed**
- **IP (simple)**
- **my\_firi**

4. Browse to an appropriate output directory.
5. Click **Generate**.
6. Edit the basic parameter values in `setup_my_firi.m` to match the equivalent settings in `setup_demo_firi.m`.

Simulate the design in MATLAB.

#### Related Information

- [Starting DSP Builder in MATLAB](#) on page 17
- [Starting DSP Builder in MATLAB](#) on page 17

- [DSP Builder Menu Options](#) on page 20
- [Simulating the IP Design in Simulink](#) on page 76
- [Starting DSP Builder in MATLAB](#) on page 17
- [DSP Builder Menu Options](#) on page 20
- [Starting DSP Builder in MATLAB](#)
- [DSP Builder Menu Options](#)  
Simulink includes a **DSP Builder** menu on any Simulink model window. Use this menu to easily start all the common tasks you need to perform on your DSP Builder model.
- [Simulating the IP Design](#)

## 5.2. Simulating the IP Design in Simulink

Create an IP design.

1. In the **demo\_firi** window, click **Start > Simulation** .  
MATLAB generates output HDL for the design.
2. Click **DSP Builder Resource Usage Design**.  
You can view the resources of the whole design and the subsystems.
3. Click **Close**.
4. Double click the **FilterSystem** subsystem, right-click on the **filter1 InterpolatingFIR** block, and click **Help**.  
After simulation, DSP Builder updates this help to include the following information:
  - The latency of the filter
  - The port interface
  - The input and output data format
  - The memory interface for the coefficients.
5. To display the latency of the filter on the schematic, right-click on **InterpolatingFIR** block and click **Block Properties**.
6. On the **Block Annotation** tab, in block property tokens double-click on **%<latency>**.
7. In **Enter text and tokens for annotation**, type `Latency = before %<latency>`. Click **OK**.  
DSP Builder shows the latency beneath the block.

Verify the design in MATLAB. Compile the design in the Quartus Prime software

### Related Information

- [Simulating, Verifying, Generating, and Compiling Your DSP Builder Design](#) on page 21
- [DSP Builder Generated Files](#) on page 63
- [Simulating, Verifying, Generating, and Compiling Your DSP Builder Design](#) on page 21

- [DSP Builder Menu Options](#)  
Simulink includes a **DSP Builder** menu on any Simulink model window. Use this menu to easily start all the common tasks you need to perform on your DSP Builder model.

### 5.3. Viewing Timing Closure and Viewing Resource Utilization for the DSP Builder IP Design

Compile the IP design in the Quartus Prime software

1. View timing closure:
  - a. In the **Task** pane, expand **TimeQuest Timing Analyzer**.
  - b. Double-click **View Report**.
  - c. In the **Table of Contents** pane expand **Slow 900mV 85C Model** and click **Fmax Summary**.
2. View the resource utilization:
  - a. On the **Task** pane expand **Fitter (Place & Route)**.
  - b. Double-click **View Report**.
  - c. In the **Table of Contents** pane expand **Resource Section** and click **Resource Usage Summary**, which shows the number of DSP block 18-bit elements.

#### Related Information

[Simulating, Verifying, Generating, and Compiling Your DSP Builder Design](#) on page 21

### 5.4. Reparameterizing the DSP Builder FIR Filter to Double the Number of Channels

Create an IP design.

1. Double-click the **EditParams** block to open **setup\_my\_firi.m** in the MATLAB Editor. Change **my\_firi\_param.ChanCount** to **32** and click **Save**.
2. Simulate the design.  
On the **filter1 InterpolatingFIR** block the latency and the number of multipliers increases
3. On the DSP Builder menu, click **Verify Design**, and click **Clear Results** to clear the output pane. Turn on **Verify at device level** and **Run Quartus Prime Software only**, turn off **Verify at subsystem level**, and click **Run Verification**  
The Simulink simulation matches a ModelSim simulation of the generated HDL. The design meets timing but the number of multipliers and logic increases. The number of channels doubles, but the number of multipliers does not double, because the design shares some multipliers.

The design now closes timing above 480 MHz. At the higher clock rate, the design shares multiplier resources, and the multiplier count decreases back to 6.

### Related Information

- [DSP Builder Menu Options](#)  
Simulink includes a **DSP Builder** menu on any Simulink model window. Use this menu to easily start all the common tasks you need to perform on your DSP Builder model.
- [Creating an IP Design](#)

## 5.5. Doubling the Target Clock Rate for a DSP Builder IP Design

Create an IP design.

1. Double-click the **EditParams** block to open `my_fir1.m` in the MATLAB Editor. Change **my\_fir1\_param.ClockRate** to **480.0** and click **Save**.
2. Simulate the design.
3. Click **DSP Builder > Verify Design**, and click **Clear Results** to clear the output pane.
4. Click **Run Verification**.

### Related Information

- [DSP Builder Menu Options](#)  
Simulink includes a **DSP Builder** menu on any Simulink model window. Use this menu to easily start all the common tasks you need to perform on your DSP Builder model.
- [Creating an IP Design](#)
- [Simulating the IP Design](#)

## 6. DSP Builder for Intel FPGAs (Advanced Blockset) Design Examples and Reference Designs

---

DSP Builder provides a variety of design examples, which you can learn from or use as a starting point for your own design.

All the design examples have the same basic structure: a top-level testbench containing an instantiated functional subsystem, which represents the hardware design.

The testbench typically includes Simulink source blocks that generate the stimulus signals and sink blocks that display simulation results. You can use other Simulink blocks to define the testbench logic.

The testbench also includes the following blocks from the DSP Builder advanced blockset:

- The **Control** block specifies information about the hardware generation environment, and the top-level memory-mapped bus interface widths.
- The **ChanView** block in a testbench allows you to visualize the contents of the <valid, channel, data> time-division multiplex (TDM) protocol. This block generates synthesizable HDL and can therefore also be useful in a functional subsystem.

The functional subsystem in each design contains a **Device** block that marks the top-level of the FPGA device and controls the target device for the hardware.

1. [DSP Builder Design Configuration Block Design Examples](#) on page 80
2. [DSP Builder FFT Design Examples](#) on page 80
3. [DSP Builder DDC Design Example](#) on page 86
4. [DSP Builder Filter Design Examples](#) on page 100
5. [DSP Builder Finite State Machine Design Example](#) on page 108
6. [DSP Builder Folding Design Examples](#) on page 108
7. [DSP Builder Floating Point Design Examples](#) on page 115
8. [DSP Builder Flow Control Design Examples](#) on page 120
9. [DSP Builder HDL Import Design Example](#) on page 124
10. [DSP Builder Host Interface Design Examples](#) on page 137
11. [DSP Builder Platform Design Examples](#) on page 138
12. [DSP Builder Primitive Block Design Examples](#) on page 140
13. [DSP Builder Reference Designs](#) on page 150
14. [DSP Builder Waveform Synthesis Design Examples](#) on page 171

## 6.1. DSP Builder Design Configuration Block Design Examples

1. [Scale](#) on page 80
2. [Local Threshold](#) on page 80

### 6.1.1. Scale

This design example demonstrates the **Scale** block.

The testbench allows you to see a vectorized block in action. Displays in the testbench track the smallest and largest values to be scaled and verify the correct behavior of the saturation modes.

The top-level testbench includes **Control**, **Signals**, **Run ModelSim**, and **Run Quartus Prime** blocks, plus a **ChanView** block that deserializes the output bus.

The **ScaleSystem** subsystem includes the **Device** and **Scale** blocks.

The model file is `demo_scale.mdl`.

### 6.1.2. Local Threshold

This design example has two identical NCOs—one is in a subsystem with a **LocalThreshold** block that is set to force soft rather than hard multipliers.

After simulation, in the resource table, you can compare the resources for NCO and NCO1. NCO1 uses no multipliers at the expense of extra logic. The resource table also contains resources for the **ChannelViewer** blocks—synthesizable blocks, that the design example uses outside the device system.

The model file is `demo_nco_threshold.mdl`.

## 6.2. DSP Builder FFT Design Examples

1. [FFT](#) on page 81
2. [FFT without BitReverseCoreC Block](#) on page 81
3. [IFFT](#) on page 82
4. [IFFT without BitReverseCoreC Block](#) on page 82
5. [Floating-Point FFT](#) on page 82
6. [Floating-Point FFT without BitReverseCoreC Block](#) on page 82
7. [Floating-Point iFFT](#) on page 83
8. [Floating-Point iFFT without BitReverseCoreC Block](#) on page 83
9. [Multichannel FFT](#) on page 83
10. [Multiwire Transpose](#) on page 83
11. [Parallel FFT](#) on page 83
12. [Parallel Floating-Point FFT](#) on page 83
13. [Single-Wire Transpose](#) on page 83
14. [Switchable FFT/iFFT](#) on page 84



15. [Variable-Size Fixed-Point FFT](#) on page 84
16. [Variable-Size Fixed-Point FFT without BitReverseCoreC Block](#) on page 84
17. [Variable-Size Fixed-Point iFFT](#) on page 84
18. [Variable-Size Fixed-Point iFFT without BitReverseCoreC Block](#) on page 84
19. [Variable-Size Floating-Point FFT](#) on page 84
20. [Variable-Size Floating-Point FFT without BitReverseCoreC Block](#) on page 84
21. [Variable-Size Floating-Point iFFT](#) on page 84
22. [Variable-Size Floating-Point iFFT without BitReverseCoreC Block](#) on page 85
23. [Variable-Size Low-Resource FFT](#) on page 85
24. [Variable-Size Low-Resource Real-Time FFT](#) on page 85
25. [Variable-Size Supersampled FFT](#) on page 85

### 6.2.1. FFT

This design example implements a 2,048 point, radix  $2^2$  FFT. This design example accepts natural order data at the input and produces natural order data at the output. The design example includes a **BitReverseCoreC** block, which converts the input data stream from natural order to bit-reversed order, and an **FFT** block, which performs an FFT on bit-reversed data and produces its output in natural order.

*Note:* The FFT designs do not inherit the width in bits and scaling information. The design example specifies these values with the *Wordlength* and *FractionLength* variables in the setup script, which are 16 and 19 for this design example. You can also set the maximum width in bits by setting the *MaxOut* variable. Most applications do not need the maximum width in bits. To save resources, set a threshold value for this variable. The default value of **inf** allows worst case bit growth.

The model file is `demo_fft.mdl`.

### 6.2.2. FFT without BitReverseCoreC Block

This design example implements a 2,048 point, radix  $2^2$  FFT. This design example accepts natural order or bit-reversed data at the input and produces bit-reversed or natural order data at the output, respectively. The design example is identical to the FFT design example, but it does not include a **BitReverseCoreC** block, which converts the input data stream from natural order to bit-reversed order.

*Note:* The FFT designs do not inherit width in bits and scaling information. The design example specifies these values with the *Wordlength* and *FractionLength* variables in the setup script, which are 16 and 19 for this design example. You can also set the maximum width in bits by setting the *MaxOut* variable. Most applications do not need the maximum width in bits. To save resources, set a threshold value for this variable. The default value of **inf** allows worst case bit growth.

The model file is `demo_fft_core.mdl`.

### 6.2.3. IFFT

This design example implements a 2,048 point, radix  $2^2$  iFFT. This design example accepts natural order data at the input and produces natural order data at the output. The design example includes a **BitReverseCoreC** block, which converts the input data stream from natural order to bit-reversed order, and an **FFT** block, which performs an FFT on bit-reversed data and produces its output in natural order.

*Note:* The FFT designs do not inherit the width in bits and scaling information. The design example specifies these values with the *Wordlength* and *FractionLength* variables in the setup script, which are 16 and 19 for this design example. To set the maximum width in bits, set the *MaxOut* variable. Most applications do not need the maximum width in bits. To save resources, set a threshold value for this variable. The default value of **inf** allows worst case bit growth.

The model file is `demo_ifft.mdl`.

### 6.2.4. IFFT without BitReverseCoreC Block

This design example implements a 2,048 point, radix  $2^2$  iFFT. This design example accepts natural order data at the input and produces natural order data at the output. The design example is identical to the iFFT design example, but it does not include a **BitReverseCoreC** block, which converts the input data stream from natural order to bit-reversed order.

*Note:* The FFT designs do not inherit width in bits and scaling information. The design example specifies these values with the *Wordlength* and *FractionLength* variables in the setup script, which are 16 and 19 for this design example. To set the maximum width in bits, set the *MaxOut* variable. Most applications do not need the maximum width in bits. To save resources, set a threshold value for this variable. The default value of **inf** allows worst case bit growth.

The model file is `demo_ifft_core.mdl`.

### 6.2.5. Floating-Point FFT

This design example implements a floating-point, 512 point, radix  $2^2$  FFT. This design example accepts natural order data at the input and produces natural order data at the output. The design example includes a **BitReverseCoreC** block, which converts the input data stream from natural order to bit-reversed order, and an **FFT\_Float** block, which performs an FFT on bit-reversed data and produces its output in natural order.

The model file is `demo_fpfft.mdl`.

### 6.2.6. Floating-Point FFT without BitReverseCoreC Block

This design example implements a floating-point, 512 point, radix  $2^2$  FFT. This design example accepts natural order data at the input and produces natural order data at the output. The design example is identical to the floating-point FFT design example, but it does not include a **BitReverseCoreC** block, which converts the input data stream from natural order to bit-reversed order.

The model file is `demo_fpfft_core.mdl`.

### 6.2.7. Floating-Point iFFT

This design example implements a floating-point 512 point, radix  $2^2$  iFFT. This design example accepts bit-reversed order data at the input and produces natural order data at the output. The design example includes a **BitReverseCoreC** block, which converts the input data stream from natural order to bit-reversed order, and an **FFT\_Float** block, which performs an FFT on bit-reversed data and produces its output in natural order.

The model file is `demo_fpifft.mdl`.

### 6.2.8. Floating-Point iFFT without BitReverseCoreC Block

This design example implements a floating-point 512 point, radix  $2^2$  iFFT. This design example accepts natural order data at the input and produces natural order data at the output. The design example is identical to the floating-point iFFT design example, but it does not include a **BitReverseCoreC** block, which converts the input data stream from natural order to bit-reversed order.

The model file is `demo_fpifft_core.mdl`.

### 6.2.9. Multichannel FFT

The FFT processes the input in blocks of 2K points. Each block contains 4 interleaved sub-channels. A 512-point FFT is performed on each sub-channel.

The model file is `demo_fft_multichannel.mdl`.

### 6.2.10. Multiwire Transpose

This design example demonstrates how to use the **MultiwireTranspose** block.

The model file is `demo_multiwiretranspose.mdl`.

### 6.2.11. Parallel FFT

The model file is `demo_parallel_fft.mdl`.

### 6.2.12. Parallel Floating-Point FFT

The model file is `demo_parallel_fpfft.mdl`.

### 6.2.13. Single-Wire Transpose

This design example demonstrates how to use the **Transpose** block.

The model file is `demo_transpose.mdl`.

#### 6.2.14. Switchable FFT/iFFT

This design example demonstrates a switchable **FFT** block.

The model file is `demo_dynamic_fft.mdl`.

#### 6.2.15. Variable-Size Fixed-Point FFT

This design example demonstrates the **FFT** block.

The model file is `demo_vfft.mdl`.

#### 6.2.16. Variable-Size Fixed-Point FFT without BitReverseCoreC Block

This design example demonstrates the **FFT** block.

The model file is `demo_vfft_core.mdl`.

#### 6.2.17. Variable-Size Fixed-Point iFFT

This design example demonstrates the **iFFT** block.

The model file is `demo_vifft.mdl`.

#### 6.2.18. Variable-Size Fixed-Point iFFT without BitReverseCoreC Block

This design example demonstrates the **iFFT** block.

The model file is `demo_vifft_core.mdl`.

#### 6.2.19. Variable-Size Floating-Point FFT

This design example demonstrates the **FFT** block.

The model file is `demo_fpvfft.mdl`.

#### 6.2.20. Variable-Size Floating-Point FFT without BitReverseCoreC Block

This design example demonstrates the **FFT** block.

The model file is `demo_fpvfft_core.mdl`.

#### 6.2.21. Variable-Size Floating-Point iFFT

This design example demonstrates the **iFFT** block.

The model file is `demo_fpvifft.mdl`.

### 6.2.22. Variable-Size Floating-Point iFFT without BitReverseCoreC Block

This design example demonstrates the **iFFT** block.

The model file is `demo_fpvifft_core.mdl`.

### 6.2.23. Variable-Size Low-Resource FFT

This design example implements a low data-rate, low-resource usage FFT suitable for applications such as vibration suppression. The maximum size of the FFT is 4K points. The actual size of each FFT iteration may be any power of 2 that is smaller than 4K points. You dynamically control the size using the size input. Irrespective of the size of the FFT, a new FFT iterates whenever it receives an additional 64 inputs. Successive FFT iterations usually use overlapping input data.

The FFT accepts 16-bit fixed-point inputs. The FFT produces block floating-point output using an 18-bit mantissa and a shared 6-bit exponent.

This FFT implementation is unusual because the FFT gain is 1. Therefore the sum-of-squares of the input values is equal (allowing for rounding errors) to the sum-of-squares of the output values.

This method of scaling gives two advantages:

- The exponent can be smaller.
- The output remains consistently scaled when the FFT dynamic size changes.

To configure the design example, edit any of the parameters in the setup file.

The model file is `demo_servofft.mdl`.

### 6.2.24. Variable-Size Low-Resource Real-Time FFT

This design example is an extension of the Variable-Size Low-Resource FFT design example. This example runs the Variable-Size Low-Resource FFT in Simulink in real time alongside the DSP System Toolbox FFT block. It produces vector scope outputs that are the same (or as near as) from both.

This design example takes care of the faster sample rate needed by the DSP Builder FFT. The setup file chooses a sample rate that is fast enough for calculation but not so fast that it slows down the simulation unnecessarily. The design also adds buffering to the original MATLAB `fft` signal path to make the signal processing delays the same in both paths.

The model file is `demo_dspba_ex_fft_tut.mdl`.

### 6.2.25. Variable-Size Supersampled FFT

This DSP Builder design example implements a variable-size supersampled FFT, with sizes ranging from 256 to 2,048 points, and a parallelism of 4 wires.

The incoming data is of fixed-point type and arrives in natural order. The number of radix-2 stages assigned to the serial section of the hybrid FFT is 7.

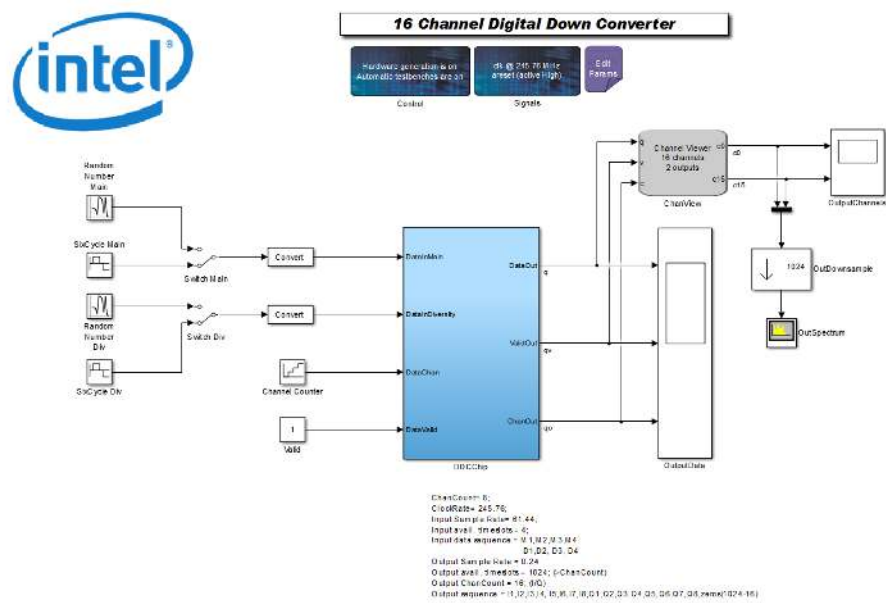
### 6.3. DSP Builder DDC Design Example

The DDC design example uses NCO/DDS, mixer, CIC, and FIR filter **IP** library blocks to build a 16-channel programmable DDC for use in a wide range of radio applications.

Intermediate frequency (IF) modem designs have multichannel, multirate filter lineups. The filter lineup is often programmable from a host processor, and has stringent demands on DSP performance and accuracy. The DDC design example is a high-performance design running at over 350 MHz in an Intel Arria 10 device. The design example is very efficient and at under 4,000 logic registers gives a low cost per channel.

View the **DDCChip** subsystem to see the components you require to build a complex, production ready system.

Figure 37. Testbench for the DDC Design



The top-level testbench includes **Control** and **Signals** blocks, and some Simulink blocks to generate source signals and visualize the output. The full power of the Simulink blocksets is available for your design.

The **DDCChip** subsystem block contains the following blocks that form the lowest level of the design hierarchy:

- The NCO and mixer
- Decimate by 16 CIC filter
- Two decimate by 4 FIR odd-symmetric filters: one with length 21, the other length with 63.

The other blocks in this subsystem perform a range of rounding and saturation functions. They also allow dynamic scaling. The **Device** block specifies the target FPGA.

### DDC Signals Block

The **Signals** block allows you to define the relationship between the sample rates and the system clock, to tell the synthesis engines how much folding or time sharing to perform. Increasing the system clock permits more folding, and therefore typically results in more resource sharing, and a smaller design.

You also need a system clock rate so that the synthesis engines know how much to pipeline the logic. For example, by considering the device and speed grade, the synthesis tool can calculate the maximum length that an adder can have. If the design exceeds this length, it pipelines the adder and adjusts the whole pipeline to compensate. This adjustment typically results in a small increase in logic size, which is usually more than compensated for by the decrease in logic size through increased folding.

The **Signals** block specifies the clock and reset names, with the system clock frequency. The bus clock or FPGA internal clock for the memory-mapped interfaces can be run at a lower clock frequency. This lets the design move the low-speed operations such as coefficient update completely off the critical path.

*Note:* To specify the clock frequency, clock margin, and bus clock frequency values in this design, use the MATLAB workspace variables `ClockRate` and `ClockMargin`, which you can edit by double-clicking on the **Edit Params** block.

### DDC Control Block

The **Control** block controls the whole DSP Builder advanced blockset environment. It examines every block in the system, controls the synthesis flow, and writes out all RTL and scripts. A single control block must be present in every top-level model.

In this design, hardware generation creates RTL. DSP Builder places the RTL and associated scripts in the directory `../rtl`, which is a relative path based on the current MATLAB directory. DSP Builder creates automatic self-checking testbenches, which saves the data that a Simulink simulation captures to build testbench stimulus for each block in your design. DSP Builder generates scripts to run these simulations.

The threshold values control the hardware generation. They control the trade-offs between hardware resources, such as hard DSP blocks or soft LE implementations of multipliers. You can perform resource balancing for your particular design needs with a few top-level controls.

### DDC Memory Maps

Many memory-mapped registers in the design exist such as filter coefficients and control registers for gains. You can access these registers through a memory port that DSP Builder automatically creates at the top-level of your design. DSP Builder can create all address decode and data multiplexing logic automatically. DSP Builder generates a memory map in XML and HTML that you can use to understand the design.

To access this memory map, after simulation, on the DSP Builder menu, point to **Resource Usage** and click **Design, Current Subsystem**, or **Selected block**. The address and data widths are set to 8 and 32 in the design.

### DDC EditParams Blocks

The **EditParams** block allows you to edit the script `setup_demo_ddc.m`, which sets up the MATLAB variables that configure your model. Use the MATLAB design example properties callback mechanism to call this script.

*Note:*

The `PreloadFcn` callback uses this script to setup the parameters when your design example opens and the `InitFcn` callback re-initializes your design example to take account of any changes when the simulation starts.

You can edit the parameters in the `setup_demo_ddc.m` script by double-clicking on the **Edit Params** block to open the script in the MATLAB text editor.

The script sets up MATLAB workspace variables. The `SampleRate` variable is set to 61.44 MHz, which is typical of a CDMA system, and represents a quarter of the system clock rate that the FPGA runs at. You can use the feature to TDM four signals onto any given wire.

### DDC Source Blocks

The Simulink environment enables you to create any required input data for your design. In the DDC design, use manual switches to select sine wave or random noise generators. DSP Builder encodes a simple six-cycle sine wave as a table in a **Repeating Sequence Stair** block from the Simulink Sources library. This sine wave is set to a frequency that is close to the carrier frequencies that you specify in the NCOs, allowing you to see the filter lineup decoding some signals. DSP Builder creates VHDL for each block as part of the testbench RTL.

### DDC Sink Blocks

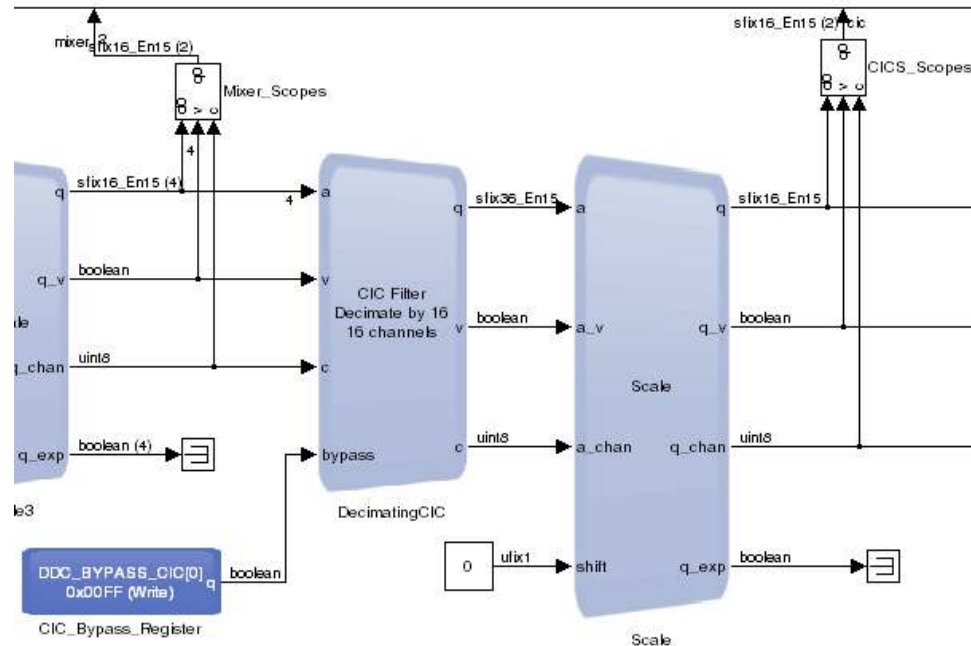
Simulink Sink library blocks display the results of the DDC simulation. The **Scope** block displays the raw output from the DDC design. The design has TDM outputs and all the data shows as **data**, **valid** and **channel** signals.

At each clock cycle, the value on the `data` wire either carries a genuine data output, or data that you can safely discard. The `valid` signal differentiates between these two cases. If the data is valid, the `channel` wire identifies the channel where the data belongs. Thus, you can use the `valid` and `channel` wires to filter the data. The **ChanView** block automates this task and decodes 16 channels of data to output channels 0 and 15. The block decimates these channels by the same rate as the whole filter line up and passes to a spectrum scope block (**OutSpectrum**) that examines the behavior in the frequency domain.



### DDC DecimatingCIC and Scale Blocks

Figure 38. DDChip Datapath (DecimatingCIC and Scale Blocks)



Two main blocks exist—the **DecimatingCIC** and the **Scale** block. To configure the CIC Filter, double click on the **DecimatingCIC** block.

The input sample rate is still the same as the data from the antenna. The `dspb_ddc.SampleRate` variable specifies the input sample rate. The number of channels, `dspb_ddc.ChanCount`, is a variable set to 16. The CIC filter has 5 stages, and performs decimation by a factor of 16. 1/16 in the dialog box indicates that the output rate is 1/16th of the input sample rate. The CIC parameter differential delay controls how many delays each CIC section uses—nearly always set to 1.


The CIC has no registers to configure, therefore no memory map elements exist.

The input data is a vector of four elements, so DSP Builder builds the decimating CIC from four separate CICs, each operating on four channels. The decimation behavior reduces the data rate at the output, therefore all 16 data samples (now at 61.44/16 MSPS each channel) can fit onto 1 wire.

The **DecimatingCIC** block multiplexes the results from each of the internal CIC filters onto a single wire. That is, four channels from vector element 1, followed by the four channels from vector element 2. DSP Builder packs the data onto a single TDM wire. Data is active for 25% of the cycles because the aggregate sample rate is now  $61.44 \text{ MSPS} \times 16 \text{ channels} / 16 \text{ decimation} = 61.44 \text{ MSPS}$  and the clock rate for the system is 245.76 MHz.

Bursts of data occur, with 16 contiguous samples followed by a gap. Each burst is tagged with the valid signal. Also the channel indicator shows that the channel order is 0..15.

**Figure 39. CIC\_All\_Scope Showing Output From the DecimatingCIC Block**



The number of input integrator sections is 4, and the number of output comb sections is 1. The lower data rate reduces the size of the overall group of 4 CICs. The Help page also reports the gain for the DCIC to be 1,048,576 or approximately  $2^{20}$ . The Help page also shows how DSP Builder combines the four channels of input data on a single output data channel. The comb section utilization (from the DSP Builder menu) confirms the 25% calculation for the folding factor.

The **Scale** block reduces the output width in bits of the CIC results.

In this case, the design requires no variable shifting operation, so it uses a Simulink constant to tie the shift input to 0. However, because the gain through the **DecimatingCIC** block is approximately  $2^{20}$  division of the output, enter a scalar value -20 for the **Number of bits to shift left** in the dialog box to perform data.

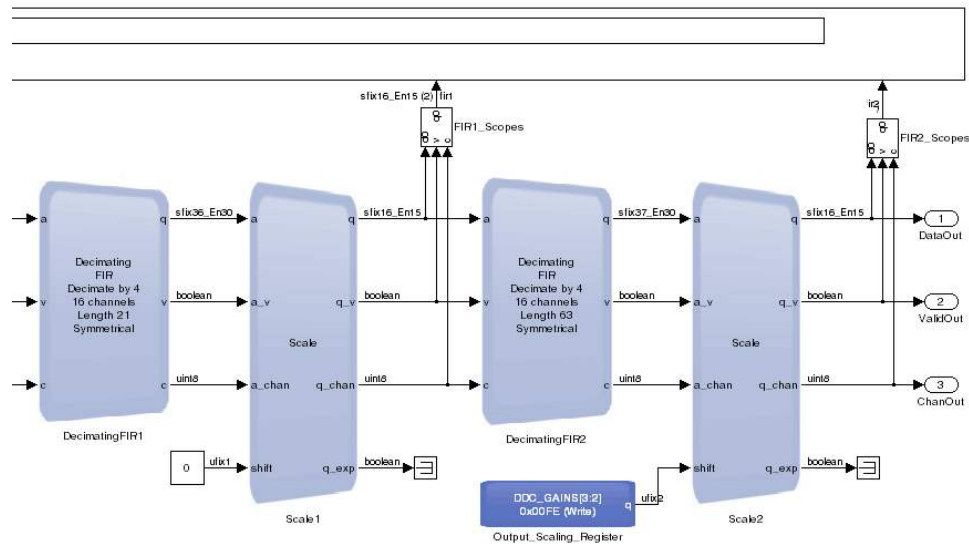
*Note:* Enter a scalar rather than a vector value to indicate that the scaling is static.

### DDC Decimating FIR Blocks

The last part of the DDC datapath comprises two decimating finite impulse response (FIR) blocks (**DecimatingFIR1** and **DecimatingFIR2**) and their corresponding scale blocks (**Scale1** and **Scale2**).

These two stages are very similar, the first filter typically compensates for the undesirable pass band response of the CIC filter, and the second FIR fine tunes the response that the waveform specification requires.

Figure 40. DDCChip Datapath (Decimating FIR and Scale Blocks)



The first decimating FIR decimates by a factor of 4.

The input rate per channel is the output sample rate of the decimating CIC, which is 16 times lower than the raw sample rate from the antenna.

*Note:* You can enter any MATLAB expression, so DSP Builder can extract the 16 out as a variable to provide additional parameterization of the whole design.

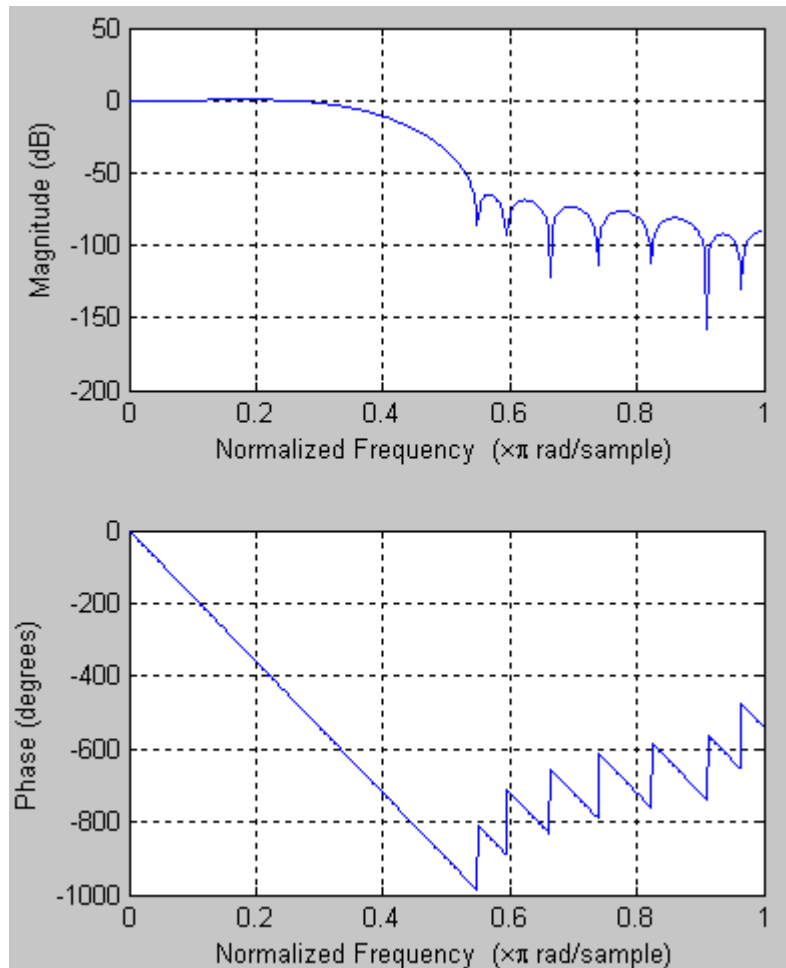
This filter performs decimation by a factor of 4 and the calculations reduce the size of the FIR filter. 16 channels exist to process and the coefficients are symmetrical.

The **Coefficients** field contains information that passes as a MATLAB fixed-point object (**fi**), which contains the data, and also the size and precision of each coefficient. Specifying an array of floating-point objects in the square brackets to the constructor to achieve this operation. The length of this array is the number of taps in the filter. At the end of this expression, the numbers 1, 16, 15 indicate that the fixed-point object is signed, and has 16-bit wide elements of which 15 are fractional bits.

For more information about **fi** objects, refer to the MATLAB Help.

This simple design uses a low-pass filter. In a real design, more careful generation of coefficients may be necessary.

Figure 41. DecimatingFIR1 Magnitude and Phase Responses



The output of the FIR filter fits onto a single wire, but because the data reduces further, there is a longer gap between frames of data.

Access a report on the generated FIR filter from the Help page.

You can scroll down in the Help page to view the port interface details. These match the hardware block, although the RTL has additional ports for clock, reset, and the bus interface.

The report shows that the input data format uses a single channel repeating every 64 clock cycles and the output data is on a single channel repeating every 256 clock cycles.

Details of the memory map include the addresses DSP Builder requires to set up the filter parameters with an external microprocessor.

You can show the total estimated resources by clicking on the DSP Builder menu, pointing to **Resources**, and clicking **Device**. Intel estimates this filter to use 338 LUT4s, 1 18×18 multiplier and 7844 bits of RAM.

The **Scale1** block that follows the **DecimatingFIR1** block performs a similar function to the **DecimatingCIC** block.

The **DecimatingFIR2** block performs a second level of decimation, in a very similar way to **DecimatingFIR1**. The coefficients use a MATLAB function. This function (**fir1**) returns an array of 63 **doubles** representing a low pass filter with cut off at 0.22. You can wrap this result in a **fi** object:

```
fi(fir1(62, 0.22), 1, 16, 15)
```

1. [DDC Design Example Subsystem](#) on page 93
2. [Building the DDC Design Example](#) on page 96

### 6.3.1. DDC Design Example Subsystem

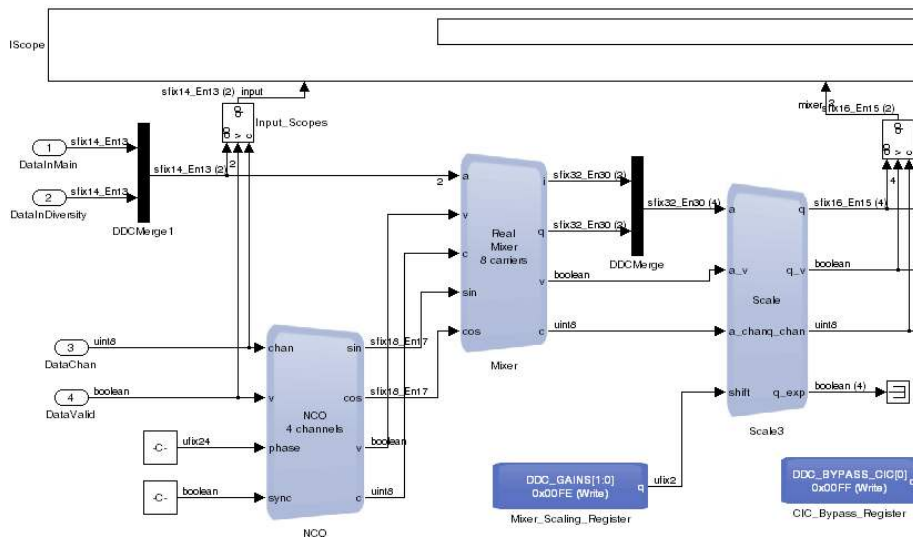
DSP Builder generates VHDL for all levels of the hierarchy, but subsystems have additional script files that build a project in the Quartus Prime software.

The **DDCChip** subsystem contains a **Device** block. This block labels this level of design hierarchy that compiles onto the FPGA. The **Device** block sets the FPGA family, design and speed grade. The family and speed grade optimize the hardware. In combination with the target clock frequency, the device determines the degree of pipelining.

The **DDCChip** subsystem has three types of block:

- The grey blocks are **IP** blocks. These represent functional IP such as black box filters, NCOs, and mixers.
- The blue blocks are processor visible registers.
- The black and white blocks are Simulink blocks.

Figure 42. DDCChip Datapath (NCO and Mixer)



Note: The inputs, NCO, and mixer stages show with Simulink signal formats turned on.

### DDCChip Primary Inputs

The primary inputs to the hardware are two parallel data signals (`DataInMain` and `DataInDiversity`), a channel signal (`DataChan`), and a valid signal (`DataValid`). The parallel data signals represent inputs from two antennas. They are of type **sfix14\_13** which is a Simulink fixed-point type of total width 14 bits. The type is signed with 13 bits of fraction, which is a typical number format that an analog-to-digital converter generates.

The data channel `DataChan` is always an 8-bit unsigned integer (**uint8**) and DSP Builder synthesizes away the top bits if not used. The valid signal `DataValid` indicates when real data transmits. The first rising edge of the valid signal starts operation of the first blocks in the chain. As the first blocks start producing outputs, their valid outputs start the next blocks in the chain. This mechanism ensures that filter chain start up is coordinated without having a global controller for the latencies of each block. The actual latencies of the blocks may change based on the clock frequency and FPGA selection.

### DDC Merge Multiplexer

The **IP** blockset supports vectors on its input and output data wires, which ensures that a block diagram is scalable when, for example, changing channel counts and operating frequencies. The merge multiplexer (**DDCMerge1**) takes two individual wires and combines them into a vector wire of width 2. This Simulink **Mux** block does not perform any multiplexing in hardware—it is just as a vectorizing block. If you examine the RTL, it contains just wires.

### DDC NCO

The NCO block generates sine and cosine waveforms to a given precision. These waveforms represent a point in the complex plane rotating around the origin at a given frequency. DSP Builder multiplies this waveform by the incoming data stream to obtain the data from the transmitted signal.

*Note:*

Four frequencies exist, because the vector in the **Phase Increment and Inversion** field is of length 4.

DSP Builder configures the NCO block to produce a signed 18-bit value with 17 bits of fraction. The internal accumulator width is set to 24 bits. This internal precision affects the spurious-free dynamic range (SFDR). DSP Builder specifies the initial frequencies for the simulation as phase increments. The phase accumulator width in bits is  $2^{24}$ , thus one complete revolution of the unit circle corresponds to a value of  $2^{24}$ . Dividing this number by 5.95, means that the design requires 5.95 cycles to perform one complete rotation. That is, the wavelength of the sine and cosine that the design produces are 5.95 cycles. The sample rate is 61.44 MHz, therefore the frequency is  $61.44/5.95$ , which is 10.32 MHz.

The input frequency in the testbench rotates every 6 cycles for a frequency of  $61.44/6=10.24$  MHz. Therefore, you can expect to recover the difference of these frequencies (0.08 MHz or 80 kHz), which fall in the low-pass filters pass bands, because DSP Builder mixes these signals.

The design exposes phase values through a memory-mapped interface at the address specified by the variable `DDC_NCO_PHASE_INCR`, which is set to address `0x0000` in the setup script. After simulation, to view resource usage for the design example, the subsystem, or a selected block, on the DSP Builder menu, point to **Resource Usage** and click **Design**, **Current Subsystem**, or **Selected block**.

DSP Builder reports for each register, the name, width, reset value, and address. This report collates all the registers from your design into a single location.

You can view the estimated results for this NCO configuration in the **Results** tab of the dialog box).

Based on the selected accumulator width and output width, DSP Builder calculates an estimated SFDR and accumulator precision. To verify this precision in a separate testbench, use `demo_nco.mdl` as a start.

### DDC Mixer

The **Mixer** block performs the superheterodyne operation by multiplying each of the two received signals (`DataInMain` and `DataInDiversity`) by each of the four frequencies. This action produces eight complex signals or 16 scalar signals (the 16 channels in the DDC design).

The mixer requires sufficient multipliers to perform this calculation. The total number of real  $\times$  complex multipliers required for each sample is  $2 \text{ signals} \times 4 \text{ frequencies} = 8$ .

Thus, 8 real  $\times$  complex multipliers require  $8 \times 2 = 16$  scalar multipliers. This processing is spread over four cycles (the folding factor given by the ratio of clock rate to sample rate), therefore DSP Builder requires four physical multipliers.

After simulation, to view resource usage for the design example, the subsystem, or a selected block, on the DSP Builder menu, point to **Resource Usage** and click **Design**, **Current Subsystem**, or **Selected block**.

You can list the input and output ports that DSP Builder creates for this block, with the data width and brief description, by right-clicking on the block and clicking **Help**. DSP Builder suffixes the vector inputs with 0 and 1 to implement the vector. This list of signals corresponds to the signals in the VHDL entity.

DSP Builder provides the results for the mixer as separate in phase and quadrature outputs—each is a vector of width 2. It performs the remaining operations on both the I and Q signals, so that DSP Builder can combine them with another Simulink multiplexer to provide a vector of width 4. This operation carries the 16 signals, with a folding factor of 4. At this point the channel counts count 0, 1, 2, 3, 0, 1, . . . .

### DDC Mixer Scale Block

At this point in the datapath, the data width is 32 bits representing the full precision output of multiplying a 14-bit data signal with an 18-bit sine or cosine signal. DSP Builder needs to reduce the data width to a lower precision to pass on to the remaining filters, which reduces the resource count considerably, and does not cause significant information loss. The **Scale3** block performs a shift-round-saturate operation to achieve this reduction. The shift is usually a 1 or 2 bit shift that you can set to adjust the gain in your design at run time.

To determine the setup, DSP Builder usually uses a microprocessor, which writes to a register to set the shift amount. This design uses a **RegField** block (**Mixer\_Scaling\_Register**). This block behaves like a constant in the Simulink simulation, but in hardware the block performs as a processor-writable register that initializes to the value in your design example.

This parameterization results in a register mapped to address `DDC_GAINS`, which is a MATLAB variable that you specify in the `setup_demo_ddc.m` script.

The register is writable from the processor, but not readable.

The register produces a 2-bit output of type **ufix(2)**—an unsigned fixed-point number. The scaling is  $2^{-0}$  so is, in effect, a 2-bit unsigned integer. These 2 bits are mapped into bits 0 and 1 of the word (another register may use other bits of this same address). The initial value for the register is set to 0. DSP Builder provides a description of the memory map in the resource usage. Sometimes, Simulink needs an explicit sample time, but you can use the default value of -1 for this tutorial.

The 2-bit unsigned integer is fed to the **Scale3** block. This block has a vector of width 4 as its data input. The **Scale3** block builds a vector of 4 internal scale units. These parameters are not visible through the user interface, but you can see them in the resource usage.

The block produces four outputs, which DSP Builder presents at the output as a vector of width 4. DSP Builder preserves the order in the vector. You can create quite a large block of hardware by passing many channels through a **IP** block. The exception output of the scale block provides signals to say when saturation occurs, which this design does not require, so this design terminates them.

The design sets the output format to 16-bit signed with 15 bits of fraction and uses the **Unbiased** rounding method. This method (convergent rounding or round-to-even) typically avoids introducing a DC bias.

The saturation method uses **Symmetric** rounding which clips values to within +0.9999 and -0.9999 (for example) rather than clipping to -1. Again this avoids introducing a DC bias.

The number of bits to shift is a vector of values that the scaling register block (**Mixer\_Scaling\_Register**) indexes. The vector has 4 values, therefore DSP Builder requires a 2-bit input.

An input of 0 uses the 0th value in the vector (address 1 in Simulink), and so on. Therefore, in this example **inout0** shifts by 0 and the result at the input has the same numerical range as the input. An input of 1 shifts left by 1, and so multiplies the input value by 2, thus increasing the gain.

### 6.3.2. Building the DDC Design Example

1. Open the model, by typing the following command in the MATLAB window:

```
demo_ddc r
```

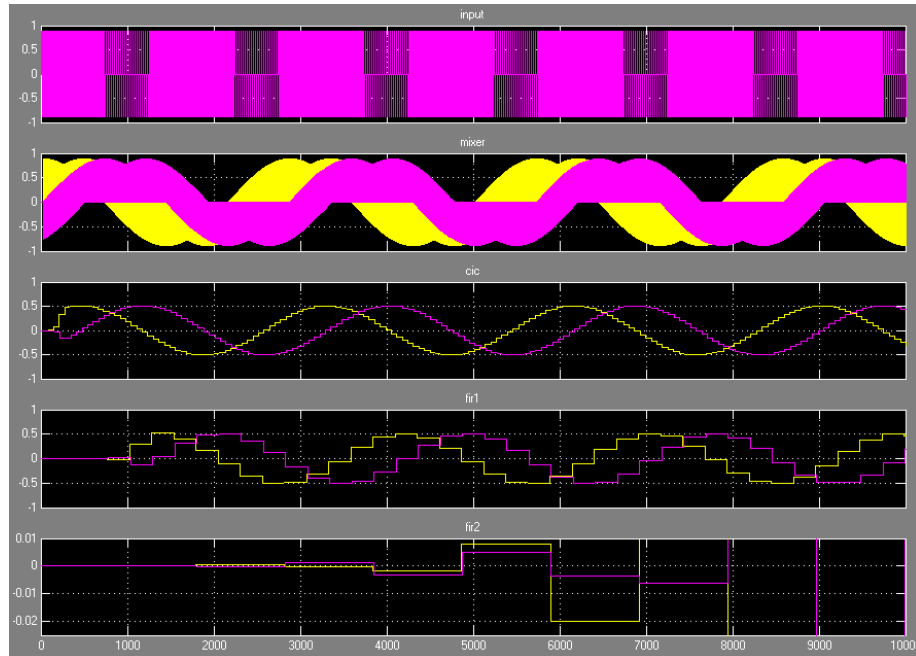
2. Simulate the design example in Simulink, by typing the following command in the MATLAB window:

```
sim('demo_ddc', 550000.0*demo_ddc.SampleTime);
```



**Figure 43. Simulation Results Shown in the IScope Block**

The **IScope** block shows the first two channels (1 real and 1 complex for the first carrier) of data (magenta and yellow) as the input signals. The first trace shows the rapidly changing input signal that the testbench generates. The second signal shows the result of the mixer. This slowly changing signal contains the information to be extracted, plus a lot of high frequency residue. Applying the series of low-pass filters and decimating results in the required data.



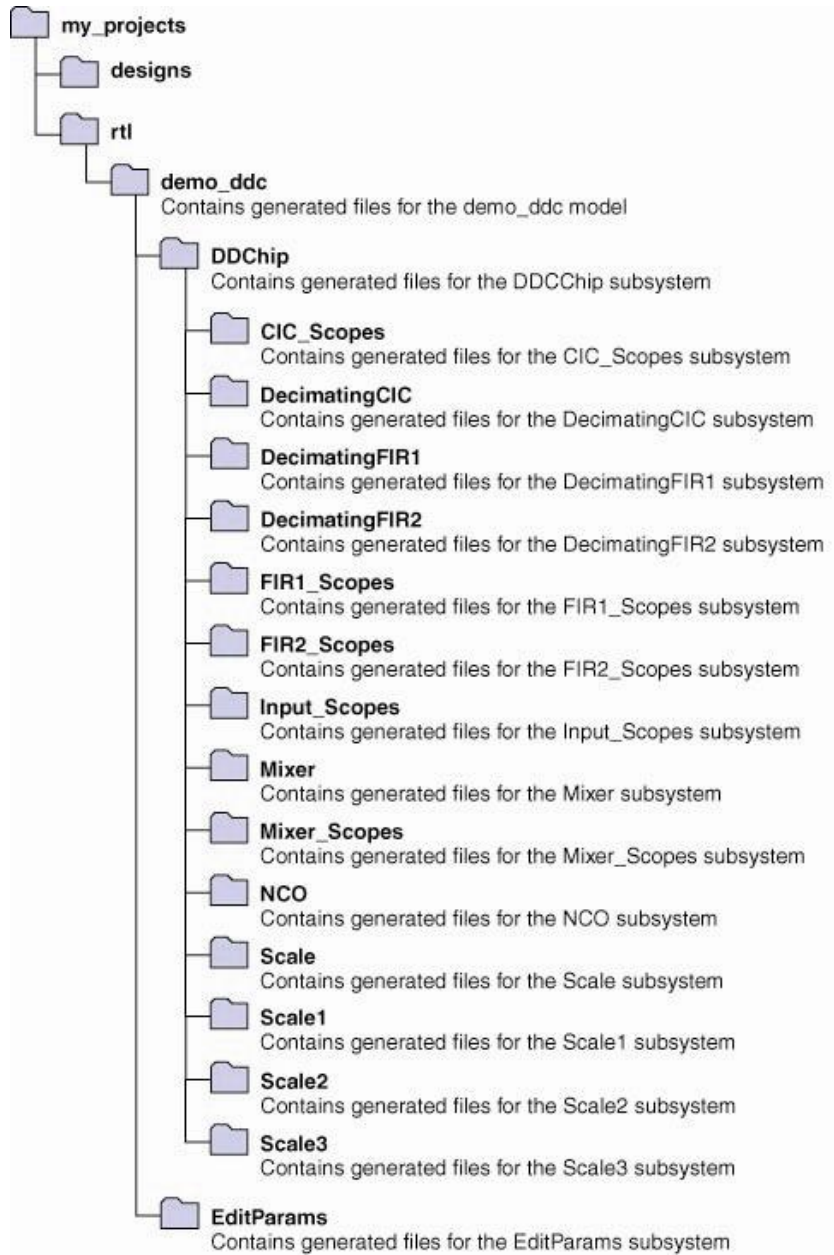
*Note:* If you turn on the **Generate Hardware** option in the parameters for the **Control** block, every time the simulation runs, DSP Builder synthesizes the underlying hardware, and writes out VHDL into the directory you specify.

3. Simulate the generated RTL in the ModelSim simulator.
4. Synthesize and fit the RTL in the Quartus Prime software.

### 6.3.2.1. DDC Design Example Generated Files

DSP Builder creates a directory structure that mirrors the structure of your model. The root to this directory can be an absolute path name or as a relative path name; for a relative path name (such as `../rtl`), the directory structure is relative to the MATLAB current directory.

**Figure 44. Generated Directory Structure for the DDC Design Example**



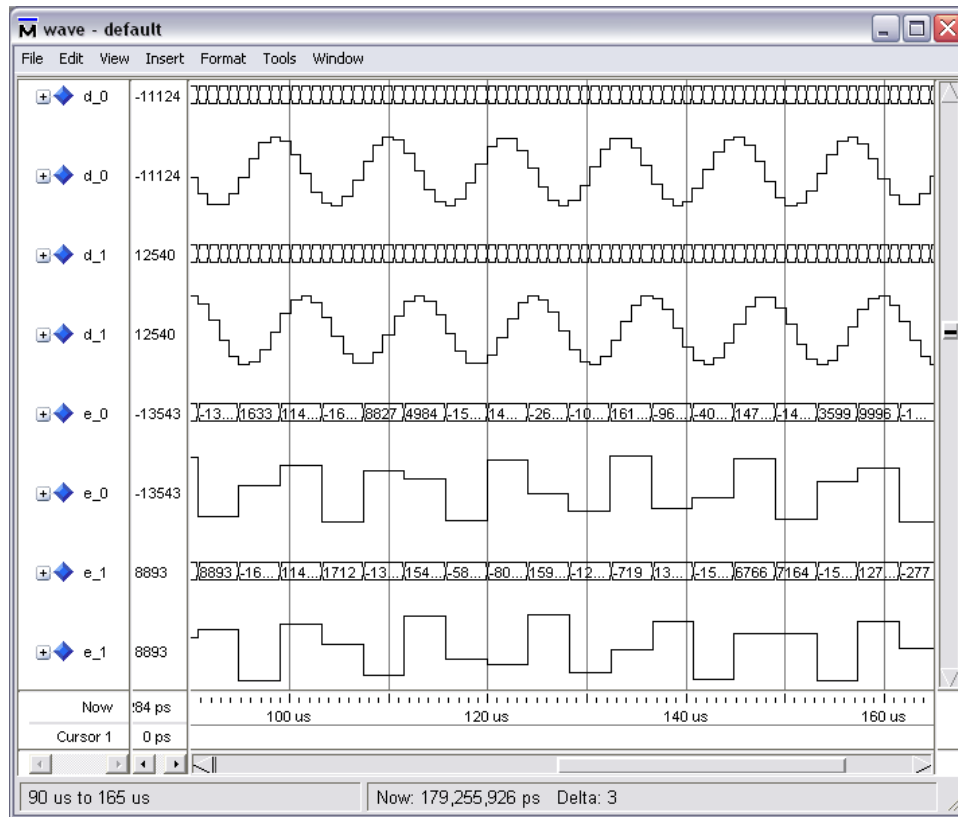
Note: Separate subdirectories exist corresponding to each hierarchical level in your design.

**Table 15. Generated Files for the DDC Design Example**

File	Description
<b>rtl</b> directory	
demo_ddc.xml	An XML file that describes the attributes of your model.
demo_ddc_entity.xml	An XML file that describes the boundaries of the system (for Signal Compiler in designs that combine blocks from the standard and advanced blocksets).
rtl\demo_ddc subdirectory	
<block name>.xml	An XML file containing information about each block in the advanced blockset which translates into HTML on demand for display in the MATLAB Help viewer and for use by the DSP Builder menu options.
demo_ddc.vhd	This is the top-level testbench file. It may contain non-synthesizable blocks, and may also contain empty black boxes for Simulink blocks that are not fully supported.
demo_ddc.add.tcl	This script loads the VHDL files in this subdirectory and in the subsystem hierarchy below it into the Quartus Prime project.
demo_ddc.qip	This file contains all the assignments and other information DSP Builder requires to process the <b>demo_ddc design</b> example in the Quartus Prime software. The file includes a reference to the .qip file in the <b>DDCChip</b> subsystem hierarchy.
<block name>.vhd	DSP Builder generates a VHDL file for each component in your model.
demo_ddc_DDCChip_entity.xml	An XML file that describes the boundaries of the <b>DDCChip</b> subsystem as a black box (for Signal Compiler in designs that combine blocks from the standard and advanced blocksets).
DDCChip.xml	An XML file that describes the attributes of the <b>DDCChip</b> subsystem.
*.stm	Stimulus files.
safe_path.vhd	Helper function that ensures a pathname is read correctly in the Quartus Prime software.
safe_path_msim.vhd	Helper function that ensures a pathname is read correctly in ModelSim.
rtl\demo_ddc\<subsystem> subdirectories	
Separate subdirectories exist for each hierarchical level in your design. These subdirectories include additional .xml .vhd, qip and .stm files describing the blocks contained in each level. Also additional .do and .tcl files exist, which it automatically calls from the corresponding files in the top-level of your model.	
<subsystem>_atb.do	Script that loads the subsystem automatic testbench into ModelSim.
<subsystem>_atb.wav.do	Script that loads signals for the subsystem automatic testbench into ModelSim.
<subsystem>/<block>/*.hex	Intel format .hex files that initialize the RAM blocks in your design for either simulation or synthesis.
<subsystem>.sdc	Design constraint file for TimeQuest support.
<subsystem>.tcl	Use this Tcl file to setup a Quartus Prime project.
<subsystem>_hw.tcl	A Tcl script that loads the generated hardware into Platform Designer.

To display a particular signal, attach a Simulink scope, regenerate, and resimulate.

**Figure 45. Simulation Results in the ModelSim Wave Window for the DDC Design Example**



## 6.4. DSP Builder Filter Design Examples

This folder contains design examples of cascaded integrator-comb (CIC) and finite impulse response (FIR) filters.

1. [Complex FIR Filter](#) on page 101
2. [Decimating CIC Filter](#) on page 101
3. [Decimating FIR Filter](#) on page 101
4. [Filter Chain with Forward Flow Control](#) on page 102
5. [FIR Filter with Exposed Bus](#) on page 102
6. [Fractional FIR Filter Chain](#) on page 102
7. [Fractional-Rate FIR Filter](#) on page 102
8. [Half-Band FIR Filter](#) on page 103
9. [IIR: Full-rate Fixed-point](#) on page 103
10. [IIR: Full-rate Floating-point](#) on page 103
11. [Interpolating CIC Filter](#) on page 104
12. [Interpolating FIR Filter](#) on page 104

13. [Interpolating FIR Filter with Multiple Coefficient Banks](#) on page 105
14. [Interpolating FIR Filter with Updating Coefficient Banks](#) on page 106
15. [Root-Raised Cosine FIR Filter](#) on page 106
16. [Single-Rate FIR Filter](#) on page 106
17. [Super-Sample Decimating FIR Filter](#) on page 107
18. [Super-Sample Fractional FIR Filter](#) on page 107
19. [Super-Sample Interpolating FIR Filter](#) on page 107
20. [Variable-Rate CIC Filter](#) on page 107

### 6.4.1. Complex FIR Filter

This design example demonstrates how to implement a complex FIR filter using three real filters. The resource efficient implementation (three real multipliers per complex multiply) maps optimally onto Intel Arria 10 DSP blocks, using the scan and cascade modes.

The model file is `demo_complex_fir.mdl`.

### 6.4.2. Decimating CIC Filter

This design example implements a decimating CIC filter.

The top-level testbench includes **Control**, **Signals**, **Run ModelSim**, and **Run Quartus Prime** blocks, plus **ChanView** block that deserialize the output buses. An **Edit Params** block allows easy access to the setup variables in the `setup_demo_dcic.m` script.

The **CICSystem** subsystem includes the **Device** and **DecimatingCIC** blocks.

The model file is `demo_dcic.mdl`.

*Note:* This design example uses the Simulink Signal Processing Blockset.

### 6.4.3. Decimating FIR Filter

This design example implements a decimating FIR filter.

This design example uses the **Decimating FIR** block to build a 20-channel decimate by 5, 49-tap FIR filter with a target system clock frequency of 240 MHz.

The top-level testbench includes **Control**, **Signals**, **Run ModelSim**, and **Run Quartus Prime** blocks, plus **ChanView** block that deserialize the output buses. An **Edit Params** block allows easy access to the setup variables in the `setup_demo_fird.m` script.

The **FilterSystem** subsystem includes the **Device** and **Decimating FIR** blocks.

The model file is `demo_fird.mdl`.

*Note:* This design example uses the Simulink Signal Processing Blockset.

#### 6.4.4. Filter Chain with Forward Flow Control

This design example builds a filter chain with forward flow control.

The top-level testbench includes **Control**, **Signals**, **Run ModelSim**, and **Run Quartus Prime** blocks, plus **ChanView** block that deserialize the output buses. An **Edit Params** block allows easy access to the setup variables in the `setup_demo_filters_flow_control.m` script.

The **FilterSystem** subsystem includes **FractionalRateFIR**, **InterpolatingFIR**, **InterpolatingCIC**, **Const** and **Scale** blocks.

The model file is `demo_filters_flow_control.mdl`.

*Note:* This design example uses the Simulink Signal Processing Blockset.

#### 6.4.5. FIR Filter with Exposed Bus

This design example is a multichannel single-rate FIR filter with rewritable coefficients. The initial configuration is a high-pass filter, but halfway through the testbench simulation, DSP Builder reconfigures it as a low-pass filter. The testbench feeds in the sum of a fast and a slow sine wave into the filter. The fast one emerges from the originally configured FIR filter; the slow one is all that is left after DSP Builder reconfigures the filter.

The model file is `demo_fir_exposed_bus.mdl`.

#### 6.4.6. Fractional FIR Filter Chain

This design example uses a chain of **InterpolatingFIR** and **DecimatingFIR** blocks to build a 16-channel fractional rate filter with a target system clock frequency of 360 MHz.

The top-level testbench includes **Control**, **Signals**, **Run ModelSim**, and **Run Quartus Prime** blocks, plus **ChanView** block that deserialize the output buses. An **Edit Params** block allows easy access to the setup variables in the `setup_demo_fir_fractional.m` script.

The **FilterSystem** subsystem includes **ChanView**, **Decimating FIR**, **InterpolatingFIR**, and **Scale** blocks.

The model file is `demo_fir_fractional.mdl`.

*Note:* This design example uses the Simulink Signal Processing Blockset.

#### 6.4.7. Fractional-Rate FIR Filter

This design example implements a fractional rate FIR filter.

The top-level testbench includes **Control**, **Signals**, **Run ModelSim**, and **Run Quartus Prime** blocks, plus **ChanView** block that deserialize the output buses. An **Edit Params** block allows easy access to the setup variables in the `setup_demo_firf.m` script.

The **FilterSystem** subsystem includes the **Device** and **FractionalRateFIR** blocks.

The model file is `demo_firf.mdl`.

*Note:* This design example uses the Simulink Signal Processing Blockset.

### 6.4.8. Half-Band FIR Filter

This design example implements a half band interpolating FIR filter.

The top-level testbench includes **Control**, **Signals**, **Run ModelSim**, and **Run Quartus Prime** blocks, plus **ChanView** block that deserialize the output buses. An **Edit Params** block allows easy access to the setup variables in the `setup_demo_firih.m` script.

The **FilterSystem** subsystem includes the **Device** block and two separate **InterpolatingFIR** blocks for the regular and interpolating filters.

The model file is `demo_firih.mdl`.

This design example uses the Simulink Signal Processing Blockset.

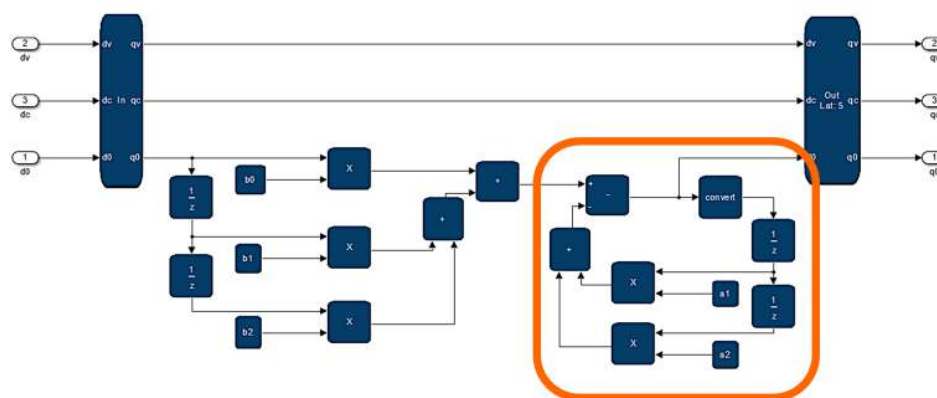
### 6.4.9. IIR: Full-rate Fixed-point

This design example implements a full-rate fixed-point IIR filter.

This design demonstrates a single-channel second-order Infinite Impulse Response (IIR) filter running at the clock rate. Usually with such designs, closing the feedback loop is difficult at high clock rates. This design recursively expands the mathematical expression from the feedback in terms of earlier samples, which gives a feed-forward scalar product and a longer feedback loop. You can make the feedback loop long enough to add any length of pipelining at the expense of more resources for the expansion.

The model file is `demo_full_rate_iir_fixed.mdl`.

**Figure 46. IIR Second-Order Biquad**



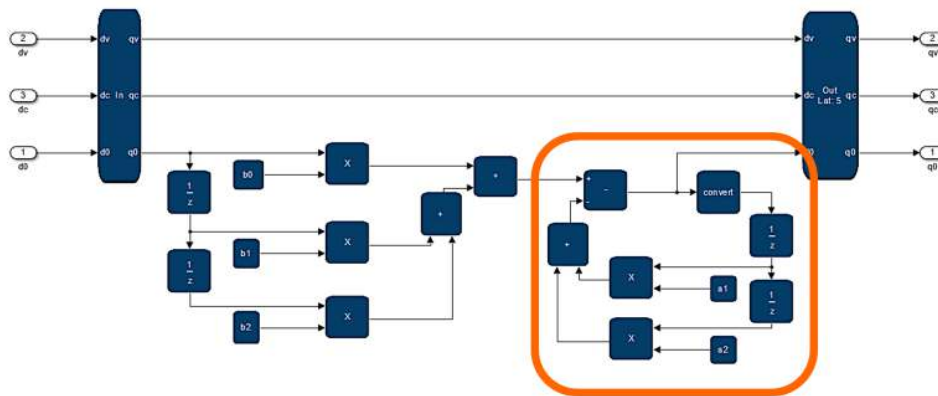
### 6.4.10. IIR: Full-rate Floating-point

This design example implements a full-rate floating-point IIR filter.

This design demonstrates a single-channel second-order Infinite Impulse Response (IIR) filter running at the clock rate. Usually with such designs, closing the feedback loop is impossible at high clock rates. This design recursively expands the mathematical expression from the feedback in terms of earlier samples, which gives a feed-forward scalar product and a longer feedback loop. You can make the feedback loop long enough to add any length of pipelining at the expense of more resources for the expansion.

The model file is `demo_full_rate_iir_floating.mdl`.

**Figure 47. IIR Second-Order Biquad**



### 6.4.11. Interpolating CIC Filter

This design example implements an interpolating CIC filter.

The top-level testbench includes **Control**, **Signals**, **Run ModelSim**, and **Run Quartus Prime** blocks, plus **ChanView** block that deserializes the output buses. An **Edit Params** block allows easy access to the setup variables in the `setup_demo_icic.m` script.

The **FilterSystem** subsystem includes the **Device** and **InterpolatingCIC** blocks.

The model file is `demo_icic.mdl`.

*Note:* This design example uses the Simulink Signal Processing Blockset.

### 6.4.12. Interpolating FIR Filter

This design example uses the **InterpolatingFIR** block to build a 16-channel interpolate by 2, symmetrical, 49-tap FIR filter with a target system clock frequency of 240 MHz.

The top-level testbench includes **Control**, **Signals**, **Run ModelSim**, and **Run Quartus Prime** blocks, plus **ChanView** block that deserializes the output buses. An **Edit Params** block allows easy access to the setup variables in the `setup_demo_firi.m` script.

The **FilterSystem** subsystem includes the **Device** and **InterpolatingFIR** blocks.

The model file is `demo_firi.mdl`.



*Note:* This design example uses the Simulink Signal Processing Blockset.

### 6.4.13. Interpolating FIR Filter with Multiple Coefficient Banks

This design example builds an interpolating FIR filter that regularly switches between coefficient banks.

Multiple sets of coefficients requires storage in memory so that the design can switch easily from one set, or bank, of coefficients in use to another in a single clock cycle.

The design must perform the following actions:

- Specify the number of coefficient banks
- Initialize the banks
- Update the coefficients in a particular bank
- Select the bank in use in the filter

You specify the coefficient array as a matrix rather than a vector—(bank rows) by (number of coefficient columns).

The addressing scheme has address offsets of base address + (bank number \* number of coefficients for each bank).

If the number of rows is greater than one, DSP Builder creates a bank select input port on the FIR filter. In a design, you can drive this input from either data or bus interface blocks, allowing either direct or bus control. The data type is unsigned integer of width  $\text{ceil}(\log_2(\text{number of banks}))$ .

The bank select is a single signal. For example, for a FIR filter with four input channels over two timeslots:

<0><1>

<2><3>

The corresponding input channel signal is:

<0><1>

Here the design receives more than one channel at a time, but can only choose a single bank of coefficients. Channels 0 and 2 use one set of coefficients and channels 1 and 3 another. Channel 0 cannot use a different set of coefficients to channel 2 in the same filter.

For multiple coefficient banks, you enter an array of coefficients sets, rather than a single coefficient set. For example, for a MATLAB array of 1 row and 8 columns [1 x 8], enter:

```
fi(fir1(7, 0.5 ),1,16,15)
```

For a MATLAB array of 2 rows and 8 columns [2 x 8] enter:

```
[fi(fir1(7, 0.5 ),1,16,15);fi(fir1(7, 0.5 ),1,16,15)]
```

Therefore, you can determine the number of banks by the number of rows without needing the number of banks. If the number of banks is greater than 1, add an additional bank select input on the block.

The model file is `demo_firi_multibank.mdl`.

#### 6.4.14. Interpolating FIR Filter with Updating Coefficient Banks

This design example is similar to the Interpolating FIR Filter with Multiple Coefficient Banks design example. While one bank is in use DSP Builder writes a new set of FIR filter coefficients to the other bank. You can see the resulting change in the filter output when the bank select switches to the updated bank.

Write to the bus interface using the **BusStimulus** block with a sample rate proportionate with the bus clock. Generally, DSP Builder does not guarantee bus interface transactions to be cycle accurate in Simulink simulations. However, in this design example, DSP Builder updates the coefficient bank while it is not in use.

The model name is `demo_firi_updatecoeff.mdl`.

#### 6.4.15. Root-Raised Cosine FIR Filter

This design example uses the Decimating FIR block to build a 4-channel decimate by 5, 199-tap root raised cosine filter with a target system clock frequency of 304 MHz.

The top-level testbench includes **Control**, **Signals**, **Run ModelSim**, and **Run Quartus Prime** blocks, plus **ChanView** block that deserialize the output buses. An **Edit Params** block allows easy access to the setup variables in the `setup_demo_fir_rrc.m` script.

The **FilterSystem** subsystem includes the **Device** and **Decimating FIR** blocks.

The model file is `demo_fir_rrc.mdl`.

*Note:* This design example uses the Simulink Signal Processing Blockset.

#### 6.4.16. Single-Rate FIR Filter

This design example uses the **SingleRateFIR** block to build a 16-channel single rate 49-tap FIR filter with a target system clock frequency of 360 MHz.

The top-level testbench includes **Control**, **Signals**, **Run ModelSim**, and **Run Quartus Prime** blocks, plus **ChanView** block that deserialize the output buses. An **Edit Params** block allows easy access to the setup variables in the `setup_demo_firs.m` script.

The **FilterSystem** subsystem includes the **Device** and **SingleRateFIR** blocks.

The model file is `demo_firs.mdl`.

*Note:* This design example uses the Simulink Signal Processing Blockset.

### 6.4.17. Super-Sample Decimating FIR Filter

This design example shows how the filters cope with data rates greater than the clock rate. The design example uses the **DecimatingFIR** block to build a single channel decimate by 2, symmetrical, 33-tap FIR filter.

The input sample rate is six times the clock rate. The filter decimates by two the input sample rate to three times the clock rate, which is visible in the vector input and output data connections. The input receives six samples in parallel at the input, and three samples are output each cycle.

After simulation, you can view the resource usage.

The model file is `demo_ssfirm.mdl`.

### 6.4.18. Super-Sample Fractional FIR Filter

This design example shows how the filters cope with data rates greater than the clock rate. The design example uses the **FractionalFIR** block to build a single channel interpolate by 3, decimate by 2, symmetrical, 33-tap FIR filter.

The input sample rate is two times the clock rate. The filter upconverts the input sample rate to three times the clock rate, which is visible in the vector input and output data connections. The input receives two samples in parallel at the input, and three samples are output each cycle.

The model file is `demo_ssfirf.mdl`.

### 6.4.19. Super-Sample Interpolating FIR Filter

This design example shows how the filters cope with data rates greater than the clock rate. The design example uses the **InterpolatingFIR** block to build a single channel interpolate by 3, symmetrical, 33-tap FIR filter.

The input sample rate is twice the clock rate and is interpolated by three by the filter to six times the clock rate, which is visible in the vector input and output data connections. The input receives two samples in parallel at the input, and six samples are output each cycle.

After simulation, you can view the resource usage.

The model file is `demo_ssfirmi.mdl`.

*Note:* This design example uses the Simulink Signal Processing Blockset.

### 6.4.20. Variable-Rate CIC Filter

CIC filters are extremely hardware efficient, as they require no multipliers. You see CIC filters commonly in applications that require large interpolation and decimation factors. Usually the interpolation and decimation factors are fixed, and you can use the CIC **IP** block. However, a subset of applications require you to change the interpolation and decimation factors at run time. This design example shows how to build a variable-rate CIC filter from primitives. It contains a variable-rate decimating CIC filter, which consists of a number of integrators and differentiators with a decimation block between them, where the rate change occurs.

You can control the rate change with a register field, which is part of the control interface. The register field controls the generation of a valid signal that feeds into the differentiators.

The design example also contains a gain compensation block that compensates for the rate change dependent gain of the CIC. It shifts the input up so that the MSB at the output is always at the same position, regardless of the rate change that you select.

The associated setup file contains parameters for the minimum and maximum decimation rate, and calculates the required internal data widths and the scaling number. To change the decimation factor for simulation, adjust variable **CicDecRate** to the desired current decimation rate.

The model file is `demo_vcic.mdl`.

## 6.5. DSP Builder Finite State Machine Design Example

The Finite State Machine example design demonstrates some of the features of the finite state machine (FSM) specification and its function in a primitive subsystem. The example first selects 20 odd numbers from the output of the counter block and then selects 8 multiples of 4 from that same counter. The model file is `demo_fsm.mdl`.

**Figure 48. Example**

```
# example 0
inputs odd four
enable ena
start go
netlist
for x 20 > 0 step -1
    transitions oddWait
    finish fin
    state _init
        if (~odd) hold
    state hold
        if (odd) _init
    end
end
end

# for loop specifies its counter as an output port
for x 0 < 8 : c
    # state transitions declare an output named q
    transitions fourWait : q
    # state transition back to _init state will
    #output a pulse on port named tko
    finish tko
    state _init
        if (~four) hold 1
    state hold
        if (four) _init 0
    end
end
end
```

### Related Information

[Finite State Machine](#) on page 364

## 6.6. DSP Builder Folding Design Examples

1. [Position, Speed, and Current Control for AC Motors](#) on page 109

2. [Position, Speed, and Current Control for AC Motors \(with ALU Folding\)](#) on page 113
3. [About FOC](#) on page 114
4. [Folded FIR Filter](#) on page 114

### 6.6.1. Position, Speed, and Current Control for AC Motors

This design example implements a field-oriented control (FOC) algorithm for AC motors such as permanent magnet synchronous machines (PMSM). Industrial servo motors, where the precise control of torque is important, commonly use these algorithms. This design example includes position and speed control, which allow the control of rotor speed and angle.

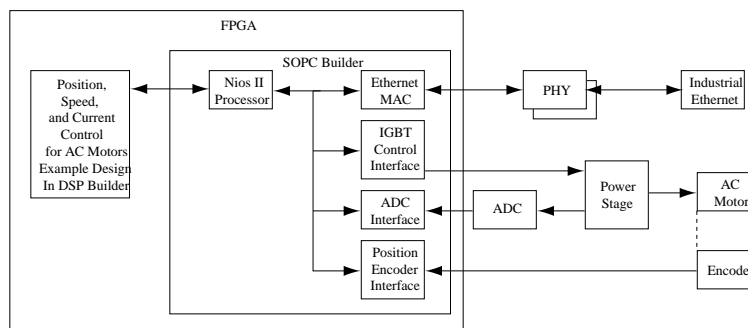
**Note:** Intel has not tested this design on hardware and Intel does not provide a model of a motor.

The model file is `psc_ctrl.mdl`. Also, an equivalent fixed-point design, `psc_ctrl_fixed.mdl`, exists. To change the precision this design uses, refer to the `setup_position_speed_current_controller_fixed.m` script.

#### Functional Description

An encoder measures the rotor position in the motor, which the FPGA then reads. An analog-to-digital converter (ADC) measures current feedback, which the FPGA then reads.

**Figure 49. AC Motor Control System Block Diagram**



Each of the FOC, speed, and position feedback loops use a simple PI controller to reduce the steady state error to zero. In a real-world PI controller, you may also need to consider integrator windup and tune the PI gains appropriately. The feedback loops for the integral portion of the PI controllers are internal to the design.

The example assumes you sample the inputs at a rate of 100 kHz and the FPGA clock rate is 100 MHz (suitable for Cyclone IV devices). ALU folding reduces the resource usage by sharing operators such as adders, multipliers, cosine. The folding factor is set to 100 to allow each operator to be timeshared up 100 times, which gives an input sample rate of 1 Msps, but as the real input sample rate is 100 ksps, only one out of every ten input timeslots are used. DSP Builder identifies the used timeslots when `valid_in` is 1. Use `valid_in` to enable the latch in the PI controller, which stores

data for use in the next valid timeslot. The `valid_out` signal indicates when the **ChannelOut** block has valid output data. You can calculate nine additional channels on the samedesign without incurring extra latency (or extra FPGA resources).

You should adjust the folding factor to see the effect it has on hardware resources and latency. To adjust, change the **Sample rate (MHz)** parameter in the **ChannelIn** and **ChannelOut** blocks of the design either directly or change the **FoldingFactor** parameter in the setup script. For example, a clock frequency of 100 MHz and sample rate of 10 MHz gives a folding factor of 10. Disabling folding, or setting the factor to 1, results in no resource sharing and minimal latency. Generally, you should not set the folding factor greater than the number of shareable operators, that is, for 24 adders and 50 multipliers, use a maximum folding factor 50.

*Note:* The testbench does not support simulations if you adjust the folding factor.

The control algorithm, with the FOC, position, speed, control loops, vary the desired position across time. The three control loops are parameterized with minimum and maximum limits, and PI values. These values are not optimized and are for demonstrations only.

### Resource Usage

**Table 16. Position, Speed, and Current Control for AC Motors Design Example Resource Usage**

Folding Factor	Add and Sub Blocks	Mult Blocks	Cos Blocks	Latency
No folding	22	22	4	170
>22	1	1	1	279

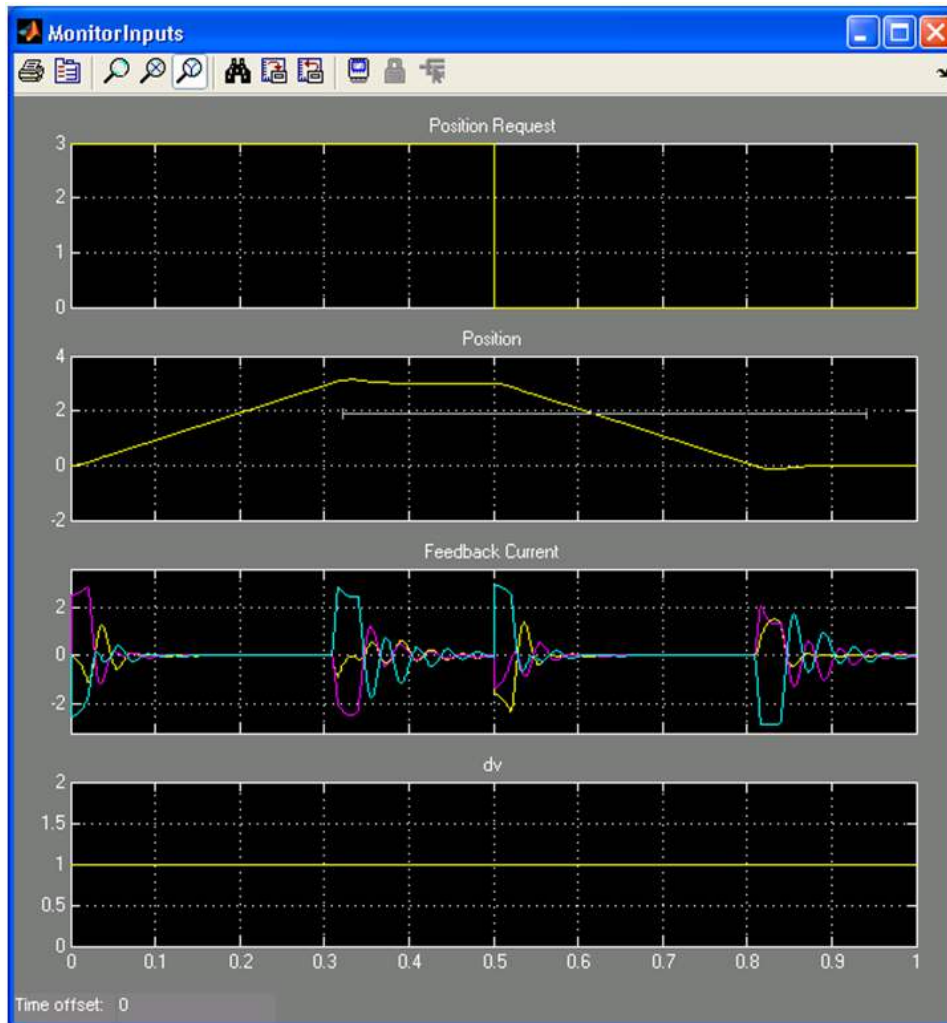
The example uses floating-point arithmetic that automatically avoids arithmetic overflow, but you can implement it in a fixed-point design and tune individual accuracies while manually avoiding overflows.

### Hardware Generation

When hardware generation is disabled, the Simulink system simulates the design at the external sample rate of 100 kHz, so that it outputs a new value once every 100 kHz. When hardware generation is enabled, the design simulates at the FPGA clock rate (100 MHz), which represents real-life latency clock delays, but it only outputs a new value every 100 kHz. This mode slows the system simulation speed greatly as the model is evaluated 1,000 times for every output. The setup script for the design example automatically detects whether hardware generation is enabled and sets the sample rates accordingly. The example is configured with hardware generation disabled, which allows fast simulations. When you enable hardware generation, set a very small simulation time (for example 0.0001 s) as simulation may be very slow.


**Figure 50. Input Position Request**

At 0 s, a position of 3 is requested and then at 0.5 s a position of 0 is requested. Also shows the actual position and motor feedback currents



**Figure 51. Output Response for Speed and Torque**

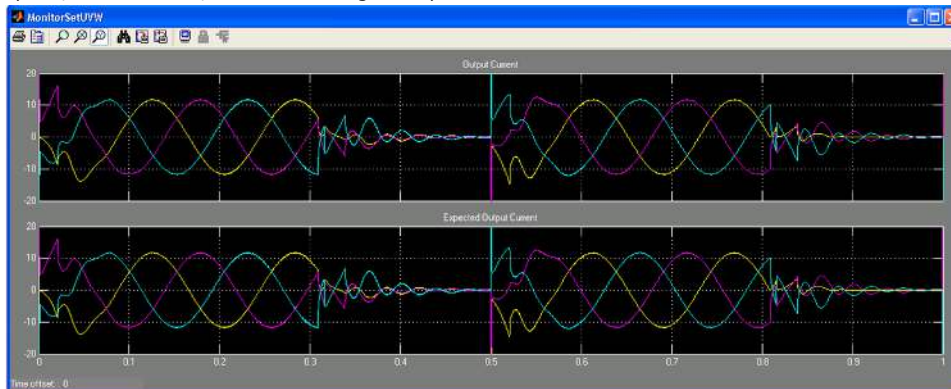
The maximum speed request saturates at 10 and the torque request saturates at 5 as set by parameters of the model. Also, some oscillation exists on the speed and torque requests because of nonoptimal settings for the PI controller causing an under-damped response.





**Figure 52. Output Current**

From 0 to 0.1, the motor is accelerating; 0.1 to 0.3, it is at a constant speed; 0.3 to 0.5, it is decelerating to stop. From 0.5 to 0.6, the motor accelerates in the opposite direction; from 0.6 to 0.8, it is at a constant speed; from 0.8 to 1, it is decelerating to stop.



### 6.6.2. Position, Speed, and Current Control for AC Motors (with ALU Folding)

The position, speed, and current control for AC motors (with ALU folding) design example is a FOC algorithm for AC motors, which is identical to the position, speed, and current control for AC motors design example. However this design example uses ALU folding.

The model file is `psc_ctrl_alu.mdl`.

The design example targets a Cyclone V device (speed grade 8). Cyclone V devices have distributed memory (MLABs). ALU folding uses many distributed memory components. ALU folding performs better in devices that have distributed memories, rather than devices with larger block memories.

The design example includes a setup script  
`setup_position_speed_current_controller_alu.m`.

**Table 17. Setup Script Variables**

Variables	Description
<code>dspb_psc_ctrl.SampleRateHz = 10000</code>	Sample rate. Default set to 10000, which is 10 kHz sample rate.
<code>dspb_psc_ctrl.ClockRate = 100</code>	FPGA clock frequency. Default set to 100, which is 100 MHz clock
<code>dspb_psc_ctrl.LatencyConstraint = 1000</code>	Maximum latency. Default 1,000 clock cycles

This design example uses a significantly large maximum latency, so resource consumption is the factor to optimize in ALU folding rather than latency.

Generally, industrial designs require a testbench that operates at the real-world sample rate. This example emulates the behavior of a motor sending current, position, and speed samples at a rate of 10 kHz.

When you run this design example without folding, the DSP Builder system operates at the same 10 kHz sample rate. Therefore, the system calculates a new packet of data for every Simulink sample. Also, the sample times of the testbench are the same as the sample times for the DSP Builder system.

The **Rate Transition** blocks translate between the Simulink testbench and the DSP Builder system. These blocks allow Simulink to manage the different sample times that the DSP Builder system requires. You need not modify the design example when you run designs with or without folding.

The **Rate Transition** blocks produce Simulink samples with a sample time of **dspb\_psc\_ctrl.SampleTime** for the testbench and **dspb\_psc\_ctrl.DSPBASampleTime** for the DSP Builder system. The samples are in the stimuli system, within the dummy motor. To hold the data consistent at the inputs to the **Rate Transition** blocks for the entire length of the output sample (**dspb\_psc\_ctrl.SampleTime**), turn on **Register Outputs**.

The data valid signal consists of a one Simulink sample pulse that signifies the beginning of a data packet followed by zero values until the next data sample, as required by ALU folding. The design example sets the period of this pulsing data valid signal to the number of Simulink samples for the DSP Builder system (at **dspb\_psc\_ctrl.DSPBASampleTime**) between data packets. This value is **dspb\_psc\_ctrl.SampleTime/dspb\_psc\_ctrl.DSPBASampleTime**.

The verification script within ALU folding uses the **To Workspace** blocks. The verification script searches for **To Workspace** blocks on the output of systems to fold. The script uses these blocks to record the outputs from both the design example with and without folding. The script compares the results with respect to valid outputs. To run the verification script, enter the following command at the MATLAB prompt:

```
Folder.Testing.RunTest('psc_ctrl_alu');
```

### 6.6.3. About FOC

FOC involves controlling the motor's sinusoidal 3-phase currents in real time, to create a smoothly rotating magnetic flux pattern, where the frequency of rotation corresponds to the frequency of the sine waves. FOC controls the amplitude of the current vector that is at 90 degrees with respect to the rotor magnet flux axis (quadrature current) to control torque.

The direct current component (0 degrees) is set to zero. The algorithm involves the following steps:

- Converting the 3-phase feedback current inputs and the rotor position from the encoder into quadrature and direct current components with the Clarke and Park transforms.
- Using these current components as the inputs to two proportional and integral (PI) controllers running in parallel to control the direct current to zero and the quadrature current to the desired torque.
- Converting the direct and quadrature current outputs from the PI controllers back to 3-phase currents with inverse Clarke and Park transforms.

### 6.6.4. Folded FIR Filter

This design example implements a simple non-symmetric FIR filter using primitive blocks, with a data sample rate much less than the system clock rate. This design example uses ALU folding to minimize hardware resource utilization.

The model file is **demo\_alu\_fir.mdl**.

## 6.7. DSP Builder Floating Point Design Examples

1. [Black-Scholes Floating Point](#) on page 115
2. [Double-Precision Real Floating-Point Matrix Multiply](#) on page 115
3. [Fine Doppler Estimator](#) on page 115
4. [Floating-Point Mandelbrot Set](#) on page 116
5. [General Real Matrix Multiply One Cycle Per Output](#) on page 117
6. [Newton Root Finding Tutorial Step 1—Iteration](#) on page 117
7. [Newton Root Finding Tutorial Step 2—Convergence](#) on page 118
8. [Newton Root Finding Tutorial Step 3—Valid](#) on page 118
9. [Newton Root Finding Tutorial Step 4—Control](#) on page 118
10. [Newton Root Finding Tutorial Step 5—Final](#) on page 118
11. [Normalizer](#) on page 118
12. [Single-Precision Complex Floating-Point Matrix Multiply](#) on page 118
13. [Single-Precision Real Floating-Point Matrix Multiply](#) on page 119
14. [Simple Nonadaptive 2D Beamformer](#) on page 119

### 6.7.1. Black-Scholes Floating Point

The DSP Builder Black-Scholes single- and double-precision floating-point design examples implement the calculation of a Black-Scholes equation and demonstrate the load exponent, reciprocal square root, logarithm and divide floating-point **Primitive** library blocks for single- or double-precision floating-point designs.

The model files are `blackScholes_S.mdl` and `blackScholes_D.mdl`.

### 6.7.2. Double-Precision Real Floating-Point Matrix Multiply

A simpler design example of a floating-point matrix multiply implementation than the complex multiply example. Each vector multiply is performed simultaneously, using many more multiply-adds in parallel.

The model file is `matmul_flash_RD.mdl`.

### 6.7.3. Fine Doppler Estimator

The fine Doppler estimator design example is an interpolator for radar applications. The example has three complex input values. It calculates the magnitude of each value, then performs a parabolic curve fit, identifies the location of the peak, and calculates the peak magnitude. The example performs all processing in single-precision floating-point data.

For more information about fine Doppler estimators, refer to *Fundamentals of Radar Signal Processing* by Mark A. Richards, McGraw-Hill, ISBN 0-07-144474-2, ch. 5.3.4.

The model file is `FineDopplerEstimator.mdl`.

### 6.7.4. Floating-Point Mandelbrot Set

This design example plots the Mandelbrot set for a defined region of the complex plane, shows many advanced blockset features, and highlights recommended design styles.

A complex number  $C$  is in the Mandelbrot set if for the following equation the value remains finite when repeatedly iterated:

$$z_{(n+1)} = z_n^2 + C$$

where  $n$  is the iteration number and  $C$  is the complex conjugate

The system takes longer to perform floating-point calculations than for the corresponding fixed-point calculations. You cannot wait around for partial results to be ready, if you want to achieve maximum efficiency. Instead, you must ensure your algorithm fully uses the floating-point calculation engines. The design contains two floating-point math subsystems: one for scaling and offsetting pixel indices to give a point in the complex plane; the other to perform the main square-and-add iteration operation.

For this design example, the total latency is approximately 19 clock cycles, depending on target device and clock speed. The latency is not excessive; but long enough that it is inefficient to wait for partial results.

FIFO buffers control the circulation of data through the iterative process. The FIFO buffers ensure that if a partial result is available for a further iteration in the  $z_{(n+1)} = z_n^2 + C$  progression, the design works on that point.

Otherwise, the design starts a new point (new value of  $C$ ). Thus, the design maintains a full flow of data through the floating-point arithmetic. This main iteration loop can exert back pressure on the new point calculation engine. If the design does not read new points off the command queue FIFO buffers quickly enough, such that they fill up, the loop iteration stalls. The design does not explicitly signal the calculation of each point when it is required (and thus avoid waiting through the latency cycles before you can use it). The design does not attempt to exactly calculate this latency in clock cycles. The design tries to issue generate point commands the exact number of clock-cycles before you need them. You must change them each time you retarget a device, or change target clock rate. Instead, the design calculates the points quickly from the start and catches them in a FIFO buffer. If the FIFO buffer starts to get full—a sufficient number of cycles ahead of full—The design stops the calculation upstream without loss of data. This selfregulating flow mitigates latency while remaining flexible.

Avoid inefficiencies by designing the algorithm implementation around the latency and availability of partial results. Data dependencies in processing can stall processing.

The design example uses the **FinishedThisPoint** signal as the valid signal. Although the system constantly produces data on the output, it marks the data as valid only when the design finishes a point. Downstream components can then just process valid data, just as the enabled subsystem in the testbench captures and plot the valid points.

In both feedback loops, you must provide sufficient delay for the scheduler to redistribute as pipelining. In feed-forward paths you can add pipelining without changing the algorithm—DSP Builder changes only the timing of the algorithm. But in feedback loops, inserting a delay can alter the meaning of an algorithm. For example,

adding N cycles of delay to an accumulator loop increments N different numbers, each incrementing every N clock cycles. The design must provide enough slack in each loop for the scheduler, which redistributes delays and pipelines operators, to be able to close timing by redistributing this slack. The scheduler must not change the total latency around the loop. The scheduler must ensure the function of the algorithm is unaltered. It must not change the total latency around the loop. It must ensure the function of the algorithm is unaltered. Such slack delays are in the top-level design of the synthesizable design in the feedback loop controlling the generation of new points, and in the **FeedBackFIFO** subsystem controlling the main iteration calculation. DSP Builder uses the minimum delay feature on the **SampleDelay** blocks to set these slack delays to the minimum possible delay that satisfies the scheduling solver. The example sets the **SampleDelay** block to the minimum latency that satisfies the schedule, which the DSP Builder solves as part of the integer linear programming problem that finds an optimum pipelining and scheduling solution. You can group delays into numbered equivalence groups to match other delays. In this design example, the single delay around the coordinate generation loop is in one equivalence group, and all the slack delays around the main calculation loop are in another equivalence group. The equivalence group field can contain any MATLAB expression that evaluates to a string. The **SampleDelay** block displays the delay that DSP Builder uses.

The FIFO buffers operate in show-ahead mode—they display the next value to be read. The read signal is a read acknowledgement, which reads the output value, discards it, and shows the next value. The design uses multiple FIFO buffers with the same control signal, which are full and give a valid output at the same time. The design only needs the output control signals from one of the FIFO buffers and can ignore the corresponding signals from the other FIFO buffers. As floating-point simulation is not bit accurate to the hardware, some points in the complex plane take fewer or more iterations to complete in hardware compared to the Simulink simulation. The results, when you are finished with a particular point, may come out in a different order. You must build a testbench mechanism that is robust to this feature. Use the testbench override feature in the **Run All Testbenches** block:

- Set the condition on mismatches to **Warning**
- Use the **Run All Testbenches** block to set an import variable, which brings the ModelSim results back into MATLAB and a custom verification function that sets the pass or fail criteria.

The model file is `Mandelbrot_S.mdl`.

### 6.7.5. General Real Matrix Multiply One Cycle Per Output

This design example implements a floating-point matrix multiply. The design performs each vector multiply simultaneously, using many multiply-adds in parallel.

The model file is `gemm_flash.mdl`.

### 6.7.6. Newton Root Finding Tutorial Step 1—Iteration

This design example is part of the Newton-Raphson tutorial. It demonstrates a naive test for convergence and exposes problems with rounding and testing equality with zero.

The model file is `demo_newton_iteration.mdl`.

### 6.7.7. Newton Root Finding Tutorial Step 2—Convergence

This design example is part of the Newton-Raphson tutorial. It demonstrates convergence criteria exposing mismatches between Simulink and ModelSim that you can correct by bit-accurate simulation. The discrepancies are worse when you use faithful rounding.

The model file is `demo_newton_convergence.mdl`.

### 6.7.8. Newton Root Finding Tutorial Step 3—Valid

This design example is part of the Newton-Raphson tutorial. It demonstrates how you avoid having the same answer multiple times on the output. It introduces a valid control signal, parallel to the datapath, to keep track of which pipeline slots the design empties. It uses equivalence groups in the minimum **SampleDelay** blocks.

The model file is `demo_newton_valid.mdl`.

### 6.7.9. Newton Root Finding Tutorial Step 4—Control

This design example is part of the Newton-Raphson tutorial. It demonstrates flow control which allows the design to buffer inputs in a FIFO buffer and insert data into pipeline slots as they become available.

The model file is `demo_newton_control.mdl`.

### 6.7.10. Newton Root Finding Tutorial Step 5—Final

This design example is part of the Newton-Raphson tutorial. It demonstrates a parallel integer datapath for counting iterations. It detects divergence in cases where the Newton method oscillates between two finite values.

The model file is `demo_newton_final.mdl`.

### 6.7.11. Normalizer

The normalizer design example demonstrates the **ilogb** block and the multifunction **ldexp** block. The parameters allow you to select the **ilogb** or **ldexp**. The design example implements a simple floating-point normalization. The magnitude of the output is always in the range 0.5 to 1.0, irrespective of the (non-zero) input.

The model file is `demo_normalizer.mdl`.

### 6.7.12. Single-Precision Complex Floating-Point Matrix Multiply

This design example uses a similar flow control style to that in the floating-point Mandelbrot set design example. The design example uses a limited number of multiply-adds, set by the vector size, to perform a complex single precision matrix multiply.

A matrix multiplication must multiply row and column dot product for each output element. For  $8 \times 8$  matrices A and B:

### Equation 1. Matrix Multiply Equation

$$AB_{ij} = \sum_{k=1}^8 A_{ik}B_{kj}$$

You may accumulate the adjacent partial results, or build adder trees, without considering any latency. However, to implement with a smaller dot product, consider resource usage folding, which uses a smaller number of multipliers rather than performing everything in parallel. Also split up the loop over k into smaller chunks. Then reorder the calculations to avoid adjacent accumulations.

A traditional implementation of a matrix multiply design is structured around a delay line and an adder tree:

$A_{11}B_{11} + A_{12}B_{21} + A_{13}B_{31}$  and so on.

The traditional implementation has the following features:

- The length and size grow with folding size (typically 8 to 12)
- Uses adder trees of 7 to 10 adders that are only used once every 10 cycles.
- Each matrix size needs different length, so you must provide for the worst case

A better implementation is to use FIFO buffers to provide self-timed control. New data is accumulated when both FIFO buffers have data. This implementation has the following advantages:

- Runs as fast as possible
- Is not sensitive to latency of dot product on devices or  $f_{MAX}$
- Is not sensitive to matrix size (hardware just stalls for small N)
- Can be responsive to back pressure, which stops FIFO buffers emptying and full feedback to control

The model file is `matmul_CS.mdl`.

### 6.7.13. Single-Precision Real Floating-Point Matrix Multiply

This design example is a simpler design example of a floating-point matrix multiply implementation than the complex multiply example. The design example uses many more multiply-adds in parallel (128 single precision multiply adds in the default parameterization), to perform each vector multiply simultaneously.

The model file is `matmul_flash_RS.mdl`.

### 6.7.14. Simple Nonadaptive 2D Beamformer

This design example demonstrates a simple nonadaptive 2D beamformer using vectors and single precision arithmetic. The parameters are the number of beams, angle, focus and intensity of each beam.

A beamformer is a key algorithm in radar and wireless and is a signal processing technique that sensor arrays use for directional signal transmission or reception. In transmission, a beamformer controls the phase and amplitude of the individual array elements to create constructive or destructive interference in the wavefront. In



reception, information from different elements are combined such that the expected pattern of radiation is preferentially observed. A number of different algorithms exist. An efficient scheme combines multiple paths constructively.

The simulation calculates the phases in MATLAB code (as a reference), simulates the beamformer 2D design to calculate the phases in DSP Builder Advanced Blockset, compares the reference to the simulation results and plots the beam pattern.

The design example uses vectors of single precision floating-point numbers, with state-machine control from two for loops.

The model file is `beamform_2d.mdl`.

## 6.8. DSP Builder Flow Control Design Examples

1. [Avalon-ST Interface \(Input and Output FIFO Buffer\) with Backpressure](#) on page 120
2. [Avalon-ST Interface \(Output FIFO Buffer\) with Backpressure](#) on page 120
3. [Kronecker Tensor Product](#) on page 121
4. [Parallel Loops](#) on page 121
5. [Primitive FIR with Back Pressure](#) on page 121
6. [Primitive FIR with Forward Pressure](#) on page 122
7. [Primitive Systolic FIR with Forward Flow Control](#) on page 123
8. [Rectangular Nested Loop](#) on page 123
9. [Sequential Loops](#) on page 124
10. [Triangular Nested Loop](#) on page 124

### 6.8.1. Avalon-ST Interface (Input and Output FIFO Buffer) with Backpressure

This example demonstrates the Avalon-ST input interface with FIFO buffers and the AvalonST output interface blocks. This example has FIFO buffers in the input and output interfaces. Use the manual switches in the testbench to change when downstream is ready for data or to turn off input. The simulation ends by turning off incoming data and ensures that it writes out as many valid data cycles as it receives.

The model file is `demo_avalon_st_input_fifo.mdl`.

### 6.8.2. Avalon-ST Interface (Output FIFO Buffer) with Backpressure

This example demonstrates the Avalon-ST input interface and the Avalon-ST output interface blocks. This example has FIFO buffers in the output interface only. Use manual switches in the testbench to change when downstream is ready for data or to turn off input. The simulation ends by turning off incoming data and ensures that it writes out as many valid data cycles as it receives.

The model file is `demo_avalon_st.mdl`.



### 6.8.3. Kronecker Tensor Product

This design example generates a Kronecker tensor product. The design example shows how to use the **Loop** block to generate datapaths that operate on regular data.

The top-level testbench includes **Control**, **Signals**, **Run ModelSim**, and **Run Quartus Prime** blocks.

The **Chip** subsystem includes the **Device** block and a lower-level **KroneckerSubsystem** subsystem.

The **KroneckerSubsystem** subsystem includes **ChannelIn**, **ChannelOut**, **Loop**, **Const**, **DualMem**, **Mult**, and **SynthesisInfo** blocks.

In this design example, the top level of the FPGA device (marked by the **Device** block) and the synthesizable **KroneckerSubsystem** subsystem (marked by the **SynthesisInfo** block) are at different hierarchy levels.

The model file is `demo_kronecker.mdl`.

### 6.8.4. Parallel Loops

This design example has two inner loops nested within the outer loop. The inner loops execute in parallel rather than sequentially. The two inner loops are started simultaneously by duplicating the control token but finish at different times. The **Rendezvous** block waits until both of them finish and then passes the control token back to the outer loop.

The model file is `forloop_parloop.mdl`.

### 6.8.5. Primitive FIR with Back Pressure

This DSP Builder design example uses **Primitive** library blocks to implement a FIR design with flow control and back pressure. The design example shows how you use the Primitive **FIFO** block to implement back pressure and flow control.

The top-level testbench includes **Control** and **Signals** blocks.

The **FirChip** subsystem includes the **Device** block and a lower-level primitive FIR subsystem.

The primitive FIR subsystem includes **ChannelIn**, **ChannelOut**, **FIFO**, **Not**, **And**, **Mux**, **SampleDelay**, **Const**, **Mult**, **Add**, and **SynthesisInfo** blocks.

In this design example, the top level of the FPGA device (marked by the **Device** block) and the synthesizable Primitive FIR subsystem (marked by the **SynthesisInfo** block) are at different hierarchy levels.

The model file is `demo_back_pressure.mdl`.

This design example shows how back pressure from a downstream block can halt upstream processing. This design example provides three FIR filters. A FIFO buffer follows each FIR filter that can buffer any data that is flowing through the FIFO buffer. If the FIFO buffer becomes half full, the design asserts the ready signal back to the upstream block. This signal prevents any new input (as flagged by valid) entering the FIR block. The FIFO buffers always show the next data if it is available and the valid

signal is asserted high. You must **AND** this FIFO valid signal with the ready signal to consume the data at the head of the FIFO buffer. If the **AND** result is high, you can consume data because it is available and you are ready for it.

You can chain several blocks together in this way, and no ready signal has to feed back further than one block, which allows you to use modular design techniques with local control.

The delay in the feedback loop represents the lumped delay that spreads throughout the FIR filter block. The delay must be at least as big as the delay through the FIR filter. This delay is not critical. Experiment with some values to find the right one. The FIFO buffer must be able to hold at least this much data after it asserts full. The full threshold must be at least this delay amount below the size of the FIFO buffer (64 – 32 in this design example).

The final block uses an external ready signal that comes from a downstream block in the system.

### 6.8.6. Primitive FIR with Forward Pressure

This DSP Builder design example uses **Primitive** library blocks to implement a FIR design with forward flow control. The design example shows how you can add a simple forward flow control scheme to a FIR design so that it can handle invalid source data correctly.

The top-level testbench includes **Control** and **Signals** blocks.

The **FirChip** subsystem includes the **Device** block and a lower-level Primitive FIR subsystem.

The primitive FIR subsystem includes **ChannelIn**, **ChannelOut**, **Mux**, **SampleDelay**, **Const**, **Mult**, **Add**, and **SynthesisInfo** blocks.

In this design example, the top level of the FPGA device (marked by the **Device** block) and the synthesizable primitive FIR subsystem (marked by the **SynthesisInfo** block) are at different hierarchy levels.

The model file is `demo_forward_pressure.mdl`.

The design example has a sequence of three FIR filters that stall when the valid signal is low, preventing invalid data polluting the datapath. The design example has a regular filter structure, but with a delay line implemented in single-cycle latches—effectively an enabled delay line.

You need not enable everything in the filter (multipliers, adders, and so on), just the blocks with state (the registers). Then observe the output valid signal, which DSP Builder pipelines with the logic, and observe the valid output data only.

You can also use vectors to implement the constant multipliers and adder tree, which also speeds up simulation.

You can improve the design example further by using the **TappedDelayLine** block.

### 6.8.7. Primitive Systolic FIR with Forward Flow Control

This DSP Builder design example uses **Primitive** library blocks to implement a systolic FIR design with forward flow control. The design example shows how you can add a simple forward flow control scheme to a FIR design so that it can handle invalid source data correctly.

The top-level testbench includes **Control** and **Signals** blocks.

The **FirChip** subsystem includes the **Device** block and a lower-level Primitive FIR subsystem.

The Primitive FIR subsystem includes **ChannelIn**, **ChannelOut**, **Mux**, **SampleDelay**, **Const**, **Mult**, **Add**, and **SynthesisInfo** blocks.

In this design example, the top level of the FPGA device (marked by the **Device** block) and the synthesizable primitive FIR subsystem (marked by the **SynthesisInfo** block) are at different hierarchy levels.

The design example has a sequence of three FIR filters that stall when the valid signal is low, preventing invalid data polluting the datapath. The design example has a regular filter structure, but with a delay line implemented in single-cycle latches—effectively an enabled delay line.

You need not enable everything in the filter (multipliers, adders, and so on), just the blocks with state (the registers). Then observe the output valid signal, which DSP Builder pipelines with the logic, and observe the valid output data only.

You can also use vectors to implement the constant multipliers and adder tree, which also speeds up simulation. You can improve the design example further with the **TappedDelayLine** block.

The model file is `demo_forward_pressure.mdl`.

### 6.8.8. Rectangular Nested Loop

In this design example all initialization, step, and limit values are constant. At the corners (at the end of loops) there may be cycles where the count value goes out of range, then the output valid signal from the loop is low.

The token-passing structure is typical for a nested-loop structure. The `bs` port of the innermost loop (**ForLoopB**) connects to the `bd` port of the same loop, so that the next loop iteration of this loop starts immediately after the previous iteration.

The `bs` port of the outer loop (**ForLoopA**) connects to the `ls` port of the inner loop; the `ld` port of the inner loop loops back to the `bd` port of the outer loop. Each iteration of the outer loop runs a full activation of the inner loop before continuing on to the next iteration.

The `ls` port of the outer loop connect to external logic and the `ld` port of the outer loop is unconnected, which is typical of applications where the control token is generated afresh for each activation of the outermost loop.

The model file is `forloop_rectangle.mdl`.

### 6.8.9. Sequential Loops

This design example nests two inner loops (**InnerLoopA** and **InnerLoopB**) within the outer loop. The design example daisy chains the `ld` port of **InnerLoopA** to the `ls` port of **InnerLoopB** rather than connecting it directly to the `bd` port of **OuterLoop**. Thus each activation of **InnerLoopA** is followed by an activation of **InnerLoopB**.

The model file is `forloop_seqloop.mdl`.

### 6.8.10. Triangular Nested Loop

The initialization, step, and limit values do not have to be constants. By using the count value from an outer loop as the limit of an inner loop, the counter effectively walks through a triangular set of indices.

The token-passing structure for this loop is identical to that for the rectangular loop, except for the parameterization of the loops.

The model file is `forloop_triangle.mdl`.

## 6.9. DSP Builder HDL Import Design Example

This digital up-converter resamples 20 MSPS complex base-band data to 80 MHz intermediate frequency, mixes it to center on +25 MHz, and applies some simple digital predistortion (DPD). This design example takes FIR and DPD VHDL components to create a complete up-conversion chain by importing existing IP and adding the up-conversion, mixer and pre-DPD scaling.

The digital upconverter includes: input memory, upconverter, FIR filter, scaler, mixer and digital predistortion (DPD).

**Table 18. Example Design Files**

<code>hdl_import_duc.mdl</code>	The DSP Builder design.
<code>hdl_import_duc_params.xml</code>	The design's parameter file.
<code>hdl_import_calc_fir_coefs.m</code>	A script to generate the FIR coefficients using MATLAB's <code>cfirpm</code> function. DSP Builder prints the coefficients to MATLAB's Command Window and you can copy and paste them into <code>coefficients.vhd</code> .
<code>calc_dpd_coefs.m</code>	A script to generate the DPD coefficients using a simple polynomial model of a power amplifier. DSP Builder prints the coefficients MATLAB's Command Window and you can copy and paste them into <code>lut_dpd.vhd</code> .
<code>to_import</code>	This directory contains 12 VHDL source files.

### VHDL Components

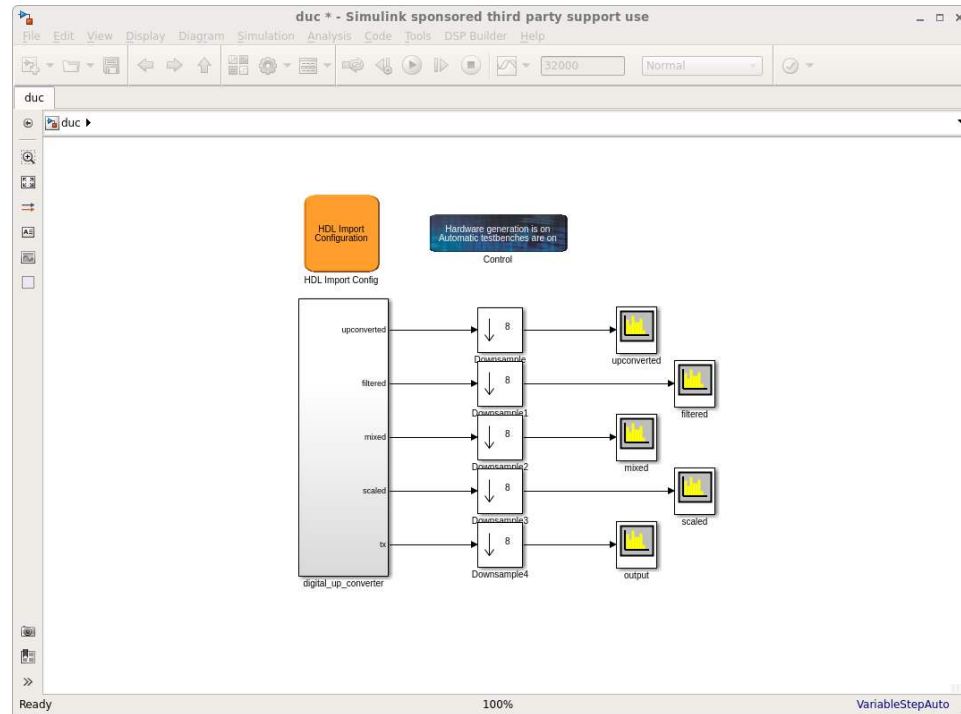
The design example includes a complex FIR filter in VHDL optimized for Intel Stratix 10 devices. This FIR filter has one valid data sample every eight clock cycles. See *Designing Filters for High Performance*.

The simple LUT-based DPD is initialized with a third-order polynomial.

## Top-Level Design

The top-level design contains the device-level subsystem and five downsample and spectrum analyzer blocks from MathWork's DSP System Toolbox. These blocks show the spectral output from the various stages of the up-conversion chain.

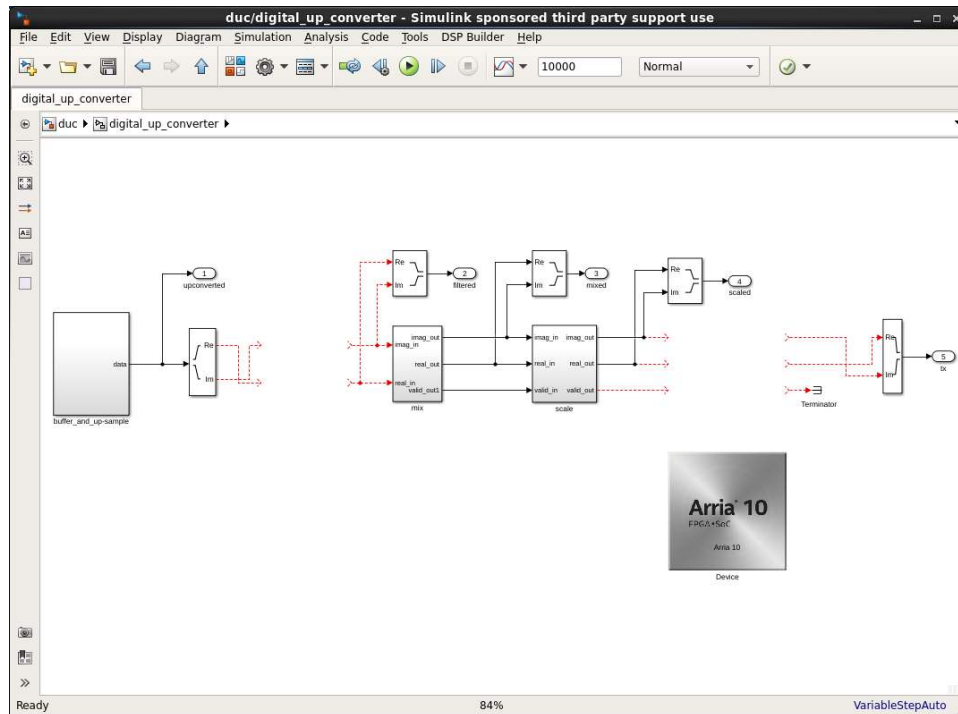
Figure 53. Top-Level Design



## Digital Up Converter

The **digital\_up\_converter** subsystem is the device-level subsystem. It contains all of the design's DSP Builder-based components and two gaps for **HDL Import** blocks.

Figure 54. Digital Upconverter

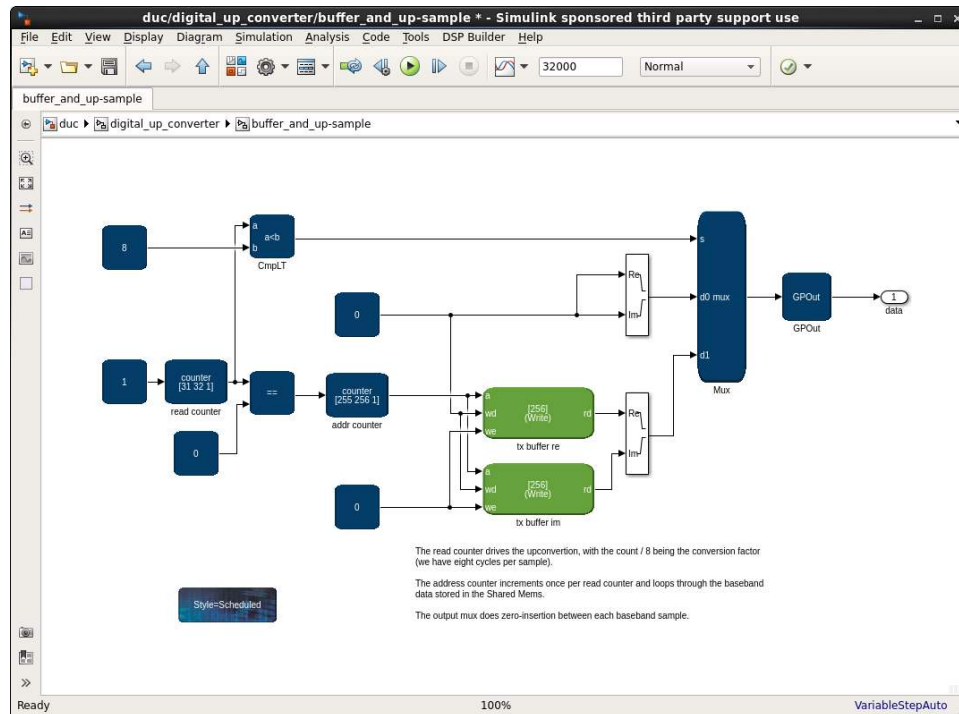


### Buffer and Upsample

This scheduled subsystem contains two **SharedMem** blocks, which contain the 20 MSPS baseband source: one for the real part of the signal and one for the imaginary part. You can write to the blocks via the bus or use the preloaded tones.

The **read\_counter** block drives the upconversion. It counts modulo 32 because it upsamples the 20 MSPS baseband by 4 to 80 MSPS and then holds each sample for 8 clock cycles at a clock rate of 640 MHz. The FIR filter accepts one sample every eight cycles. By holding the samples, the FIR does not need synchronization logic.

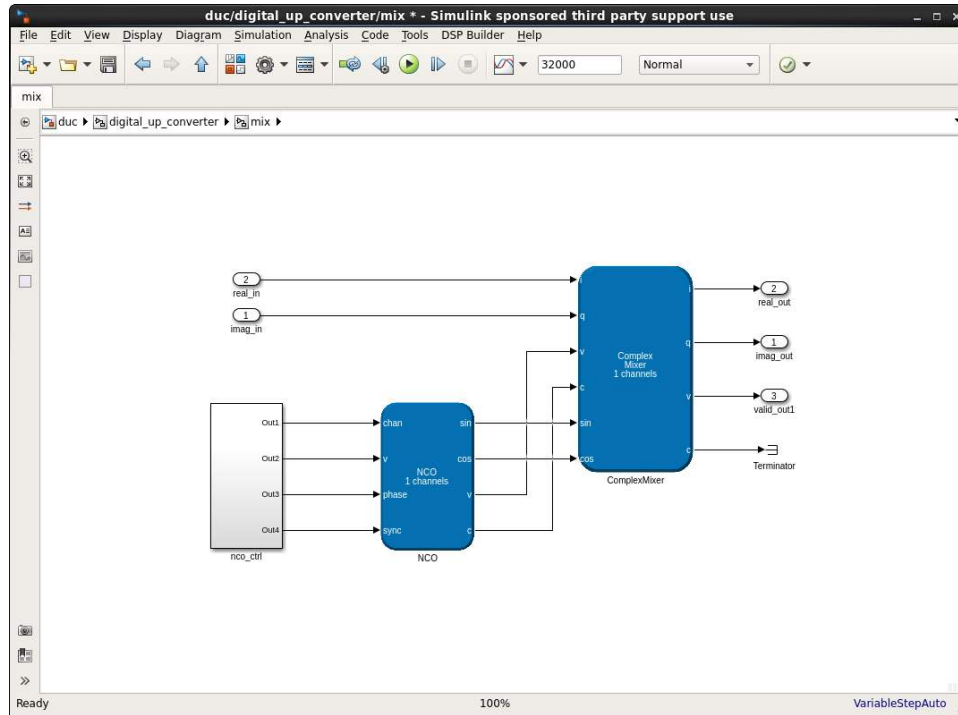
Figure 55. Buffer and Up-Sample



### Mixer

This single-channel mixer consists of **NCO** and **ComplexMixer** IP blocks and a scheduled subsystem for controlling the **NCO**. The control subsystem asserts the valid signal once every eight cycles. The **NCO** generates a 16 MHz complex tone, which the **ComplexMixer** uses to mix the filtered signal.

Figure 56. Mixer

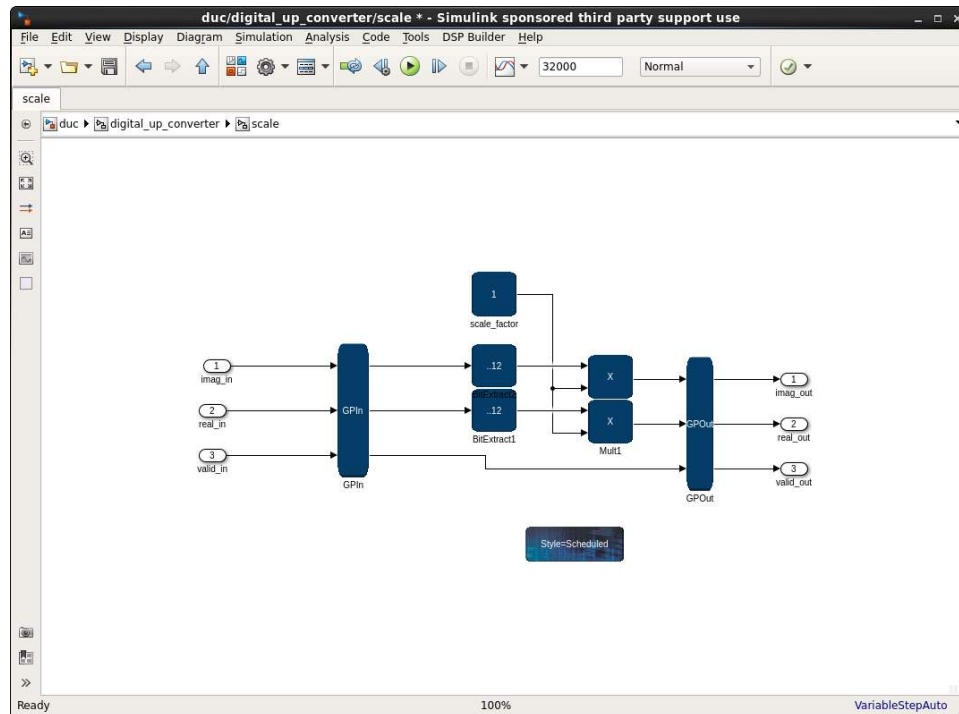


### Scale

The scale scheduled subsystem scales the data so that it fits within the DPD's range of operation by bit-shifting from the mixer's output. You can use the optional multiplier for increasing the signal level if bit-shifting is insufficient.



Figure 57. Scale



### FIR Coefficients

The FIR coefficients are defined in `coefficients.vhd`.

The coefficients are calculated in `hdl_import_calc_fir_coefs.m`. This script uses MATLAB's `cfirpm` command to create complex coefficients.

### DPD


The file `lut_dpd.vhd` contains the DPD for this design example. The DPD consists of an address generator that indexes a LUT. The output of the LUT is then multiplied with the complex input data. The LUT contents are calculated in `hdl_import_calc_dpd_coefs.m`. This script uses a simple, real-numbered, third-order model of an amplifier to calculate predistortion coefficients. DSP Builder uses these coefficients to calculate the LUT contents.

### Simulink Simulations Results

**Figure 58. Simulation**


The first four waveforms are the real and imaginary input and output of the the FIR. The FIR smooths the zero-padded signals.

The next four waveforms are the real and imaginary input and output of the the DPD.




**Figure 59. Upconverted**

The two preloaded memory signals are clearly visible about 0, as are their four aliases because of the zero-insert upsampling.




**Figure 60. Filtered**

The aliased signals are attenuated by 40dB, as expected from the analysis in `calc_fir_coefs.m`.



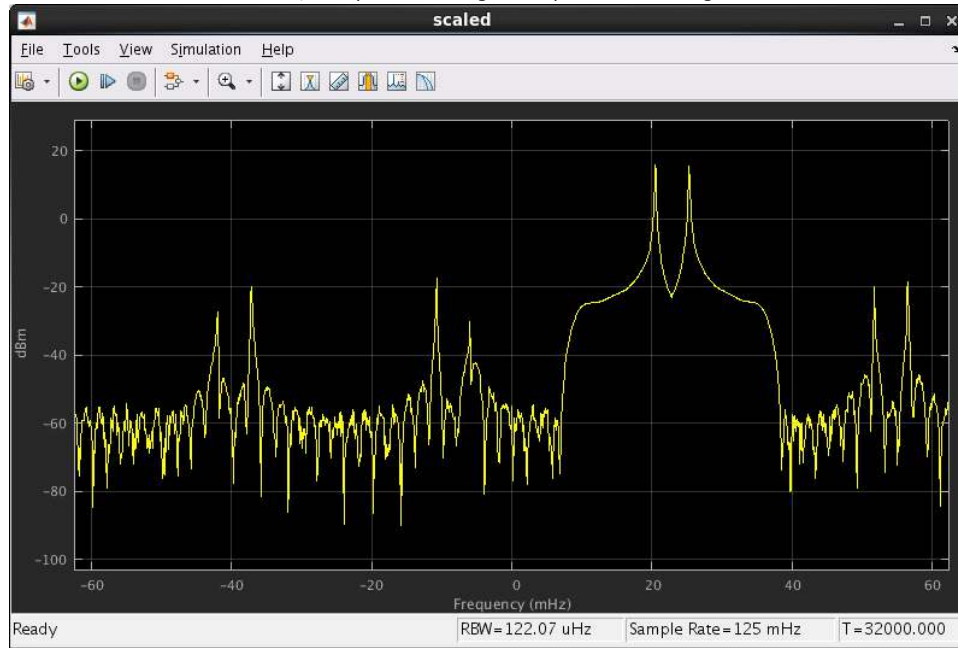
**Figure 61. Mixed**

The mixed spectrum shows the baseband signal moving over to be centered on 16 MHz. This view shows the Simulink clock rate of 1 Hz rather than the FPGA clock rate of 640 MHz, so 16 MHz becomes 25 mHz.



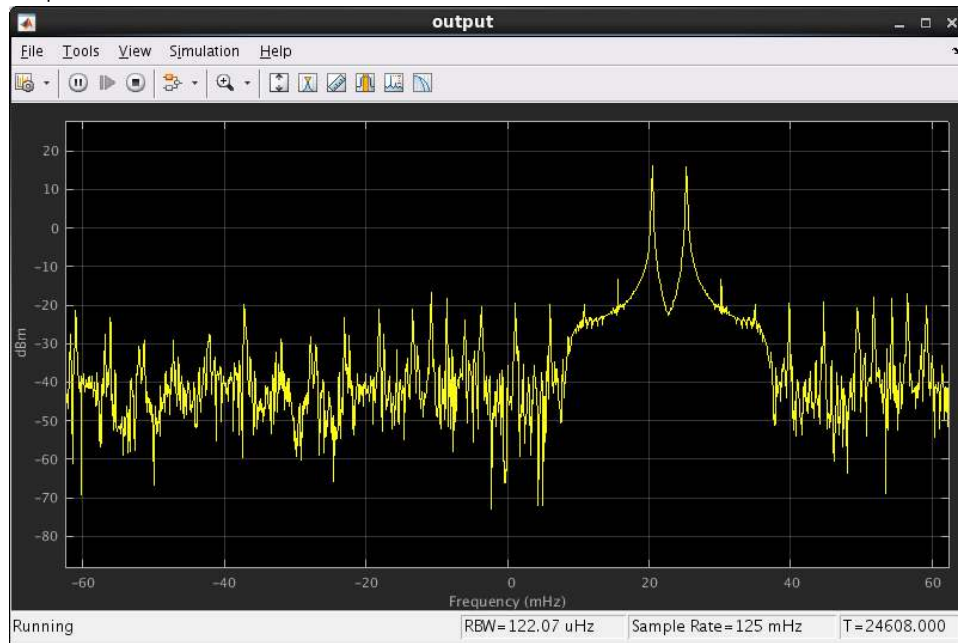
**Figure 62. Scaled**

Scaled looks identical to mixed, except that the signal amplitude is much greater.



**Figure 63. Output**

The post-DPD output signal is a noisier version of the scaled signal. Observe the two third-order harmonics in the pass-band.



**Related Information**

[Designing Filters for High Performance](#)

### 6.9.1. Performing a Cosimulation

This tutorial uses the DSP Builder HDL import design example.

The design example has two HDL entities: the DPD (`lut_dpd.vhd`) and the FIR (`complex_fir.vhd`).

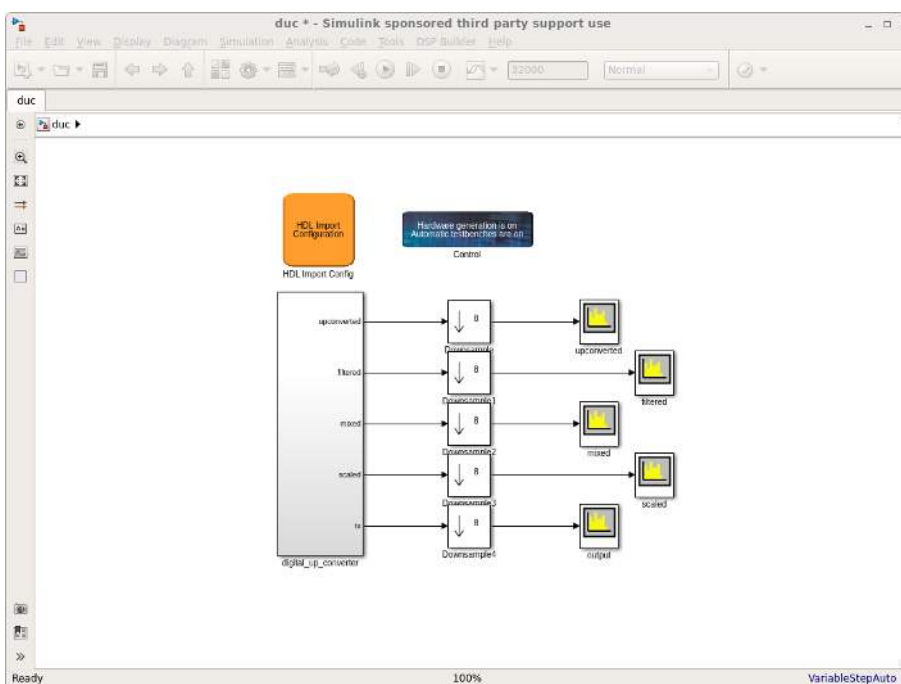
In DSP Builder cosimulation, each **HDL Import** block represents an HDL instance. You must instantiate both of these entities in a top-level VHDL file. For this design example, Intel provides `top.vhd`.

In addition, the FIR filter uses a signed data type with a generic for the data width. When DSP Builder instantiates the FIR filter, it uses its own paradigm (i.e. `std_logic_vector` and no generics). This design example adds a wrapper entity: `complex_fir_wrapper.vhd`. This entity instantiates `complex_fir`, including setting the generic to the appropriate value, and converts signed to `std_logic_vector`.

These two files, `top.vhd` and `complex_fir_wrapper.vhd` are in the `to_import` directory.

1. Add a HDL Import Config block to the top-level design.

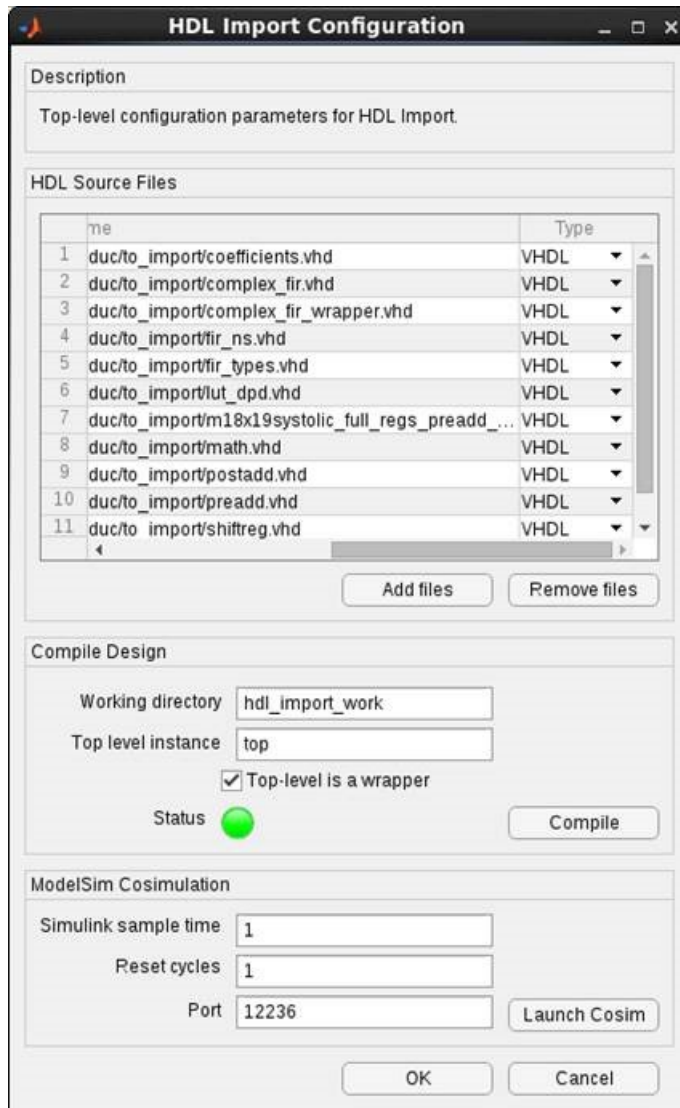
**Figure 64. Top-level Design with HDL Import Config Block**



2. Parameterize the **HDL Import Config** block.
  - a. Click **Add** to add all of the files from the `to_import` directory.  
The order of the files does not matter. DSP Builder determines the type of HDL file by the extension, but you can change the type manually.
  - b. Enter `top` in the **Top level instance**.
  - c. Turn on **Top-level is a wrapper**.

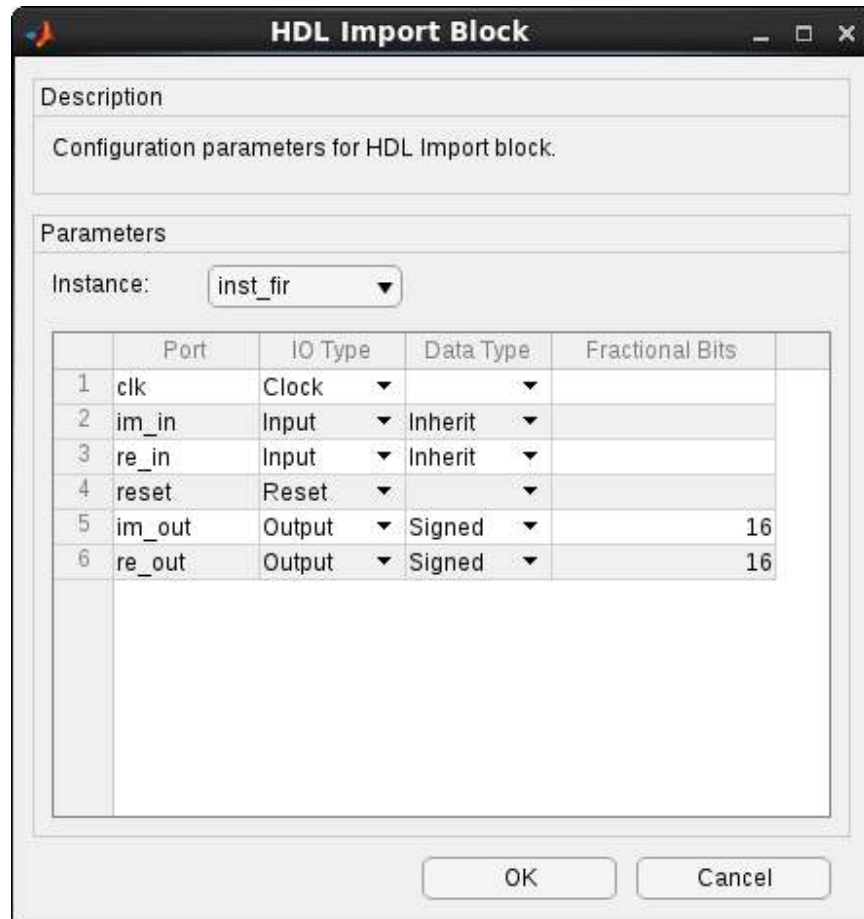
- d. Click the **Compile** button.
- e. Set the **Simulink sample time** field to 1.
- f. When the status light is green, click **Launch Cosim**.

Figure 65. HDL Import Configuration



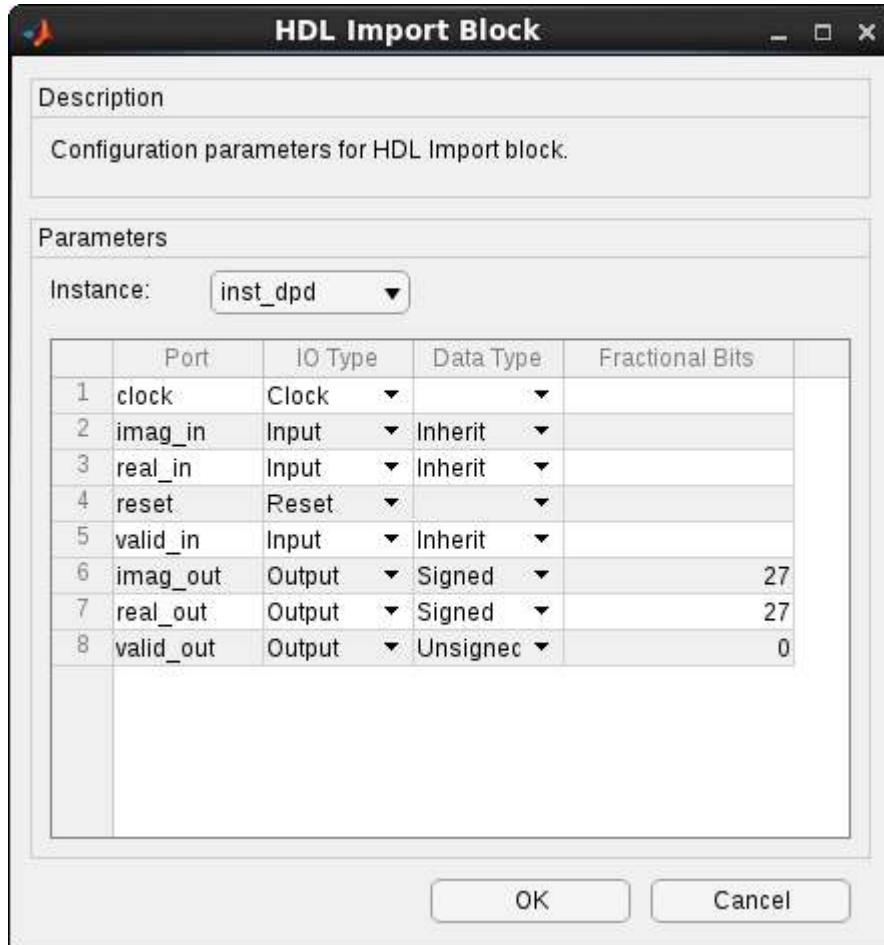
3. Add a **HDL Import** block to the digital\_up\_converter subsystem.
  - a. Double click the **HDL Import** block
  - b. Click **Instance** and select **inst\_fir**.
  - c. Set the fractional bits of the two output signals to 16.

Figure 66. HDL Import Block inst\_fir Parameters



4. Add a second **HDL Import** block to the digital\_up\_converter subsystem.
  - a. Double click the **HDL Import** block
  - b. Click **Instance** and select **inst\_dpd**.
  - c. Set the fractional bits of the two output signals to 27.
  - d. Set the valid output to unsigned.

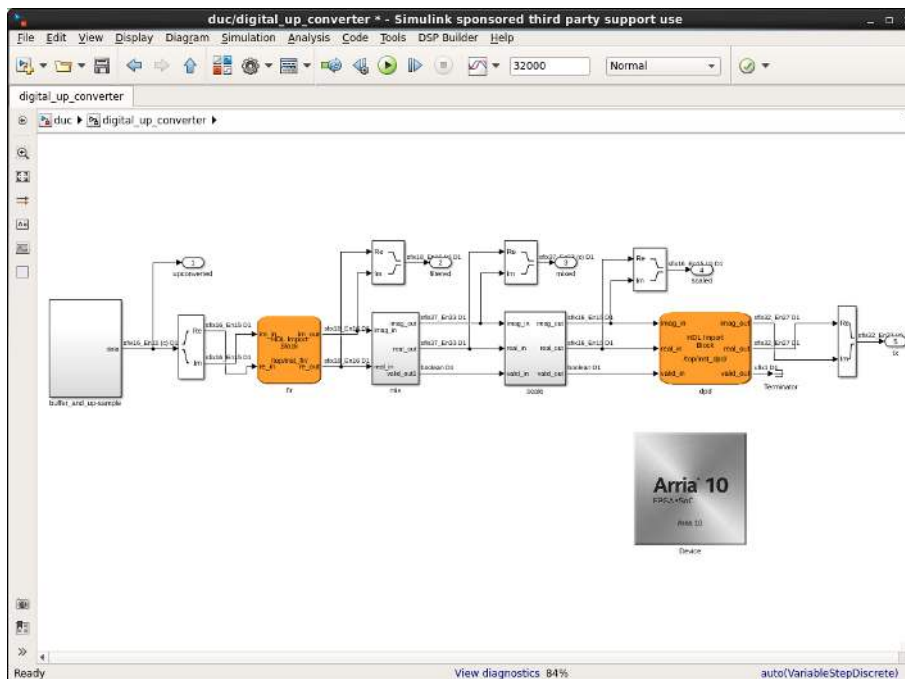
Figure 67. HDL Import inst\_dpd Parameters



- Wire up **HDL import** blocks.  
The HDL Import block port names are in alphabetical order.



Figure 68. Wire up HDL Import Blocks



6. Press the play button or advance through the simulation a cycle at a time.
7. Verify HDL import with the ModelSim simulator, in DSP Builder, select **DSP Builder > Run ModelSim > Device**.  
The cosimulation turns any non-high state (e.g. U or X) to a zero.
8. Compile the design in Intel Quartus Prime, by selecting **DSP Builder > Run Quartus Prime Software**.

## 6.10. DSP Builder Host Interface Design Examples

1. [Memory-Mapped Registers](#) on page 137

### 6.10.1. Memory-Mapped Registers

This design example is an extreme example of using the processor registers to implement a simple calculator. Registers and shared memories write arguments and read results.

The top-level testbench includes **Control**, **Signals**, **Run ModelSim**, and **Run Quartus Prime** blocks.

This design also includes **BusStimulus** and **BusStimulusFileReader** blocks.

The **RegChip** subsystem includes **RegField**, **RegBit**, **RegOut**, **SharedMem**, **Const**, **Add**, **Sub**, **Mult**, **Convert**, **Select**, **BitExtract**, **Shift**, and **SynthesisInfo** blocks.

The model file is `demo_regs.mdl`.

## 6.11. DSP Builder Platform Design Examples

This folder contains design examples that illustrate how you can implement a DDC or digital up converter (DUC) for use in a radio basestation. Use these designs as a starting point to build your own filter chain that meets your exact needs.

1. [16-Channel DDC](#) on page 138
2. [16-Channel DUC](#) on page 138
3. [2-Antenna DUC for WiMAX](#) on page 139
4. [2-Channel DUC](#) on page 140
5. [Super-Sample Rate Digital Upconverter](#) on page 140

### 6.11.1. 16-Channel DDC

This design example shows how to use using **IP** and **Interface** blocks to build a 16-channel digital-down converter for modern radio systems.

Decimating CIC and FIR filters down convert eight complex carriers (16 real channels) from 61.44 MHz. The total decimation rate is 64. A real mixer and NCO isolate the eight carriers. The testbench isolates two channels of data from the TDM signals using a channel viewer.

The top-level testbench includes **Control**, **Signals**, **Run ModelSim**, and **Run Quartus Prime** blocks, plus a **ChanView** block that deserializes the output bus. An **Edit Params** block allows easy access to the setup variables in the `setup_demo_ddc.m` script.

The **DDCChip** subsystem includes **Device**, **Decimating FIR**, **DecimatingCIC**, **Mixer**, **NCO**, **Scale**, **RegBit**, and **RegField** blocks.

The model file is `demo_ddc.mdl`.

*Note:* This design example uses the Simulink Signal Processing Blockset.

### 6.11.2. 16-Channel DUC

This design example shows how to build a 16-channel DUC as found in modern radio systems using **Interface**, **IP**, and **Primitive** blocks.

This design example shows an interpolating filter chain with interpolating CIC and FIR filters that up convert eight complex channels (16 real channels). The total interpolation rate is 50. DSP Builder integrates several **Primitive** subsystems into the datapath. This design example shows how you can integrate **IP** blocks with **Primitive** subsystems:

- The programmable **Gain** subsystem, at the start of the datapath, shows how you can use processor-visible register blocks to control a datapath element.
- The **Sync** subsystem is a **Primitive** subsystem that shows how to manage two data streams coming together and synchronizing. The design writes the data from the NCOs to a memory with the channel as an address. The data stream uses its channel signals to read out the NCO signals, which resynchronizes the data correctly. Alternatively, you can simply delay the NCO value by the correct number of cycles to ensure that the NCO and channel data arrive at the **Mixer** on the same cycle.

Extensive use is made of Simulink multiplexer and demultiplexer blocks to manage vector signals.

The top-level testbench includes **Control**, **Signals**, **Run ModelSim**, and **Run Quartus Prime** blocks, plus a **ChanView** block that deserializes the output bus. An **Edit Params** block allows easy access to the setup variables in the `setup_demo_duc.m` script.

The **DUCChip** subsystem includes a **Device** block and a lower level **DUC16** subsystem.

The **DUC16** subsystem includes **InterpolatingFIR**, **InterpolatingCIC**, **ComplexMixer**, **NCO**, and **Scale** blocks.

It also includes lower level **Gain**, **Sync**, and **CarrierSum** subsystems which make use of other **Interface** and Primitive blocks including **AddSLoad**, **And**, **BitExtract**, **ChannelIn**, **ChannelOut**, **CompareEquality**, **Const**, **SampleDelay**, **DualMem**, **Mult**, **Mux**, **Not**, **Or**, **RegBit**, **RegField** blocks, and **SynthesisInfo** blocks.

The model file is `demo_duc.mdl`.

*Note:* This design example uses the Simulink Signal Processing Blockset.

### 6.11.3. 2-Antenna DUC for WiMAX

This design example shows how to build a 2-antenna DUC to meet a WiMAX specification.

The top-level testbench includes **Control**, **Signals**, **Run ModelSim**, and **Run Quartus Prime** blocks, plus a **ChanView** block that deserializes the output bus.

The **DUCChip** subsystem includes a **Device** block and a lower level **DUC2Antenna** subsystem.

The **DUC2Antenna** subsystem includes **InterpolatingFIR**, **SingleRateFIR**, **Const**, **ComplexMixer**, **NCO**, and **Scale** blocks.

The model file is `demo_wimax_duc.mdl`.

*Note:* This design example uses the Simulink Signal Processing Blockset.

#### 6.11.4. 2-Channel DUC

This design example shows how to build a 2-channel DUC.

Interpolating CIC and FIR filters up convert a single complex channel (2 real channels). A NCO and **Mixer** subsystem combine the complex input channels into a single output channel.

This design example shows how quick and easy it is to emulate the contents of an existing datapath. A **Mixer** block implements the mixer in this design example as the data rate is low enough to save resource using a time-shared hardware technique.

The top-level testbench includes **Control**, **Signals**, **Run ModelSim**, and **Run Quartus Prime** blocks, plus a **ChanView** block that deserializes the output bus. An **Edit Params** block allows easy access to the setup variables in the `setup_demo_AD9856.m` script.

The **AD9856** subsystem includes a **Device** block and a lower level **DUCIQ** subsystem.

The **DUCIQ** subsystem includes **Const**, **InterpolatingFIR**, **SingleRateFIR**, **InterpolatingCIC**, **NCO**, **Scale** blocks, and a lower level **Mixer** subsystem.

The **Mixer** subsystem includes **ChannelIn**, **ChannelOut**, **Mult**, **Const**, **BitExtract**, **CompareEquality**, **And**, **Delay**, **Sub**, and **SynthesisInfo** blocks.

The model file is `demo_AD9856.mdl`.

*Note:* This design example uses the Simulink Signal Processing Blockset.

#### 6.11.5. Super-Sample Rate Digital Upconverter

The model file is `demo_ssduc.mdl`.

### 6.12. DSP Builder Primitive Block Design Examples

1. [8×8 Inverse Discrete Cosine Transform](#) on page 141
2. [Automatic Gain Control](#) on page 141
3. [Bit Combine for Boolean Vectors](#) on page 142
4. [Bit Extract for Boolean Vectors](#) on page 142
5. [Color Space Converter](#) on page 142
6. [CORDIC from Primitive Blocks](#) on page 143
7. [Digital Predistortion Forward Path](#) on page 143
8. [Fibonacci Series](#) on page 143
9. [Folded Vector Sort](#) on page 144
10. [Fractional Square Root Using CORDIC](#) on page 144
11. [Fixed-point Maths Functions](#) on page 144
12. [Gaussian Random Number Generator](#) on page 144

13. [Hello World](#) on page 145
14. [Hybrid Direct Form and Transpose Form FIR Filter](#) on page 145
15. [Loadable Counter](#) on page 145
16. [Matrix Initialization of LUT](#) on page 146
17. [Matrix Initialization of Vector Memories](#) on page 146
18. [Multichannel IIR Filter](#) on page 147
19. [Quadrature Amplitude Modulation](#) on page 147
20. [Reinterpret Cast for Bit Packing and Unpacking](#) on page 147
21. [Run-time Configurable Decimating and Interpolating Half-Rate FIR Filter](#) on page 148
22. [Square Root Using CORDIC](#) on page 148
23. [Test CORDIC Functions with the CORDIC Block](#) on page 148
24. [Uniform Random Number Generator](#) on page 148
25. [Vector Sort—Sequential](#) on page 149
26. [Vector Sort—Iterative](#) on page 149
27. [Vector Initialization of Sample Delay](#) on page 149
28. [Wide Single-Channel Accumulators](#) on page 150

### 6.12.1. 8×8 Inverse Discrete Cosine Transform

This design example uses the Chen-Wang algorithm to implement a fully pipelined 8×8 inverse discrete cosine transform (IDCT).

Separate subsystems perform the row transformation (**Row**), corner turner (**CornerTurn**), and column transformation (**Col**) functions. The design example synthesizes each separate subsystem separately. The **Row** and **Col** subsystems have additional levels of hierarchy for the different stages. The **SynthesisInfo** block is at the row or column level, so the design example flattens these subsystems before synthesis.

The **CornerTurn** turn block makes extensive use of Simulink **Goto/From** blocks to reduce the wiring complexity. The top-level testbench includes **Control** and **Signals** blocks. The **IDCTChip** subsystem includes the **Device** block and a lower level **IDCT** subsystem. The **IDCT** subsystem includes lower level subsystems that it describes with the **ChannelIn**, **ChannelOut**, **Const**, **BitCombine**, **Shift**, **Mult**, **Add**, **Sub**, **BitExtract**, **SampleDelay**, **OR** Gate, **Not**, **Sequence**, and **SynthesisInfo** blocks.

The model file is `demo_idct8x8.mdl`.

### 6.12.2. Automatic Gain Control

This design example implements an automatic gain control.

This design example shows a complex loop with several subloops that it schedules and pipelines without inserting registers. The design example spreads a lumped delay around the circuit to satisfy timing while maintaining correctness. Processor visible registers control the thresholds and gains.

In complex algorithmic circuits, the zero-latency blocks make it easy to follow a data value through the circuit and investigate the algorithm without offsetting all the results by the pipelining delays.

The top-level testbench includes **Control**, **Signals**, **Run ModelSim**, and **Run Quartus Prime** blocks.

The **AGC\_Chip** subsystem includes the **Device** block, a **RegField** block and a lower level **AGC** subsystem.

The **AGC** subsystem includes **RegField**, **ChannelIn**, **ChannelOut**, **Mult**, **SampleDelay**, **Add**, **Sub**, **Convert**, **Abs**, **CmpGE**, **Lut**, **Const**, **SharedMem**, **Shift**, **BitExtract**, **Select**, and **SynthesisInfo** blocks.

The model file is `demo_agc.mdl`.

### 6.12.3. Bit Combine for Boolean Vectors

This design example demonstrates different ways to use the **BitCombine** primitive block to create signals of different widths from a vector of Boolean signals.

The one input **BitCombine** block is a special case that concatenates all the components of the input vector and produces one wide scalar output signal. You can apply 1-bit reducing operators to vectors of Boolean signals. The **BitCombine** block supports multiple input concatenation. When vectors of Boolean signals are input on multiple ports, corresponding components from each vector are combined so that the output is a vector of signals.

The model file is `demo_bitcombine.mdl`.

### 6.12.4. Bit Extract for Boolean Vectors

This design example demonstrates different ways to use the **BitExtract** block to split a wide signal into a vector of narrow signal components.

This block converts a scalar signal into a vector of Boolean signals. You use the initialization parameter to arbitrarily order the components of the vector output by the **BitExtract** block. If the input to a **BitExtract** block is a vector, different bits can be extracted from each of the components. The output does not always have to be a vector of Boolean signals. You may split a 16-bit wide signal into four components each 4-bits wide.

The model file is `demo_bitextract.mdl`.

### 6.12.5. Color Space Converter

This design example demonstrates DSP Builder **Primitive** subsystems with simple RGB to Y'CbCr color space conversion

- $Y = 0.257R + 0.504G + 0.098B + 16$
- $Cb = -0.148R - 0.291G + 0.439B + 128$
- $Cr = 0.439R - 0.368G - 0.071B + 128$

The RGB data arrives as three parallel signals each clock cycle. The model file is `demo_csc.mdl`.

### 6.12.6. CORDIC from Primitive Blocks

This design example demonstrates building a CORDIC out of basic operators. This design has the same functionality as the CORDIC library block in the **demo\_cordic\_lib\_block** example

The model file is `demo_cordic_primitives.mdl`.

### 6.12.7. Digital Predistortion Forward Path

This design example demonstrates forward paths that implement digital predistortion (DPD).

Forward paths compensate for nonlinear power amplifiers by applying the inverse of the distortion that the power amplifier generates, such that the pre-distortion and the distortion of the power amplifier cancel each other out. The power amplifier's non-linearity may change over time, therefore such systems are typically adaptive.

This design example is based on "A robust digital baseband pre-distorter constructed using memory polynomials," L. Ding, G. T. Zhou, D. R. Morgan, et al., *IEEE Transactions on Communications*, vol. 52, no. 1, pp. 159-165, 2004.

This design example only implements the forward path, which is representative of many systems where you implement the forward path in FPGAs, and the feedback path on external processors. The design example sets the predistortion memory,  $Q$ , to 8; the highest nonlinearity order  $K$  is 5 in this design example. The file `setup_demo_dpd_fwdpath` initializes the complex valued coefficients, which are stored in registers. During operation, the external processor continuously improves and adapts these coefficients with a microcontroller interface.

The model file is `demo_dpd_fwdpath.mdl`.

### 6.12.8. Fibonacci Series

This DSP Builder design example generates a Fibonacci sequence.

This design example shows that even for circuitry with tight feedback loops and 120-bit adders, designs can achieve high data rates by the pipelining algorithms. The top-level testbench includes **Control**, **Signals**, **Run ModelSim**, and **Run Quartus Prime** blocks. The **Chip** subsystem includes the **Device** block and a lower level **FibSystem** subsystem. The **FibSystem** subsystem includes **ChannelIn**, **ChannelOut**, **SampleDelay**, **Add**, **Mux**, and **SynthesisInfo** blocks.

*Note:* In this design example, the top-level of the FPGA device (marked by the **Device** block) and the synthesizable **Primitive** subsystem (marked by the **SynthesisInfo** block) are at different hierarchy levels.

The model file is `demo_fibonacci.mdl`.

### 6.12.9. Folded Vector Sort

This design sorts the values on the input vector from largest to smallest. The design is a masked subsystem that allows for sorting with either a comparator and mux block, or a minimum and a maximum block. The first implementation is more efficient. Both use the reconfigurable subsystem to choose between implementations using the **BlockChoice** parameter.

Folded designs repeatedly use a single dual sort stage. The throughput of the design is limited in the number of channels, vector width, and data rate. The data passes through the dual sort stage (vector width)/2 times. The vector sort design example uses full throughput with (vector width)/2 dual sort stages in sequence.

Look under the mask to view the implementation of reconfigurable subsystem templates and the blocks that reorder and interleave vectors.

The model file is `demo_foldedsort.mdl`.

### 6.12.10. Fractional Square Root Using CORDIC

This design example demonstrates CORDIC techniques, but does not use the **CORDIC** block. This design example is fully iterative.

The design example allows you to generate a valid signal. The design example only generates output and can only accept input every  $N$  cycles, where  $N$  depends on the number of stages, the data output format, and the target  $f_{MAX}$ . The valid signal goes high when the output is ready. You can use this output signal to trigger the next input, for example, a FIFO buffer read for bursty data.

The model file is `demo_cordic_fracsqrt.mdl`.

### 6.12.11. Fixed-point Maths Functions

This design example demonstrates how the **Math**, **Trig** and **Sqrt** functions support fixed-point types and the fixed-point **Divide** function. You can use fixed-point types of width up to and including 32 bits.

DSP Builder generates results using the same techniques as in the floating point functions but at generally reduced resource usage, depending on data bit width. Outputs are faithfully rounded. If the exact result is between two representable numbers within the data format, DSP Builder uses either of them. In some instances you see a difference in output result between simulation and hardware by one LSB. To get bit-accurate results at the subsystem level, this example uses the **Bit Exact** option on the SynthesisInfo block.

The model file is `demo_fixed_math.mdl`.

### 6.12.12. Gaussian Random Number Generator

This DSP Builder design example demonstrates a random number generator (CLT component method) that produces random numbers with normal distribution and standard deviation that you specify using the input `sigma_input`.



You can also specify the seed value for the random sequence using the `seed_value` input. The reset input resets the sequence to the initial state defined by the `seed_value`. The output is a 32-bit single-precision floating-point number.

### 6.12.13. Hello World

This DSP Builder design example produces a simple text message that it stores in a look-up table.

An external input enables a counter that addresses a lookup-table (LUT) that contains some text. The design example writes the result to a MATLAB array. You can examine the contents with a `char(message)` command in the MATLAB command window.

This design example does not use any **ChannelIn**, **ChannelOut**, **GPIn**, or **GPOut** blocks. The design example uses Simulink ports for simplicity although they prevent the automatic testbench flow from working.

The top-level testbench includes **Control**, **Signals**, **Run ModelSim**, and **Run Quartus Prime** blocks.

The **Chip** subsystem includes **Device**, **Counter**, **Lut**, and **SynthesisInfo** blocks.

*Note:* In this design example, the top-level of the FPGA device (marked by the **Device** block) and the synthesizable **Primitive** subsystem (marked by the **SynthesisInfo** block) are at the same level.

The model file is `helloWorld.mdl`.

### 6.12.14. Hybrid Direct Form and Transpose Form FIR Filter

The design example uses small, four-tap direct form filters to use the structure inside the DSP block efficiently. The design example combines these direct form minifilters into a transpose structure, which minimizes the logic and memory that the sample pipe uses. This FIR filter shows a FIR architecture that is a hybrid between the direct form and transpose form FIR filter. It combines the advantages of both.

The model file is `demo_hybrid_fir_mc.mdl`.

### 6.12.15. Loadable Counter

This design example demonstrates the **LoadableCounter** block.

The testbench reloads the counter with new parameters every 64 cycles. A manual switch allows you to control whether the counter is permanently enabled, or only enabled on alternate cycles. You can view the signals input and output from the counter with the provided scope.

The model file is `demo_ld_counter.mdl`.

### 6.12.16. Matrix Initialization of LUT

This design example feeds a vector of addresses to the **Primitive** block such that DSP Builder gives each vector component a different address. This design example also shows **Lut** blocks working with complex data types. You can initialize **Lut** blocks in exactly the same way.

Using this design example avoids demultiplexing, connecting, and multiplexing, so that you can build parameterizable systems.

You can use one of the following ways to specify the contents of the **Lut** block:

- Specify table contents as single row or column vector. The length of the 1D row or column vector determines the number of addressable entries in the table. If DSP Builder reads vector data from the table, all components of a given vector share the same value.
- When a look-up table contains vector data, you can provide a matrix to specify the table contents. The number of rows in the matrix determines the number of addressable entries in the table. Each row specifies the vector contents of the corresponding table entry. The number of columns must match the vector length, otherwise DSP Builder issues an error.

*Note:*

The default initialization of the LUT is a row vector **round([0:255]/17)**. This vector is inconsistent with the default for the **DualMem** block, which is a column vector **[zeros(16, 1)]**. The latter form is consistent with the new matrix initialization form in which the number of rows determines the addressable size.

The model file is **demo\_lut\_matrix\_init.mdl**.

### 6.12.17. Matrix Initialization of Vector Memories

Use this feature in DSP Builder designs that handle vector data and require individual components of each vector in the dual memory to be initialized uniquely.

The design example file is **demo\_dualmem\_matrix\_init.mdl**.

You can initialize both the dual memory and LUT **Primitive** library blocks with matrix data.

The number of rows in the 2D matrix that you provide for initialization determines the addressable size of the dual memory. The number of columns must match the width of the vector data. So the  $n$ th column specifies the contents of the  $n$ th dual memory. Within each of these columns the  $i$ th row specifies the contents at the  $(j - 1)$ th address (the first row is address zero, second row address 1, and so on).

The exception for this row and column interpretation of the initialization matrix is for 1D data, where the initialization matrix consists of either a single column or single row. In this case, the interpretation is flexible and maps the vector (row or column) into the contents of each dual memory. In the previous behavior all dual memories have identical initial contents.

The **demo\_dualmem\_matrix\_init** design example uses complex values in both the initialization and the data that it later writes to the dual memory. You set up the contents matrix in the model's set-up script, which runs on model initialization.

### 6.12.18. Multichannel IIR Filter

This DSP Builder design example implements a masked multi-channel infinite impulse response (IIR) filter with a masked subsystem that it builds from **Primitive** library blocks.

This design example has many feedback loops. The design example implements all the pipelined delays in the circuit automatically. The multiple channels provide more latency around the circuit to ensure a high clock frequency result. Lumped delays allow you to easily parameterize the design example when changing the channel counts. For example, masking the subsystem provides the benefits of a black-box IP block but with visibility.

The top-level testbench includes **Control** and **Signals** blocks, plus **ChanView** block that deserialize the output buses.

The **IIRChip** subsystem includes the **Device** block and a masked **IIRSubsystem** subsystem. The coefficients for the filter are set from  $[b, a] = \text{ellip}(2, 1, 10, 0.3)$ ; in the callbacks for the masked subsystem. You can look under the mask to see the implementation details of the **IIRSubsystem** subsystem which includes **ChannelIn**, **ChannelOut**, **SampleDelay**, **Const**, **Mult**, **Add**, **Sub**, **Convert**, and **SynthesisInfo** blocks.

The model file is `demo_iir.mdl`.

### 6.12.19. Quadrature Amplitude Modulation

This design example implements a simple quadrature amplitude modulation (QAM256) design example with noise addition. The testbench uses various Simulink blocks.

The top-level testbench includes **Control**, **Signals**, **Run ModelSim**, and **Run Quartus Prime** blocks.

The **QAM256Chip** subsystem includes **Add**, **GPIn**, **GPOut**, **BitExtract**, **Lut**, **BitCombine**, and **SynthesisInfo** blocks.

The model file is `demo_QAM256.mdl`.

*Note:* This design example uses the Simulink Communications Blockset.

### 6.12.20. Reinterpret Cast for Bit Packing and Unpacking

This design example demonstrates the **ReinterpretCast** block, which packs signals into a long word and extracts multiple signals from a long word.

The first datapath reinterprets a single precision complex signal into raw 32-bit components that separate into real and imaginary parts. A **BitCombine** block then merges it into a 64-bit signal. The second datapath uses the **BitExtract** block to split a 64-bit wide signal into a two component vectors of 32-bit signals. The **ReinterpretCast** block then converts the raw bit pattern into single-precision IEEE format. The HDL that the design synthesizes is simple wire connections, which performs no computation.

The model file is `demo_reinterpret_cast.mdl`.

### 6.12.21. Run-time Configurable Decimating and Interpolating Half-Rate FIR Filter

This design example contains a half-rate FIR filter, which can perform either decimation or interpolation by a factor of two during run time.

In decimation mode, the design example accepts a new sample every clock cycle, and produces a new result every two clock cycles. When interpolating, the design example accepts a new input every other clock cycle, and produces a new result every clock cycle. In both cases, the design example fully uses multipliers, making this structure very efficient compared to parallel instantiations of interpolate and decimate filters, or compared to a single rate filter with external interpolate and decimate stages.

The coefficients are set to [1 0 3 0 5 6 5 0 3 0 1] to illustrate the operation of the filter in `setup_demo_fir_tdd.m`.

The model file is `demo_fir_tdd.mdl`.

### 6.12.22. Square Root Using CORDIC

This design example demonstrates the **CORDIC** block. It configures the **CORDIC** block for **uint(32)** input and **uint(16)** output. The example is partially parallelized (four stages).

The design example allows you to generate a valid signal. The design example only generates output and can only accept input every  $N$  cycles, where  $N$  depends on the number of stages, the data output format, and the target  $f_{MAX}$ . The valid signal goes high when the output is ready. You can use this output signal to trigger the next input, for example, a FIFO buffer read for bursty data.

The model file is `demo_cordic_sqrt.mdl`.

### 6.12.23. Test CORDIC Functions with the CORDIC Block

This design example demonstrates how to use the DSP Builder **Primitive CORDIC** block to implement the coordinate rotation digital (CORDIC) algorithm.

The **Mode** input can either rotate the input vector by a specified angle, or rotate the input vector to the x-axis while recording the angle required to make that rotation. You can experiment with different size of inputs to control the precision of the CORDIC output.

The top-level testbench includes **Control** and **Signals** blocks.

The **SinCos** and **AGC** subsystem includes **ChannelIn**, **ChannelOut**, **CORDIC**, and **SynthesisInfo** blocks.

The model file is `demo_cordic_lib_block.mdl`.

### 6.12.24. Uniform Random Number Generator

This DSP Builder design example demonstrates a random number generator (Tausworthe-88) that produces uniformly distributed random numbers.

You can specify the seed value for the random sequence using the `seed_value` input. The reset input resets the sequence to the initial state defined by the `seed_value`. The output is a 32-bit random number, which can be interpreted as a random integer sampled from the uniform distribution.

### 6.12.25. Vector Sort—Sequential

This design example sorts the values on the input vector from largest to smallest. The sorting is a configurable masked subsystem: `sortstages`.

For sorting, the `sortstages` subsystem allows either a comparator and mux based block, or one based on a minimum and a maximum block. The first is more efficient. Both use the reconfigurable subsystem to choose between implementations using the **BlockChoice** parameter.

The design repeatedly uses a dual sort stage in series. The data passes through the dual sort stage  $(\text{vector width})/2$  times.

Look under the mask to view the implementation of reconfigurable subsystem templates and the blocks that reorder and interleave vectors.

The model file is `demo_vectorsort.mdl`.

### 6.12.26. Vector Sort—Iterative

This design sorts the values on the input vector from largest to smallest. The design is a masked subsystem that allows for sorting with either a comparator and mux block, or a minimum and a maximum block. The first implementation is more efficient. Both use the reconfigurable subsystem to choose between implementations using the **BlockChoice** parameter.

Folded designs repeatedly use a single dual sort stage. The throughput of the design is limited in the number of channels, vector width, and data rate. The data passes through the dual sort stage  $(\text{vector width})/2$  times. The vector sort design example uses full throughput with  $(\text{vector width})/2$  dual sort stages in sequence.

Look under the mask to view the implementation of reconfigurable subsystem templates and the blocks that reorder and interleave vectors.

The model file is `demo_foldedsort.mdl`.

### 6.12.27. Vector Initialization of Sample Delay

This DSP Builder design example shows that one sample delay can replace what usually requires a **Demultiplex**, **SampleDelay**, and **Multiplex** combination.

When the **SampleDelay Primitive** library block receives vector input, you can independently specify a different delay for each of the components of the vector.

You may give individual components zero delay resulting in a direct feed through of only that component. Avoid algebraic loops if you select some components to be zero delays.

This rule only applies when DSP Builder is reading and outputting vector data. A scalar specification of delay length still sets all the delays on each vector component to the same value. You must not specify a vector that is not the same length as the vector on the input port. A negative delay on any one component is also an error. However, as in the scalar case, you can specify a zero length delay for one or more of the components.

The model file is `demo_sample_delay_vector.mdl`.

### 6.12.28. Wide Single-Channel Accumulators

This example design shows various ways to connect up an adder, sample delay (depth=1), and optional multiplexer to implement reset or load.

The output type of the adder is propagated from one of the inputs. You must select the correct input, otherwise the accumulator fails to schedule. You may add a **Convert** block to ensure the accumulator also maintains sufficient precision.

The wide single-channel accumulator consists of a two-input adder and sample-delay feedback with one cycle of latency. If you use a fixed-point input to this accumulator, you can make it arbitrarily wide provided the types of the inputs match with a data type prop duplicate block. The output type of the **Add** block can be with or without word growth. Alternatively, you can propagate the input type to the output of the adder.

The optional use of a two-to-one multiplexer allows the accumulator to load values according to a Boolean control signal. The inputs differ in precision, so the type with wider fractional part must be propagated to the output type of the adder, otherwise the accumulator fails to schedule. Converting both inputs to the same precision ensures that the single-channel accumulator can always be scheduled even at high  $f_{MAX}$  targets.

If neither input has a fixed-point type that is suitable for the adder to output, use a **Convert** block to ensure that the precision of both inputs to the **Add** block are the same. Scheduling of this accumulator at high  $f_{MAX}$  fails.

The model file is `demo_wide_accumulators.mdl`.

## 6.13. DSP Builder Reference Designs

DSP Builder also includes reference designs that demonstrate the design of DDC and DUC systems for digital intermediate frequency (IF) processing.

This folder accesses groups of reference designs that illustrate the design of DDC and DUC systems for digital intermediate frequency (IF) processing.

The first group implements IF modem designs compatible with the Worldwide Interoperability for Microwave Access (WiMAX) standard. Intel provides separate models for one and two antenna receivers and transmitters.

The second group implement IF modem designs compatible with the wideband Code Division Multiple Access (W-CDMA) standard.

This folder also contains reference designs.

STAP for radar systems applies temporal and spatial filtering to separate slow moving targets from clutter and null jammers. Applications demand highprocessing requirements and low latency for rapid adaptation. High-dynamic ranges demand floating-point datapaths.

1. [1-Antenna WiMAX DDC](#) on page 152
2. [2-Antenna WiMAX DDC](#) on page 152
3. [1-Antenna WiMAX DUC](#) on page 153
4. [2-Antenna WiMAX DUC](#) on page 153
5. [4-Carrier, 2-Antenna W-CDMA DDC](#) on page 154
6. [1-Carrier, 2-Antenna W-CDMA DDC](#) on page 155
7. [4-Carrier, 2-Antenna W-CDMA DUC](#) on page 155
8. [4-Carrier, 4-Antenna DUC and DDC for LTE](#) on page 156
9. [1-Carrier, 2-Antenna W-CDMA DDC](#) on page 157
10. [4-Carrier, 2-Antenna High-Speed W-CDMA DUC at 368.64 MHz with Total Rate Change 32](#) on page 158
11. [4-Carrier, 2-Antenna High-Speed W-CDMA DUC at 368.64 MHz with Total Rate Change 48](#) on page 158
12. [4-Carrier, 2-Antenna High-Speed W-CDMA DUC at 307.2 MHz with Total Rate Change 40](#) on page 159
13. [Cholesky-based Matrix Inversion](#) on page 160
14. [Cholesky Solver Multiple Channels](#) on page 164
15. [Crest Factor Reduction](#) on page 165
16. [Direct RF with Synthesizable Testbench](#) on page 165
17. [Dynamic Decimating FIR Filter](#) on page 165
18. [Multichannel QR Decomposition](#) on page 166
19. [QR Decomposition](#) on page 166
20. [QRD Solver](#) on page 167
21. [Reconfigurable Decimation Filter](#) on page 168
22. [Single-Channel 10-MHz LTE Transmitter](#) on page 168
23. [STAP Radar Forward and Backward Substitution](#) on page 169
24. [STAP Radar Steering Generation](#) on page 169
25. [STAP Radar QR Decomposition 192x204](#) on page 169
26. [Time Delay Beamformer](#) on page 170
27. [Transmit and Receive Modem](#) on page 170
28. [Variable Integer Rate Decimation Filter](#) on page 171

### Related Information

[AN 544: Digital IF Modem Design with the DSP Builder Advanced Blockset](#)  
For more information about these designs

### 6.13.1. 1-Antenna WiMAX DDC

This reference design uses **IP** and **Interface** blocks to build a 2-channel, 1-antenna, single-frequency modulation DDC for use in an IF modem design compatible with the WiMAX standard.

The top-level testbench includes **Control**, **Signals**, and **Run Quartus Prime** blocks. The design includes an **Edit Params** block to allow easy access to the setup variables in the `setup_wimax_ddc_lrx.m` script.

The **DDCChip** subsystem includes **Device**, **Decimating FIR**, **Mixer**, **NCO**, **SingleRateFIR**, and **Scale** blocks. Also, an **Interleaver** subsystem extracts the correct I and Q channel data from the demodulated data stream.

The FIR filters implement a decimating filter chain that down convert the two channels from a frequency of 89.6 MSPS to a frequency of 11.2 MSPS (a total decimation rate of eight). The real mixer, NCO, and **Interleaver** subsystem isolate the two channels. The design configures the NCO with a single-channel to provide one sine and one cosine wave at a frequency of 22.4 MHz. The NCO has the same sample rate (89.6 MSPS) as the input data sample rate.

A system clock rate of 179.2 MHz drives the design on the FPGA that the **Device** block defines inside the **DDCChip** subsystem.

The model file is `wimax_ddc_lrx.mdl`.

*Note:* This reference design uses the Simulink Signal Processing Blockset.

### 6.13.2. 2-Antenna WiMAX DDC

This reference design uses **IP** and **Interface** blocks to build a 4-channel, 2-antenna, 2-frequency modulation DDC for use in an IF modem design compatible with the WiMAX standard.

The top-level testbench includes **Control**, **Signals**, and **Run Quartus Prime** blocks. The design includes an **Edit Params** block to allow easy access to the setup variables in the `setup_wimax_ddc_2rx_iiqq.m` script.

The **DDCChip** subsystem includes **Device**, **Decimating FIR**, **Mixer**, **NCO**, **SingleRateFIR**, and **Scale** blocks.

The FIR filters implement a decimating filter chain that down convert the two channels from a frequency of 89.6 MSPS to a frequency of 11.2 MSPS (a total decimation rate of 8). The real mixer and NCO isolate the two channels. The design configures the NCO with two channels to provide two sets of sine and cosine waves at the same frequency of 22.4 MHz. The NCO has the same sample rate of (89.6 MSPS) as the input data sample rate.

A system clock rate of 179.2 MHz drives the design on the FPGA, which the **Device** block defines inside the **DDCChip** subsystem.

The model file is `wimax_ddc_2rx_iiqq.mdl`.

*Note:* This reference design uses the Simulink Signal Processing Blockset.



### 6.13.3. 1-Antenna WiMAX DUC

This reference design uses **IP**, **Interface**, and **Primitive** library blocks to build a 2-channel, 1-antenna, single-frequency modulation DUC for use in an IF modem design compatible with the WiMAX standard.

The top-level testbench includes **Control**, **Signals**, and **Run Quartus Prime** blocks. The design includes an **Edit Params** block to allow easy access to the setup variables in the `setup_wimax_duc_1tx.m` script.

The **DUCChip** subsystem includes a **Device** block to specify the target FPGA device, and a **DUC2Channel** subsystem which contains **SingleRateFIR**, **Scale**, **InterpolatingFIR**, **NCO**, and **ComplexMixer** blocks. The **deinterleaver** subsystem contains a series of **Primitive** blocks including delays and multiplexers that deinterleave the two I and Q channels.

The FIR filters implement an interpolating filter chain that up converts the two channels from a frequency of 11.2 MSPS to a frequency of 89.6 MSPS (a total interpolating rate of 8). The complex mixer and NCO modulate the two input channel baseband signals to the IF domain. The design configures the NCO with a single channel to provide one sine and one cosine wave at a frequency of 22.4 MHz. The NCO has the same sample rate (89.6 MSPS) as the input data sample rate.

A system clock rate of 179.2 MHz drives the design on the FPGA, which the **Device** block defines inside the **DUCChip** subsystem.

The model file is `wimax_duc_1tx.mdl`.

*Note:* This reference design uses the Simulink Signal Processing Blockset.

### 6.13.4. 2-Antenna WiMAX DUC

This reference design uses **IP**, **Interface**, and **Primitive** library blocks to build a 4-channel, 2-antenna, single-frequency modulation DUC for use in an IF modem design compatible with the WiMAX standard.

The top-level testbench includes **Control**, **Signals**, and **Run Quartus Prime** blocks. The design includes an **Edit Params** block to allow easy access to the setup variables in the `setup_wimax_duc_2tx_iiqq.m` script.

The **DUCChip** subsystem includes a **Device** block to specify the target FPGA device, and a **DUC2Channel** subsystem which contains **SingleRateFIR**, **Scale**, **InterpolatingFIR**, **NCO**, **ComplexMixer**, and **Const** blocks. It also contains a **Sync** subsystem, which shows how to manage two data streams coming together and synchronizing. The design writes the data from the NCOs to a memory with the channel index as an address. The data stream uses its channel signals to read out the NCO signals, which resynchronizes the data correctly. (Alternatively, you can simply delay the NCO value by the correct number of cycles to ensure that the NCO and channel data arrive at the **Mixer** on the same cycle). The **deinterleaver** subsystem contains a series of **Primitive** blocks including delays and multiplexers that deinterleave the four I and Q channels.

The FIR filters implement an interpolating filter chain that up converts the two channels from a frequency of 11.2 MSPS to a frequency of 89.6 MSPS (a total interpolating rate of 8).

A complex mixer and NCO modulate the two input channel baseband signals to the IF domain. The design configures the NCO to provide two sets of sine and cosine waves at a frequency of 22.4 MHz. The NCO has the same sample rate (89.6 MSPS) as the input data sample rate.

The **Sync** subsystem shows how to manage two data streams coming together and synchronizing. The design writes the data from the NCOs to a memory with the channel as an address. The data stream uses its channel signals to read out the NCO signals, which resynchronizes the data correctly.

A system clock rate of 179.2 MHz drives the design on the FPGA, which the **Device** block defines inside the **DUCChip** subsystem.

The model file is `wimax_duc_2tx_iiqq.mdl`.

*Note:* This reference design uses the Simulink Signal Processing Blockset.

### 6.13.5. 4-Carrier, 2-Antenna W-CDMA DDC

This reference design uses **IP** and **Interface** blocks to build a 16-channel, 2-antenna, multiple-frequency modulation DDC for use in an IF modem design compatible with the W-CDMA standard.

The top-level testbench includes **Control**, **Signals**, and **Run Quartus Prime** blocks, plus a **ChanView** block that isolates two channels of data from the TDM signals.

The **DDCChip** subsystem includes **Device**, **DecimatingCIC**, **Decimating FIR**, **Mixer**, **NCO**, and **Scale** blocks. It also contains a **Sync** subsystem which provides the synchronization of the channel data to the NCO carrier waves.

The CIC and FIR filters implement a decimating filter chain that down converts the eight complex carriers (16 real channels from two antennas with four pairs of I and Q inputs from each antenna) from a frequency of 122.88 MSPS to a frequency of 7.68 MSPS (a total decimation rate of 16). The real mixer and NCO isolate the four channels. The design configures the NCO with four channels to provide four pairs of sine and cosine waves at frequencies of 12.5 MHz, 17.5 MHz, 22.5 MHz, and 27.5 MHz, respectively. The NCO has the same sample rate (122.88 MSPS) as the input data sample rate.

The **Sync** subsystem shows how to manage two data streams that come together and synchronize. The data from the NCOs writes to a memory with the channel as an address. The data stream uses its channel signals to read out the NCO signals, which resynchronizes the data correctly.

A system clock rate of 245.76 MHz drives the design on the FPGA, which the **Device** block defines inside the **DDCChip** subsystem.

The model file is `wcdma_multichannel_ddc_mixer.mdl`.

*Note:* This reference design uses the Simulink Signal Processing Blockset.

### 6.13.6. 1-Carrier, 2-Antenna W-CDMA DDC

This reference design uses **IP** and **Interface** blocks to build a 4-channel, 2-antenna, single-frequency modulation DDC for use in an IF modem design compatible with the W-CDMA standard.

The top-level testbench includes **Control**, **Signals**, and **Run Quartus Prime** blocks, plus a **ChanView** block that isolates two channels of data from the TDM signals.

The **DDCChip** subsystem includes **Device**, **DecimatingCIC**, **Decimating FIR**, **Mixer**, **NCO**, and **Scale** blocks.

The CIC and FIR filters implement a decimating filter chain that down converts the two complex carriers (4 real channels from two antennas with one pair of I and Q inputs from each antenna) from a frequency of 122.88 MSPS to a frequency of 7.68 MSPS (a total decimation rate of 16). The real mixer and NCO isolate the four channels. The design configures the NCO with a single channel to provide one sine and one cosine wave at a frequency of 17.5 MHz. The NCO has the same sample rate (122.88 MSPS) as the input data sample rate.

A system clock rate of 122.88 MHz drives the design on the FPGA, which the **Device** block defines inside the **DDCChip** subsystem.

The model file is `wcdma_picocell_ddc_mixer.mdl`.

*Note:* This reference design uses the Simulink Signal Processing Blockset.

### 6.13.7. 4-Carrier, 2-Antenna W-CDMA DUC

This reference design uses **IP** and **Interface** blocks to build a 16-channel, 2-antenna, multiple-frequency modulation DUC for use in an IF modem design compatible with the W-CDMA standard.

The top-level testbench includes **Control**, **Signals**, and **Run Quartus Prime** blocks. A **Spectrum Scope** block computes and displays the periodogram of the outputs from the two antennas.

The **DUCChip** subsystem includes a **Device** block to specify the target FPGA device, and a **DUC** subsystem that contains **InterpolatingFIR**, **InterpolatingCIC**, **NCO**, **ComplexMixer**, and **Scale** blocks.

The FIR and CIC filters implement an interpolating filter chain that up converts the 16-channel input data from a frequency of 3.84 MSPS to a frequency of 122.88 MSPS (a total interpolation factor of 32). The complex mixer and NCO modulate the four channel baseband input signal onto the IF region. The design configures the NCO with four channels to provide four pairs of sine and cosine waves at frequencies of 12.5 MHz, 17.5 MHz, 22.5 MHz, and 27.5 MHz, respectively. The NCO has the same sample rate (122.88 MSPS) as the final interpolated output sample rate from the last CIC filter in the interpolating filter chain.

The subsystem **SyncMixSumSel** uses Primitive blocks to implement the synchronization, mixing, summation, scaling, and signal selection. This subsystem separates each operation into further subsystems. The **Sync** subsystem shows how to manage two data streams that come together and synchronize. The data from the NCOs writes to a memory with the channel as an address. The data stream uses its channel signals to read out the NCO signals, which resynchronizes the data correctly.

The **Sum** and **SampSelectr** subsystems sum up the correct modulated signals to the designated antenna.

A system clock rate of 245.76 MHz drives the design on the FPGA, which the **Device** block defines inside the **DUC** subsystem.

The model file is `wcdma_multichannel_duc_mixer.mdl`.

*Note:* This reference design uses the Simulink Signal Processing Blockset.

### 6.13.8. 4-Carrier, 4-Antenna DUC and DDC for LTE

These DUC and matching DDC designs connect to 4 antennas and can process 4 channels per antenna. With a sample rate of 61.44 MHz and a clock rate of 491.52 MHz, these designs represent up- and downconverters used in LTE.

#### DUC

The top-level design of the upconverter contains a **TEST\_BENCH** block with signal sources, the upconverter, and a **SINKS** block that stores the datastreams coming out of the upconverter in MATLAB variables. Depending on which simulation you run, the **TEST\_BENCH** block uses either real LTE sample streams or specialized debugging patterns. The upconverter consists of the LDUC module, the lower DUC, which contains a channel filter and two interpolating filters, each interpolating by a factor of 2. The filtered sample stream feeds into the **COMPLEX MIXER** block, where a NCO generates separate frequencies for each of the four channels, and multiplies the generated sinewaves with the filtered sample stream. A delay match block ensures that the sample stream and the generated frequencies align correctly. After the **COMPLEX MIXER** block is an **antenna summer** block, which adds up the different channels for each antenna, multiplies each with a different frequency, and outputs them to the four separate antennas.

The model file is `duc_4c4ant.mdl`.

#### DDC

The top-level design of the DDC also contains a **TESTBENCH** block, which contains source blocks that read from workspace. It uses the data that DSP Builder generates during the simulation of the DUC. The **SINKS** block again traces the outputs of the design in MATLAB variables, which you can analyze and manipulate in MATLAB. The DDC consists of a complex mixer that matches the complex mixer of the DUC, and the LDDC (Lower DownConverter), which contains two decimate-by-2 filters and a channel filter.

The model file is `ddc_4c4ant.mdl`.

#### Simulation Scripts

The design, which is in the `Examples\ReferenceDesigns\DDC4c4ant\4C4T4R_echodemo\4C4T4R\Design` directory, contains two separate parts: `duc_4c4ant.mdl` contains the upconverter, and `ddc_4c4ant.mdl` contains the downconverter. The directory also contains two scripts that allow you to run the simulation of both designs: Both `Run_DUC_DDC_demo.m` and `Test_DUC_DDC_demo.m` create test vectors, run the upconverter first, which

generates the input vectors for the downconverter, - then run the downconverter and analyze the outputs. The designs contains no channel model, but you can add your own channel model and apply it to the output data of the DUC before running the DDC to simulate more realistic operating conditions. `Run_DUC_DDC_demo.m` uses typical LTE waveforms; `Test_DUC_DDC_demo.m` works with ramps that help visualizing which data goes into which channel and which antenna it transmits on. In the test pattern, an impulse is set first, followed by a ramp on channel 1 on antenna 1. All other channels and antenna are 0. The next section transmits channel 1 on antenna 1, channel 2 on antenna 2 ... channel 4 on antenna 4. The last section transmits all 4 channels on all 4 antennas, using the full capacity of the system. Use this debug pattern, if you want to modify or extend the design. Run the scripts using the `echodemo` command, to step through the script section by section, by typing `echodemo Run_DUC_DDC_demo.m` at the MATLAB command prompt, and then clicking **Next** several times to step through the simulation script. Alternatively, you can run the entire script by typing `Run_DUC_DDC_demo.m` at the MATLAB command prompt. The last step of the script calls up a plot function that generates input vs output plots for each channel, with overlaid input and output plots. These plots should match closely, displaying only a small quantization error. The script also produces channel scopes, which show each channel's data in time and frequency domains.

### 6.13.9. 1-Carrier, 2-Antenna W-CDMA DDC

This reference design uses **IP** and **Interface** blocks to build a 4-channel, 2-antenna, single-frequency modulation DUC for an IF modem design compatible with the W-CDMA standard.

The top-level testbench includes **Control**, **Signals**, and **Run Quartus Prime** blocks. A **Spectrum Scope** block computes and displays the periodogram of the outputs from the two antennas.

The **DUCChip** subsystem includes a **Device** block to specify the target FPGA device, and a **DUC** subsystem that contains **InterpolatingFIR**, **InterpolatingCIC**, **NCO**, **ComplexMixer**, and **Scale** blocks.

The FIR and CIC filters implement an interpolating filter chain that up convert the four channel input data from a frequency of 3.84 MSPS to a frequency of 122.88 MSPS (a total interpolation factor of 32). The complex mixer and NCO modulate the four channel baseband input signal onto the IF region.

The design example configures the NCO with a single channel to provide one sine and one cosine wave at a frequency of 17.5 MHz. The NCO has the same sample rate (122.88 MSPS) as the final interpolated output sample rate from the last CIC filter in the interpolating filter chain.

A system clock rate of 122.88 MHz drives the design on the FPGA, which the **Device** block defines inside the **DDC** subsystem.

The model file is `wcdma_picocell_duc_mixer.mdl`.

*Note:* This reference design uses the Simulink Signal Processing Blockset.

### 6.13.10. 4-Carrier, 2-Antenna High-Speed W-CDMA DUC at 368.64 MHz with Total Rate Change 32

This reference design uses **IP** and **Interface** blocks to build a high-speed 16-channel, 2-antenna, multiple-frequency modulation DUC for use in an IF modem design compatible with the W-CDMA standard.

The top-level testbench includes **Control**, **Signals**, and **Run Quartus Prime** blocks. A **Spectrum Scope** block computes and displays the periodogram of the outputs from the two antennas.

The **DUCChip** subsystem includes a **Device** block to specify the target FPGA device, and a **DUC** subsystem that contains **InterpolatingFIR**, **InterpolatingCIC**, **NCO**, **ComplexMixer**, and **Scale** blocks.

The FIR and CIC filters implement an interpolating filter chain that up converts the 16-channel input data from a frequency of 3.84 MSPS to a frequency of 122.88 MSPS (a total interpolation factor of 32). This design example uses dummy signals and carriers to achieve the desired rate up conversion, because of the unusual FPGA clock frequency and total rate change combination. The complex mixer and NCO modulate the four channel baseband input signal onto the IF region. The design example configures the NCO with four channels to provide four pairs of sine and cosine waves at frequencies of 12.5 MHz, 17.5 MHz, 22.5 MHz and 27.5 MHz, respectively. The NCO has the same sample rate (122.88 MSPS) as the final interpolated output sample rate from the last CIC filter in the interpolating filter chain.

The **Sync** subsystem shows how to manage two data streams that come together and synchronize. The data from the NCOs writes to a memory with the channel as an address. The data stream uses its channel signals to read out the NCO signals, which resynchronizes the data correctly.

The **GenCarrier** subsystem manipulates the NCO outputs to generate carrier signals that can align with the datapath signals.

The **CarrierSum** and **SignalSelector** subsystems sum up the right modulated signals to the designated antenna.

A system clock rate of 368.64 MHz, which is 96 times the input sample rate, drives the design on the FPGA, which the **Device** block defines inside the **DUC** subsystem. The higher clock rate can potentially allow resource re-use in other modules of a digital system implemented on an FPGA.

The model file is `mcducmix96x32R.mdl`.

*Note:* This reference design uses the Simulink Signal Processing Blockset.

### 6.13.11. 4-Carrier, 2-Antenna High-Speed W-CDMA DUC at 368.64 MHz with Total Rate Change 48

This reference design uses **IP** and **Interface** blocks to build a high-speed 16-channel, 2-antenna, multiple-frequency modulation DUC for use in an IF modem design compatible with the W-CDMA standard.

The top-level testbench includes **Control**, **Signals**, and **Run Quartus Prime** blocks. A **Spectrum Scope** block computes and displays the periodogram of the outputs from the two antennas.

The **DUCChip** subsystem includes a **Device** block to specify the target FPGA device, and a **DUC** subsystem that contains **InterpolatingFIR**, **InterpolatingCIC**, **NCO**, **ComplexMixer**, and **Scale** blocks.

The FIR and CIC filters implement an interpolating filter chain that up converts the 16-channel input data from a frequency of 3.84 MSPS to a frequency of 184.32 MSPS (a total interpolation factor of 48).

The complex mixer and NCO modulate the four channel baseband input signal onto the IF region. The design configures the NCO with four channels to provide four pairs of sine and cosine waves at frequencies of 12.5 MHz, 17.5 MHz, 22.5 MHz, and 27.5 MHz, respectively. The NCO has the same sample rate (184.32 MSPS) as the final interpolated output sample rate from the last CIC filter in the interpolating filter chain.

The **Sync** subsystem shows how to manage two data streams that come together and synchronize. The data from the NCOs writes to a memory with the channel as an address. The data stream uses its channel signals to read out the NCO signals, which resynchronizes the data correctly.

The **CarrierSum** and **SignalSelector** subsystems sum up the right modulated signals to the designated antenna.

A system clock rate of 368.64 MHz, which is 96 times the input sample rate, drives the design on the FPGA, which the **Device** block defines inside the **DUC** subsystem. The higher clock rate can potentially allow resource re-use in other modules of a digital system implemented on an FPGA.

The model file is `mcducmix96x48R.mdl`.

*Note:* This reference design uses the Simulink Signal Processing Blockset.

### 6.13.12. 4-Carrier, 2-Antenna High-Speed W-CDMA DUC at 307.2 MHz with Total Rate Change 40

This reference design uses **IP** and **Interface** blocks to build a high-speed 16-channel, 2-antenna, multiple-frequency modulation DUC for use in an IF modem design compatible with the W-CDMA standard

The top-level testbench includes **Control**, **Signals**, and **Run Quartus Prime** blocks. A **Spectrum Scope** block computes and displays the periodogram of the outputs from the two antennas.

The **DUCChip** subsystem includes a **Device** block to specify the target FPGA device, and a **DUC** subsystem that contains **InterpolatingFIR**, **InterpolatingCIC**, **NCO**, **ComplexMixer**, and **Scale** blocks.

The FIR and CIC filters implement an interpolating filter chain that up converts the 16-channel input data from a frequency of 3.84 MSPS to a frequency of 153.6 MSPS (a total interpolation factor of 40).



The complex mixer and NCO modulate the four channel baseband input signal onto the IF region. The design configures the NCO with four channels to provide four pairs of sine and cosine waves at frequencies of 12.5 MHz, 17.5 MHz, 22.5 MHz, and 27.5 MHz, respectively. The NCO has the same sample rate (153.6 MSPS) as the final interpolated output sample rate from the last CIC filter in the interpolating filter chain.

The **Sync** subsystem shows how to manage two data streams that come together and Synchronize. The design writes data from the NCOs to a memory with the channel as an address. The data stream uses its channel signals to read out the NCO signals, which resynchronizes the data correctly.

The **CarrierSum** and **SignalSelector** subsystems sum up the right modulated signals to the designated antenna.

A system clock rate of 307.2 MHz, which is 80 times the input sample rate, drives the design on the FPGA, which the **Device** block defines inside the **DUC** subsystem. The higher clock rate can potentially allow resource re-use in other modules of a digital system implemented on an FPGA.

The model file is `mducmix80x40R.mdl`.

*Note:* This reference design uses the Simulink Signal Processing Blockset.


### 6.13.13. Cholesky-based Matrix Inversion

Matrix inversion has many applications in wireless communications, e.g. digital predistortion (DPD) for RF linearization and multiple-input multiple-output (MIMO) detection. Matrix inversion algorithms typically require high-resolution numerics to guarantee accuracy and numerical stability. The implementation is normally resource demanding in particular if the matrix dimension grows. The DSP Builder Cholesky-based Matrix Inversion reference design offers an efficient implementation of matrix inversion for minimized resource utilization and improved latency and throughput. The Cholesky decomposition technique inverts a positive-definite real or complex square matrix. Cholesky decomposition-based matrix inversion is more efficient than direct matrix inversion.

**Figure 69. Matrix inversion based on Cholesky decomposition**

The figure shows the three steps of implementing a Hermitian matrix inversion using Cholesky decomposition:

1. Cholesky decomposition
2. Triangular matrix inversion through forward substitution
3. Triangular matrix multiplication



The Cholesky decomposition calculates the reciprocal values of the diagonal elements of  $L$ ,  $\frac{1}{L_{kk}}$  which the triangular matrix inversion requires. The design propagates those values to the output interface of the Cholesky decomposition reducing resource usage and latency.



Assuming matrix  $A$  is an  $N \times N$  positive-definite square matrix, Cholesky decomposition of  $A$  into lower and upper triangular matrices,  $L$ , and  $L^H$  is given by:

$$A = L^H L$$

The inverse of Hermitian  $A$ ,  $A^{-1}$  is:

$$A^{-1} = (L^{-1})^H \cdot L^{-1}$$

The design performs Cholesky decomposition and calculates the inverse of  $L$ ,  $J = L^{-1}$ , through forward substitution.  $J$  is a lower triangle matrix. The inverse of the input matrix requires a triangular matrix multiplication, followed by a Hermitian matrix multiplication:

$$A^{-1} = J^H \cdot J$$

The Cholesky-based matrix inversion reference design comprises a Cholesky decomposition design and a triangular matrix inversion design. Both designs are fully pipelined, with multichannel input and output streaming to maximize throughput. The size of dot-product engines in both designs are compile-time configurable according to the size of the input matrices. The datapath and control logic are split.

**Figure 70. Cholesky Decomposition Top-level Design**

Input = Size\* (size + 1)\* channel/2

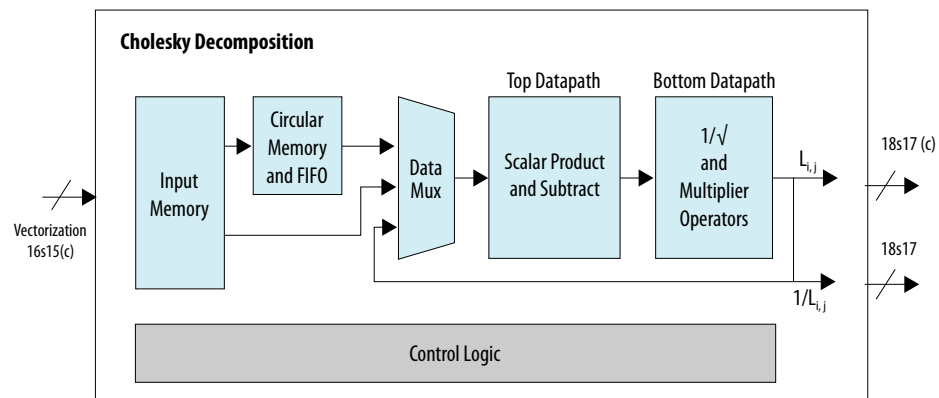
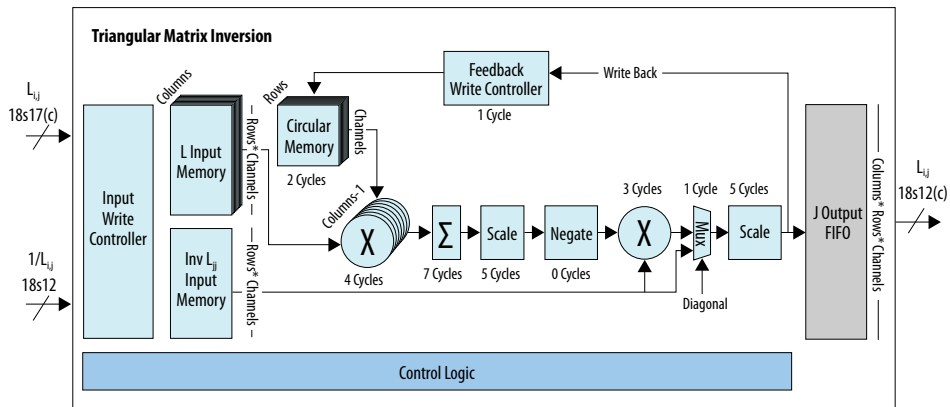


Figure 71. Triangular Matrix Inversion Top-level Design



This design supports single-precision floating-point Cholesky matrix inversion. DSP Builder requires a single-precision floating-point input for the floating point inversion.

Matrix inversion takes multiple matrices and interleaves the inverse computations for all matrices. This method hides the latency in computing each element by pipelining inversion of a completely different channel. Multichannel designs use the idle cycles in the computation chain to process the next channel. Two buffers at the input and output of the design create channels for streaming matrices into multichannel interfaces.

Table 19. Top-level matrix inversion input and output ports

The input and output interfaces follow Avalon™ streaming (Avalon-ST) standard.

Signal	Direction	Type	Width	Description
Sink_Valid	Input	Boolean	1	Avalon streaming sink valid signal for the input matrix interface. Number of valid input = (matrix size*(matrix size + 1))/2
Sink_Channel	Input	unsigned integer	8	Avalon streaming sink channel bus for the input matrix interface.
Sink_Data	Input	Single floating-point complex	64 bit I/Q	Avalon streaming sink data bus for the input matrix interface. Lower matrix elements are streamed in column major order.
Source_Valid	Output	Boolean	1	Avalon streaming source valid signal for output interface. This signal is asserted for (size*(size+1))/2 clocks
Source_Channel	Output	unsigned integer	8	Avalon streaming source channel bus for output interface.
Source_Data	Output	Single floating-point complex	64 bit I/Q	Avalon streaming source data bus for output interface. Lower matrix elements are streamed in column major order.

## Parameters

**Table 20. Parameters of the matrix inversion design**

The parameters are compile-time configurable using the setup file. .

Parameter	Description
<b>Size of Matrix</b>	The size of matrix to invert.
<b>Channels</b>	Number of matrices inverted in a burst. Minimum of 16 channel.
<b>Latency</b>	The period in cycles the module waits before receiving the next set of matrices.


DSP Builder calculates the throughput of the design by setting the latency value and the system clock:

$$\text{Throughput (matrix inversion per second)} = \text{System clock} / \text{Latency}$$

Although elements of input matrices arrive in streaming format, the internal channelizer vectorizes the input matrices into several channels (the default is 16). This vectorization significantly improves the throughput.

**Figure 72. Input streaming interface for 8x8 Hermitian input matrix**

The figure shows the latency configuration parameter in the input interface including data, valid, and channel signals. In this example of 8x8 matrix inversion, the valid signal remains high for 36 clock cycles (total number of lower triangle elements of the Hermitian matrix of 8x8) and remains low for (latency - 36) cycles before inserting the next matrix elements. The minimum duration to remain low and hence the minimum latency period may vary depending on the matrix size and the pipelining required to meet timing constraints.



**Table 21. Recommended Values for the Minimum Latency (maximum throughput)**

In Intel Stratix 10 and Intel Arria 10 devices, speed grade -1 and -2, for three different matrix sizes.

Matrix Dimension	Latency in clock cycles	
	Intel Arria 10 Devices	Intel Stratix 10 Devices
4x4	≥ 30	≥ 30
8x8	≥ 75	≥ 74
16x16	≥ 230	≥ 220

### Performance and Resource Usage

**Table 22. Floating-point implementation resource utilization targeting GX/SX/TX 280 FPGA**

The table shows the resource count of the floating-point Cholesky-based matrix inversion design including the channelizing input and output buffers.

Matrix Dimension	Number of channels	Logic Elements (ALMs)	DSP Blocks	Memory bits	RAM blocks	Registers
4x4	16	8,236	55	548,448	55	22,066
8x8	16	16,665	103	2,001,664	194	45,463
16x16	16	35,025	199	7,085,088	521	95,079

**Table 23. Performance of the floating-point matrix inversion module for different matrix dimensions**

This table shows the  $f_{MAX}$  performance of the floating-point design for different matrix sizes with a system clock of 368.64 MHz and targeting a FPGA device. The maximum throughput is in millions of matrix inversions per second.

Matrix Dimension	Number of channels	Target System clock (MHz)	$f_{MAX}$ (MHz)	Throughput <sub>MAX</sub>
4x4	16	368.64	468.06	12.2
8x8	16	368.64	403.88	5.0
16x16	16	368.64	392.77	1.67

### 6.13.14. Cholesky Solver Multiple Channels

The Cholesky Solver Multiple Channels reference design performs Cholesky decomposition to solve column vector  $x$  in  $Ax = b$

$A$  is a Hermitian, positive definite matrix (for example covariance matrix) and  $b$  is a column vector.

The design uses forward and backward substitution to solve  $x$ .

The design decomposes  $A$  into  $L*L'$ , therefore  $L*L'*x = b$ , or  $L*y = b$ , where  $y = L'*x$ . The design solves  $y$  with forward substitution and  $x$  with backward substitution.

This design uses cycle stealing and command FIFO techniques to enhance performance. Although it targets multiple channels, it also works well with single channels.

To input the lower triangular elements of matrix A and b with the input bus, specify the column, row, and channel index of each element. The design transposes and appends the column vector b to the bottom of A and treats it as an extension of A in terms of column and row addressing.

The output is column vector x with the bottom element output first.

A multiple channel design optimizes performance by prioritizing diagonal element calculation over non-diagonal ones.

The model file is `cholseky_solver_mc.mdl`.

### 6.13.15. Crest Factor Reduction

This reference design implements crest factor reduction, based on the peak cancelling algorithm.

For further information refer to the web page.

You can change the simulation length by clicking on the Simulink Length block.

The model file is `demo_cfr.mdl`.

#### Related Information

[Crest factor reduction for wireless systems](#)

### 6.13.16. Direct RF with Synthesizable Testbench

This very large reference design implements a digital upconversion to RF and digital predistortion, with a testbench that you can synthesize to hardware for easier on-chip testing.

The model file is `DirectRFTest_and_DPD_SV.mdl`.

### 6.13.17. Dynamic Decimating FIR Filter

The dynamic decimating FIR reference design offers multichannel run-time decimation ratios in integer power of 2 and run-time control of channel count (in trading with bandwidth). The design supports dynamic channel count to signal bandwidth trade off (if you halve the channel count, the input sample rate doubles).

The FIR filter length is  $2 \times (D_{\max} / D_{\min}) \times N + 1$  where  $D_{\max}$  and  $D_{\min}$  are the maximum and minimum decimation ratios and N is the number of (1 sided) symmetric coefficients at  $D_{\min}$ .

All channels must have the same decimation ratio. The product of the number of channels and the minimum decimation ratio must be 4 or more. The design limits the wire count to 1 and:

number of channels x sample rate = clock rate.

The model file is `demo_dyndeci.mdl`

### 6.13.18. Multichannel QR Decomposition

This reference design is a complete linear equations system solution that uses QR decomposition.

To optimize the overall throughput the solver can interleave multiple data instances at the same time. The inputs of the design are system matrices  $A [n \times m]$  and input vectors.

The reference design uses the Gram-Schmidt method to decompose system matrix  $A$  to  $Q$  and  $R$  matrices. It calculates the solution of the system by completing backward substitution.

The reference design is fully parametrizable: system dimensions  $n$  and  $m$ , the processing vector size, which defines the parallelization ratio of the dot product engine, and the number of channels that the design processes in parallel. This design uses single-precision **Multiply** and **Add** blocks that perform most of the floating-point calculations to implement a parallel dot product engine. The design uses a processor, which executes a fixed set of micro-instructions and generates operation indexes, to route different phases of the calculation through these blocks. The design uses for-loop macro blocks, which allow very efficient, flexible, and high-level implementation of iterative operations, to implement the processor.

The model file is `demo_mcqrd.mdl`.

### 6.13.19. QR Decomposition

This reference design is a complete linear equations system solution that uses QR decomposition.

The input of the design is a system matrix  $A [n \times m]$  and input vector.

The reference design uses the Gram-Schmidt method to decompose system matrix  $A$  to  $Q$  and  $R$  matrices, and calculates the solution of the system by completing backward substitution.

The reference design is fully parametrizable—system dimensions  $n$  and  $m$ , and the processing vector size, which defines the parallelization ratio of the dot product engine. This design uses single-precision **Multiply** and **Add** blocks that perform most of the floating-point calculations to implement a parallel dot product engine. The design uses a processor, which executes a fixed set of microinstructions and generates operation indexes, to route different phases of the calculation through these blocks. The design uses for-loop macro blocks, which allow very efficient, flexible and high-level implementation of iterative operations, to implement the processor.

This design uses the **Run All Testbenches** block to access enhanced features of the automatically-generated testbench. An application-specific m-function verifies the simulation output, to correctly handle the complex results and the numerical approximation because of the floating-point format.

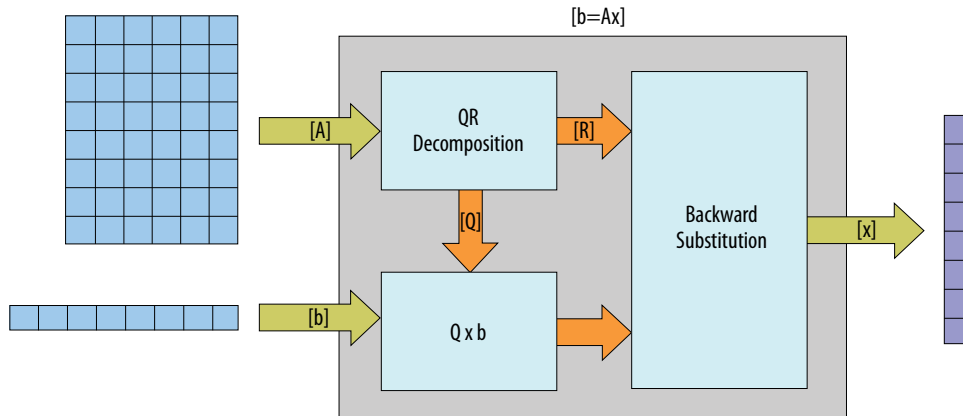
The model file is `demo_qrd.mdl`.

### 6.13.20. QRD Solver

The QRD Solver reference design is a complete linear equations system solution using QR decomposition. The input of the design is a system matrix  $A$  [ $n \times m$ ] and input vector  $[b]$ .

**Figure 73. QRD Solver**

The design decomposes the system matrix  $A$  to  $Q$  and  $R$  matrices using the Gram-Schmidt method. The design calculates the solution of the system by completing backward substitution.



The reference design is fully parameterizable over system dimensions  $n$  and  $m$  and the processing vector size, which defines the parallelization ratio of the dot product engine. This design implements parallel dot product engine using single-precision **Multiply** and **Add** blocks that perform most of the floating-point calculations. The design routes different phases of the calculation through these blocks with a controlling processor that executes a fixed set of microinstructions and generates operation indexes. The design implements the controlling processor using for-loop macro blocks, which allow very efficient, flexible, and high-level implementation of iterative operations.

This design uses the **Run All Testbenches** block to access enhanced features of the automatically generated testbench. An application-specific m-function verifies the simulation output, to correctly handle the complex results and the numerical approximation because of the floating-point format. Intel optimized the design for Intel Stratix 10 FPGAs. The design implements hardened floating-point operators in the FPGA DSP blocks.

**Table 24. Performance**

Intel tested the design with Intel Quartus Prime v18.1.1 build 259, targeting a 1SG280LN3F43E2VG device

Matrix Size	Parallel Processing Vector Size	fMAX (MHz)	Resources			Throughput		Latency	
			ALM	DSPs	M20K	Cycles	Matrices/s	Cycles	ms
512x256	512	320	461K (49%)	4,370 (76%)	1,313 (11%)	71,232	4,492	137,545	0.43
64x64	64	418	60.5 (6%)	562 (10%)	160 (1%)	7,920	52,777	12,392	0.03

The model file is demo\_qrd\_s10.mdl.

### 6.13.21. Reconfigurable Decimation Filter

The reconfigurable decimation filter reference design uses primitive blocks to build a variable integer rate decimation FIR filter.

The reference design has the following features:

- Supports arbitrary integer decimation rate (including the cases without rate change), arbitrary number of channels and arbitrary clock rate and input sample rate, if the clock rate is high enough to process all channels in a single data path (i.e. no hardware duplication).
- Supports run-time reconfiguration of decimation rate.
- Uses two memory banks for filter coefficients storage instead of prestoring coefficients for all rates in memory. Updates one memory bank while the design is reading coefficients from the other bank.
- Implements real time control of scaling in the FIR datapath.

You can modify the parameters in the `setup_var Downsampler.m` file, which you access from the Edit Params icon.

The model file is `var Downsampler.mdl`.

### 6.13.22. Single-Channel 10-MHz LTE Transmitter

This reference design uses **IP**, **Primitive**, and blocks from the FFT Blockset library to build a single-channel 10-MHz LTE transmitter.

The top-level testbench includes blocks to access control and signals, and to run the Quartus Prime software. It also includes an **Edit Params** block to allow easy access to the configuration variables in the `setup_sc_LTEtxr.m` script. A discrete-time scatter plot scope displays the constellation of the modulated signal in inphase versus quadrature components.

The **LTE\_txr** subsystem includes a **Device** block to specify the target FPGA device, and **64QAM**, **1K\_IFFT**, **ScaleRnd**, **CP\_bReverse**, **Chg\_Data\_Format**, and **DUC** blocks.

The **64QAM** subsystem uses a lookup table to convert the source input data into 64 QAM symbol mapped data. The **1K\_IFFT** subsystem converts the frequency domain quadrature amplitude modulation (QAM) modulated symbols to the time domain. The **ScaleRnd** subsystem follows the conversion, which scales down the output signals and converts them to the specified fixed-point type.

The **bit CP\_bReverse** subsystem adds extended cycle prefix (CP) or guard interval for each orthogonal frequency-domain multiplexing (OFDM) symbol to avoid intersymbol interference (ISI) that causes multipaths. The **CP\_bReverse** block reorders the output bits of IFFT subsystems, which are in bit-reversed order, so that they are in the correct order in the time domain. The design adds the cyclic prefix bit by copying the last 25% of the data frame, then appends to the beginning of it.

The **Chg\_Data\_Format** subsystem changes the output data format of **CP\_bReverse** subsystem to match the working protocol format of **DUC** subsystem.



The **DUC** subsystem uses an interpolating filter chain to achieve an interpolation factor of 16, such that the design interpolates the 15.36 Msps input channel to 245.76 Msps. In this design, an interpolating finite impulse response (FIR) filter interpolates by 2, followed by a cascaded integrator-comb (CIC) filter with an interpolation rate of 8. An NCO generates orthogonal sinusoids at specified carrier frequency. The design mixes the signals with complex input data with a **ComplexMixer** block. The final SINC compensation filter compensates for the digital analog converter (DAC) frequency response roll-off.

A system clock rate of 245.76 MHz drives the design on the FPGA. The **Signals** block of the design defines this clock. The input random data for the **64QAM** symbol mapping subsystem has a data rate of 15.36 Msps.

The model file is `sc_LTEtxr.mdl`.

### 6.13.23. STAP Radar Forward and Backward Substitution

The QR decomposition reference design produces an upper triangular matrix and a lower triangular matrix.

The design applies this linear system of equations to the steering vector in the following two steps:

- Forward substitution with the lower triangular matrix
- Backward substitution with the lower triangular matrix

A command pipeline controls the routing of floating-point vectors. Nested **ForLoop** blocks generate these commands. Another FIFO unit queues the commands. This decoupled system of FIFO buffers maximizes the usage of the shared vector floating-point block while automatically throttling the rate of the **ForLoop** system.

This design uses advanced settings from the **DSP Builder > Verify Design** menu to access enhanced features of the automatically generated testbench. An application specific m-function verifies the simulation output, to correctly compare complex results and properly handle floating-point errors that arise from the ill-conditioning of the QRD output.

The model file is `STAP_ForwardAndBackwardSubstitution.mdl`.

### 6.13.24. STAP Radar Steering Generation

The STAP radar steering generation reference design uses **ForLoop** blocks and floating-point primitives to generate the steering vector. You input the angle of arrival and Doppler frequency.

The model file is `STAP_steeringGen.mdl`.

### 6.13.25. STAP Radar QR Decomposition 192x204

The QR decomposition reference design implements a sequence of floating-point vector operations.

Single-precision **Multiply** and **Add** blocks perform most of the floating-point calculations. The design routes different phases of the calculation through these blocks with a controlling processor that executes a fixed set of microinstructions. FIFO units ensure this architecture maximizes the usage of the **Multiply** and **Add** blocks.

This design uses the **Run All Testbenches** block to access enhanced features of the automatically generated testbench. An application specific m-function verifies the simulation output, to correctly handle the complex results and the numerical approximation due to the floating-point format.

The model file is `STAP_qrd192x204.mdl`. The parallel version model file is **STAP\_qrd192x204\_p.mdl**.

### 6.13.26. Time Delay Beamformer

The time delay beamformer reference design implements a time-delay beamformer that has many advantages over traditional phase-shifted beamformer. It uses a (full-band) Nyquist filter and Farrow-like structure for optimal performance and resource usages.

The design includes the following features so you can simulate and verify the transmit and receive beamforming operations:

- Waveform (chirp) generation
- Target emulation
- Receiver noise emulation
- Aperture tapering
- Pulse compression

### 6.13.27. Transmit and Receive Modem

The transmit and receive modem design contains a QAM transmitter, a synthesizable channel model and a receiver, working at sample rates that match or exceed the clock rate. The design works at different sample rates, and can provide up to 16 parallel data streams between transmitter and receiver.

The transmitter can produce random data, which is useful for generating a hardware demo, or you can feed it with data from the MATLAB environment. You can modulate the data, where the modulation order can be QAM4 or QAM64. The design filters the signal, and then feeds it into optional crest factor reduction (CFR) and digital predistortion (DPD) blocks. Intel assumes you have a control processor that configures modulation scheme and CFR and DPD parameters.

The channel model contains a random noise source, and a channel model, which you can configure through the setup script. This channel model allows you to build a hardware demonstrator on a standard FPGA development platform, without DA or AD converters and analogue components. Following the channel model is the model of a decimating ADC, which emulates the behavior of some existing ADC components that provide this functionality.

The receiver contains an RRC filter, followed by an equalizer. Intel assumes that a control processor calculates the equalizer coefficients. The equalizer feeds into an AGC block, which feeds into a demapper. You can configure the demapper to different modulation orders.

The model file is `tx_ch_rx.mdl`

### 6.13.28. Variable Integer Rate Decimation Filter

The variable integer rate decimation filter reference design implements a 16-channel interpolate-by-2 symmetrical 49-tap FIR filter. The target system clock frequency is 320 MHz.

You can modify the parameters in the `setup_vardecimator_rt.m` file, which you access from the **Edit Params** icon.

The model file is `vardecimator_rt.mdl`.

## 6.14. DSP Builder Waveform Synthesis Design Examples

This folder contains design examples that synthesize waveforms with a NCO or direct digital synthesis (DDS).

1. [Complex Mixer](#) on page 171
2. [Four Channel, Two Banks NCO](#) on page 171
3. [Four Channel, Four Banks NCO](#) on page 173
4. [Four Channel, Eight Banks, Two Wires NCO](#) on page 173
5. [Four Channel, 16 Banks NCO](#) on page 174
6. [IP](#) on page 175
7. [NCO](#) on page 175
8. [NCO with Exposed Bus](#) on page 175
9. [Real Mixer](#) on page 175
10. [Super-sample NCO](#) on page 176

### 6.14.1. Complex Mixer

This design example shows how to mix complex signals.

The top-level testbench includes **Control**, **Signals**, **Run ModelSim**, and **Run Quartus Prime** blocks, plus **ChanView** block that deserialize the output buses. An **Edit Params** block allows easy access to the setup variables in the `setup_demo_complex_mixer.m` script.

The **FilterSystem** subsystem includes the **Device** and **ComplexMixer** blocks.

The model file is `demo_complex_mixer.mdl`.

*Note:* This design example uses the Simulink Signal Processing Blockset.

### 6.14.2. Four Channel, Two Banks NCO

This design example implements an NCO with four channels and two banks.

This design example demonstrates frequency-hopping with the NCO block to generate four channels of sinusoidal waves that you can switch from one set (bank) of frequencies to another.

The phase increment values are set directly into the **NCO Parameter** dialog box as a 2 (rows) × 4 (columns) matrix. The input for the bank index is set up so that it alternates between the two predefined banks with each one lasting 2000 steps.

A **BusStimulus** block sets up an Avalon-MM interface that writes into the phase increment memory registers. It shows how you can use the Avalon-MM interface to dynamically change the frequencies of the NCO-generated sinusoidal signals at run time. This design example uses a 16-bit memory interface (as the **Control** block specifies) and a 24-bit the accumulator in the **NCO** block. The design example requires two registers for each phase increment value. With the base address of the phase increment memory map set to 1000 in this design example, the addresses [1000 1001 1002 1003 1012 1013 1014 1015] write to the phase increment memory registers of channels 1 and 2 in bank 1, and to the registers of channels 3 and 4 in bank 2. The write data is also made up of two parts with each part writing to one of the registers feeding the selected phase increment accumulators.

This design example has two banks of frequencies with each bank processes 2,000 steps before switching to the other. You should write a new value into the phase increment memory register for each bank to change the NCO output frequencies after 8,000 steps during simulation. To avoid writing new values to the active bank, the design example configures the write enable signals in the following way:

```
[zeros(1,7000) 1 1 1 1 zeros(1,2000) 1 1 1 1 zeros(1,8000)]
```

This configuration ensures that a new phase increment value for bank 0 is written at 7000 steps when the NCO is processing bank 1; and a new phase increment value for bank 1 is written at 9000 steps when the NCO is processing bank 0.

Four writes for each bank exist to write new values for channel 1 and 2 into bank 0, and new values for channel 3 and 4 into bank 1. Each new phase value needs two registers due to the size of the memory interface.

The **Spectrum Scope** block shows three peaks for a selected channel with the first two peaks representing the two banks and the third peak showing the frequency that you specify through the memory interface. The scope of the select channel shows the sinusoidal waves of the channel you select. You can zoom in to see the smooth and continuous sinusoidal signals at the switching point. You can also see the frequency changes after 8000 steps where the phase increment value alters through the memory interface.

The top-level testbench includes **Control**, **Signals**, **BusStimulus**, **Run ModelSim**, and **Run Quartus Prime** blocks, plus **ChanView** blocks that deserialize the output buses. An **Edit Params** block allows easy access to the setup variables in the `setup_demo_mc_nco_2banks_mem_interface.m` script.

The **NCOSubSystem** subsystem includes the **Device** and **NCO** blocks.

The model file is `demo_mc_nco_2banks_mem_interface.mdl`.

**Note:** This design example uses the Simulink Signal Processing Blockset.

### 6.14.3. Four Channel, Four Banks NCO

This design example implements a NCO with four channels and four banks.

This design example is similar to the Four Channel, Two Banks NCO design, but it has four banks of frequencies defined for the phase increment values. Each spectrum plot has five peaks: the fifth peak shows the changes the design example writes through the memory interface.

The design example uses a 32-bit memory interface with a 24-bit accumulator. Hence, the design example requires only one phase increment memory register for each phase increment value—refer to the address and data setup on the **BusStimulus** block inside this design example.

This design example has four banks of frequencies with each bank processed for 2,000 steps before switching to the other. You should write a new value into the phase increment memory register for each bank to change the NCO output frequencies after 16,000 steps during simulation. To avoid writing new values to the active bank, the design example configures the write enable signals in the following way:

```
[zeros(1,15000) 1 zeros(1,2000) 1 zeros(1,2000) 1 zeros(1,2000) 1 zeros(1,8000)]
```

This configuration ensures that a new phase increment value for bank 0 is written at 15000 steps when the NCO is processing bank 3; a new phase increment value for bank 1 is written at 17000 steps when the NCO is processing bank 0; a new phase increment value for bank 2 is written at 19000 steps when the NCO is processing bank 1; and a new phase increment value for bank 3 is written at 21000 steps when the NCO is processing bank 2.

There is one write for each bank to write a new value for channel 1 into bank 0; a new value for channel 2 into bank 1; a new value for channel 3 into bank 2; and a new value for channel 4 into bank 3. Each new phase value needs only one register due to the size of the memory interface.

The top-level testbench includes **Control**, **Signals**, **BusStimulus**, **Run ModelSim**, and **Run Quartus Prime** blocks, plus **ChanView** blocks that deserialize the output buses. An **Edit Params** block allows easy access to the setup variables in the `setup_demo_mc_nco_4banks_mem_interface.m` script.

The **NCOSubSystem** subsystem includes the **Device** and **NCO** blocks.

The model file is `demo_mc_nco_4banks_mem_interface.mdl`.

*Note:* This design example uses the Simulink Signal Processing Blockset.

### 6.14.4. Four Channel, Eight Banks, Two Wires NCO

This design example implements a NCO with four channels and eight banks.

This design example is similar to the Four Channel, 16 Banks NCO design, but has only eight banks of phase increment values (specified in the setup script for the workspace variable) feeding into the NCO. Furthermore, the sample time for the NCO requires two wires to output the four channels of the sinusoidal signals. Two wires exist for the NCO output, each wire only contains two channels. Hence, the channel indicator is from 0 .. 3 to 0 .. 1.

You can inspect the eight peaks on the spectrum graph for each channel and see the smooth continuous sinusoidal waves on the scope display.

This design example uses an additional subsystem (**Select\_bank\_out**) to extract the NCO-generated sinusoidal signal of a selected bank on a channel.

The design example outputs the data to the workspace and plots through with the separate `demo_mc_nco_extracted_waves.mdl`, which demonstrates that the output of the bank you select does represent a genuine sinusoidal wave. However, from the scope display, you can see that the sinusoidal wave is no longer smooth at the switching point, because the design example uses the different values of phase increment values between the selected banks. You can only run the **demo\_mc\_nco\_extracted\_waves.mdl** model after you run `demo_mc_nco_8banks_2wires.mdl`.

The top-level testbench includes **Control**, **Signals**, **BusStimulus**, **Run ModelSim**, and **Run Quartus Prime** blocks, plus **ChanView** blocks that deserialize the output buses. An **Edit Params** block allows easy access to the setup variables in the `setup_demo_mc_nco_8banks_2wires.m` script.

The **NCOSubSystem** subsystem includes the **Device** and **NCO** blocks.

The **Select\_bank\_out** subsystem contains **Const**, **CompareEquality**, and **AND Gate** blocks.

The model file is `demo_mc_nco_8banks_2wires.mdl`.

*Note:* This design example uses the Simulink Signal Processing Blockset.

#### 6.14.5. Four Channel, 16 Banks NCO

This design example implements a NCO with four channels and 16 banks. This design example demonstrates frequency-hopping with the **NCO** block to generate 4 channels of sinusoidal waves, which you can switch from one set (bank) of frequencies to another in the 16 predefined frequency sets.

A workspace variable `phaseIncr` defines the 16 (rows) × 4 (columns) matrix for the phase increment input with the phase increment values that the setup script calculates.

The input for the bank index is set up so that it cycles from 0 to 15 with each bank lasting 1200 steps.

The spectrum display shows clearly 16 peaks for the selected channel indicating that the design example generates 16 different frequencies for that channel. The scope of the selected channel shows the sinusoidal waves of the selected channel. You can zoom in to see that the design example generates smooth and continuous sinusoidal signals at the switching point.

The top-level testbench includes **Control**, **Signals**, **Run ModelSim**, and **Run Quartus Prime** blocks, plus **ChanView** blocks that deserialize the output buses. An **Edit Params** block allows easy access to the setup variables in the `setup_demo_mc_nco_16banks.m` script.

The **NCOSubSystem** subsystem includes the **Device** and **NCO** blocks.

The model file is `demo_mc_nco_16banks.mdl`.

*Note:* This design example uses the Simulink Signal Processing Blockset.

### 6.14.6. IP

The IP design example describes how you can build a NCO design with the **NCO** block from the Waveform Synthesis library.

*Note:* This design example uses the Simulink Signal Processing Blockset.

### 6.14.7. NCO

This design example uses the **NCO** block from the Waveform Synthesis library to implement an NCO. A Simulink double precision sine or cosine wave compares the results.

The top-level testbench includes **Control**, **Signals**, **Run ModelSim**, and **Run Quartus Prime** blocks, plus **ChanView** blocks that deserialize the output buses. An **Edit Params** block allows easy access to the setup variables in the `setup_demo_nco.m` script.

The **NCOSubSystem** subsystem includes the **Device** and **NCO** blocks.

The model file is `demo_nco.mdl`.

*Note:* This design example uses the Simulink Signal Processing Blockset.

#### Related Information

[NCO on page 248](#)

### 6.14.8. NCO with Exposed Bus

This design example is a multichannel NCO that outputs four waveforms with slightly different frequencies. Halfway through the simulation, DSP Builder reconfigures the NCO for smaller increments, which gives a waveform with a longer period.

The model file is `demo_nco_exposed_bus.mdl`.

### 6.14.9. Real Mixer

This design example shows how to mix non-complex signals.

The top-level testbench includes **Control**, **Signals**, **Run ModelSim**, and **Run Quartus Prime** blocks, plus **ChanView** block that deserialize the output buses. An **Edit Params** block allows easy access to the setup variables in the `setup_demo_mix.m` script.

The **MixerSystem** subsystem includes the **Device** and **Mixer** blocks.

The model file is `demo_mix.mdl`.

*Note:* This design example uses the Simulink Signal Processing Blockset.

### 6.14.10. Super-sample NCO

This design example uses the **NCO** block from the Waveform Synthesis library to implement a super-sample NCO. The design demonstrates run-time reconfiguring of the frequency using a register bus.

A super-sample NCO uses multiple NCOs that each have an initial phase offset. When you combine the parallel outputs into a serial stream, they can describe frequencies  $N$  times the Nyquist frequency of a single NCO. Where  $N$  is the total number of NCOs that the design uses.

The NCO block produces four outputs, which all have the same phase increment but each have a different, evenly distributed initial phase offset. With the four parallel outputs in series they describe frequencies up to four times higher than the Nyquist frequency of an individual NCO.

To change the frequency of the super-sample NCO using the bus, write a new phase increment and offset to each of the four constituent NCOs and then strobe the synchronization register. The NCO block includes the phase increment register; a separate primitive subsystem implements the phase offset and synchronization registers.

The `setup_demo_nco_super_sample` scripts allows you to configure the clock rate, number of NCOs, NCO accumulator size, and many other parameters. This script calculates the required phase increment and offsets required to sweep the super-sample NCO through five frequencies. The script defines the memory map and creates the bus stimulus.

DSP Builder writes the output of the super-sample NCO into a MATLAB workspace variable and compares it with a MATLAB-generated waveform in the script `test_demo_nco_super_sample`.

DSP Builder schedules the bus in HDL but not in Simulink, so bus writes occur at different clock cycles. Therefore, the function `verify_demo_nco_super_sample` function verifies the design, which checks that the Simulink and ModelSim frequency distributions match within a tolerance.

The output of the **Spectrum Analyser** block show the simulation initializes to the last frequency in `dspb_super_nco.frequencies` and then rotates through the list.

The model file is `demo_nco_super_sample.mdl`.

**Note:** This design example uses the Simulink Signal Processing Blockset.



## 7. DSP Builder Design Rules, Design Recommendations, and Troubleshooting

---

1. [DSP Builder Design Rules and Recommendations](#) on page 177
2. [Troubleshooting DSP Builder Designs](#) on page 179

### 7.1. DSP Builder Design Rules and Recommendations

Use the design rules and recommendations to ensure your design performs correctly.

#### Design Rules for the Top-Level Design

- Ensure the top-level design has a **Control** block and a **Signals** block.
- Ensure the synthesizable part of your design is a subsystem or contained within a subsystem of the top-level design.
- Ensure testbench stimulus data types that feed into the synthesizable design are correct, as DSP Builder propagates them.
- Ensure you place **Interface ► ExternalMemory** and **Avalon Memory-MappedSettings** blocks only in the top-level design

#### Design Rules for the Synthesized Top-Level Design

- Ensure your synthesized hardware top-level subsystem has a **Device** block.
- Ensure you place some non-synthesizable blocks (from the **Interface ► MemoryMapped ► Stimulus** and **Utilities ► Testbench** libraries outside the synthesized system.

#### Design Rules for the Primitive Top-Level Design

- Ensure the primitive top-level subsystem contain a **SynthesisInfo** block with style set to **Scheduled**.
- Ensure the Primitive subsystems do not contain IP blocks.
- Only use primitive blocks in primitive subsystems and delimit them by primitive boundary blocks.
- If using ALU folding, ensure the **ALU Folding** block is in the primitive top-level subsystem.
- Route all subsystem inputs with associated `valid` and `channel` signals that are to be scheduled together through the same **ChannelIn** blocks immediately following the subsystem inputs. Route any other subsystem inputs through **GPIn** blocks.
- Route all subsystem outputs with associated `valid` and `channel` signals that are to be scheduled together through the same **ChannelOut** blocks immediately before the subsystem outputs. Route any other subsystem outputs through **GPOut** blocks.

- Ensure all primitive subsystem input boundary blocks (**GPIn** or **ChannelIn**) or output boundary blocks (**GPOut** or **ChannelOut**) are in primitive top-level subsystem.

*Note:* Also Avalon-MM interface blocks can be subsystem schedule boundaries

- Ensure the `valid` signal is a scalar Boolean signal or `ufix(1)`.
- Ensure the `channel` signal is a scalar `uint(8)`

### Design Rules for Avalon-MM Interface Blocks

- Place shared memory blocks inside primitive scheduled subsystem.
- Ensure the **RegField** and **RegBit** blocks output type width exactly match the range you specify for these blocks through MSB and LSB parameters.
- Ensure the specified ranges through MSB and LSB parameters fit within Avalon-MM word width set from **Avalon Interfaces** > **Avalon Memory-Mapped Settings**.
- Ensure different instances of register blocks (**RegBit**, **RegField**, or **RegOut**) that map to the same Avalon-MM address specify disjoint ranges.
- For shared memory blocks, ensure output data width matches or is twice the size of Avalon-MM data width set from **Avalon Interfaces** > **Avalon Memory-Mapped Settings**.
- Locate the **BusStimulus** and **BusStimulusFileReader** blocks in the testbench, which is outside the synthesizable system.

### Recommendations for your Top-Level Design

- Create a Simulink project for your model file, libraries, and scripts.
- Use workspace variables to set parameters, which allows you to globally organize and change them.
- Use set-up scripts to set the workspace variables and clear-up scripts to clear them from the workspace afterwards.
- Run set-up, analysis, and clear-up scripts automatically by adding them to the model callbacks.
- Build a testbench that is parameterizable with system parameters such as sample rate, clock rate, and number of channels. Use the **Channelizer** block to create data in the valid-channel-data protocol.
- Hierarchically structure your design into subsystems. A modular design with well-defined subsystem boundaries allows you to precisely manage latency and speed of different modules and achieve timing closure.
- Save repeated subsystems as library blocks. Replace the design blocks with copies from the library.
- Make library blocks configurable and self-modifying.
- Create and use your own libraries of reusable components. Organize them into separate library files.
- Use configurable subsystem blocks in libraries to switch implementations in place.

- Build separate testbenches for library blocks
- Keep block and subsystem names short, but descriptive. Do not use names with special characters, slashes, or that begin with numbers.
- Use vectors to build parameterizable designs. DSP Builder does not need to redraw them when parameters such as the number of channels change. A design that uses a vector input of width  $N$  is the same as connecting  $N$  copies of the block with a single scalar connection to each.

### Recommendations for Loops in Primitive Subsystems

- Ensure sufficient sample delays (**SampleDelay** blocks) exist around loops to allow for pipelining
- To determine the minimum loop latency, turn on **Minimum Delay** on the **SampleDelay** block
- Simulink performs data type, complexity, and vector width propagation. Sometimes Simulink does not successfully resolve propagation around loops, particularly multiple nested loops.
- If Simulink is unsuccessful, look for where data types are not annotated.
- You may have to explicitly set data types. Simulink provides a library of blocks to help in such situations, which duplicate data types. For example, the data type prop duplicate block, **fixpt\_dtprop**, (type `open fixpt_dtprop` from the MATLAB command prompt), which the control library latches use.
- Avoid primitive subsystems with logic that clocked inputs do not drive, because either reset behavior determines hardware behavior or the hardware is inefficient.
- Avoid examples that start from reset, as the design simulation in Simulink may not match that of the generated hardware. You should start a counter from the valid signal, rather than the constant. If the counter repeats without stopping after the first valid, add a zero-latency latch into this connection.
- Avoid loops that DSP Builder drives without clocked inputs.

### Related Information

- [Control](#) on page 223
- [Avalon Memory-Mapped Agent Settings \(AvalonMemoryMappedAgentSettings\)](#) on page 221
- [External Memory, Memory Read, Memory Write](#) on page 265
- [Channel In \(ChannelIn\)](#) on page 359
- [Channel Out \(ChannelOut\)](#) on page 360
- [Synthesis Information \(SynthesisInfo\)](#) on page 362
- [Setting DSP Builder Design Parameters with MATLAB Scripts](#) on page 184

## 7.2. Troubleshooting DSP Builder Designs

You might see errors when you build, test, update, simulate, or verify your DSP Builder design.

1. Check your design construction:

- Follow the recommendations for structuring and managing your model.
  - Follow the Simulink setup guidelines.
  - Follow the design rules
  - Follow the rules for **Primitive** and **IP** library blocks and specific blocks like **SampleDelays** blocks
2. Check for common Simulink errors including algebraic loops and unresolved data types.
  3. Ensure your DSP Builder does not use **Primitive** library blocks in unsupported modes – either outside of primitive subsystems or in loops without sufficient start to end of loop timing offset.
  4. Read DSP Builder error messages to see the root cause.
  5. Click **DSP Builder** ► **Design Checker**, to check your design for common mistakes.
  6. Select individual steps and click **Check**.

The output only matches the hardware when `valid` is high.

If your design uses FIFO buffers within multiple feedback loops, while the data throughput and frequency of invalid cycles is the same, their distribution over a frame of data might vary (because of the final distribution of delays around the loop). If you find a mismatch, step past errors.

1. [About Loops](#) on page 180
2. [DSP Builder Timed Feedback Loops](#) on page 181
3. [DSP Builder Loops, Clock Cycles, and Data Cycles](#) on page 182

### Related Information

[DSP Builder Design Rules and Recommendations](#) on page 177

## 7.2.1. About Loops

Your design can contain many loops that can interact with or be nested inside each other. DSP Builder uses standard mathematical linear programming techniques to solve a set of simultaneous timing constraints.

Consider the following two main cases:

- The simpler case is feed-forward. When no loops exist, feed-forward datapaths are balanced to ensure that all the input data reaches each functional unit in the same cycle. After analysis, DSP Builder inserts delays on all the non-critical paths to balance out the delays on the critical path.
- The case with loops is more complex. Loops cannot be combinational—all loops in the Simulink design must include delay memory. Otherwise Simulink displays an 'algebraic loop' error. In hardware, the signal has to have a specified number of clock cycles latency round the feedback loop. Typically, one or more lumped delays exist with **SampleDelay** blocks specifying the latency around some or all of the loop. DSP Builder preserves the latency around the loop to maintain correct functional operation. To achieve the target clock frequency, the total delay of the sum of **SampleDelay** blocks around the loop must be greater or equal to the required pipelining.

If the pipelining requirements of the functional units around the loop are greater than the delay specified by the **SampleDelay** blocks on the loop path, DSP Builder generates an error message. The message states that distribution of memory failed as there was insufficient delay to satisfy the  $f_{MAX}$  requirement. DSP Builder cannot simultaneously satisfy the pipelining to achieve the given  $f_{MAX}$  and the loop criteria to re-circulate the data in the number of clock cycles specified by the **SampleDelay** blocks.

DSP Builder automatically adjusts the pipeline requirements of every **Primitive** block according to these factors

- The type of block
- The target  $f_{MAX}$
- The device family and speedgrade
- The inputs of inputs
- The bit width in the data inputs

**Note:** DSP Builder implements multipliers on DSP blocks. The DSP block latency affects these multipliers and varies with target device family and clock-frequency. Very wide fixed-point multipliers incur higher latency when DSP Builder splits them into smaller multipliers and adders. You cannot count the multiplier and adder latencies separately because DSP Builder may combine them into a single DSP block. The latency of some blocks depends on what pipelining you apply to surrounding blocks. DSP Builder avoids pipelining every block but inserts pipeline stages after every few blocks in a long sequence of logical components, if  $f_{MAX}$  is sufficiently low that timing closure is still achievable.

In the **SynthesisInfo** block, you can optionally specify a latency constraint limit that can be a workspace variable or expression, but must evaluate to a positive integer. However, only use this feature to add further latency. Never use the feature to reduce latency to less than the latency required to pipeline the design to achieve the target  $f_{MAX}$ .

After you run a simulation in Simulink, the help page for the **SynthesisInfo** block shows the latency, port interface, and estimated resource utilization for the current **Primitive** subsystem.

When no loops exist, feed-forward datapaths are balanced to ensure that all the input data reaches each functional unit in the same cycle. After analysis, DSP Builder inserts delays on all the non-critical paths to balance out the delays on the critical path.

In designs with loops, DSP Builder advanced blockset must synthesize at least one cycle of delay in every feedback loop to avoid combinational loops that Simulink cannot simulate. Typically, one or more lumped delays exist. To preserve the delay around the loop for correct operation, the functional units that need more pipelining stages borrow from the lumped delay.

### 7.2.2. DSP Builder Timed Feedback Loops

Take care with feedback loops generally, in particular provide sufficient delay around the loop.

Designs that have a cycle containing two adders with only a single sample delay are not sufficient. In automatically pipelining designs, DSP Builder creates a schedule of signals through the design. From internal timing models, DSP Builder calculates how

fast certain components, such as wide adders, can run and how many pipelining stages they require to run at a specific clock frequency. DSP Builder must account for the required pipelining while not changing the order of the schedule. The single sample delay is not enough to pipeline the path through the two adders at the specific clock frequency. DSP Builder is not free to insert more pipelining, as it changes the algorithm, accumulating every  $n$  cycles, rather than every cycle. The scheduler detects this change and gives an appropriate error indicating how much more latency the loop requires for it to run at the specific clock rate. In multiple loops, this error may be hit a few times in a row as DSP Builder balances and resolves each loop.

### 7.2.3. DSP Builder Loops, Clock Cycles, and Data Cycles

Never confuse clock cycles and data cycles in relation to feedback loops. For example, you may want to accumulate previous data from the same channel. The DSP Builder multichannel IIR filter design example (**demo\_iir**) shows feedback accumulators processing multiple channels. In this example, consecutive data samples on any particular channel are 20 clock cycles apart. DSP Builder derives this number from clock rate and sample rate.

The folded IIR filter design example (**demo\_iir\_fold2**) demonstrates one channel, at a low data rate. This design example implements a single-channel infinite impulse response (IIR) filter with a subsystem built from Primitive blocks folded down to a serial implementation.

The design of the IIR is the same as the IIR in the multichannel example, **demo\_iir**. As the channel count is one, the lumped delays in the feedback loops are all one. If you run the design at full speed, there is a scheduling problem. With new data arriving every clock cycle, the lumped delay of one cycle is not enough to allow for pipelining around the loops. However, the data arrives at a much slower rate than the clock rate, in this example 32 times slower (the clock rate in the design is 320 MHz, and the sample rate is 10 MHz), which gives 32 clock cycles between each sample.

You can set the lumped delays to 32 cycles long—the gap between successive data samples—which is inefficient both in terms of register use and in underused multipliers and adders. Instead, use folding to schedule the data through a minimum set of fully used hardware.

Set the **SampleRate** on both the **ChannelIn** and **ChannelOut** blocks to 10 MHz, to inform the synthesis for the **Primitive** subsystem of the schedule of data through the design. Even though the clock rate is 320 MHz, each data sample per channel is arriving only at 10 MHz. The RTL is folded down—in multiplier use—at the expense of extra logic for signal multiplexing and extra latency.

## 8. About DSP Builder for Intel FPGAs Optimization

---

Improve your designs and learn about folding and floating-point data types.

1. [Associating DSP Builder with MATLAB](#) on page 183
2. [Setting Up Simulink for DSP Builder Designs](#) on page 183
3. [The DSP Builder Windows Shortcut](#) on page 184
4. [Setting DSP Builder Design Parameters with MATLAB Scripts](#) on page 184
5. [Managing your Designs](#) on page 187
6. [How to Manage Latency](#) on page 188
7. [Flow Control in DSP Builder Designs](#) on page 195
8. [Reset Minimization](#) on page 197
9. [About Importing HDL](#) on page 199

### 8.1. Associating DSP Builder with MATLAB

If you install another version of MATLAB or you install DSP Builder without associating it with a version of MATLAB, you can associate DSP Builder with MATLAB

1. Type the following command into a command window:  

```
<path to dsp_builder.bat> -m "<path to matlab executable>" For example:c:\intel_fpga_pro\quartus\dspba\dsp_builder.bat -m "c:\tools\matlab\R2015a\windows64\bin\matlab.exe"
```

### 8.2. Setting Up Simulink for DSP Builder Designs

1. [Setting Up Simulink Solver](#) on page 183
2. [Setting Up Simulink Signal Display Option](#) on page 184

#### 8.2.1. Setting Up Simulink Solver

1. On the File menu, click **Preferences**.
2. Expand **Configuration defaults** and click **Solver**.
3. For Type, select **Fixed-step solver**, unless you have folding turned on in some part of your design. In that case, you need to select **Variable-step solver**.
4. For Solver select **Discrete (no continuous states)**.
5. Click on Display Defaults and turn on **Show port data types**.

### 8.2.2. Setting Up Simulink Signal Display Option

Display various port and signal properties to aid debugging and visualization.

1. In Simulink, click **Format > Port/Signal Displays**.
2. Click **Sample Time Colors** to change the color of blocks and wires in particular clock domain—useful when creating multirate designs.
3. Click **Port Data Type** option to display the data type of the blocks. You can only connect ports of same data type.
4. Click **Signal Dimensions** to display the dimensions of particular signal wire.
5. Make show data types and wide non-scalar lines the default for new models:
  - a. Click **File > Preferences**.
  - a. Select **Display Defaults for New Models**
  - b. Turn on **Wide nonscalar lines** and **Show port data types**.

### 8.3. The DSP Builder Windows Shortcut

Create a shortcut to set the file paths to DSP Builder and run a batch file with an argument for the MATLAB executable to use

The shortcut target is:

```
<dsp_builder.bat from the DSP Builder release to use> -m "<path to the MATLAB executable to use>"
```

For example

```
C:\Altera\16.0\quartus\dspba\dsp_builder.bat -m "C:\tools\matlab\R2013a\windows64\bin\matlab.exe"
```

You can copy the shortcut from the **Start** menu and paste it to your desktop to create a desktop shortcut. You can edit the properties to use different installed DSP Builder releases, different MATLAB releases, or different start directories.

#### Related Information

[Starting DSP Builder in MATLAB](#)

### 8.4. Setting DSP Builder Design Parameters with MATLAB Scripts

1. Set block and design parameters using MATLAB workspace variables with names unique to your model.
2. Define the MATLAB workspace variables in a MATLAB script or set of scripts, where you can manage them.
3. Run the scripts run automatically on opening the model and again before simulation.  
DSP Builder evaluates and updates all parameters before it generates hardware.
4. Clean up the workspace using a separate script when you close the design.



1. [Running Setup Scripts Automatically](#) on page 185
2. [Defining Unique DSP Builder Design Parameters](#) on page 185
3. [Example DSP Builder Custom Scripts](#) on page 185

### 8.4.1. Running Setup Scripts Automatically

1. In a Simulink model file **.mdl**, click **File > Model properties**.
2. Select **Callbacks** tab.
3. Select **PreLoadFcn** and type the setup script name in the window on the right hand side. When you open your Simulink design file, the setup script runs.
4. Select **InitFcn** and type the setup script name in the window on the right hand side. Simulink runs your setup script first at the start of each simulation before it evaluates the model design file **.mdl**.

### 8.4.2. Defining Unique DSP Builder Design Parameters

Define unique parameters to avoid parameters clashing with other open designs and to help clear the workspace.

1. Create named structures and append a common root to all parameter names.

For example;

```
my_design_params.clockrate = 200;  
my_design_params.samplerate = 50;  
my_design_params.inputChannels = 4;
```

2. Clear the specific workspace variables you create with a clear-up script that run when you close the model. Do not use `clear all`.

For example, if you use the named structure `my_design_params`, run `clear my_design_params`. You may have other temporary workspace variables to clear too.

### 8.4.3. Example DSP Builder Custom Scripts

You can write scripts that directly change parameters (such as the hardware destination directory) on the **Control** and **Signals** blocks.

For example, in a script that passes the design name (without **.mdl** extension) as model you can use:

```
%% Load the model  
load_system(model);  
%% Get the Signals block  
signals = find_system(model, 'type', 'block', 'MaskType', 'DSP Builder Advanced  
Blockset Signals Block');  
if (isempty(signals))  
    error('The design must contain a Signals Block. ');  
end;  
%% Get the Controls block  
control = find_system(model, 'type', 'block', 'MaskType', 'DSP Builder Advanced  
Blockset Control Block');  
if (isempty(control))  
    error('The design must contain a Control Block. ');  
end;%%
```

```
Example: set the RTL destination directory
dest_dir = ['./rtl' num2str(freq)];
dspba.SetRTLDestDir(model, rtlDir);
```

Similarly you can get and set other parameters. For example, on the **Signals** block you can set the target clock frequency:

```
fmax_freq = 300.0; dspba.set_param(signals{1}, 'freq', fmax_freq);
```

You can also change the following threshold values that are parameters on the **Control** block:

- **distRamThresholdBits**
- **hardMultiplierThresholdLuts**
- **mlabThresholdBits**
- **ramThresholdBits**

You can loop over changing these values, change the destination directory, run the Quartus Prime software each time, and perform design space exploration. For example:

```
%% Run a simulation; which also does the RTL generation.
t = sim(model);
%% Then run the Quartus Prime compilation flow.
[success, details] = run_hw_compilation(<model>, './')%%
where details is a struct containing resource and timing information
    details.Logic,
    details.Comb_Aluts,
    details.Mem_Aluts,
    details.Reggs,
    details.ALM,
    details.DSP_18bit,
    details.Mem_Bits,
    details.M9K,
    details.M144K,
    details.IO,
    details.FMax,
    details.Slack,
    details.Required,
    details.FMax_unres,
    details.timingpath,
    details.dir,
    details.command,
    details.pwd
such that >> disp(details) gives output something like:
Logic: 4915
Comb_Aluts: 3213
Mem_Aluts: 377
Reggs: 4725
ALM: 2952
DSP_18bit: 68
Mem_Bits: 719278
M9K: 97
M144K: 0 IO: 116
FMax: 220.1700
Slack: 0.4581
Required: 200
FMax_unres: 220.1700
timingpath: [1x4146 char]
dir: './quartus_demo_iff_t_4096_for_SPR_FFT_4K_n_2'
command: [1x266 char]
pwd: 'D:\test\script'
```

**Note:**

The Timing Report is in the timingpath variable, which you can display by **disp(details.timingpath)**. Unused resources may appear as -1, rather than 0.

You must previously execute **load\_system** before commands such as **find\_system** and **run\_hw\_compilation** work.

A useful set of commands to generate RTL, compile in the Quartus Prime software and return the details is:

```
load_system(<model>);  
sim(<model>);  
[success, details] = run_hw_compilation(<model>, './')
```

## 8.5. Managing your Designs

DSP Builder supports parameterization through scripting.

1. To define many of DSP Builder advanced blockset parameters as a MATLAB workspace variables, such as clock rate, sample rate, and bit width, define these variables in a **.m** file.
2. Run this setup script before running your design.
3. Explore different values for various parameters, without having to modify the Simulink design.  
  
For instance, you can evaluate the performance impact of varying bit width at different stages of your design.
4. Define the data type and width of **Primitive** library blocks in the script
5. Experiment with different values. DSP Builder advanced blockset vector signal and ALU folding support allows you to use the same design file to target single and multiple channels designs.
6. Use a script for device options in your setup script, which eases design migration, whether you are targeting a new device or you are upgrading the design to support more data channels.
7. Use advanced scripting to fine tune Quartus Prime settings and to build automatic test sweeping, including parameter changes and device changes.
  1. [Managing Basic Parameters](#) on page 187
  2. [Creating User Libraries and Converting a Primitive Subsystem into a Custom Block](#) on page 188
  3. [Revision Control](#) on page 188

### 8.5.1. Managing Basic Parameters

Before you start implementing your design, you should define key parameters in a script.

Based on the FPGA clock rate and data sample rates, you can derive how many clock cycles are available to process unique data samples. This parameter is called *Period* in many of the design examples. For example, for a period of three, a new sample for the same channel appears every three clock cycles. For multiplication, you have three clock cycles to compute one multiplication for this channel. In a design with multiple channels, you can accommodate three different channels with just one multiplier. A resource reuse potential exists when the period is greater than one.

1. Define the following parameters:
  - FPGA clock rate
  - Data sample rates at various stages of your design
  - Number of channels or data sources
  - Bit widths of signals at various stages of your design, including possible bit growth throughout the computational datapath
  - Coefficients of filters

### 8.5.2. Creating User Libraries and Converting a Primitive Subsystem into a Custom Block

You can group frequently used custom blocks into libraries for future use.

1. Mask the block to hide the block's contents and provide a custom block dialog.
2. Place the block in a library to prohibit modifications and allow you to easily update copies of the block.

*Note:* This procedure is similar to creating a Simulink custom block and custom library. You can also add a custom library to the Simulink library browser.

### 8.5.3. Revision Control

Use Simulink revision control to manage your DSP Builder advanced blockset design revision control.

The Simulink **Model Info** block displays revision control information about a model as an annotation block in the model's block diagram. It shows revision control information embedded in the model and information maintained by an external revision control or configuration management system.

You can customize some revision control tools to use the Simulink report generator XML comparison, which allows you to compare two versions of the same file.

You must add the following files to revision control:

- Your setup script (.m file)
- Model design files .mdl.
- All the customized library files.
- `_params.xml` file

*Note:* You do not need to archive autogenerated files such as Quartus Prime project files or synthesizable RTL files.

## 8.6. How to Manage Latency

The **Primitive** library blocks are untimed circuits, so they are not cycle accurate. A one-to-one mapping does not exist between the blocks in the Simulink model and the blocks you implement in your design in RTL. This decoupling of design intent from design implementation gives productivity benefits. The **ChannelOut** block is the boundary between the untimed section and the cycle accurate section. This block creates the additional delay that the RTL introduces, so that data going in to the

**ChannelOut** block delays internally, before DSP Builder presents it externally. The latency of the block shows on the **ChannelOut** mask. You may want to fix or constrain the latency after you complete part of a DSP Builder design, for example on an **IP** library block or for a **Primitive** subsystem. In other cases, you may want to limit the latency in advance, which allows future changes to other subsystems without causing undesirable effects upon the overall design.

To accommodate extra latency, insert registers. This feature applies only to **Primitive** subsystems. To access, use the **Synthesis Info** block.

Latency is the number of delays in the `valid` signal across the subsystem. The DSP Builder advanced blockset balances delays in the `valid` and `channel` path with delays that DSP Builder inserts for autopipelining in the datapath.

*Note:* User-inserted sample delays in the datapath are part of the algorithm, rather than pipelining, and are not balanced. However, any uniform delays that you insert across the entire datapath optimize out. If you want to constrain the latency across the entire datapath, you can specify this latency constraint in the **SynthesisInfo** block.

1. [Reading the Added Latency Value for an IP Block](#) on page 189
2. [Zero Latency Example](#) on page 189
3. [Implicit Delays in DSP Builder Designs](#) on page 190
4. [Distributed Delays in DSP Builder Designs](#) on page 191
5. [Latency and fMAX Constraint Conflicts in DSP Builder Designs](#) on page 193
6. [Control Units Delays](#) on page 193

### 8.6.1. Reading the Added Latency Value for an IP Block

1. Select the block and type the following command:

```
get_param(gcb, 'latency')
```

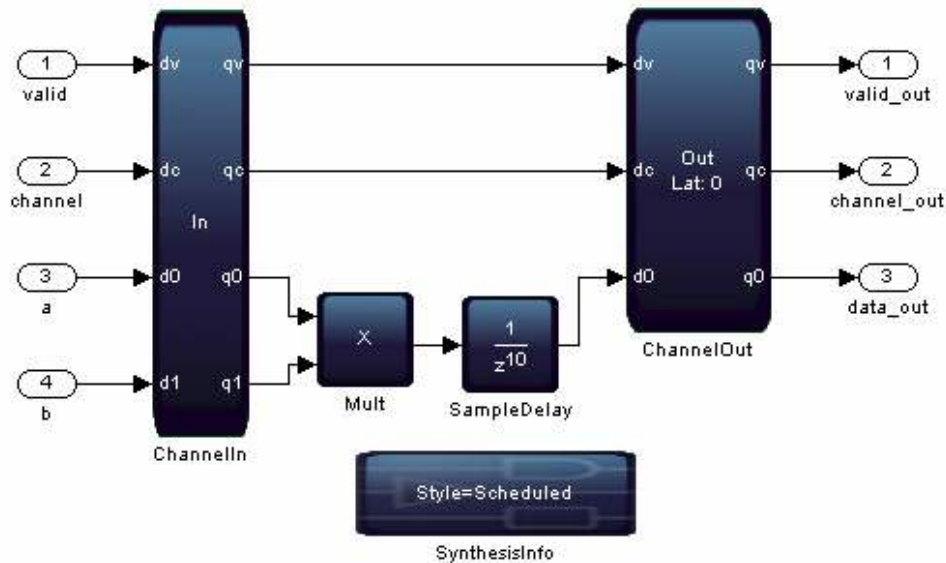
You can also use this command in an M-script. For example when you want to use the returned latency value to balance delays with external circuitry.

*Note:* If you use an M-script to get this parameter and set latency elsewhere in your design, by the time it updates and sets on the **IP** block, it is too late to initialize the delays elsewhere. You must run your design twice after any changes to make sure that you have the correct latency. If you are scripting the whole flow, you must run once with end time 0, and then run again immediately with the desired simulation end time.

### 8.6.2. Zero Latency Example

In this example, sufficient delays in the design ensure that DSP Builder requires no extra automatic pipelining to reach the  $f_{MAX}$  target (although DSP Builder distributes this user-added delay through the datapath). Thus, the reported latency is zero. DSP Builder inserts no extra pipelining registers in the datapath to meet  $f_{MAX}$  and thus inserts no balancing registers on the **channel** and **valid** paths. The delay of the **valid** signal across the subsystem is zero clock cycles, as the **Lat: 0** latency value on the **ChannelOut** block shows.

**Figure 74. Latency Example with a User-Specified Delay**



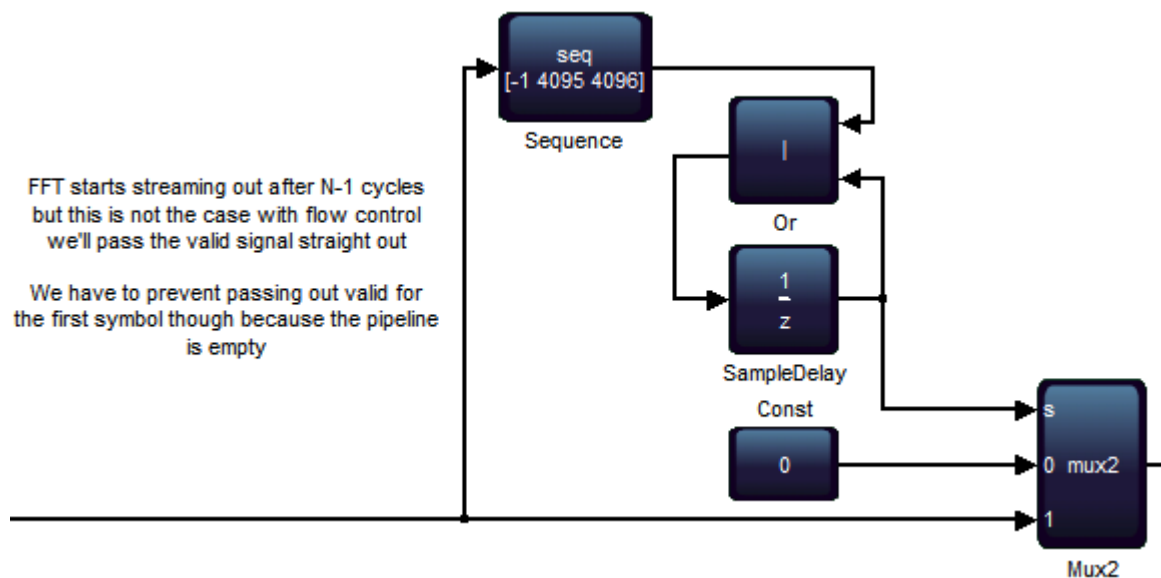
### 8.6.3. Implicit Delays in DSP Builder Designs

The DSP Builder scheduler may add extra delays on paths between the **ChannelIn** and **ChannelOut** blocks. The extra latency is the same for all such paths and is displayed on the **ChannelOut** block.

If the valid input drives directly the valid output, the delay on the valid signal matches the latency displayed on the **ChannelOut** block. It doesn't, if the valid output is generated in any other way, for example by using a Sequence block.

For example, the 4K FFT design example uses a **Sequence** block to drive the `valid` signal explicitly.

Figure 75. Sequence Block in the 4K FFT Design Example



The latency that the **ChannelOut** block reports is therefore not 4096 + the automatic pipelining value, but just the pipelining value.

#### 8.6.4. Distributed Delays in DSP Builder Designs

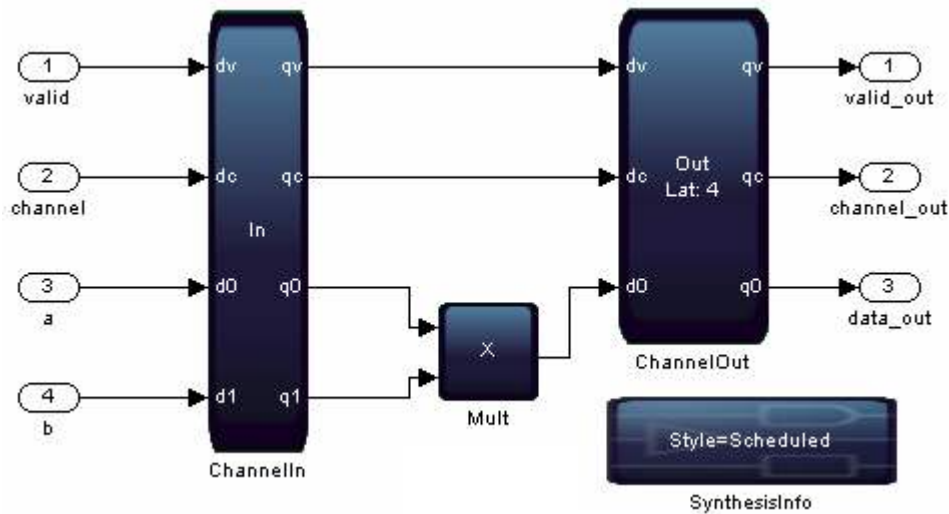
Distributed delays are not cycle-accurate inside a primitive subsystem, because DSP Builder distributes and optimizes the user-specified delay. To consistently apply extra latency to a primitive subsystem, use latency constraints.

In this example, the **Mult** block has a direct feed-through simulation model, and the following **SampleDelay** block has a delay of 10. The **Mult** block has zero delay in simulation, followed by a delay of 10. In the generated hardware, DSP Builder distributes part of this 10-stage pipelining throughout the multiplier optimally, such that the **Mult** block has a delay (in this case, four pipelining stages) and the **SampleDelay** block a delay (in this case, six pipelining stages). The overall result is the same—10 pipelining stages, but if you try to match signals in the primitive subsystem against hardware, you may find DSP Builder shifts them by several cycles.

Similarly, if you have insufficient user-inserted delay to meet the required  $f_{MAX}$ , DSP Builder automatically pipelines and balances the delays, and then corrects the cycle-accuracy of the primitive subsystem as a whole, by delaying the output signals in simulation by the appropriate number of cycles at the **ChannelOut** block.

If you specify no pipelining, the simulation design example for the multiplier is direct-feed-through, and the result appears on the output immediately.

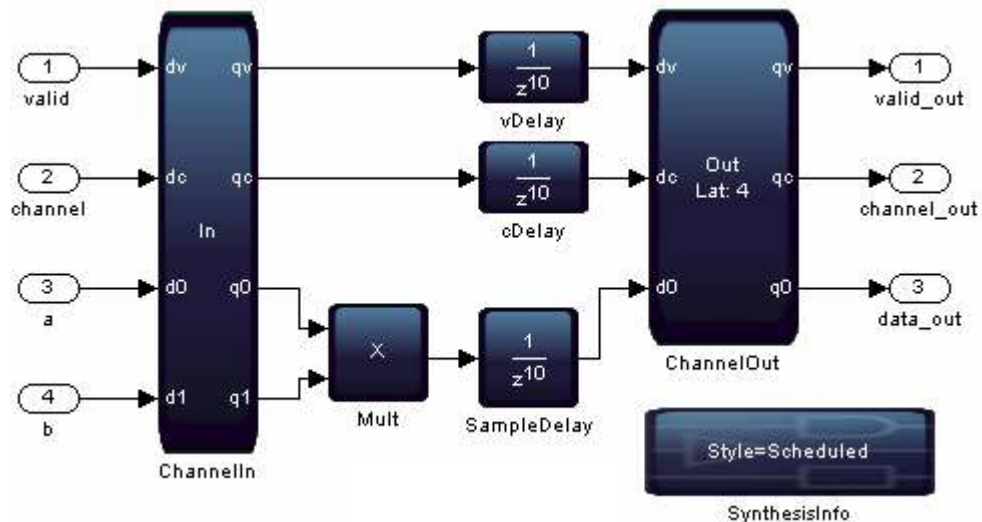
**Figure 76. Latency Example without a User-Specified Delay**



To reach the desired  $f_{MAX}$ , DSP Builder then inserts four pipelining stages in the multiplier, and balances these with four registers on the `channel` and `valid` paths. To correct the simulation design example to match hardware, the **ChannelOut** block delays the outputs by four cycles in simulation and displays **Lat: 4** on the block. Thus, if you compare the output of the multiplier simulation with the hardware it is now four cycles early in simulation; but if you compare the primitive subsystem outputs with hardware they match, because the **ChannelOut** block provides the simulation correction for the automatically inserted pipelining.

If you want a consistent 10 cycles of delay across the `valid`, `channel` and `datapath`, you may need latency constraints.

**Figure 77. Latency Example with Consistent Delays**





This example has a consistent line of **SampleDelay** blocks inserted across the design. However, the algorithm does not use these delays. DSP Builder recognizes that designs do not require them and optimizes them away, leaving only the delay that designs require. In this case, each block requires a delay of four, to balance the four delay stages to pipeline the multiplier sufficiently to reach the target  $f_{MAX}$ . The delay of 10 in simulation remains from the non-direct-feed-through **SampleDelay** blocks. In such cases, you receive the following warning on the MATLAB command line:

```
DSP Builder optimizes away some user inserted SampleDelays. The latency on the valid path across primitive subsystem design name in hardware is 4, which may differ from the simulation model. If you need to preserve extra SampleDelay blocks in this case, use the Constraint Latency option on the SynthesisInfo block.
```

**Note:** **SampleDelay** blocks reset to unknown values ('X'), not to zero. Designs that rely on **SampleDelays** output of zero after reset may not behave correctly in hardware. Use the `valid` signal to indicate valid data and its propagation through the design.

### 8.6.5. Latency and $f_{MAX}$ Constraint Conflicts in DSP Builder Designs

Some blocks need to have a minimum latency, either because of logical or silicon limitations. In these cases, you can create an abstracted design that cannot be realized in hardware. While these cases can generally be addressed, in some cases like IIRs, find algorithmic alternatives.

Generally, problems occur in feedback loops. You can solve these issues by lowering the  $f_{MAX}$  target, or by restructuring the feedback loop to reduce the combinatorial logic or increasing the delay. You can redesign some control structures that have feedback loops to make them completely feed forward.

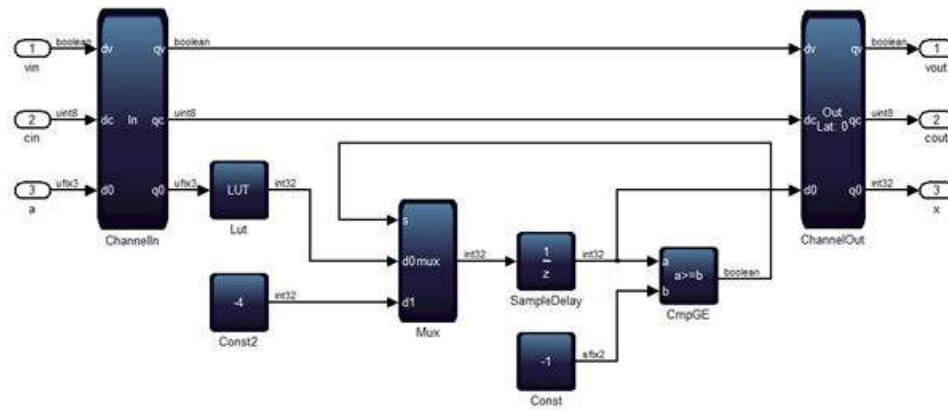
You cannot set a latency constraint that conflicts with the constraint that the  $f_{MAX}$  target implies. For example, a latency constraint of **< 2** may conflict with the  $f_{MAX}$  implied pipelining constraint. The multiplier may need four pipelining stages to reach the target  $f_{MAX}$ . The simulation fails and issues an error, highlighting the Primitive subsystem.

DSP Builder gives this error because you must increase the constraint limit by at least 3 (that is, to **< 5**) to meet the target  $f_{MAX}$ .

### 8.6.6. Control Units Delays

Commonly, you may use an FSM to design control units. An FSM uses DSP Builder **SampleDelay** blocks to store its internal state. DSP Builder automatically redistributes these **SampleDelay** blocks, which may alter the functional behavior of the control unit subsystem. Then the generated hardware no longer matches the simulation. Also, redistribution of **SampleDelay** blocks throughout the design may change the behavior of the FSM by altering its initial state. Classically, you exploit the reset states of the constituent components to determine the initial state; however this approach may not work. DSP Builder may not preserve any given component because it automatically pipelines **Primitive** subsystems. Also it can leave some components combinatorial based on  $f_{MAX}$  target, device family, speed grade, and the locations of registers immediately upstream or downstream.

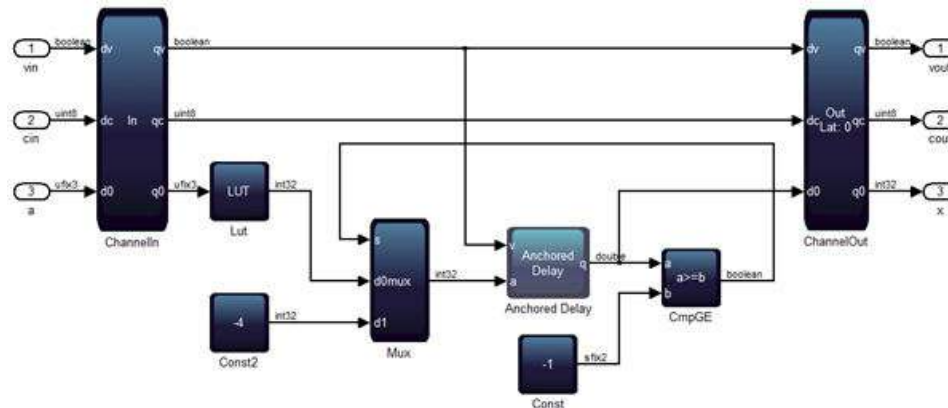
Figure 78. SampleDelay Block Example



DSP Builder relocates the sample delay, to save registers, to the Boolean signal that drives the s-input of the 2-to-1 **Mux** block. You may see a mismatch in the first cycle and beyond, depending on the contents of the LUT.

When you design a control unit as an FSM, the locations of **SampleDelay** blocks specify where DSP Builder expects zero values during the first cycle. In [Figure 78](#) on page 194, DSP Builder expects the first sample that the a-input receives of the **CmpGE** block to be zero. Therefore, the first output value of that compare block is high. Delay redistribution changes this initialization. You cannot rely on the reset state of that block, especially if you embed the Primitive subsystem within a larger design. Other subsystems may drive the feedback loop whose pipeline depth adapts to  $f_{MAX}$ . The first valid sample may only enter this subsystem after some arbitrary number of cycles that you cannot predetermine. To avoid this problem, always ensure you anchor the **SampleDelay** blocks to the valid signal so that the control unit enters a well-defined state when valid-in first goes high.

Figure 79. SampleDelay Block Example 2



To make a control unit design resistant to automated delay redistribution and to solve most hardware FSM designs that fail to match simulation, replace every **SampleDelay** block with the **Anchored Delay** block from the **Control** folder in the Additional libraries. When the valid-in first goes high, the **Anchored Delay** block outputs one (or more) zeros, otherwise it behaves just like an ordinary **SampleDelay** block.

Synthesizing the example design ( $f_{MAX} = 250\text{MHz}$ ) on Arria V (speedgrade 4), shows that DSP Builder is still redistributing the delays contained inside of the **Anchored Delay** block to minimize register utilization. DSP Builder still inserts a register initialized to zero before the s-input of the 2-to-1 **Mux** block. However, the hardware continues to match Simulink simulation because of the anchoring. If you place highly pipelined subsystems upstream so that the control unit doesn't enter its first state until several cycles after device initialization, the FSM still provides correct outputs. Synchronization is maintained because DSP Builder inserts balancing delays on the valid-in wire that drives the **Anchored Delay** and forces the control unit to enter its initial state the correct number of cycles later.

Control units that use this design methodology are also robust to optimizations that alter the latency of components. For example, when a LUT block grows sufficiently large, DSP Builder synthesizes a **DualMem** block in its place that has a latency of at least one cycle. Automated delay balancing inserts a sufficient number of one bit wide delays on the valid signal control path inside every **Anchored Delay**. Hence, even if the **CmpGE** block is registered, its reset state has no influence on the initial state of the control unit when the valid-in first goes high.

Each **Anchored Delay** introduces a 2-to-1 **Mux** block in the control path. When targeting a high  $f_{MAX}$  (or slow device) tight feedback loops may fail to schedule or meet timing. Using **Anchored Delay** blocks in place of **SampleDelay** blocks may also use more registers and can also contribute to routing congestion.

## 8.7. Flow Control in DSP Builder Designs

Use DSP Builder `valid` and `channel` signals with data to indicate when data is valid for synchronizing. You should use these signals to process valid data and ignore invalid data cycles in a streaming style to use the FPGA efficiently. You can build designs that run as fast as the data allows and are not sensitive to latency or devices  $f_{MAX}$  and that can be responsive to backpressure.

This style uses FIFO buffers for capturing and flow control of valid outputs, loops, and for loops, for simple and complex nested counter structures. Also add latches to enable only components with state—thus minimizing enable line fan-out, which can otherwise be a bottleneck to performance.

### Flow Control Using Latches

Generally hardware designers avoid latches. However, these subsystems synthesize to flip-flops.

Often designs need to stall or enable signals. Routing an enable signal to all the blocks in the design can lead to high fan-out nets, which become the critical timing path in the design. To avoid this situation, enable only blocks with state, while marking output data as invalid when necessary.

DSP Builder provides the following utility functions in the **Additional Blocks Control library**, which are masked subsystems.

- **Zero-Latency Latch (latch\_0L)**
- **Single-Cycle Latency Latch (latch\_1L)**
- **Reset-Priority Latch (SRLatch\_PS)**
- **Set-Priority Latch (SRLatch)**

Some of these blocks use the Simulink **Data Type Prop Duplicate** block, which takes the data type of a reference signal `ref` and back propagates it to another signal `prop`. Use this feature to match data types without forcing an explicit type that you can use in other areas of your design.

### Forward Flow Control Using Latches

The **demo\_forward\_pressure** example design shows how to use latches to implement forward flow control.

### Flow Control Using FIFO Buffers

You can use FIFO buffers to build flexible, self-timed designs insensitive to latency. They are an essential component in building parameterizable designs with feedback, such as those that implement back pressure.

### Flow Control and Backpressure Using FIFO Buffers

The **demo\_back\_pressure** design example shows how to use latches to implement back pressure flow control.

You must acknowledge reading of invalid output data. Consider a FIFO buffer with the following parameters:

- Depth = 8
- Fill threshold = 2
- Fill period = 7

A three cycle latency exists between the first write and valid going high. The  $q$  output has a similar latency in response to writes. The latency in response to read acknowledgements is only one cycle for all output ports. The valid out goes low in response to the first read, even though the design writes two items to the FIFO buffer. The second write is not older than three cycles when the read occurs.

With the fill threshold set to a low value, the  $t$  output can go high even though the  $v$  out is still zero. Also, the  $q$  output stays at the last value read when valid goes low in response to a read.

Problems can occur when you use no feedback on the read line, or if you take the feedback from the  $t$  output instead with fill threshold set to a very low value ( $< 3$ ). A situation may arise where a read acknowledgement is received shortly following a write but before the valid output goes high. In this situation, the internal state of the FIFO buffer does not recover for many cycles. Instead of attempting to reproduce this behavior, Simulink issues a warning when a read acknowledgement is received while valid output is zero. This intermediate state between the first write to an empty FIFO buffer and the valid going high, highlights that the input to output latency across the FIFO buffer is different in this case. This situation is the only time when the FIFO buffer behaves with a latency greater than one cycle. With other primitive blocks, which have consistent constant latency across each input to output path, you never have to consider these intermediate states.

You can mitigate this issue by taking care when using the FIFO buffer. The model needs to ensure that the read is never high when valid is low using the simple feedback. If you derive the read input from the  $t$  output, ensure that you use a sufficiently high threshold.

You can set fill threshold to a low number ( $<3$ ) and arrive at a state where output  $t$  is high and output  $v$  is low, because of differences in latency across different pairs of ports—from  $w$  to  $v$  is three cycles, from  $x$  to  $t$  is one cycle, from  $w$  to  $t$  is one cycle. If this situation arises, do not send a read acknowledgement signal to the FIFO buffer. Ensure that when the  $v$  output is low, the  $x$  input is also low. A warning appears in the MATLAB command window if you ever violate this rule. If you derive the read acknowledgement signal with a feedback from the  $t$  output, ensure that the fill threshold is set to a sufficiently high number (3 or above). Similarly for the  $f$  output and the full period.

If you supply vector data to the  $d$  input, you see vector data on the  $q$  output. DSP Builder does not support vector signals on the  $w$  or  $x$  inputs, as the behavior is unspecified. The  $v$ ,  $t$ , and  $f$  outputs are always scalar.

### Flow Control using Simple Loop

Designs may require counters, or nested counters to implement indexing of multidimensional data. The **Loop** block provides a simple nested counter—equivalent to a simple software loop.

The enable input and **demo\_kronecker** design example demonstrate flow control using a loop.

### Flow Control Using the ForLoop Block

You can use either **Loop** or **ForLoop** blocks for building nested loops.

The **Loop** block has the following advantages:

- A single **Loop** block can implement an entire stack of nested loops.
- No wasted cycles when the loop is active but the count is not valid.
- The implementation cost is lower because no overhead for the token-passing scheme exists.

The **ForLoop** block has the following advantages:

- Loops may count either up or down.
- You may specify the initial value and the step, not just the limit value.
- The token-passing scheme allows the construction of control structures that are more sophisticated than just nesting rectangular loops.

When a stack of nested loops is the appropriate control structure (for example, matrix multiplication) use a single **Loop** block. When a more complex control structure is required, use multiple **ForLoop** blocks.

## 8.8. Reset Minimization

Reset minimization reduces the amount of reset logic in your design. A reduction in reset logic can give an area decrease and potential  $f_{MAX}$  increase. Reset minimization removes resets on the datapath. You can apply reset minimization globally to floating-point operators and to your synthesizable subsystems. By default, DSP Builder turns on reset minimization for HyperFlex™ architectures and off for all other devices.

DSP Builder distinguishes control flow from data flow: control flow is the logic you connect to the **ChannelIn and ChannelOut** valid signal path. DSP Builder applies little or no reset minimization to control logic and aggressive minimization to data flow.

By default, DSP Builder chooses reset minimization options for you automatically. It automatically applies reset minimization if your target device includes the HyperFlex architecture.

You may override the default automatic reset minimization options, for example as part of design space optimization.

When you globally apply reset minimization, DSP Builder determines a local reset minimization setting for each of your synthesizable subsystems. DSP Builder applies this local reset minimization conditionally, if your subsystem contains **ChannelIn or ChannelOut** blocks.

**Table 25. Reset Minimization Summary**

If your synthesizable subsystem uses a mixture of Channel and GP blocks, choose **Conditional** for **Local reset minimization**.

Global Enable	Local Setting	Synthesizable Subsystem	Reset Minimization
Off	Any	Any	No
On	Off	Any	No
On	Conditional	ChannelIn and ChannelOut	Yes
On	Conditional	GPIn and GPOut	No
On	On	ChannelIn and ChannelOut	Yes
On	On	GPIn and GPOut	Yes

DSP Builder does not apply reset minimization to blocks with innate state, user-constructed cycles, and enable logic in your design, as that can give undefined initial values.

Reset minimization only detects local cycles within a subsystem. You should avoid broader feedback cycles.

Reset minimization may affect the behavior of your design during Simulink simulation and on hardware.

### Simulink Simulation

The DSP Builder simulation engine within Simulink is unaware of the reset minimization optimization and therefore always simulates your design behavior with reset present.

In general there is no difference in behavior, and this is aided by the testbench inputs defaulting typically to zero and a longer minimum reset pulse-width allowing such defaults to propagate through the datapath register stages.

However in some cases mismatches may occur, because data entering a **Sample Delay** in your design during reset is non-zero.

If an input does not default to zero or the internal behavior is incompatible with **Sample Delay** blocks resetting to zeros (or the minimum reset-pulse width is less than the design latency), the Simulink simulation might be different than the HDL simulation.

### Implementation on Hardware

Removing a reset on the datapath means that when DSP Builder releases a reset, your data flow logic may contain values clocked in during reset, which might affect the initial post-reset behavior of your system.

Reset minimization detects and avoids optimizing cycles in your synthesizable subsystem. It does not detect cycles constructed outside of a single synthesizable subsystem. Do not enable it for such designs.

### Related Information

- [Control](#) on page 223
- [Synthesis Information \(SynthesisInfo\)](#) on page 362

## 8.9. About Importing HDL

Importing HDL enables you to cosimulate existing HDL as a subsystem within your DSP Builder designs.

Importing HDL has the following software requirements:

- HDL Verifier toolbox
- An HDL Verifier compatible version of the ModelSim simulator (importing HDL does not support ModelSim AE)

Additionally, your HDL must conform to DSP Builder design rules and must:

- Have only one clock domain
- Match reset level with DSP Builder
- Use the `std_logic` data type for clock and reset ports
- Use `std_logic_vector` for all other ports
- Have no top-level generics
- Contain no bus components

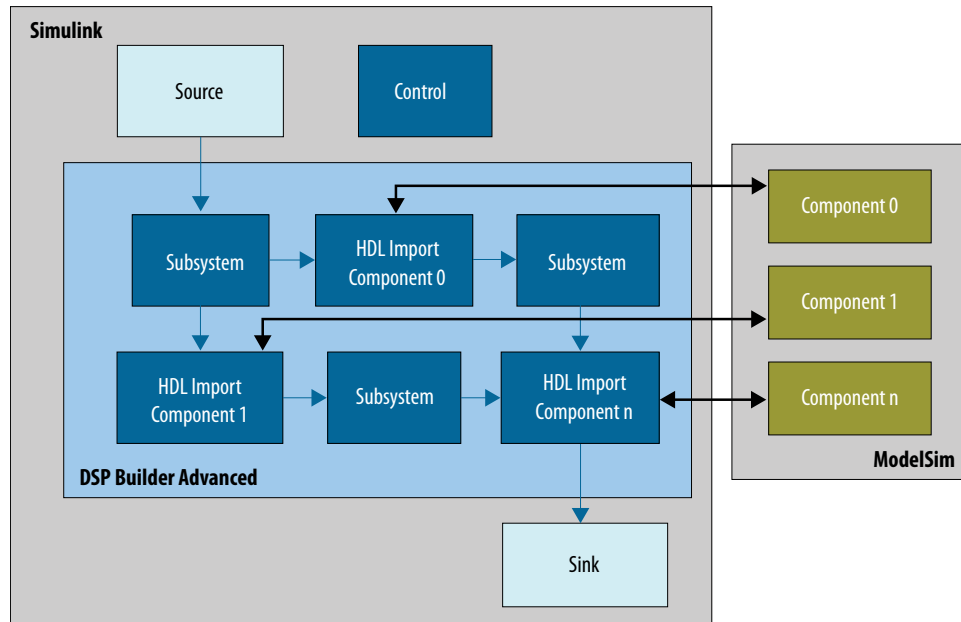
You may need to write a wrapper HDL file that instantiates your HDL, which might configure generics, convert from other data types to `std_logic_vector`, or invert the reset signal.

DSP Builder can import any number of instantiated entities. To import multiple copies of an entity or multiple distinct entities, instantiate the entities in a top-level wrapper file.

Simulink does not model all the signal states that ModelSim uses (e.g. 'U'). Simulink interprets all non-'1' states as a '0'.

Importing HDL uses the HDL Verifier toolbox to communicate with an HDL simulation running in ModelSim. You can have as many components in your ModelSim simulation as you like; each component communicates with a separate DSP Builder **HDL Import** block. Your top-level design must include an **HDL Import Config** block.

**Figure 80. HDL Import Block Placement**



You cannot place **HDL Import** blocks inside a primitive scheduled subsystem.

DSP Builder creates the appropriate instantiation of the component represented by the **HDL Import** block.

DSP Builder sees imported HDL as a scheduled system. DSP Builder does not try to schedule your imported HDL. You cannot import HDL into a scheduled subsystem. Imported HDL acts like other DSP Builder IP blocks (e.g. NCO, FFT). You must manually delay-balance any parallel datapaths and turn on **Generate Hardware** in the **Control** block.



## 9. About Folding

---

Folding optimizes hardware usage for low throughput systems, which have many clock cycles between data samples. Low throughput systems often inefficiently use hardware resources. When you map designs that process data as it arrives every clock cycle to hardware, many hardware resources may be idle for the clock cycles between data.

Folding allows you to create your design and generate hardware that reuses resources to create an efficient implementation.

The folding factor is the number of times you reuse a single hardware resource, such as a multiplier, and it depends on the ratio of the data and clock rates:

Folding factor = clock rate/data rate

DSP Builder offers ALU folding for folding factors greater than 500. With ALU folding, DSP Builder arranges one of each resource in a central arithmetic logic unit (ALU) with a program to schedule the data through the shared operation.

1. [ALU Folding](#) on page 201
2. [Removing Resource Sharing Folding](#) on page 208

### 9.1. ALU Folding

ALU folding generates an ALU architecture specific to the DSP Builder design. The functional units in the generated ALU architecture depend on the blocks and data types in your design. DSP Builder maps the operations performed by connecting blocks in Simulink to the functional units on the generated architecture.

ALU folding reduces the resource consumption of a design by as much as it can while still meeting the latency constraint. The constraint specifies the maximum number of clock cycles a system with folding takes to process a packet. If ALU folding cannot meet this latency constraint, or if ALU folding cannot meet a latency constraint internal to the DSP Builder system due to a feedback loop, you see an error message stating it is not possible to schedule the design.

1. [ALU Folding Limitations](#) on page 202
2. [ALU Folding Parameters](#) on page 202
3. [ALU Folding Simulation Rate](#) on page 202
4. [Using ALU Folding](#) on page 206
5. [Using Automated Verification](#) on page 207
6. [Ready Signal](#) on page 207
7. [Connecting the ALU Folding Ready Signal](#) on page 207

8. [About the ALU Folding Start of Packet Signal](#) on page 208

### 9.1.1. ALU Folding Limitations

Avoid using ALU folding with designs that use many data types. ALU folding is ideal for large designs with a uniform data type, such as single-precision floating-point. The design uses less logic when creating a single hardware resource for an operation in the ALU that it can share across the design.

For designs that use more than one data type, a **Convert** block between two data types uses more resources if the design requires saturation and rounding. An unbiased rounding operation uses more resources than a biased rounding mode.

Some DSP Builder blocks store state, for example:

- **Sample Delay**
- **Counter**
- **DualMem**
- **FIFO**

With ALU folding, any blocks that store state have a separate state for each channel. DSP Builder only updates the state for a channel when the system processes the channel. Thus, a sample delay delays a signal until processing the next data sample. For 200 clock cycles to a data period, DSP Builder delays the signal for the 200 clock cycles. Also, data associated with one channel cannot affect the state associated with any other channel. Changing the number of channels does not affect the behavior of the design.

*Note:* For designs without ALU folding, state is associated with a block, which you can update in any clock cycle. Data input with channel 0 can affect state that then affects a computation with data input with channel 1.

### 9.1.2. ALU Folding Parameters

**Table 26. ALU Folding Parameters**

Parameter	Description
Sample Rate	Data sample rate.
Number of Channels	Supports single or multiple channels
Maximum latency	Maximum latency for the system.
Register outputs	The format of data outputs
Simulation rate	Specify <b>clock rate</b> or <b>data rate</b> to control how Simulink models the system

### 9.1.3. ALU Folding Simulation Rate

In the ALU folding parameters, you can specify **Data rate** or **Clock rate** for **Simulation rate**. the Simulation rate only controls the simulink simulation; the hardware is identical.

Date rate simulation offers the following features:

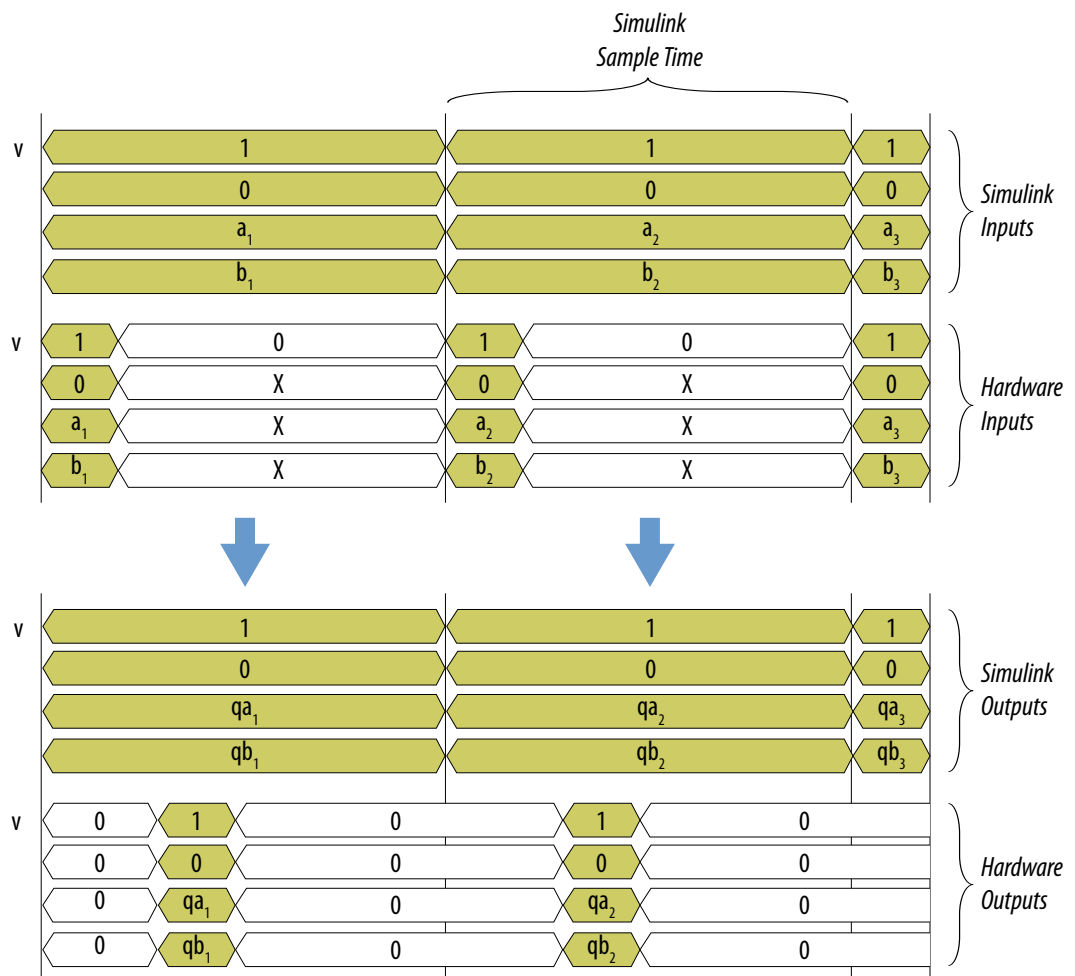
- Simulates faster.
- Simulates original unfolded model.
- Each Simulink sample represents a data sample.
- Generates automatic ModelSim testbench (if turned on in the Control block).

Clock rate simulation offers:

- Simulink sample rates identical to the clock rate.
- Simulation matches the hardware interface.
- Modelling of clock level timings and jitter in the data inputs.

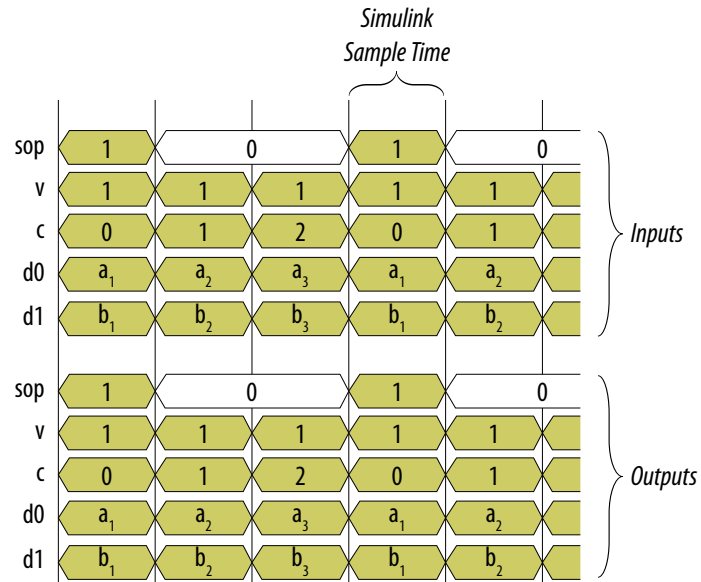
### Data Rate

**Figure 81. Single Channel Data Rate Simulation with no Register Outputs**



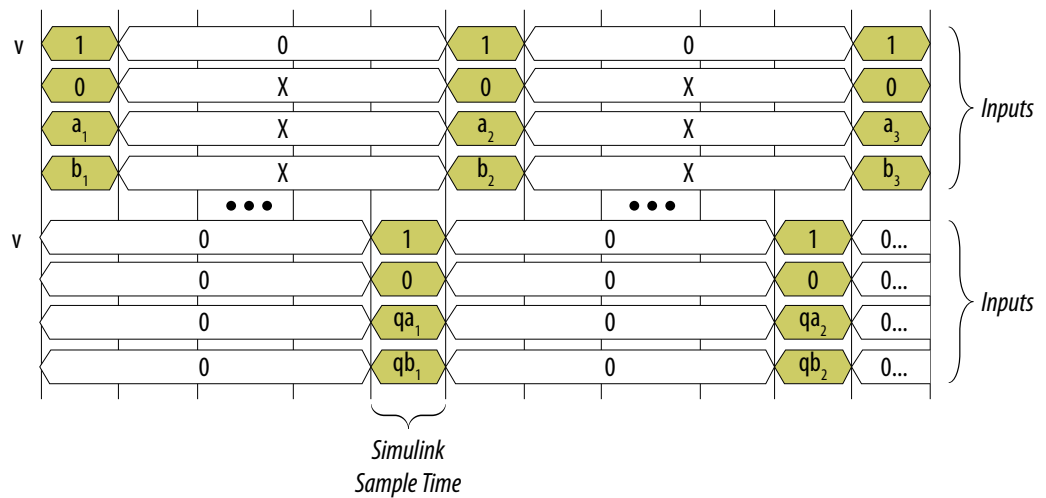
**Figure 82. Multichannel Data Rate Simulation with no Register Outputs**

The Simulink sample time is a third of the data sample period

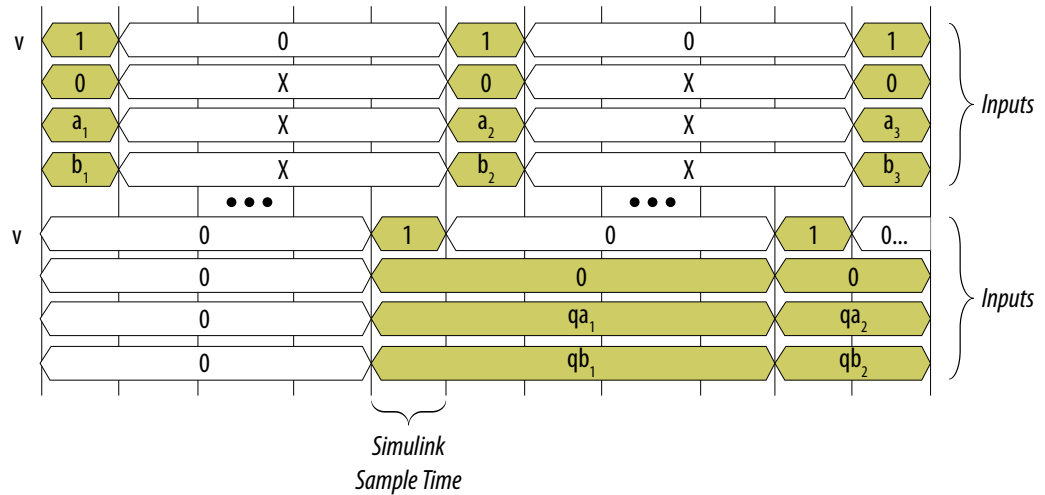


**Clock Rate**

**Figure 83. Single Channel Clock Rate Simulation with no Register Outputs**



**Figure 84. Single Channel Clock Rate Simulation with Register Outputs**



**Figure 85. Multichannel Clock Rate Simulation with no Register Outputs**

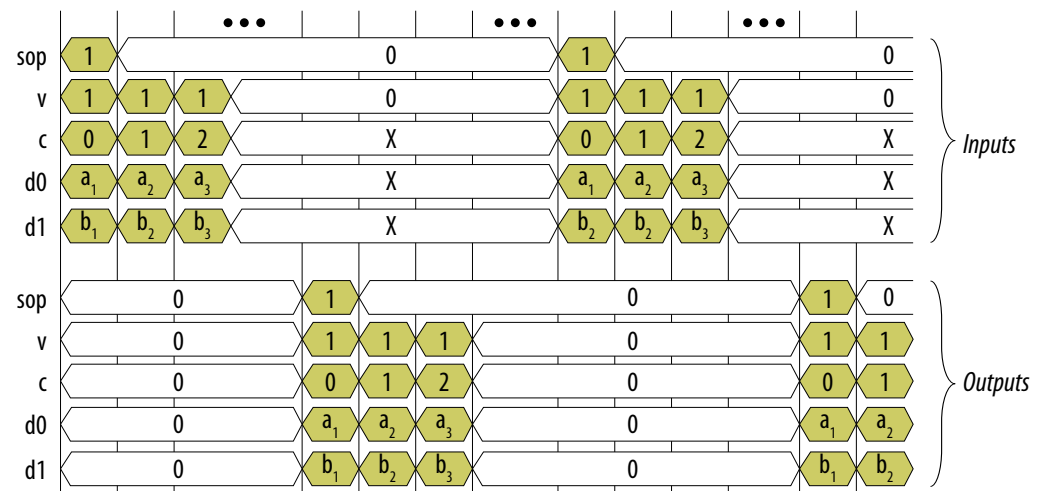
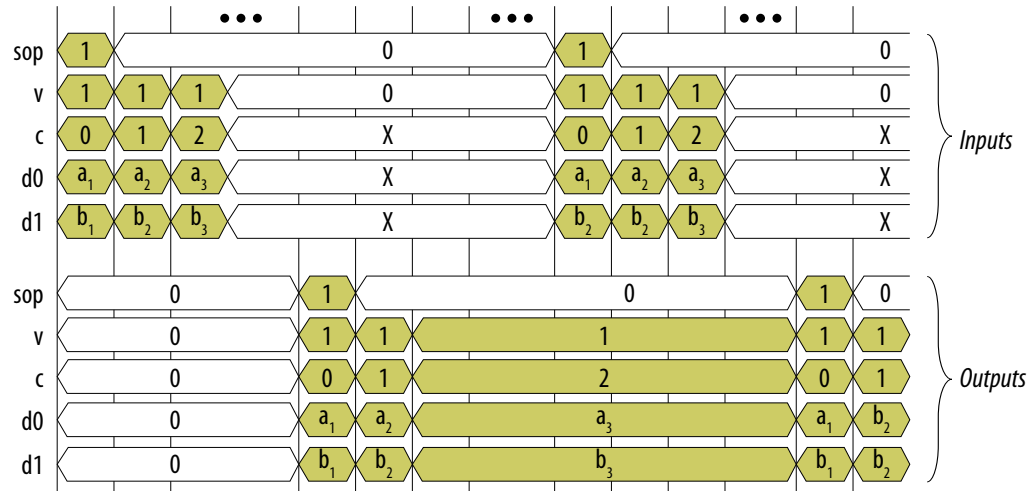


Figure 86. Multichannel Clock Rate Simulation with Register Outputs



### 9.1.4. Using ALU Folding

**Note:** In the **ChannelIn** and **ChannelOut** blocks, before you use ALU folding, ensure you turn off **Folding enabled**.

1. Open the top-level design that contains the primitive subsystem you want to add ALU folding to.
2. Save a backup of your original design.
3. Replace:
  - Constant multiplier blocks with multipliers blocks.
  - Reciprocal blocks with Divide blocks
  - Sin( $\pi x$ ) blocks with sin( $x$ ) blocks.
4. Avoid low-level bit manipulation
5. Open the primitive subsystem (which contains the **ChannelIn** and **ChannelOut** blocks) and add an **ALU Folding** block from the DSP Builder **Utilities** library.
6. Double click the **ALU Folding** block to open the **Block Parameters** window.
7. Enter a value for **Sample rate (MHz)**.
8. Enter a value for **Maximum latency (cycles)**
9. Turn off **Register Outputs** to make the output format the same as the input format. Turn on **Register Outputs**, so that the outputs hold their values until the next data sample output occurs.
10. Select the **Simulation rate**
11. Simulate your design.

DSP Builder generates HDL for the folded implementation of the subsystem and a testbench. The testbench verifies the sample rate Simulink simulation against a clock rate ModelSim simulation of the generated HDL.

### 9.1.5. Using Automated Verification


To use automated verification, on the DSP Builder menu click **Verify Design**.

The testbench uses captured test vectors from the Simulink simulation and plays through the clock rate simulation of the generated hardware at the data rate. DSP Builder checks the order and bit-accuracy of the hardware simulation outputs against the Simulink simulation.

### 9.1.6. Ready Signal

The ready signal is an output that goes high to indicate when you can input data into your design. It provides flow control that allows you to reduce jitter in your design. The ready signal output is high when the internal architecture is idle.


**Figure 87. Ready Signal Timing**



### 9.1.7. Connecting the ALU Folding Ready Signal

1. Connect a Ready block from the **Primitive** library to the **ChannelOut** block.

Figure 88. Connecting Ready Block



### 9.1.8. About the ALU Folding Start of Packet Signal

DSP Builder uses a start of packet signal for systems using ALU folding. The start of packet signal is an extra signal on the **ChannelIn** and **ChannelOut** blocks. To use the start of packet signal turn on **Has Start of Packet Signal** on the **ChannelIn** and **ChannelOut** blocks. You must use the start of packet signal for multichannel designs.

With the Start of Packet signal:

- The system is in an idle state after reset and after it finishes processing a data sample.
- The system indicates the first clock cycle of a packet of data when the start of packet signal goes high.
- The system processes the data packet if it is in an idle state when it receives the start of packet signal.
- The system is not idle the clock cycle after the start of packet signal until it finishes processing a data sample.

You may use the valid signal instead of the start of packet signal, which does not allow the folded system to process a non-valid data sample.

## 9.2. Removing Resource Sharing Folding

With DSP Builder v14.0 you could use resource sharing folding, which is now removed in v14.1 or later. When you open pre v14.1 designs in v14.1 or later, you must remove resource sharing folding, which you originally selected on the **ChannelIn** block:



1. Open the design in v14.1 or later.
2. Replace **ChannelIn** and **ChannelOut** blocks with new **ChannelIn** and **ChannelOut** blocks.
3. Change any part of your design that uses TDM vectors.
4. Change any aspects of your design that uses a sample delays rather than clock delays.

*Note:* If you do not remove resource sharing folding, and you simulate your design you see a MATLAB system error.

## 10. Floating-Point Data Types

---

Most **Primitive** library blocks support floating-point data types. DSP Builder generates a parallel datapath optimized for Intel FPGAs from the Simulink model.

Floating-point designs are useful in:

- Scientific applications
- Numerical algorithms
- High-dynamic range data designs
- Statistical modelling

Fixed-point designs often cannot support data with a high dynamic range unless the design explicitly uses a high precision type. Floating-point designs can represent data over a high dynamic range with limited precision. A compact representation makes efficient use of memory and minimizes data widths. The lowest precision type that DSP Builder supports is float16\_m10, otherwise known as half-precision float, which occupies 16 bits of storage. It can represent a range between  $-2^{16}$  to  $+2^{16}$  (exclusive) and non-zero magnitudes as small as  $2^{-14}$ .

Typically, fixed-point designs may include fixed-point types of various bit widths and precisions. When you create fixed-point designs, keep variations in word growth and word precision within acceptable limits. When you create floating-point designs, you must limit rounding error to ensure an accurate result. A floating-point design typically has only one or two floating-point data types.

DSP Builder provides a comprehensive library of elementary mathematical functions with complete support for all floating-point types. Each core is parameterized by precision, clock frequency, and device family.

1. [DSP Builder Floating-Point Data Type Features](#) on page 211
2. [DSP Builder Supported Floating-Point Data Types](#) on page 211
3. [DSP Builder Round-Off Errors](#) on page 212
4. [Trading Off Logic Utilization and Accuracy in DSP Builder Designs](#) on page 212
5. [Upgrading Pre v14.0 Designs](#) on page 213
6. [Floating-Point Sine Wave Generator Tutorial](#) on page 213
7. [Newton-Raphson Root Finding Tutorial](#) on page 218
8. [Forcing Soft Floating-point Data Types with the Advanced Options](#) on page 220

## 10.1. DSP Builder Floating-Point Data Type Features

- Variable precision. Nine floating-point precisions including half-precision, single-precision, and double-precision data.
- High throughput. Architecture adapts to  $f_{MAX}$ .
- Bit-accurate simulation.
- Efficient DSP block usage in all device families .
- Standards compliant. Fundamental components (`add`, `sub`, `mult`, and `div`) are IEEE compliant. Elementary mathematical functions are OpenCL compliant.
- Configurable. Trade off logic utilization against accuracy. Significant DSP block reduction using faithful rounding.

## 10.2. DSP Builder Supported Floating-Point Data Types

The supported floating-point types are either IEEE 754 formats (half, single and double precision) or custom IEEE 754-like formats with user-specified exponent and fraction-field widths .

Type Name	Sign Width $s$	Exponent Width $e$	Exponent Bias $b$	Mantissa Width $m$	Description
float16_m7	1	8	127	7	Bfloat16
float16_m10		5	15	10	Half-precision IEEE 754-2008)
float19_m10		8	127	10	Also known as TF32
float26_m17		8	127	17	
float32_m23		8	127	23	Single-precision IEEE 754
float35_m26		8	127	26	
float46_m35		10	511	35	
float55_m44		10	511	44	
float64_m52		11	1023	52	Double-precision IEEE 754

DSP Builder represents the special values positive zero, negative zero, subnormals, and non-numbers in the standard IEEE 754 manner, namely:

- zero is  $m=0$  and  $e=0$  with  $s$  giving the sign.
- subnormal is  $m \neq 0$  and  $e=0$  with  $s$  giving the sign.
- infinity is  $m=0$  and  $e=\text{all ones}$  with  $s$  giving the sign.
- not a number (NaN) is  $m \neq 0$  and  $e=\text{all ones}$ .

Except for the preceding special values, the numerical value of a float type is given in terms of its bit-wise representation by:

$$f = (-1)^s \times 2^{(e-b)} \times (1 + (m / (2^{m\_width})))$$

where:

- $e$ ,  $b$ , and  $m$  are the base-10 equivalents of the respective bit sequences
- the field widths for each of  $s$ ,  $e$  and  $m$  and the value of  $b$  are given for each format in the table

For example, for a 32-bit single precision floating point number with a bit-wise representation of 0x40300000:

$s = 0b$	$= 0$
$e = 10000000b$	$= 128$
$m = 01100000000000000000000b$	$= 3145728$

then:

$$\begin{aligned}
 f &= (-1)^0 \times 2^{(128-127)} \times (1+(3145728/(2^{23}))) \\
 &= 1 \times 2 \times (1+0.375) \\
 &= 2.75
 \end{aligned}$$

### 10.3. DSP Builder Round-Off Errors

Every mathematical operation on floating-point data incurs a round-off error.

For the fundamental operations (add, subtract, multiple, divide) this error is determined by the rounding mode:

- Correct. A typical relative error is half the magnitude of the LSB in the mantissa.
- Faithful. A typical relative error is equal to the magnitude of the LSB in the mantissa.

The relative error for float16\_m10 is approximately 0.1% for faithful rounding, and 0.05% for correct rounding. The rounding mode is a configurable mask parameter.

The elementary mathematical functions conform to the error tolerances specified in the OpenCL standard. In practice, the relative error exhibited by the DSP Builder mathematical library lies comfortably within the specified tolerances.

Bit cancellations can occur when subtracting two floating-point numbers that are very close in value, which can introduce very large relative errors. You need to take the same precautions with floating-point designs as with numerical software to prevent bit cancellations.

### 10.4. Trading Off Logic Utilization and Accuracy in DSP Builder Designs

1. If your design exceeds the relative accuracy but is using too much hardware:
  - a. Use the next lowest precision
  - b. Use faithful rounding instead of correct rounding
  - c. Enable the fused datapath option

Each of these changes reduces logic utilization at the expense of accuracy. A design may use more than one floating-point precision for different sections of the circuit, however if there are too many different precisions you will need to have more type conversion blocks. Each convert block increases logic utilization.

## 10.5. Upgrading Pre v14.0 Designs

DSP Builder designs in v14.0 onwards have floating-point data turned on by default, which provides access to all floating-point data types. Floating-point designs from before v14.0 have a dedicated **EnhancedPrecision** block in the primitive subsystem. These designs work correctly in v14.0, but Intel recommends you find and remove any **EnhancedPrecision** blocks from your primitive subsystems.

## 10.6. Floating-Point Sine Wave Generator Tutorial

1. [Creating a Sine Wave Generator in DSP Builder](#) on page 213
2. [Using Data Type Variables to Parameterize Designs](#) on page 214
3. [Using Data-Type Propagation in DSP Builder Designs](#) on page 215
4. [DSP Builder Testbench Verification](#) on page 215

### Related Information

[DSP Builder Floating Point Design Examples](#) on page 115

### 10.6.1. Creating a Sine Wave Generator in DSP Builder

1. In Simulink, click **DSP Builder** > **New Model wizard** and create a floating-point primitive (simple) model.
2. Create a primitive subsystem with the following components:
  - **ChannelIn** and **ChannelOut**
  - **SynthesisInfo** (configure to **Scheduled**)
  - **Counter** with a large period (e.g. 32,768) and incrementing in steps of 1
  - **Convert** block (set the mask parameter **Output data type mode** to single)
  - **Mult** block
  - **Trig** block (configure the **Mask** parameter function to **sin(x)**)
3. Connect the blocks.
4. Simulate the design.
5. Connect the single-precision input of the subsystem to a Simulink built-in source. For example, repeating sequence stair.
6. Set the repeating sequence stair block parameter **Output Data Type** to **single**. Hence, both inputs to the **Mult** block are single. This data type propagates through to the **Trig** block
7. Simulate this design with hardware generation turned on. DSP Builder generates HDL files for a single-precision IEEE multiply, a single-precision sine function, and a fixed-to-float conversion component.

8. Click **DSP Builder > Resource Usage > Resource Usage Report** and record the DSP and LUT usage for the design.
9. To change the floating-point precision of the synthesized design, insert a **Convert** block on the floating-point input wire.
10. Parameterize the **Convert** block:
  - a. Set **Output data type mode** to **Variable precision floating point**.
  - b. In the **Floating Point Precision** drop-down menu select **float26\_m17**.
11. Apply the same parameters to the **Convert** block in the primitive subsystem.
12. To connect the floating-point output port of the subsystem to a scope, or some other built-in Simulink sink block:
  - a. Insert a **Convert** block on the floating-point output wire.
  - b. Set the **Output data type mode** to **double**.

*Note:* If you do not connect a **Convert** block, you cannot simulate your design. Simulink scopes do not recognize the DSP Builder custom floating-point type.
13. Simulate the design to generate HDL for the reduced precision multiplier, sine function, and fixed-to-float convert.
14. Re-examine the resource usage report. The DSP and LUT utilization is significantly lower than for the single-precision design.

### 10.6.2. Using Data Type Variables to Parameterize Designs

Commonly, you change the precision of a DSP Builder design by using scripts. However, writing scripts to update the floating-point precision drop-down menu for all blocks in the design is tedious.

Use data-type variables to parameterize designs by data type.

1. At the MATLAB console, initialize a variable with the following command:

```
>> inputType = dspba.vpfloat(26,17)
inputType =
Class: 'FLOAT'
ExpBits: 8
FracBits: 17
```

This MATLAB structure specifies the floating-point precision similar to how `fixdt()` specifies fixed-point precisions.

2. For the top-level **Convert** block on the input wire, open the parameter dialog box:
  - a. Set the **Output data type mode** to **Specify via dialog**.

- b. Delete the **Output data type** field and type the variable name `inputType`.
3. Repeat for the **Convert** block in the primitive subsystem, so that the same data type propagates to both inputs of the **Mult** blocks.
4. Change the floating-point precision of the design, by assigning a different type to the variable `inputType`. You can also initialize the type variable using the data type name:

```
>> inputType = dspba.namedVPFloat('float26_m17')
```

```
>> inputType = dspba.namedVPFloat('single')
```

### 10.6.3. Using Data-Type Propagation in DSP Builder Designs

To simplify setup scripts for parameterized designs, use data-type propagation.

1. Add the DSP Builder custom **SameDT** block to your design. Do not use the built-in Simulink same-DT block, which does not propagate data types.
2. Use any of the following blocks to allow you to back propagate DSP Builder's floating-point data types via their output ports:
  - **Const**
  - **Lut**
  - **Convert**
  - **ReinterpetCast**
3. Set the **Output data type mode** parameter for these blocks to **Inherit via back propagation**. Using this option and the custom **SameDT** block minimizes scripting for setting up data types in your design.  
The data type propagates via the built-in multiplex to three different wires, and then back propagates via the respective output ports of the **Convert** block, (coefficients) **LUT** block, and **Const** block.

### 10.6.4. DSP Builder Testbench Verification

The Simulink simulation model generates the stimulus files for the automated testbench. A multiple precision floating-point library processes the variable precision floating-point signal values. All functions round the output results to the nearest representable value, even if you configure the fundamental operators to use faithful rounding. Also, the elementary mathematical functions do not need to round to nearest to comply with the IEEE754 standard. Hence, the hardware does not always output the same bit pattern as the Simulink simulation. The benefit of this minor non-compliance with the IEEE standard is a potential improvement in quality of results, namely some combination of increased  $f_{MAX}$ , reduced latency, and reduced area. DSP Builder provides tools to help analyze floating-point signals when simulating your designs.

1. [Tuning ATB Thresholds](#) on page 216
2. [Writing Application Specific Verification](#) on page 216
3. [Performing Bit-Accurate Simulation](#) on page 217
4. [Adder Trees and Scalar Products](#) on page 217
5. [Creating Floating-Point Accumulators for Designs that Use Iteration](#) on page 218

### 10.6.4.1. Tuning ATB Thresholds

If your design uses floating-point components, the autogenerated testbench uses a special floating-point comparison when detecting mismatches. Two thresholds influence the sensitivity of the mismatch detection:

- Floating-point mismatch tolerance, which is the largest relative error that is not flagged as a mismatch
- Floating-point mismatch zero tolerance, which is the largest magnitude for a signal value to be considered as equivalent to zero.

1. Click **DSP Builder > Verify Design > Advanced**.
2. Enter new values in **Floating-point testbench settings**.

### 10.6.4.2. Writing Application Specific Verification

Generally, use the threshold method for detecting mismatches in hardware for most designs. For more sophisticated designs you can write your own application specific verification function.

1. Click **DSP Builder > Verify Design > Advanced**.
2. Turn on **Import ModelSim results to MATLAB**.
3. Enter a MATLAB variable name in **Import device output to variable**.
4. Optionally, enter a different variable name in **Import valid map to variable** field.
5. Enter the name of the verification m-function in **Verification function**.
6. Simulate and generate the hardware. DSP Builder modifies the ATB to write to a file the output signal values that ModelSim simulates:
  - a. Click **DSP Builder > Verify Design**
  - b. Turn on **Run simulation** and **Verify at device level**,
  - c. Click **OK**

MATLAB stores the simulation results using field names derived from the names of the output ports in your design.

```
>> atbPaths = vsimOut.keys; vsimOut(atbPaths{1})
ans =
  vout: [2000x1 embedded.fi]
  vout_stm: [2000x1 embedded.fi]
  cout: [2000x1 embedded.fi]
  cout_stm: [2000x1 embedded.fi]
  x: [2000x1 double]
  x_stm: [2000x1 double]
```

The fields ending in `_stm` are from the stimulus files that Simulink normally writes out during simulation. You can use these as the golden standard against which to compare the simulated hardware output. The verification function you specified is started, passing this struct as the first parameter



### 10.6.4.3. Performing Bit-Accurate Simulation

For most floating-point designs, and any fixed-point designs that use the **ConstMult** block, the default Simulink simulation only approximates the behavior of the generated hardware. Therefore, DSP Builder supports bit-accurate simulation.

1. Open the **SynthesisInfo** block parameter dialog box.
2. Turn on **Bit Accurate Simulation**.

When you turn on bit-accurate simulation and you simulate the design, the Simulink simulation is no longer based on the multiprecision floating-point library. Instead, the Simulink simulation signal values are generated from the same model of the floating-point operations that generate the RTL. Therefore, the simulation values exactly match the values that the hardware produces.

*Note:* Do not turn on **Bit Accurate Simulation** when your design includes **Memory-Mapped** library blocks, otherwise the simulation is all zeros.

The bit-accurate simulation correctly reproduces the signal values of the data paths in the design, but the timing of those signals may not necessarily reflect any register retiming that the scheduler has carried out. To make the bit-accurate simulation also reproduce the correct timing:

- a. Click **DSP Builder > Verify Design**
- b. Click on the **Advanced** tab.
- c. Turn on **Bit accuracy is also cycle accurate**

### 10.6.4.4. Adder Trees and Scalar Products

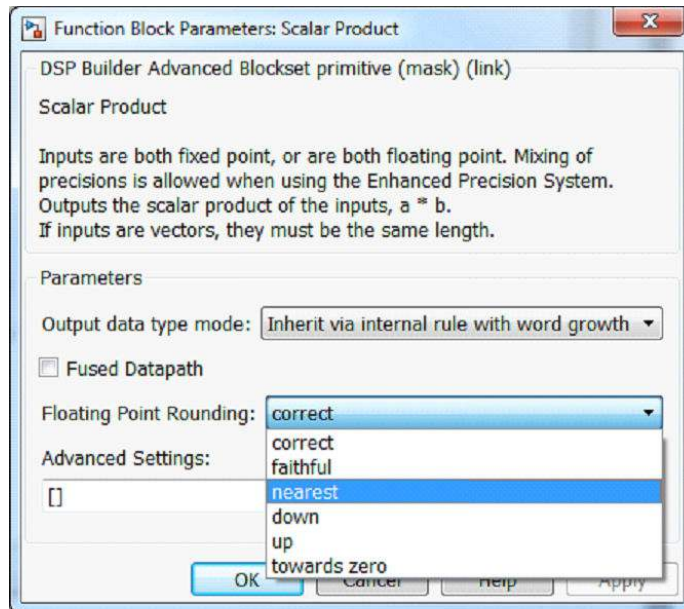
The **matmul\_flash\_RS** and **matmul\_flash\_RD** design examples use vector signals with floating-point components.

The **gemm\_flash** design example is a generalized matrix multiplication design that uses the **Scalar Product** block to calculate an inner product. The adder tree and the **Scalar Product** block have similar parameters:

- Fused datapath. Enable this option to reduce logic utilization at the expense of IEEE compliance
- Rounding modes:
  1. Nearest
  2. Down (towards negative infinity)
  3. Up (towards positive infinity)
  4. Towards zero

When you turn on **Fused datapath**, you can select only the rounding modes **Nearest** and **Towards zero**. Logic utilization is highest when your design uses rounding mode **Nearest**.

Figure 89. Floating-Point Rounding



#### 10.6.4.5. Creating Floating-Point Accumulators for Designs that Use Iteration

In DSP Builder, you can create fixed-point accumulators with one or more channels, but you cannot create floating-point equivalents.

1. If your design requires a single channel accumulator, use the **Acc** block.
2. Use multiple channels where the number of channels is at least as large as the latency of the floating-point adder. The lower limit on the number of channels depends on  $f_{MAX}$ , device family, and speed grade:
  - a. DSP Builder redistributes sample delays in your multiple channel accumulator. You cannot rely on DSP Builder preserving the reset state of zero. Use a 2:1 **Mux** block in your accumulator to make the initial state explicit.
  - b. Use the custom **SameDT** block to avoid any data type propagation problems in feedback cycles.

## 10.7. Newton-Raphson Root Finding Tutorial

This DSP Builder tutorial implements a floating-point iterative algorithm. The tutorial also demonstrates how to exploit pipeline parallelism.

Consider an application that finds the intersections of two equations:

$$y = e^x$$

$$y = mx + c$$

The design finds the roots of the equation,  $e^x - mx - c = 0$ , using Newton-Raphson iteration. The Newton-Raphson part of the design derives an improved approximation to the root from the previous guess.

*Note:* The following design examples show the various stages of the Newton-Raphson root finding tutorial:

- `demo_newton_iteration.mdl`
  - `demo_newton_convergence.mdl`
  - `demo_newton_valid.mdl`
  - `demo_newton_control.mdl`
  - `demo_newton_final.mdl`
1. [Implementing the Newton Design](#) on page 219
  2. [Improving DSP Builder Floating-Point Designs](#) on page 219

### 10.7.1. Implementing the Newton Design

1. Add and connect the blocks in the Newton design.
2. Reduce logic usage by configuring the **Mult**, **Add**, **Sub**, and **Divide** blocks to use faithful rounding.
3. Create the iteration loop by feeding back the output guess to the input guess through a **SampleDelay** block. The design detects when a sample finishes iterating by comparing the residue with zero.
4. Ensure that the length of this delay is sufficiently large so that the scheduling succeeds.
5. Turn on the **SampleDelay** block **Minimum delay** parameter so that DSP Builder determines this length automatically.

### 10.7.2. Improving DSP Builder Floating-Point Designs

In floating-point designs when comparing against zero, many of the samples never terminate and circulate through the feedback loop indefinitely. Because of the round-off errors accumulating in the **NewtonStep** subsystem, the residue may never reach exactly zero for many of the data samples.

1. Use a subsystem to detect convergence and divergence of the iterative feedback loop
2. Simulate the design. The number of valid samples on the output far exceeds the number of valid samples on the input.
3. To track which pipeline slots contain a valid sample, add a control signal path that lies parallel to the datapath feedback.
4. This 1-bit wide control path must be the same latency as the datapath.
  - a. Ensure that for both **SampleDelay** blocks you turn on **Minimum delay enabled**.
  - b. Set the **Equivalence Group** to the string, `newton`. When a sample converges to a root, DSP Builder outputs it as a valid solution and marks the pipeline slot as empty. If a sample diverges, DSP Builder marks the pipeline slot as empty but keeps valid low.

5. Simulate this version of the design and verify that the number of valid samples output equals the number of valid samples input.

The design may exceed the pipeline capacity if you provide too many valid samples on the input. The scheduled size of the sample delays indicates the maximum number of pipeline slots that are available for iterative computation. If you input more than this number, you risk overwriting previous valid samples before their iteration converges.

6. To overcome this limitation, introduce a FIFO buffer for the input samples. When an empty pipeline slot becomes available at the front of the pipeline, DSP Builder removes a sample from the queue from the FIFO buffer and inserts it into the free slot to begin iterating around the **NewtonStep** feedback loop.
7. Simulate the design and verify that you can safely input more valid samples than the pipeline depth of the iterative loop.
8. Set the size of the FIFO buffer to adjust the capacity and ensure that it is as large as your application requires.
9. The rounding errors of the floating-point blocks can interact in such a way that a sample forever oscillates between two values: never converging and never diverging. To detect this oscillation add another control path in parallel to the datapath feedback to count the number of iterations each sample passes through.

*Note:* The **TooMany** subsystem compares the iteration count against a threshold to detect divergence.

*Note:* The **Mandelbrot\_S** design example implements another iterative algorithm that shows parallel feedback loops for floating-point data and control paths.

*Note:* The **matmul\_CS** design example exploits both vector and pipeline parallelism. This example also shows how to incorporate memories in the feedback path to store intermediate results.

## 10.8. Forcing Soft Floating-point Data Types with the Advanced Options

The **Add**, **Sub**, **Mult**, **Sum of Elements** and **Scalar Product** blocks have default hard floating-point data types (for devices that implement hard floating-point designs and single precision only).

To force soft floating-point data type, you must perform this task on all applicable blocks in your design.

1. Type `struct('forceSoftFP', 1)` in the **Advanced Options** dialog box.

### Related Information

- [Add](#) on page 305
- [Multiply \(Mult\)](#) on page 343
- [Subtract \(Sub\)](#) on page 354
- [Scalar Product](#) on page 0
- [Sum of Elements \(SumOfElements\)](#) on page 355

## 11. Design Configuration Library

The DSP Builder **Design Configuration** library blocks control your design flow and run external synthesis and simulation tools. These blocks set the design parameters, such as device family, target  $f_{MAX}$ , and bus interface signal width.

1. [Avalon Memory-Mapped Agent Settings \(AvalonMemoryMappedAgentSettings\)](#) on page 221
2. [Control](#) on page 223
3. [Device](#) on page 226
4. [Edit Params](#) on page 227
5. [LocalThreshold](#) on page 227

### 11.1. Avalon Memory-Mapped Agent Settings (AvalonMemoryMappedAgentSettings)

The **AvalonMemoryMappedAgentSettings** block specifies information about the top-level memory-mapped bus interface widths.

*Note:* You can either use this block in your design or view the Avalon-MM agent interface settings on the **DSP Builder** ► **Avalon Interface** menu.

**Table 27. Parameters for the AvalonMemoryMappedAgentSettings Block**

Parameter	Description
Bus interface name	Specifies the prefix for the address, data and control signals in the generated control bus.
Address width	Specifies the width in bits of the memory-mapped address bus (1-32, default=10).
Data width	Specifies the width in bits of the memory-mapped data bus (16, 32, or 64, default=16). DSP Builder does not support byte enables for Avalon memory-mapped agent interface. Only connect hosts to this interface that have the same or a smaller data width. For example, to attach a JTAG host, set the data width to 32 bits or less. When using <b>SharedMem</b> block ensure the output data width matches the <b>AvalonMemoryMappedAgentSettings</b> bus data width or is exactly twice the bus data width.
Bus is:	Specifies whether the memory-mapped address bus is <b>Big Endian</b> or <b>Little Endian</b> .
Separate bus clock	Turn on so any processor-visible control registers are clocked by a separate control bus clock to ease timing closure.
Bus clock frequency (MHz)	Specifies the frequency of the separate processor interface bus clock (when enabled).
Bus clock synchronous with system clock	Turn on so the bus clock is synchronous with the system clock.

## Word and Byte Addressing

Use word addressing when accessing memory-mapped blocks in DSP Builder from the design logic or through DSP Builder processor interface (using the **BusStimuli** block). Use byte addressing when you access the same locations through DSP Builder MATLAB API. To change the word address to byte address, multiply it by the number of bytes in **AvalonMemoryMappedAgentSettings** block data width. If you use the **BusStimuliFileReader** block to drive the **BusStimulus** block, ensure values for **Data Width** and **Address Width** parameters exactly match the address and data width you set in **Avalon Interfaces** > **Avalon Memory-Mapped Settings**

*Note:* Ensure your access permissions are correct, when using the **RegBit**, **RegField**, and **SharedMem** blocks from the **Interface** library.

*Note:* If you read from a nonreadable address, the output data is not valid.

When using the **SharedMem** block the output data width is twice the bus data width. In the DSP Builder processor interface, the block appears to have twice the number of entries compared with the design view. Also DSP Builder interprets each element in an initialization array to be of output data width. Use the System Console MATLAB API in DSP Builder to access the memory-mapped locations in DSP Builder designs on the FPGA. Use byte addressing when using this interface:

```
dspba_design_base_address_in_qsys +  
(block_address_in_dspba_design* dspba_bus_data_width_bytes)
```

Read and write requests time out in 1 minute if the device shows no response for the initiated request. For example:

- Read or write requests to an address that is not assigned to any slave in the top-level system.
- Read requests to a memory-mapped location that does not have read access (i.e. write only).

DSP Builder responds to read requests to non-readable or unassigned addresses with invalid data, because unanswered read requests may block the interconnect, so further valid requests don't go through. DSP Builder accepts write requests, but ignores them if the address is non-writable.

If the subsequent requests to valid addresses and locations continue to time out, the initial request disables the bus interconnect. You must then reset the system or reprogram the board.

Additionally, close all your host connections in MATLAB before switching off or reprogramming the board, because MATLAB corrupts the existing connection. If you cannot start a new debugging session, restart MATLAB.

## Clock Crossing

DSP Builder designs use a separate clock for all processor visible control registers if you select **Separate bus clock** in **Avalon Interfaces** > **Avalon Memory-Mapped Settings**. This clock is asynchronous to a main system clock if you turn off **Bus clock synchronous with system clock**.

DSP Builder inserts simple two-stage synchronizers between the bus and system clock domains. DSP Builder adds the synchronization to the autogenerated bus slave logic if you use any of the **Interface** blocks (e.g. **RegField**) and to NCO IP if you enable writable access to configuration registers inside the NCO.

The DSP Builder-generated timing constraints set maximum and minimum delays for paths between two different clocks to a big enough range, so timing analyzer doesn't show an error. Using this method allows you to overwrite constraints for concrete paths if required. However, specifying a false path constraint takes precedence over other constraints.

You can use similar constraints for all such paths in DSP Builder blocks for the higher level projects.

When you add synchronizers to DSP Builder designs, the Quartus Prime timing analyzer also provides a metastability report.

#### Related Information

- [Register Bit \(RegBit\)](#) on page 269
- [Register Field \(RegField\)](#) on page 270
- [Shared Memory \(SharedMem\)](#) on page 271

## 11.2. Control

The **Control** block specifies information about the hardware generation environment and the top-level memory-mapped bus interface widths.

*Note:* DSP Builder applies globally the options in the **Control** block to your design.

*Note:* You must include a **Control** block in the top-level model.

**Table 28. Control Block General Parameters**

Parameter	Description
Generate hardware	Turn on to generate output file.
Hardware description language	Specify VHDL or Verilog HDL.
Hardware destination directory	Specify the root directory in which to write the output files. This location can be an absolute path or a relative path (for example, <code>./rtl</code> ). A directory tree is created under this root directory that reflects the names of your model hierarchy.
Use separate working directory for Quartus Prime project	Turn on to create separate working directory.
Generate a single Avalon Conduit interface for the Platform Designer	In v18.1 and earlier, DSP Builder designs that you import and generate in Platform Designer have a single Avalon interface for data, valid, and channel signals. In v19.1 or later, if you regenerate an existing design, turn on (default) this option to preserve the single Avalon Conduit interface. Turn off if you want to connect multiple interfaces in Platform Designer.
Small memory minimum fill	This threshold controls whether the design uses registers or small memories (MLABs) to implement delay lines. DSP Builder uses a small memory only if it fills it with at least the threshold number of bits. On device families that don't support small memories, DSP Builder ignores this threshold.
<i>continued...</i>	

Parameter	Description
Medium memory minimum fill	This threshold controls when the design uses a medium memory (M9K, M10K or M20K) instead of a small memory or registers. DSP Builder uses the medium memory only if it fills it with at least the threshold number of bits.
Large memory minimum fill	This threshold controls whether the design uses a large memory (M144K) instead of multiple medium memories. DSP Builder uses the large memory only when it can fill it with at least the threshold number of bits. <b>Default</b> prevents the design using any M144Ks. On device families that don't support large memories, DSP Builder ignores this threshold.
Multiplier: logic and DSP threshold	Specifies the number of logic elements you want to use to save a multiplier. If the estimated cost of implementing a multiplier in logic is no more than this threshold, DSP Builder implements that multiplier in logic. Otherwise DSP Builder uses a hard multiplier. <b>Default</b> means the design always uses hard multipliers.

**Table 29. Control Block Clock Parameters**

Parameter	Description
Clock signal name	Specifies the name of the system clock signal that DSP Builder uses in the RTL generation, in the <code>_hw.tcl</code> file, and that you see in Platform Designer.
Clock frequency (MHz)	Specifies the system clock rate for the system.
Clock margin (MHz)	Specifies the margin requested to achieve a high system frequency in the fitter. The specified margin does not affect the folding options because the system runs at the rate specified by the <b>Clock frequency</b> parameter setting. Specify a positive clock margin if you need to pipeline your design more aggressively (or specify a negative clock margin to save resources) when you do not want to change the ratio between the clock speed and the bus speed.
Reset signal name	Specifies the name of the reset signal that DSP Builder uses in the RTL generation, the <code>_hw.tcl</code> file, and that you see in Platform Designer.
Reset active	Specifies whether the logic generated is reset with an active high or active low reset signal.
Use default minimum reset pulse width	Turn on to enter a minimum reset value pulse width.
Minimum reset pulse width	Enter a value for the minimum number of system clock cycles for which you assert the reset signal in your target hardware. This setting does not enforce that your design correctly resets in the number of cycles you specify, in particular when you apply reset minimization. You should simulate your design with this value (which DSP Builder applies in the simulation testbench) to confirm that your design works. DSP Builder reset minimization uses a longer minimum reset pulse width to remove resets on the control path. Applying a reset value at an earlier register propagates to later registers during the reset period, without them needing an explicit reset. When you turn <b>Global enable On</b> , DSP Builder enters a large, minimum reset pulse width according to the reset-minimization. When you turn <b>Global enable Off</b> it selects a small minimum reset pulse width as in previous versions of DSP Builder. DSP Builder reports the actual minimum reset pulse width value when it generates your design.

**Table 30. Control Block Testbenches Tab Parameters**

Parameter	Description
Create automatic testbenches	Turn on to generate additional automatic testbench files. These files capture the input and output of each block in a <code>.stm</code> file. DSP Builder creates a test harness ( <code>_atb.vhd</code> ) that simulates the generated RTL alongside the captured data. DSP Builder generates a script ( <code>&lt;model&gt;_atb.do</code> ) that you can use to simulate the design in ModelSim and ensure bit and cycle accuracy between the Simulink model and the generated RTL.
Action on ChannelOut mismatch	Select <b>Error</b> or <b>Warning</b> .

*continued...*



Parameter	Description
Import ModelSim results to MATLAB	Turn on to import the ModelSim simulator results back into MATLAB. Specify: <ul style="list-style-type: none"> <li>The name of the variable to which you assign the device output</li> <li>The name of the variable to which you assign the valid map</li> <li>The name of the function to process the results.</li> </ul>
Floating-point mismatch tolerance	Specify the floating-point mismatch tolerance, which is the largest difference in magnitude for a signal value to be equivalent. Use this field to record cumulative differences in floating-point precision between the Simulink model and the RTL.
Floating-point zero mismatch tolerance	Specify the floating-point mismatch zero tolerance, which is the largest magnitude for a signal value to be equivalent to zero. Use this field to record cumulative differences in floating point precision between the Simulink model and the RTL.
ModelSim input values during reset	Defines the behavior of ModelSim input signals during the period in which DSP Builder holds your design in reset. Use this parameter to test whether the values of the input data during reset affect the behavior of a design with reset-minimization. The Simulink simulation assumes all inputs are zero during reset. Using a non-zero value causes a mismatch if input values during reset affect the RTL behavior. To eliminate this behavior, place a mux on the inputs, which sets them to zero during reset. Only valid if the <b>Minimum reset pulse width</b> is greater than <b>1</b> clock cycle. This parameter can be: <b>0</b> , (default) holds all input data at zero during reset. This matches the Simulink simulation and does not identify any RTL dependence on input values during reset. <b>1</b> , sets all bits of all input data to 1. <b>X</b> , sets all bits of all input data to X. <b>01</b> , alternates all bits of all input data between 0 and 1 on each clock cycle. <b>X01</b> , alternates all bits of all input data between X, 0 and 1 on each clock cycle.

**Table 31. Optimization Tab Parameters**

Sets the reset minimization parameters, which default to **Auto**. When set to **Auto**, DSP Builder turns on **Global enable** and **Floating-point**, if the design targets a HyperFlex device.

Parameter	Value	Description
<b>Reset - Minimization</b>		
Global enable	Auto, on, or off	Globally enables reset minimization. DSP Builder also applies local settings on the <b>SynthesisInfo</b> blocks. Reset minimization applies to all subsystems in your design that include <b>ChannelIn</b> and <b>ChannelOut</b> blocks. DSP Builder does not apply reset minimization to subsystems that include <b>GPIIn</b> and <b>GPOut</b> blocks. When global enable is on, use the <b>SynthesisInfo</b> block to change its behavior at the subsystem level.
Floating-point	Auto, on, or off	Applies reset minimization to all floating-point operators. Use this feature when your design employs floating-point operators that are not control flow.
<b>Uninitialized Memory Content</b>		
Uninitialized means	<b>0 - Initialized to Zeroes</b> or <b>X - Truly Uninitialized</b>	Specify what happens when you turn off <b>Initialize Hardware Memory Blocks with Initial Data Contents</b> on any memory block ( <b>DualMem</b> or <b>SharedMem</b> ) in your design. This parameter does not apply to your memory blocks when you turn on <b>Initialize Hardware Memory Blocks with Initial Data Contents</b> for a memory block. <b>0 - Initialized to Zeroes</b> means the content of each uninitialized memory is zeroes in: <ul style="list-style-type: none"> <li>Simulink simulation</li> <li>HDL simulation</li> <li>Hardware</li> </ul>

*continued...*

Parameter	Value	Description
		<p><b>X - Truly Uninitialized</b> means the content of each uninitialized memory is:</p> <ul style="list-style-type: none"> <li>• Zeroes in Simulink simulation</li> <li>• Xs in HDL simulation</li> <li>• Undefined in hardware.</li> </ul> <p>DSP Builder cannot represent undefined or X initial values it reads from the memory in Simulink simulation. Unless you eliminate such values with suitable logic, you may see a mismatch in HDL simulation.</p> <p>The default value is <b>0 - Initialized to Zeroes</b>, which is benign in simulation.</p>
<b>FIFO Optimizations</b>		
Report FIFO fill level	-	Turn on to report the maximum FIFO depth reached during simulation to the Simulink Diagnostic Viewer. You can use the information to optimize the FIFO buffers.

Options in the **Control** block specify whether hardware generates for your design example and the location of the generated RTL. You can also create automatic RTL testbenches for each subsystem in your design example and specify the depth of signals that DSP Builder includes when your design example simulates in the ModelSim simulator.

You can specify the address and data bus widths that the memory-mapped bus interface use and specify whether DSP Builder stores the high-order byte of the address in memory at the lowest address and the low-order byte at the highest address (big endian), or the high-order byte at the highest address and the low-order byte at the lowest address (little endian).

**Related Information**

- [Reset Minimization](#) on page 197
- [Shared Memory \(SharedMem\)](#) on page 271
- [Dual Memory \(DualMem\)](#) on page 327

**11.2.1. DSP Builder Memory and Multiplier Trade-Off Options**

When your design synthesizes to logic, DSP Builder creates delay blocks, whether explicitly from primitive delays, or in the **IP** library blocks. DSP Builder tries to balance the implementation between logic elements (LEs) and block memories (M9K, M20K, M20K, or M144K). The trade-off depends on the target FPGA family, but as a guideline the default trade-off is set to minimize the absolute silicon area the design uses. You can influence this trade-off.

DSP Builder converts multipliers with a single constant input into balanced adder trees, which occurs automatically where the depth of the tree is not greater than 2. If the depth is greater than 2, DSP Builder compares the hard multiplier threshold with the estimated size of the adder tree, which is generally much lower than the size of a full soft multiplier. If DSP Builder combines two non-constant multipliers followed by an adder into a single DSP block, DSP Builder does not convert the multiplier into LEs, even if a large threshold is present.

**11.3. Device**

The **Device** block indicates a particular Simulink subsystem as the top-level design of an FPGA device. It also specifies a particular device and allows you to specify the target device and speed grade for the device.

**Note:** All blocks in subsystems below this level of hierarchy, become part of the RTL design. All blocks above this level of hierarchy become part of the testbench.

You can hierarchically separate parts of the design into synthesizable systems. You must use a **Device** block, which sets the device family, part number, speed grade, and so on, to indicate the top-level synthesizable system.

You can further hierarchically split the synthesizable system into Primitive subsystems for Primitive blocks and **IP** blocks.

You can optionally include **LocalThreshold** blocks to override threshold settings defined higher up the hierarchy.

DSP Builder generates project files and scripts that relate to this level of hierarchy. All blocks in subsystems below this level become part of the RTL design. All blocks above this level of hierarchy become part of the testbench.

You can insert multiple **Device** blocks in non-overlapping subsystems to use multiple FPGAs in the same design. You can mix device families freely.

**Table 32. Parameters for the Device Block**

Parameter	Description
Device family	Select the required target device family.
Device	Select the specific device.
Family member	Specify the device member as free-form text or enter <b>AUTO</b> for automatic selection.
Speed grade	Select the speed grade for the FPGA target device.

## 11.4. Edit Params

The **Edit Params** block opens the `setup_<model_name>.m` file. If you call your set up script `setup_<model_name>.m` you can use the **Edit Params** block in your design as a shortcut to open the set-up file.

The **Edit Params** block is available as a functional block in the Simulink library browser. To view it open the library, by right clicking on the **Design Configuration Blocks** library in the Simulink Library Browser and select **Open Design Configuration Blocks Library**.

Examples of **Edit Params** blocks are in many of the design examples.

To call your script automatically:

- When your model opens, add a `PreloadFcn` reference to your script in the **Callbacks** tab of your Model Properties in Simulink.
- At the start of a simulation run, add a `InitFcn` reference to your script in the **Callbacks** tab of your Model Properties in Simulink.

## 11.5. LocalThreshold

The **LocalThreshold** block allows hierarchical overrides of the global clock margin and threshold settings set on the **Control** and **Signals** blocks.

You can place the **LocalThreshold** block anywhere in your design to define over-ride values for the margin and threshold settings for that subsystem and any embedded subsystems. You can over-ride these values further down in the hierarchy by implementing more **LocalThreshold** blocks.

For example, you can specify different clock margins for different regions of your design.

**Table 33. Parameters for the LocalThreshold Block**

Parameter	Description
Clock margin (MHz)	Specifies the margin to influence the tradeoff between performance and resources. The specified margin does not affect the folding options because the system runs at the rate specified by the <b>Clock frequency</b> parameter setting. Specify a positive clock margin if you need to pipeline your design more aggressively (or specify a negative clock margin to save resources) when you do not want to change the ratio between the clock speed and the bus speed.
Generation Thresholds	
Small memory minimum fill	This threshold controls whether registers or small memories (MLABs) implement delay lines. DSP Builder uses a small memory only if it fills it with at least the threshold number of bits. On device families that don't support small memories, DSP Builder ignores this threshold.
Medium memory minimum fill	This threshold controls when the design uses a medium memory (M9K, M10K or M20K) in place of a small memory or registers. DSP Builder uses the medium memory only if it fills it with at least the threshold number of bits.
Large memory minimum fill	This threshold controls whether the design uses a large memory (M144K) instead of multiple medium memories. DSP Builder uses the large memory only when it can fill it with at least the threshold number of bits. <b>Default</b> prevents the design using any M144Ks. On device families that don't support large memories, DSP Builder ignores this threshold.
Multiplier: logic and DSP threshold	Specifies the number of logic elements you want to use to save a multiplier. If the estimated cost of implementing a multiplier in logic is no more than this threshold, DSP Builder implements that multiplier in logic. Otherwise DSP Builder uses a hard multiplier. <b>Default</b> means the design always uses hard multipliers.
Apply Karatsuba method to complex multiply blocks	Implements this equation: $(a+jb) * (c+jd) = (a-b)*(c+d) - a*d + b*c + j(a*d + b*c)$ . DSP Builder includes internal preadder steps into DSP blocks but you see bit growth in the multipliers.

This block has no inputs or outputs.

**Related Information**

[DSP Builder Memory and Multiplier Trade-Off Options](#) on page 226

## 12. IP Library

---

Use the DSP Builder advanced blockset **IP** library blocks to implement full IP functions. Only use these blocks outside of primitive subsystems.

1. [Channel Filter and Waveform Library](#) on page 229
2. [Dependent Delay Library](#) on page 257
3. [FFT IP Library](#) on page 258

### 12.1. Channel Filter and Waveform Library

The DSP Builder advanced blockset **Channel Filter and Waveform** library contains several decimating and interpolating cascaded integrator comb (CIC), and finite impulse response (FIR) filters including single-rate, multirate, and fractional rate FIR filters.

Multirate filters are essential to the up and down conversion tasks that modern radio systems require. Cost effective solutions to many other DSP applications also use multirate filters to reduce the multiplier count.

FIR filter memory-mapped interfaces allow you to read and write coefficients directly, easing system integration.

1. [DSP Builder FIR and CIC Filters](#) on page 230
2. [DSP Builder FIR Filters](#) on page 233
3. [Channel Viewer \(ChanView\)](#) on page 235
4. [Complex Mixer \(ComplexMixer\)](#) on page 236
5. [Decimating CIC](#) on page 238
6. [Decimating FIR](#) on page 239
7. [Fractional Rate FIR](#) on page 241
8. [Interpolating CIC](#) on page 244
9. [Interpolating FIR](#) on page 245
10. [NCO](#) on page 248
11. [Real Mixer \(Mixer\)](#) on page 252
12. [Scale](#) on page 254
13. [Single-Rate FIR](#) on page 255

### 12.1.1. DSP Builder FIR and CIC Filters

Finite impulse response (FIR) and cascaded integrator comb (CIC) filters share many common features and use advanced high-level synthesis techniques to generate filters with higher clock speeds, lower logic, multiplier, and memory counts. Using these high clock rates allows you to reduce your costs by choosing smaller FPGAs.

1. [Common CIC and FIR Filter Features](#) on page 230
2. [Updated Help](#) on page 231
3. [Half-Band and L-Band Nyquist FIR Filters](#) on page 232
4. [Parameterization of CIC and FIR Filters](#) on page 232
5. [Setting and Changing FIR Filter Coefficients at Runtime in DSP Builder](#) on page 233

#### 12.1.1.1. Common CIC and FIR Filter Features

- Filter length from 1 to unlimited taps
- Data input width from 2 to 32 bits
- Data output width from 4 to 64 bits
- Multichannel (up to 256 channels)
- Powerful MATLAB integration
- Simulink fixed-point integration
- Automatic pipelining
- Plug and play connectivity
- Simplified timing closure
- Generates updated help for your parameters

*Note:* Each channel is an independent data source. In an IF modem design, two channels are required for the complex pair from each antenna.

*Note:* This library does not support complex data.

*Note:* All input data and coefficients must be fixed-point data.

#### Automatic Pipelining

The required system clock frequency, and the device family and speed grade determine the maximum logic depth permitted in the output RTL. DSP Builder pipelines functions such as adders by splitting them into multiple sections with a registered carry between them. This pipelining decreases the logic depth allowing higher frequency operation.

#### High-Speed Operation

The DSP Builder filter generator is responsive to the system clock frequency, therefore timing closure is much easier to achieve. The generator uses heuristics that ensure the logic can run at the desired system clock frequency on the FPGA. You can help timing closure by adding more clock margin, resulting in additional pipelining that shortens the critical paths. The FPGA structures such as internal multiplier and memory delays determine the maximum clock frequencies.

### Scalability

In some cases, the aggregate sample rate for all channels may be higher than the system clock rate. In these cases, the filter has multiple input or output buses to carry the additional data, so DSP Builder implements this requirement in the Simulink block by increasing the vector width of the data signals.

### Coefficient Generation

You can generate filter coefficients using a MATLAB function that reloads at run time with the memory-mapped interface registers. For example, the Simulink fixed-point object **fi(fir1(49, 0.3),1,18,19)**

### Channelization

The generated help page for the block shows the input channel data format and output data channel format that a FIR or CIC filter uses, after you run a Simulink simulation.

#### 12.1.1.2. Updated Help

After you run a simulation, DSP Builder updates the help pages with specific information about each instance of a block. This updated help overrides the default help link. To find the updated help click on the help link on the block after simulation.

This updated help includes a link back to the help for the general block and the following information about the generated FIR instance:

- Date and time of generation
- The version number and revision for the FIR
- Number of physical input and output data buses
- Bit width of data output.
- Number of different phases
- Implementation folding. The number of times that the design uses each multiplier per sample to reduce the implementation size.
- Filter utilization. For some sample rates and some interpolation/decimation settings, the filter may stall internally one or more cycles. The filter utilization is the percentage of time that the filter is actively working, assuming that the input arrives at the specified data rate.
- Tap utilization. When some filters are folded, the design may have extra unused taps. The extra taps increase the filter length with no hardware resource increase.
- Latency. The depth of pipelining added to the block to meet the target clock frequency on the chosen target device.
- Parameters table that lists the system clock, clock margin, and all FIR input parameters.
- Port interface table.
- Input and output data format. An ASCII rendering of the input and output channelized data ordering.

The updated help includes the following information about the CIC instance:

- Date and time of generation
- The version number and revision for the CIC
- Number of integrators. Depending on the input data rate and interpolation factor the number of integrator stages DSP Builder needs to process the data may be more than 1. In these instances, the integrator sections of the filter duplicate (vectorize) to satisfy the data rate requirement.
- Calculated output bit width. The width in bits of the (vectorized) data output from the filter.
- Calculated stage bit widths. Each stage in the filter has precise width in bits requirements—N comb sections followed by N integrator sections.
- The gain through the CIC filter. CIC filters usually have large gains that you must scale back.
- Comb section utilization. In the comb section, the data rate is lower, so that you can perform more resource sharing. This message indicates the efficiency of the subtractor usage.
- Integrator section utilization. In the integrator section, the data rate is higher, so that you can perform less resource sharing. This message indicates the efficiency of the adder usage.
- The latency that this block introduces.
- Parameters table that lists the decimation rate, number of stages, differential delay, number of channels, clock frequency, and input sample rate parameters.
- Port interface table.
- Input and output data format.

### 12.1.1.3. Half-Band and L-Band Nyquist FIR Filters

Some filtering functions can use a half-band filter where nearly half of the coefficients are zero. The half-band support uses these extra zeros to further reduce the number of multipliers, and thereby reduce the filter cost.

The generalized form of these filters is L-band Nyquist filters, in which every Lth coefficient is zero counting out from the center tap. DSP Builder also supports these structures and can often reduce the number of multipliers required in a filter.

### 12.1.1.4. Parameterization of CIC and FIR Filters

The system specification, including the channel count and sample rates, determines the main parameters for a filter. The enclosing Simulink mode infers the remaining parameters such as data widths and system clock rates. Any changes to these parameters ripple through your design, changing the system performance without you having to update all the components. You can express any of the parameters as MATLAB expressions, to rapidly parameterize a whole system.

The hardware generation techniques create efficient filters with combinations of parameters, such as a symmetric 3-band FIR filter with seven channels and 100 cycles to process a sample from each channel. Hardware generation is fast and can change at runtime with every Simulink simulation, so that the edit simulation loop time is much reduced, improving productivity.



### 12.1.1.5. Setting and Changing FIR Filter Coefficients at Runtime in DSP Builder

1. Set the base address of the memory-mapped coefficients with the **Base address** parameter.
2. Set the filter coefficients by entering a Simulink fixed-point array into the **Coefficients** parameter.
3. Generate a vector of coefficients either by entering an array of numbers, or using one of the many MATLAB functions to build the required coefficients.
4. Update the parameters through a processor interface during run time using the **BusStimulus** block. Alternatively, update the parameters from you model by exposing hidden processor interface ports (turn on **Expose Bus Ports**).

### 12.1.2. DSP Builder FIR Filters

1. [FIR Filter Avalon-MM Interfaces](#) on page 233
2. [Reconfigurable FIR Filters](#) on page 234

#### 12.1.2.1. FIR Filter Avalon-MM Interfaces

All DSP Builder FIR blocks can provide Avalon memory-mapped interfaces to coefficients, allowing you to change the coefficient values at run time.

- To allow read, write, or readwrite access to coefficients from system bus interfaces, select **Read**, **Write**, or **Readwrite** for the FIR block **Bus mode**. Select **Constant** to disable this interface.
- Specify the bus address width and data type on the FIR block and on the **Avalon memory-mapped Settings** block. FIR blocks automatically generate the appropriate bus logic and interface for your design.
- To place the bus logic on a separate clock domain specify **Separate bus clock** in the **Avalon memory-mapped Settings** block.
- DSP Builder hides the agent interface ports by default in the simulation model. Use the **BusStimulus** block to access this interface during simulation (similar to **Interface** blocks). Use the base address specified on the FIR block for accessing the coefficients.
- In generated RTL, DSP Builder adds these ports to FIR blocks and routes them to the Avalon memory-mapped agent interface. Set the width of data and address ports (**Avalon Data** and **Address Width**) in the **Avalon memory-mapped Settings** block.

*Note:* If the FIR coefficient is wider than Avalon memory-mapped data width, the design requires several accesses to write or read a single coefficient.

In your higher level system, access FIR coefficients through the agent interface at the base address you specified on the FIR block.

*Note:* The FIR base address is now an offset from the base address assigned to the agent interface in your Platform Designer system.

When you expose bus interface ports in the Simulink design (turn on **Expose Bus Interface**), a valid sub-set of Avalon memory-mapped agent interface ports appears on the block based on the selected bus mode. You can now make direct connections to these ports in the Simulink model for accessing the coefficients. The FIR coefficient width sets the data ports (write and read). DSP Builder places bus agent logic on the system clock domain.

**Table 34. FIR Filter Avalon-MM Ports**

Your design contains `address` when **Read/Write Mode** is not **Constant**, `write` when **Read/Write Mode** is **Write** or **ReadWrite**, `read` and `valid` when **Read/Write Mode** is **Read** or **ReadWrite**

Name	Direction	Description
<code>address</code>	Input	Address of the request. DSP Builder adds <code>address</code> to your design when <b>Bus Mode</b> is not set to <b>Constant</b> . The port width depends on the <b>Bus Address Width</b> in the <b>Avalon Memory-Mapped Settings</b> block. For the first coefficient use the <b>Base Address</b> you specify for the block and for the last one use: <b>Base Address + Number Of Coefficients -1</b>
<code>data</code>	Input	Write data. The port width depends on the coefficient width of the FIR block. Set <code>data</code> and <code>address</code> and assert <code>write</code> port simultaneously to initiate a write request.
<code>read</code>	Input	Read enable. Set the <code>address</code> to a valid address and assert this single-bit input simultaneously every time you want to initiate a read request. After sending a read request, wait for <code>valid</code> to be asserted indicating that read data is available on <code>readdata</code> . You don't need to wait for the completion of the first read request to initiate a second read request. The agent supports pipelined reads. DSP Builder provides the responses in the exact same order you send the read requests.
<code>write</code>	Input	Write enable. Assert this single-bit input every time you need to initiate a write request. Do not assert <code>read</code> and <code>write</code> ports at the same time, otherwise, you see undefined behavior.
<code>readdata</code>	Output	Read data. DSP Builder sets the port width the same as the FIR coefficient width. This output provides data for read responses. Only capture this output if you assert <code>valid</code> output.
<code>valid</code>	Output	Read data valid. This single-bit output indicates that valid data is available on <code>readdata</code> .

**Related Information**

[Avalon Memory-Mapped Agent Settings \(AvalonMemoryMappedAgentSettings\)](#) on page 221

**12.1.2.2. Reconfigurable FIR Filters**

Trades off the bandwidth of different channels at runtime.

The input rate determines the bandwidth of the FIR. If you turn off **Reconfigurable carrier** (nonreconfigurable FIR), the IP core allocates this bandwidth equally amongst each channel. The reconfigurable FIR feature allows the IP core to allocate the bandwidth manually. You set these allocations during parameterization and you can change which allocation the IP core uses at run-time using the mode signal. You can use one channel's bandwidth to process a different channel's data. You specify the allocation by listing the channels you want the IP core to process in the mode mapping. For example, a mode mapping of 0,1,2,2 gives channel 2 twice the bandwidth of channel 0 and 1, at the cost of not processing channel 3.

### 12.1.3. Channel Viewer (ChanView)


The **ChanView** block deserializes the bus on its inputs to produce a configurable number of output signals that DSP Builder does not apply TDM protocol.

You can use a **ChanView** block in a testbench to visualize the contents of the TDM protocol. It produces synthesizable RTL, so you can use it anywhere in your design.

When a single channel is input, the **ChanView** block strips out all the non-valid samples, thus cleaning up the display in the Simulink scope.

The channel outputs are not aligned. For example, if you have input channels *c0* and *c1* on a single wire and view both channels, the output is not aligned.

**Figure 90. Channel Viewer Output for Two Channels on a Single Wire**



*Note:* You can add delays after **ChanView** blocks if you want to realign the output channels.

**Table 35. Parameters for the ChanView Block**

Parameter	Description
Number of input channels	Specifies the number of unique channels the block can process. The design does not use this parameter unless the data bus is a vector or the folding factor is greater than the number of channels. If the data bus is a vector, this value determines which vector element contains the correct channel.
Output channels	A vector that controls the input channels to decode and present as outputs. The number of outputs equals the length of this vector, and each output corresponds to one channel in order.

**Table 36. Port Interface for the ChanView Block**

Signal	Direction	Description
q	Input	The data input to the block. This signal may be a vector. This block does not support floating-point types.
v	Input	Indicates validity of data input signals. If v is high, the data on the wire is valid.
c	Input	Indicates the channel of the data input signals. If v is high, c indicates which channel the data corresponds to.
cn	Output	Each output is a deserialized version of the channel contained on the TDM bus. The output value is updated on each clock cycle that has valid data when the channel matches the required channel.
ov	Output	Optional. Pulses 1 at last cycle of a frame (when all held channel output signals have correct value for the frame) provided valid is high throughout the frame data.

After DSP Builder runs a simulation, it updates the help pages with specific information about each instance of a block. For resource usage, on the DSP Builder menu, point to **Resources**, and click **Design**.

**Table 37. Messages for the ChanView Block**

Message Example	Description
Written on Tue Feb 19 11:25:27 2008	Date and time when you ran this file.
Latency is 2	The latency that this block introduces.
Port interface table	Lists the port interfaces to the <b>ChanView</b> block.

### 12.1.4. Complex Mixer (ComplexMixer)

The **ComplexMixer** block performs a complex by complex multiply on streams of data and it splits the inputs and outputs into their real and imaginary components. This function can shift the frequency of a data stream in a digital up converter, where the first complex data is the *i* and *q* data and the second complex data is the cosine and sine data provided by an NCO.

The **ComplexMixer** block multiplies a complex input stream by a synchronized complex data stream, sample by sample.

You can use this block in a digital up converter for a radio system or a general purpose DSP application. The data has fixed-point types, and the output is the implied full precision fixed-point type.

**Note:** You can easily replicate the **ComplexMixer** block with a **Multiply** block that takes complex inputs within a **Primitive** subsystem.

The **ComplexMixer** performs element-by-element multiplication on  $n$  channels and  $m$  frequencies.

The system specification, including such factors as the channel count and sample rates, determines the main parameters for this block. The input sample rate of the block determines the number of channels present on each input wire and the number of wires:

$$\text{Number of Channels per wire} = \text{Clock\_Rate} / \text{Sample\_Rate}$$

$$\text{Number of Wires} = \text{ceiling}(\text{Chan\_Count} \times \text{Sample\_Rate} / \text{Clock\_Rate})$$

For example, a sample rate of 60 MSPS and system clock rate of 240 MHz gives four samples to be TDM on to each input wire.

If a wire has more channels than TDM slots available, the input wire is a vector of sufficient width to hold all the samples. Similarly, the number of frequencies (the number of complex numbers) determines the width of the sine and cosine inputs. The number of results produced by the **ComplexMixer** is the product of the sample input vector and the frequency vector. The results are TDM on to the  $i$  and  $q$  outputs in a similar manner to the inputs.

**Table 38. Parameters for the ComplexMixer Block**

Parameter	Description
Input Rate Per Channel (MSPS)	The data rate per channel measured in millions of samples per second.
Number of Complex Channels	The number of complex input channels.
Number of Frequencies	The number of complex frequencies in the multiplier.

**Table 39. Port Interface for the ComplexMixer Block**

Signal	Direction	Description
$i$	Input	The real (in phase) component of the complex data input. If you request more channels than can fit on a single bus, this signal is a vector. The width in bits inherits from the input wire.
$q$	Input	The imaginary (quadrature phase) component of the complex data input. If you request more channels than can fit on a single bus, this signal is a vector. The width in bits inherits from the input wire.
$v$	Input	Indicates validity of data input signals. If $v$ is high, the data on the $a$ wire is valid.
$c$	Input	Indicates the channel of the data input signals. If $v$ is high, $c$ indicates the data channel data.
$\sin$	Input	The imaginary part of the complex number. For example, the NCO's sine output.
$\cos$	Input	The real part of the complex number. For example, the NCO's cosine output.
$i$	Output	The in-phase (real) output of the mixer, which is $(i \times \cos - q \times \sin)$ . If you request more channels than can fit on a single bus, this signal is a vector. The width in bits is wide enough for the full precision result.
$q$	Output	The quadrature phase (imaginary) output of the mixer, which is $(i \times \sin + q \times \cos)$ . If you request more channels than can fit on a single bus, this signal is a vector. The width in bits is wide enough for the full precision result.
$v$	Output	Indicates validity of data output signals.
$c$	Output	Indicates the channel of the data output signals.

The **ComplexMixer** block performs the multiplication on corresponding components; the **RealMixer** block does not. The **ComplexMixer** block uses modulo indexing if one vector is shorter than another. Hence, the output vector width is the maximum of the widths of the input vectors. The **RealMixer** block performs a full outer product on the input vectors. The number of components in the output vector is the product of the width of the input vectors for sin and cos (must be the same) and the width of the input vector for a.

### 12.1.5. Decimating CIC

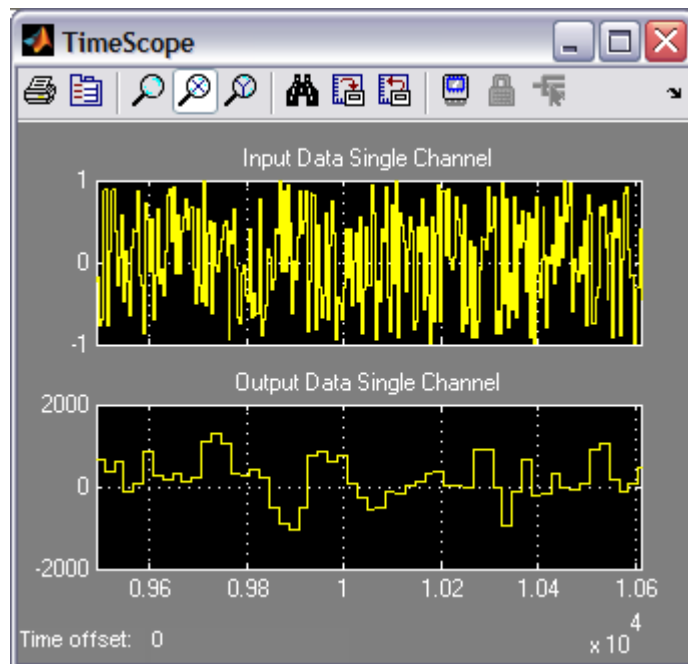
The **DecimatingCIC** block implements a highly efficient multichannel CIC filter across a broad range of parameters directly from a Simulink model. The **DecimatingCIC** block performs filtering on a stream of multichannel input data and produces a stream of output data with decreased sampling frequency.

You can use the **DecimatingCIC** block in a digital down converter for a radio system or a general purpose DSP application. The coefficients and input data are fixed-point types, and the output is the implied full-precision fixed-point type. You can reduce the precision with a separate **Scale** block, which can perform rounding and saturation to provide the required output precision.

The **DecimatingCIC** block supports rate changes from two upwards.

The **DecimatingCIC** has a lower output sample rate than the input sample rate by a factor  $D$ , where  $D$  is the decimation factor. Usually, the **DecimatingCIC** discards  $(D-1)$  out of  $D$  output samples thus lowering the sample rate by a factor  $D$ . The physical implementation avoids performing additions leading to these discarded samples, reducing the filter cost.

**Figure 91. Decimate by 5 Filter Decreasing Sample Rate of a Random Noise Input**



**Table 40. Parameters for the DecimatingCIC Block**

Parameter	Description
Input rate per channel	Specifies the sampling frequency of the input data per channel measured in millions of samples per second (MSPS).
Number of channels	Specifies the number of unique channels to process.
Number of stages	Specifies the number of comb and integrator stages.
Decimation factor	Specifies the decimation factor 1/(integer). (An integer greater than 1 implies interpolation.)
Differential delay	Specifies the differential delay.

**Table 41. Port Interface for the DecimatingCIC Block**

Signal	Direction	Description
a	Input	The fixed-point data input to the block. If you request more channels than can fit on a single bus, this signal is a vector. The width in bits is inherited from the input wire.
v	Input	Indicates validity of the data input signals. If v is high, the data on the a wire is valid.
c	Input	Indicates channel of data input signals. If v is high, c indicates which channel the data corresponds to.
bypass	Input	When this input asserts, the input data is zero-stuffed and scaled by the gain of the filter, which is useful during hardware debugging.
q	Output	The data output from the block. If you request more channels than can fit on a single bus, this signal is a vector. The width in bits is a function of the input width in bits and the parameterization.
v	Output	Indicates validity of data output signals.
c	Output	Indicates channel of data output signals.

**Related Information**

[DSP Builder FIR and CIC Filters](#) on page 230


### 12.1.6. Decimating FIR

The **DecimatingFIR** block implements a highly efficient multichannel FIR filter across a broad range of parameters directly from a Simulink model. A memory-mapped interface allows you to read and write coefficients directly, easing system integration. The **Decimating FIR** block performs filtering on a stream of multichannel input data and produces a stream of output data with increased sampling frequency.

Use the **Decimating FIR** block in a digital down converter for a radio system or a general purpose DSP application. The coefficients and input data are fixed-point types, and the output is the implied full precision fixed-point type. You can reduce the precision by using a separate **Scale** block, which can perform rounding and saturation to provide the required output precision.

The **Decimating FIR** block supports rate changes from two upwards, coefficient width in bits from 2 to 32 bits, half-band and L-band Nyquist filters, real and complex filters, symmetry and anti(negative)-symmetry.

**Figure 92. Decimating by 5 Filter Decreasing Sample Rate of a Sine Wave Input**



At each sample time  $k$ , the new output  $y$ , is calculated by multiplying coefficients  $a$ , by the recent past values of the input  $x$ .

The **Decimating FIR** has a lower output sample rate than the input sample rate by a factor,  $D$ , the decimation factor. The decimating FIR discards  $D-1$  out of  $D$  output samples, thus lowering the sample rate by a factor  $D$ .

The physical implementation avoids performing multiplications with these zero samples, reducing the filter cost.

**Table 42. Parameters for the DecimatingFIR Block**

Parameter	Description
Input rate per channel	Specifies the sampling frequency of the input data per channel measured in millions of samples per second (MSPS).
Decimation	Specifies the decimation rate. Must be an integer.
Number of channels	Specifies the number of unique channels to process.
Symmetry	You can select <b>Symmetrical</b> or <b>Anti-Symmetrical</b> coefficients. Symmetrical coefficients can result in hardware resource savings over the asymmetrical version.
Coefficients	You can specify the filter coefficients using a Simulink fixed-point object <b>fi(0)</b> . The data type of the fixed-point object determines the width and format of the coefficients. The length of the array determines the length of the filter. For example, <b>fi(fir1(49, 0.3),1,18,19)</b>
Base address	You can memory map the filter's coefficients into the address space of the system. This field determines the starting address for the coefficients. It is specified as a MATLAB <b>double</b> type (decimal integer) but you can use a MATLAB expression to specify a hexadecimal or octal type if required.

*continued...*



Parameter	Description
Read/Write mode	You can allow <b>Read</b> , <b>Write</b> , or <b>Read/Write</b> access from the system interface. Turn on <b>Constant</b> to map coefficients to the system address space.
Filter structure	You can select <b>Use All Taps</b> , <b>Half Band</b> , or other specified band (from <b>3rd Band</b> to <b>46th Band</b> ).
Expose Avalonmemory-mapped agent in Simulink	Allows you to reconfigure coefficients without Platform Designer. Also, it allows you to reprogram multiple FIR filters simultaneously. Turn on to show the Avalon-MM inputs and outputs as normal ports in Simulink. The <b>Read/Write mode</b> decides the valid subset of Avalon memory-mapped agent ports that appear on the block. If you select <b>Constant</b> , the block shows no Avalon-MM ports.
Reconfigurable channels	Turn on for a reconfigurable FIR filter.
Channel mapping	Enter parameters as a MATLAB 2D array for reconfigurable FIR filter. Each row represents a mode; each entry in a row represents the channel input on that time slot. For example, [0,0,0,0;0,1 2,3] gives the first element of the second row as 0, which means DSP Builder processes channel 0 on the first cycle when the FIR is set to mode 1.

For more information about Simulink fixed-point objects and MATLAB functions, refer to the MATLAB Help.

**Table 43. Port Interface for the DecimatingFIR Block**

Signal	Direction	Description
a	Input	The fixed-point data input to the block. If you request more channels than can fit on a single bus, this signal is a vector. The width in bits is inherited from the input wire.
v	Input	Indicates validity of the data input signals. If v is high, the data on the a wire is valid.
c	Input	Indicates the channel of the data input signals. If v is high, then c indicates which channel the data corresponds to.
m	Input	Indicates a reconfigurable filter.
b	Input	Indicates multibank filter. This input appears when you add a second filter definition to the <b>Coefficients</b> parameter in the parameters dialog box.
q	Output	The fixed-point filtered data output from the block. If you request more channels than can fit on a single bus, this signal is a vector. The width in bits is a function of the input width in bits and the parameterization.
v	Output	Indicates validity of data output signals.
c	Output	Indicates the channel of the data output signals. The output data can be non-zero when v is low.

#### Related Information

- [DSP Builder FIR and CIC Filters](#) on page 230
- [DSP Builder FIR Filters](#) on page 233

### 12.1.7. Fractional Rate FIR

The **FractionalRateFIR** block implements a highly efficient multichannel FIR filter across a broad range of parameters directly from a Simulink model. A memory-mapped interface allows you to read and write coefficients directly, easing system integration. The **FractionalRateFIR** block performs filtering on a stream of multichannel input data and produces a stream of output data with increased sampling frequency.

You can use the **FractionalRateFIR** block in a digital down converter for a radio system or a general purpose DSP application. The coefficients and input data are fixed-point types, and the output is the implied full precision fixed-point type. You can reduce the precision by using a separate **Scale** block, which can perform rounding and saturation to provide the required output precision.

The **FractionalRateFIR** block supports:

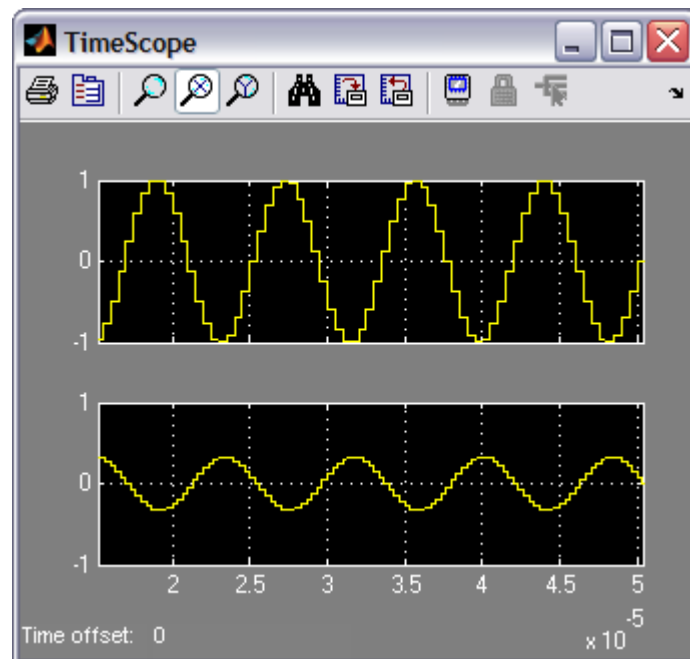
- Interpolation rate changes and decimation rate changes from two upwards
- Rational fractional rate changes
- Coefficient width in bits from 2 to 32 bits
- Half-band and L-band Nyquist filters
- Symmetry and anti(negative)-symmetry.

In the basic filter operation, at each sample time,  $k$ , the new output  $y$ , is calculated by multiplying coefficients  $a$ , by the recent past values of the input  $x$ .

The **FractionalRateFIR** has a modified output sample rate that differs from the input sample rate by a factor,  $I/D$ , where  $I$  is the interpolation rate and  $D$  is the decimation factor. Usually, the fractional rate interpolates by a factor  $I$  by inserting  $(I-1)$  zeros before performing the filter operation. Then the FIR discards  $D-1$  out of  $D$  output samples, thus lowering the sample rate by a factor  $D$ .

The physical implementation avoids performing multiplications with these zero samples, reducing the filter cost.

**Figure 93. Sample Rate of a Sine Wave Input Interpolated by 3 and Decimated by 2**



**Table 44. Parameters for the FractionalRateFIR Block**

Parameter	Description
Input rate per channel	Specifies the sampling frequency of the input data per channel measured in millions of samples per second (MSPS).
Interpolation	Specifies the interpolation rate. Must be an integer.
Decimation	Specifies the decimation rate. Must be an integer.
Number of channels	Specifies the number of unique channels to process.
Symmetry	You can select <b>Symmetrical</b> or <b>Anti-Symmetrical</b> coefficients. Symmetrical coefficients can result in hardware resource savings over the asymmetrical version.
Coefficients	You can specify the filter coefficients using a Simulink fixed-point object <b>fi(0)</b> . The data type of the fixed-point object determines the width and format of the coefficients. The length of the array determines the length of the filter. For example, <b>fi(fir1(49, 0.3),1,18,19)</b> .
Base address	You can memory map the filter's coefficients into the address space of the system. This field determines the starting address for the coefficients. It is specified as a MATLAB <b>double</b> type (decimal integer) but you can use a MATLAB expression to specify a hexadecimal or octal type if required.
Read/Write mode	You can allow <b>Read</b> , <b>Write</b> , or <b>Read/Write</b> access from the system interface. Turn on <b>Constant</b> to map coefficients to the system address space.
Filter structure	You can select <b>Use All Taps</b> , <b>Half Band</b> , or a specified band (from <b>3rd Band</b> to <b>46th Band</b> ).
Expose Avalon memory-mapped agent in Simulink	Allows you to reconfigure coefficients without Platform Designer. Also, it allows you to reprogram multiple FIR filters simultaneously. Turn on to show the Avalon-MM inputs and outputs as normal ports in Simulink. The <b>Read/Write mode</b> decides the valid subset of Avalon memory-mapped agent ports that appear on the block. If you select <b>Constant</b> , the block shows no Avalon-MM ports.
Reconfigurable channels	Turn on for a reconfigurable FIR filter.
Channel mapping	Enter parameters as a MATLAB 2D array for reconfigurable FIR filter. Each row represents a mode; each entry in a row represents the channel input on that time slot. For example, [0,0,0,0;0,1 2,3] gives the first element of the second row as 0, which means DSP Builder processes channel 0 on the first cycle when the FIR is set to mode 1.

**Table 45. Port Interface for the FractionalRateFIR Block**

Signal	Direction	Description
a	Input	The fixed-point data input to the block. If you request more channels than can fit on a single bus, this signal is a vector. The width in bits is inherited from the input wire.
v	Input	Indicates validity of the data input signals. If v is high, the data on the a wire is valid.
c	Input	Indicates the channel of the data input signals. If v is high, c indicates which channel the data corresponds to.
b	Input	Indicates multibank filter. This input appears when you add a second filter definition to the <b>Coefficients</b> parameter in the parameters dialog box.
m	Input	Indicates reconfigurable filter.
q	Output	The fixed-point filtered data output from the block. If you request more channels than can fit on a single bus, this signal is a vector. The width in bits is a function of the input width in bits and the parameterization.
v	Output	Indicates validity of data output signals. The output data can be non-zero when v is low.
c	Output	Indicates the channel of the data output signals. The output data can be non-zero when v is low.

### Related Information

- [DSP Builder FIR and CIC Filters](#) on page 230
- [DSP Builder FIR Filters](#) on page 233

## 12.1.8. Interpolating CIC

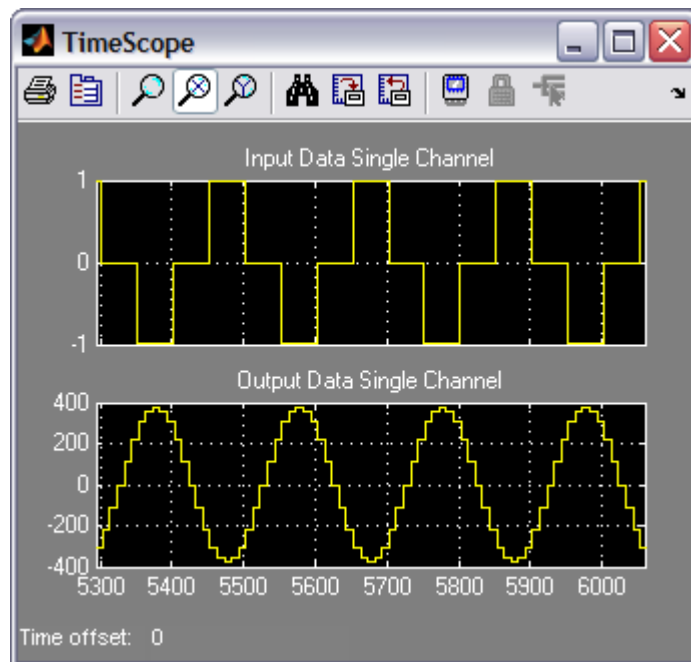
The **InterpolatingCIC** block implements a highly efficient multichannel cascaded integrator-comb filter across a broad range of parameters directly from a Simulink model. The **InterpolatingCIC** block performs filtering on a stream of multichannel input data and produces a stream of output data with increased sampling frequency.

You can use the **InterpolatingCIC** block in a digital up converter for a radio system or a general purpose DSP application. The coefficients and input data are fixed-point types, and the output is the implied full precision fixed-point type. You can reduce the precision by using a separate **Scale** block, which can perform rounding and saturation to provide the required output precision.

The **InterpolatingCIC** block supports rate changes from two upwards.

The **InterpolatingCIC** has a higher output sample rate than the input sample rate by a factor  $I$ , where  $I$  is the interpolation rate. Usually, the **InterpolatingCIC** inserts  $(I-1)$  zeros for every input sample, thus raising the sample rate by a factor  $I$ .

**Figure 94. Interpolate by 5 Filter Increasing Sample Rate of a Sine Wave Input**



**Table 46. Parameters for the InterpolatingCIC Block**

Parameter	Description
Input rate per channel	Specifies the sampling frequency of the input data per channel measured in millions of samples per second (MSPS).
Number of channels	Specifies the number of unique channels to process.
Number of stages	Specifies the number of comb and integrator stages.
Interpolation factor	Specifies the interpolation factor. Must be an integer.
Differential delay	Specifies the differential delay.
Final decimation	You can optionally specify a final decimation by 2 to allow interpolation rates which are multiples of 0.5. The decimation works by simply throwing away data values. Only use this option to reduce the number of unique outputs the CIC generates.

**Table 47. Port Interface for the InterpolatingCIC Block**

Signal	Direction	Description
a	Input	The fixed-point data input to the block. If you request more channels than can fit on a single bus, this signal is a vector. The width in bits is inherited from the input wire.
v	Input	Indicates validity of the data input signals. If v is high, the data on the a wire is valid.
c	Input	Indicates the channel of the data input signals. If v is high, c indicates which channel the data corresponds to.
bypass	Input	When this input is asserted, the input data is zero-stuffed and scaled by the gain of the filter. This option can be useful during hardware debug.
q	Output	The fixed-point filtered data output from the block. If you request more channels than can fit on a single bus, this signal is a vector. The width in bits is a function of the input width in bits and the parameterization.
v	Output	Indicates validity of data output signals. The output data can be non-zero when v is low.
c	Output	Indicates the channel of the data output signals. The output data can be non-zero when v is low.

### Related Information

[DSP Builder FIR and CIC Filters](#) on page 230

## 12.1.9. Interpolating FIR

The **InterpolatingFIR** block implements a highly efficient multichannel FIR filter across a broad range of parameters directly from a Simulink model. A memory-mapped interface allows you to read and write coefficients directly, easing system integration. The **InterpolatingFIR** block performs filtering on a stream of multichannel input data and produces a stream of output data with increased sampling frequency.

You can use the **InterpolatingFIR** block in a digital up converter for a radio system or a general purpose DSP application. The coefficients and input data are fixed-point types, and the output is the implied full precision fixed-point type. You can reduce the precision by using a separate **Scale** block, which can perform rounding and saturation to provide the required output precision.

The **InterpolatingFIR** block supports:


- Rate changes from two upwards
- Coefficient width in bits from 2 to 32 bits
- Data output width in bits from 4 to 64 bits
- Half-band and L-band Nyquist filters
- Symmetry and anti(negative)-symmetry
- Real filters

In the basic equation, at each sample time  $k$ , the new output  $y$ , is calculated by multiplying coefficients  $a$ , by the recent past values of the input  $x$ .

The **InterpolatingFIR** has a higher output sample rate than the input sample rate by a factor,  $I$ , the interpolation factor. Usually, the interpolating FIR inserts  $I-1$  zeroes for every input sample, thus raising the sample rate by a factor  $I$ .

The physical implementation avoids performing multiplications with these zero samples, reducing the filter cost.

**Figure 95. Interpolate by 2 Filter Increasing Sample Rate of a Sine Wave Input**



**Table 48. Parameters for the InterpolatingFIR Block**

Parameter	Description
Input rate per channel	Specifies the sampling frequency of the input data per channel measured in millions of samples per second (MSPS).
Interpolation	Specifies the interpolation rate. Must be an integer.
Number of channels	Specifies the number of unique channels to process.
<i>continued...</i>	

Parameter	Description
Symmetry	You can select <b>Symmetrical</b> or <b>Anti-Symmetrical</b> coefficients. Symmetrical coefficients can result in hardware resource savings over the asymmetrical version.
Coefficients	You can specify the filter coefficients using a Simulink fixed-point object <b>fi(0)</b> . The data type of the fixed-point object determines the width and format of the coefficients. The length of the array determines the length of the filter. For example, <b>fi(fir1(49, 0.3),1,18,19)</b> .
Base address	You can memory map the filter's coefficients into the address space of the system. This field determines the starting address for the coefficients. It is specified as a MATLAB <b>double</b> type (decimal integer) but you can use a MATLAB expression to specify a hexadecimal or octal type if required.
Read/Write mode	You can allow <b>Read</b> , <b>Write</b> , or <b>Read/Write</b> access from the system interface. Turn on <b>Constant</b> to map coefficients to the system address space.
Filter structure	You can select <b>Use All Taps</b> , <b>Half Band</b> , or a specified band (from <b>3rd Band</b> to <b>46th Band</b> ).
Expose Avalonmemory-mapped agent in Simulink	Allows you to reconfigure coefficients without Platform Designer. Also, it allows you to reprogram multiple FIR filters simultaneously. Turn on to show the Avalon-MM inputs and outputs as normal ports in Simulink. The <b>Read/Write mode</b> decides the valid subset of Avalon memory-mapped agent ports that appear on the block. If you select <b>Constant</b> , the block shows no Avalon-MM ports.
Reconfigurable channels	Turn on for a reconfigurable FIR filter.
Channel mapping	Enter parameters as a MATLAB 2D array for reconfigurable FIR filter. Each row represents a mode; each entry in a row represents the channel input on that time slot. For example, [0,0,0,0;0,1 2,3] gives the first element of the second row as 0, which means DSP Builder processes channel 0 on the first cycle when the FIR is set to mode 1.

**Table 49. Port Interface for the InterpolatingFIR Block**

Signal	Direction	Description
a	Input	The fixed-point data input to the block. If you request more channels than can fit on a single bus, this signal is a vector. The width in bits is inherited from the input wire.
v	Input	Indicates validity of the data input signals. If v is high, the data on the a wire is valid.
c	Input	Indicates the channel of the data input signals. If v is high, c indicates which channel the data corresponds to.
m	Input	Indicates reconfigurable filter.
b	Input	Indicates multibank filter. This input appears when you add a second filter definition to the <b>Coefficients</b> parameter in the parameters dialog box.
q	Output	The fixed-point filtered data output from the block. If you request more channels than can fit on a single bus, this signal is a vector. The width in bits is a function of the input width in bits and the parameterization.
v	Output	Indicates validity of data output signals. The output data can be non-zero when v is low
c	Output	Indicates the channel of the data output signals. The output data can be non-zero when v is low

#### Related Information

- [DSP Builder FIR and CIC Filters](#) on page 230
- [DSP Builder FIR Filters](#) on page 233

### 12.1.10. NCO

The DSP Builder **NCO** block uses an octant-based algorithm with trigonometric interpolation. A numerically controlled oscillator (NCO) or digitally controlled oscillator (DCO) is an electronic system for synthesizing a range of frequencies from a fixed time base. Use NCOs when you require a continuous phase sinusoidal signal with variable frequency, such as when receiving the signal from an NCO-based transmitter in a communications system.

The NCO accumulates a phase angle in an accumulator. DSP Builder uses this angle as a lookup into sine and cosine tables to find a coarse sine and cosine approximation. DSP Builder implements the tables with a ROM. A Taylor series expansion of the small angle error refines this coarse approximation to produce accurate sine and cosine values. The **NCO** block uses folding to produce multiple sine and cosine values if the sample rate is an integer fraction of the system clock rate.

You can use this block in a digital up- or down-converter for a radio system or a general purpose DSP application. The coefficients and input data are fixed-point types, and the output is the implied full precision fixed-point type.

An NCO sometimes needs to synchronize its phase to an exact cycle. It uses the `phase` and `sync` inputs for this purpose. The `sync` input is a write enable for the channel (address) specified by the `chan` input when the new phase value (data) is available on the `phase` input. You may need some external logic (which you can implement as a primitive subsystem) to drive these signals. For example, you can prepare a sequence of new phase values in a shared memory and then write all the values to the NCO on a synchronization pulse. This option is particularly useful if you want an initial phase offset in the upper sinusoid.

The system specification, including such factors as the channel count, sample rates, and noise floor, determines the main parameters for this block. You can express all the parameters as MATLAB expressions, making it easy to parameterize a complete system.

The hardware generation techniques create very efficient NCOs, which are fast enough to update with every Simulink simulation. The edit-simulation loop time is much reduced, improving productivity.

**Table 50. Specification Parameters for the NCO Block**

Parameter	Description
Output Rate Per Channel (MSPS)	The sine and cosine output rate per channel measured in millions of samples per second.
Output Data Type	The output width in bits of the NCO. The bit width controls the internal precision of the NCO. The spurious-free dynamic range (SFDR) of the waves produced is approximately $6.02 \times \text{bit width}$ . The 6.02 factor comes from the definition of decibels with each added bit of precision increasing the SFDR by a factor of $20 \times \log_{10}(2)$ .
Output Scaling Value	This value interprets the output data in the Simulink environment. The power of 2 scaling provided lets you specify the range of the output value.
Accumulator Bit Width	Specifies the width of the memory-mapped accumulator bit width, which governs the NCO frequency accuracy that you can control. The width is limited to the range 15–30 for use with a 32-bit memory map (shared by other applications such as a Nios II processor). The top two bits in the 32-bit width are reserved to control the inversion of the sine and cosine outputs. Select <b>Constant</b> for the <b>Read/Write Mode</b> to increase the width to 40 bits. Frequency resolution = $\text{clock frequency} / 2^{\text{accumulator bit width}}$

*continued...*



Parameter	Description
Phase Increment and Inversion	A vector that represents the step in phase between each sample. This vector controls the frequencies generated during simulation. The length of the vector determines how many channels (frequencies) of data are generated from the NCO. The unit of the vector is one (sine or cosine) cycle.
Phase Increment and Inversion Memory Map	Specifies where in the memory-mapped space the NCO registers are mapped.
Read/Write Mode	Specifies whether the NCO phase increment and inversion registers are mapped as <b>Read</b> , <b>Write</b> , <b>Read/Write</b> , or <b>Constant</b> .
Expose Avalon Memory-Mapped Agent in Simulink	Allows you to reconfigure increments without Platform Designer. Also, it also allows you to reprogram multiple NCOs simultaneously. When you turn on this parameter, the following three additional input ports and two output ports appear in Simulink. <ul style="list-style-type: none"> <li>data, address, write</li> <li>readdata, valid</li> </ul>

[NCO Block Phase Increment and Inversion](#) on page 249

[NCO Block Phase Increment Memory Registers](#) on page 250

[NCO Block Frequency Hopping](#) on page 251

### Related Information

- [Super-sample NCO](#) on page 0
- [NCO](#) on page 175

#### 12.1.10.1. NCO Block Phase Increment and Inversion

The **Phase Increment and Inversion** parameter allows you to specify the phase increment values that control the frequencies of the sinusoidal wave signals generated during simulation. You can also specify whether to invert the generated sinusoidal signals. This parameter is closely related to the **Output Rate per Channel** and the **Accumulator Bit Width** parameters.

To achieve a desired frequency (in MHz) from the **NCO** block, you must specify a phase increment value defined by:

$$\text{Phase Increment Value} = \text{Frequency} * 2^{\text{Accumulator Bit Width}} / \text{Output Data Rate}$$

This value must fall within the range specified by the **Accumulator Bit Width** parameter. For example, for an accumulator bit width of 24 bits, you can specify a phase increment value less than  $2^{24}$ .

You can specify the phase increment values in a vector format that generates multichannel sinusoidal signals. The length of the vector determines how many channels (frequencies) of data are generated from the **NCO** block. For example, a length of 4 implies that four channels of data are generated.

When the design uses the NCO for super-rate applications (NCO frequency is higher than output data rate), for example direct RF DUC, use multiple channels (in evenly distributed phases). The phase increment value is:

$$\text{Phase increment value} = \text{mod}((\text{frequency})/(\text{output data rate}), 1) \times 2^{\text{accumulator bit width}}$$

The modulus function limits the phase value to less than 1 and prevents interfering with the inversion bits.

When the input is in matrix format (with multiple rows of vectors), the design configures the **NCO** block as a multi-bank NCO for frequency hopping for multicarrier designs. The number of rows in the matrix represents the number of banks of frequencies (of sine and cosine waves) that generate for a given channel. An additional **bank** input and **b** output port automatically add to the **NCO** block.

*Note:* No upper limit to the number of rows exists in the matrix and you can specify any number of frequency banks. However, you should carefully monitor the resource usage to ensure that the specified design fits into the target device.

You can also use the **Phase Increment and Inversion** parameter to indicate whether the generated sinusoidal signals are inverted. For an accumulator width in bits of 24 bits, you can add two bits (the 25th and 26th bits) to the phase increment value for a given frequency. These bits indicate if the sine (26th bit) and cosine (25th bit) are inverted.

### 12.1.10.2. NCO Block Phase Increment Memory Registers

Use the **Phase Increment and Inversion Memory Map** parameter to specify the base address of the memory-mapped space where the NCO registers are mapped. The **System Data Width** specified in the **DSP Builder > Avalon Interfaces > Avalon memory-mapped Agent** menu and the **Accumulator Bit Width** specified in the **NCO** block determines the number of registers required for each phase increment value. You can specify the **System Data Width** to be either 8, 16, or 32 bits. If the **Accumulator Bit Width** is larger than the **System Data Width**, two registers are required to store each phase increment value.

The **NCO** block only supports one or two registers for each phase increment value. If one register is required for each phase increment value, the phase increment value for the first frequency is written into the base address, the second value into the next address (base address + 1) and so on. If you require two registers, the design uses the base address and the next address (base address + 1) for the first value with each address storing part of the value. The next pair of addresses store the next value and so on.

For example, for a **System Data Width** of 16, **Accumulator Bit Width** of 24 and **Phase Increment and Inversion Memory Map** base address of 1000, addresses 1000 and 1001 store the phase increment value for the first frequency. Address 1001 stores the lower 16 bits (15 .. 0) and address 1000 stores the remaining 8 bits (23 .. 16). If DSP Builder generates four channels of sinusoidal signals, it uses addresses 1002 and 1003 for the second channel, addresses 1004 and 1005 for the third channel, addresses 1006 and 1007 for the fourth channel.

In summary:

<total addresses required> = <number of registers per value> × <number of channels>

When DSP Builder writes to the phase increment and inversion memory map registers (in write mode), the new value takes effect immediately.

If the application is a super-rate operation (like direct RF DUC) and multiple channels in the NCO are configured for a new center frequency, first configure the phase increment value for each channel. DSP Builder then synchronizes the phase offsets of all channels at the same time by asserting the **sync** pulse.

To minimize the duration of disruption, you may use two banks of phase increment registers. The new phase increment registers bank switches first. Then, you can apply the **sync** pulse to synchronize the new phase offsets.

### 12.1.10.3. NCO Block Frequency Hopping

Use the **NCO** block to configure multiple banks of predefined frequencies for frequency hopping. If you specify a matrix comprising multiple rows of vectors as the **Phase Increment and Inversion** values, DSP Builder configures the NCO for multiple banks and defines the number of banks by the number of rows of vectors specified by inputs to the **Phase Increment and Inversion** parameter. A `bank` input and `b` output are automatically added to the NCO block. It also allocates phase increment memory registers for the multiple banks of frequencies automatically.

You can use the Avalon-MM interface to access (read or write) the phase increment memory registers in the same way as for a single bank with the register address for the `i`th bank frequencies starting from:

$$\langle \text{base address} \rangle + (i - 1) \times \langle \text{number of registers per value} \rangle \times \langle \text{number of channels} \rangle.$$

You can use the `bank` input as the index to switch the generated sinusoidal waves to the specified set (bank) of predefined frequencies.

*Note:* Ensure you constrain the bank input to the range (0 .. <number of banks> - 1). You can expect unreliable outputs from the NCO block if the bank input exceeds the number of banks.

When using an Avalon-MM interface to access (read or write) the phase increment memory registers, ensure that you only write to the inactive banks (banks which are not equal to the index specified by the input bank port). The dual-port memory that the **NCO** block uses is in **DONT\_CARE** mode when reading and writing to the same address. The **NCO** block uses the active bank to read the phase increment value. Writing to the active bank may cause unreliable values to read out and the active bank may pass out unexpected sinusoidal signals through the memory interface.

The read data, from the address to which you write the new values to, may also be unreliable because of the memory type that the **NCO** block uses. Only use read data from banks where they do not write at the same time.

The **Results** tab shows the implications of your parameter settings.

**Table 51. Results Tab Parameters for the NCO Block**

Parameter	Description
Expected SFDR	The SFDR in decibels relative to the carrier (dBc): $(\text{Output Data Type Width}) \times 20 \times \log_{10}(2)$ .
Accumulator precision	Accumulator precision in Hz: $10^6 \times (\text{output rate}) / 2^{(\text{accumulator width in bits}+1)}$ .
Frequency	Frequency in MHz: $(\text{output rate}) \times (\text{phase increment and inversion}) / 2^{(\text{accumulator width in bits})}$ .
# outputs per cycle	The number of outputs per cycle is the width of the vector of output signals: $\text{physical channels out} = \text{ceil}(\text{length}(\text{phase increment and inversion}) / ((\text{system clock frequency}) / (\text{output rate})))$
log2 of look-up table	The number of address bits in the internal look-up tables.

**Table 52. Port Interface for the NCO Block**

Signal	Direction	Description
chan	Input	Indicates the channel. If <i>v</i> is high, <i>chan</i> indicates which channel the data corresponds to.
v	Input	Indicates validity. If <i>v</i> is high, new data generates.
phase	Input	Specifies the phase offset. The size of this port should match the wire count of the NCO. The number of sines/cosines per cycle is limited to 1–16 outputs. Use multiple NCO blocks if more outputs are required.
sync	Input	Specifies the phase synchronization. The size of this port should match the wire count of the NCO output. When asserted, the phase offsets of all channels synchronize to the phase inputs. This signal has no effect to the phase increment and inversion registers. When you use this signal, you may need to initialize the offsets upon system power-up or reset. The number of sines/cosines per cycle is limited to 1–16 outputs. Use multiple NCO blocks if more outputs are required.
bank	Input	This input is available when you specify a matrix of predefined vectors for the phase increment values. You can use this input to switch to the bank of predefined frequencies.
data	Input	The data port has unsigned integers with a width equal to the width of the accumulator plus two for the inversion bits.
address	Input	Only available when you turn on <b>Expose Avalonmemory-mapped Agent in Simulink</b> . The address port is the same width as the system address width that you configure in the <b>DSP Builder &gt; Avalon Interfaces &gt; Avalon memory-mapped Agent</b> menu. Also the base address is the same.
write	Input	Deassert the write port to make a read occur.
sin	Output	The sine data output from the block. If you request more channels than can fit on a single bus, this signal is a vector. The width in bits is a function of the input width in bits and the parameterization.
cos	Output	The cosine data output from the block. If you request more channels than can fit on a single bus, this signal is vector. The width in bits is a function of the input width in bits and the parameterization. The number of sines/cosines per cycle is limited to 1–16 outputs. Use multiple NCO blocks if more outputs are required.
v	Output	Indicates validity of the data output signals.
c	Output	Indicates channel of the data output signals.
b	Output	Indicates the bank that the output signals use. This output is available when you specify a matrix of predefined vectors for the phase increment values.
readdata	Output	The data port has unsigned integers with a width equal to the width of the accumulator plus two for the inversion bits.
valid	Output	Indicates a valid output.

### 12.1.11. Real Mixer (Mixer)

The DSP Builder **Mixer** block performs a real by complex multiply on streams of data. This function creates quadrature data from an antenna input, where the real data is the antenna data and the complex data is the cosine and sine data provided by an NCO.

The **Mixer** block multiplies a real input stream by a synchronized complex data stream, sample by sample.

You can use the **Mixer** block in a digital down converter for a radio system or a general purpose DSP application. The data has fixed-point types, and the output is the implied full precision fixed-point type.

**Note:** You can easily replicate the **Mixer** block with a **Multiply** block that takes one real and one complex input within a primitive subsystem.

The **Mixer** performs element-by-element multiplication on  $n$  channels and  $m$  frequencies.

The system specification, including such factors as the channel count and sample rates, determines the main parameters for this block. The input sample rate of the block determines the number of channels present on each input wire and the number of wires:

$$\text{Number of Channels per wire} = \text{Clock\_Rate} / \text{Sample\_Rate}$$

$$\text{Number of Wires} = \text{ceiling}(\text{Chan\_Count} \times \text{Sample\_Rate} / \text{Clock\_Rate})$$

For example, a sample rate of 60 MSPS and system clock rate of 240 MHz gives four samples to be TDM on to each input wire:

If there are more channels than TDM slots available on a wire, the input wire is a vector of sufficient width to hold all the samples. Similarly, the number of frequencies (the number of complex numbers) determines the width of the sine and cosine inputs. The number of results that the **Mixer** produces is the product of the sample input vector and the frequency vector. The results are TDM on to the  $i$  and  $q$  outputs in a similar way to the inputs.

**Table 53. Parameters for the Mixer Block**

Parameter	Description
Input Rate Per Channel (MSPS)	The data rate per channel measured in millions of samples per second.
Number of Channels	The number of real input channels.
Number of Frequencies	The number of real frequencies in the multiplier.

**Table 54. Port Interface for the Mixer Block**

Signal	Direction	Description
a	Input	The real data input to the block. If you request more channels than can fit on a single bus, this signal is a vector. The width in bits is inherited from the input wire.
v	Input	Indicates the validity of the data input signals. If $v$ is high, the data on the $a$ wire is valid.
c	Input	Indicates the channel of the data input signals. If $v$ is high, $c$ indicates the data channel.
sin	Input	The imaginary part of the complex number. For example, the NCO's sine output.
cos	Input	The real part of the complex number. For example, the NCO's cosine output.
i	Output	The in-phase (real) output of the mixer, which is $(a \times \cos)$ . If you request more channels than can fit on a single bus, this signal is a vector. The width in bits is wide enough for the full precision result.
q	Output	The quadrature phase (imaginary) output of the mixer, which is $(a \times \sin)$ . If you request more channels than can fit on a single bus, this signal is a vector. The width in bits is wide enough for the full precision result.
v	Output	Indicates the validity of the data output signals.
c	Output	Indicates the channel of the data output signals.

### 12.1.12. Scale

The **Scale** block selects part of a wide input word, performs various types of rounding, saturation and fixed-point scaling, and produces an output of specified precision.

By default, DSP Builder preserves the binary point so that the fixed-point interpretation of the result has the same value, subject to rounding, as the fixed-point interpretation of the input.

You can dynamically perform additional scaling, by specifying a variable number of bits to shift, allowing you to introduce any power of two gain.

**Note:** Always use **Scale** blocks to change data types in preference to **Convert** blocks, because they vectorize and automatically balance the delays with the corresponding valid and channel signals.

The **Scale** block provides scaling in addition to rounding and saturation to help you manage bit growth. The basic functional modules of a **Scale** block are shifts followed by rounding and saturation. The multiplication factor (default is 1) is a constant scale to apply to the input.

The number of bits to shift left allows you to select the most meaningful bits of a wide word, and discard unused MSBs. You can specify the number of shifts as a scalar or a vector. The block relies on shift input port to decide which value to use if you specified the number of shifts as a vector. The shift input signal selects which gain to use cycle-by-cycle.

In a multichannel design, changing the shift value cycle-by-cycle allows you to use a different scaling factor for different channels.

A positive number of **Number of bits to shift left** indicates that the MSBs are discarded, and the **Scale** block introduces a gain to the input. A negative number means that zeros (or 1 in the signed data case) are padded to the MSBs of the input data signal, and the output signal is attenuated.

**Table 55. Parameters for the Scale Block**

Parameter	Description
Output data type	The type of the result. For example: <b>sfix(16), uint(8)</b> .
Output scaling value	The scaling of the result if the result type is fixed-point. For example: <b>2<sup>-15</sup></b> .
Rounding method	Specifies one of the following three rounding methods for discarding the least significant bits (LSBs): <ul style="list-style-type: none"> <li>• <b>Truncate:</b> truncates the least significant bits. Has the lowest hardware usage, but introduces the worst bias.</li> <li>• <b>Biased:</b> rounds up if the discarded bits are 0.5 or above.</li> <li>• <b>Unbiased:</b> rounds up if the discarded bits are greater than 0.5, and rounds to even if the discarded bits equal 0.5.</li> </ul>
Multiplication factor	Modify the interpreted value by scaling it by this factor. This factor does not affect the hardware generated for the <b>Scale</b> block, but merely affects the interpretation of the result. For example: 1, 2, 3, 4, 8, 0.5.
Saturation method	Specifies one of the following three saturation methods for discarding the most significant bits (MSBs):

*continued...*

Parameter	Description
	<ul style="list-style-type: none"> <li>• <b>None:</b> no saturation is performed.</li> <li>• <b>Asymmetric:</b> the range of the number produced occupies the whole of the two's complement range (for example -1.0 to 0.999). One more negative number is available, which introduces a slight bias.</li> <li>• <b>Symmetric:</b> the range of the result is clipped to between symmetrical boundaries (for example -0.999 and 0.999), which ensure that no bias enters the dataflow.</li> </ul>
Number of bits to shift left	A scalar or a vector that determines the gain of the result. A positive number indicates that the scale block introduces a gain to the input. A negative number means that the output signal is attenuated. A vector of gains allows the shift input signal to select which gain to use on a cycle per cycle basis. The value of the shift input performs zero-based indexing of the vector.

**Table 56. Port Interface for the Scale Block**

Signal	Direction	Description
a	Input	The fixed-point data input to the block. If you request more channels than can fit onto a single bus, this signal is a vector. The width in bits is inherited from the input wire.
a_v	Input	Indicates the validity of the data input signals. If a_v is high, the data on the a wire is valid.
a_chan	Input	Indicates the channel of the data input signals. If a_v is high, a_chan indicates to which channel the data corresponds.
shift	Input	Indicates which element of the zero-based shift vector to use.
q	Output	The scaled fixed-point data output from the block. If you request more channels than can fit onto a single bus, this signal is a vector. The width in bits is calculated as a function of the input width in bits and the parameterization.
q_v	Output	Indicates the validity of the data output signals.
q_chan	Output	Indicates the channel of the data output signals.
q_exp	Output	Indicates whether the output sample has saturated or overflowed.

After you run a simulation, DSP Builder updates the help pages with specific information about each instance of a block.

**Table 57. Messages for the Scale Block**

Message Example	Description
Written on Tue Feb 19 11:25:27 2008	Date and time when this file ran.
Number of physical buses: 4	Depending on the input data rate, the number of data wires needed to carry the input data may be more than 1.
Calculated bit width of output stage: 16	The width in bits of the (vectorized) data output.
Latency is 2	The latency introduced by this block.
Parameters table	Lists the current rounding and saturation modes.
Port interface table	Lists the port interfaces to the <b>Scale</b> block.

### 12.1.13. Single-Rate FIR

The **SingleRateFIR** block implements a highly efficient multichannel finite impulse response filter across a broad range of parameters directly from a Simulink model. A memory-mapped interface allows you to read and write coefficients directly, easing

system integration. The **SingleRateFIR** block performs filtering on a stream of multichannel input data and produces a stream of output data with increased sampling frequency.

You can use the **SingleRateFIR** block in a digital up converter for a radio system or a general purpose DSP application. The coefficients and input data are fixed-point types, and the output is the implied full precision fixed-point type. You can reduce the precision by using a separate **Scale** block, which can perform rounding and saturation to provide the required output precision.

The **SingleRateFIR** block supports sample rates from 1 to 500, coefficient width in bits from 2 to 32 bits, half-band and L-band Nyquist filters, real and complex filters, and symmetry and anti(negative)-symmetry.

**Table 58. Parameters for the Single-Rate FIR Block**

Parameter	Description
Input rate per channel	Specifies the sampling frequency of the input data per channel measured in millions of samples per second (MSPS).
Number of channels	Specifies the number of unique channels to process.
Symmetry	You can select <b>Symmetrical</b> or <b>Anti-Symmetrical</b> coefficients. Symmetrical coefficients can result in hardware resource savings over the asymmetrical version.
Coefficients	You can specify the filter coefficients using a Simulink fixed-point object <b>fi(0)</b> . The data type of the fixed-point object determines the width and format of the coefficients. The length of the array determines the length of the filter. For example, <b>fi(fir1(49, 0.3),1,18,19)</b>
Base address	You can memory map the filter's coefficients into the address space of the system. This field determines the starting address for the coefficients. It is specified as a MATLAB <b>double</b> type (decimal integer) but you can use a MATLAB expression to specify a hexadecimal or octal type if required.
Read/Write mode	You can allow <b>Read</b> , <b>Write</b> , or <b>Read/Write</b> access from the system interface. Turn on <b>Constant</b> , to map coefficients to the system address space.
Expose Avalonmemory-mapped agent in Simulink	Allows you to reconfigure coefficients without Platform Designer. Also, it allows you to reprogram multiple FIR filters simultaneously. Turn on to show the Avalon memory-mapped inputs and outputs as normal ports in Simulink. The <b>Read/Write mode</b> decides the valid subset of Avalon memory-mapped agent ports that appear on the block. If you select <b>Constant</b> , the block shows no Avalon-MM ports.
Reconfigurable channels	Turn on for a reconfigurable FIR filter.
Channel mapping	Enter parameters as a MATLAB 2D array for a reconfigurable FIR filter. Each row represents a mode; each entry in a row represents the channel input on that time slot. For example, [0,0,0,0;0,1 2,3] gives the first element of the second row as 0, which means DSP Builder processes channel 0 on the first cycle when the FIR is set to mode 1.

**Table 59. Port Interface for the Single-Rate FIR Block**

Signal	Direction	Description
a	Input	The fixed-point data input to the block. If you request more channels than can fit on a single bus, this signal is a vector. The width in bits is inherited from the input wire.
v	Input	Indicates validity of the data input signals. If v is high, the data on the a wire is valid.
c	Input	Indicates the channel of the data input signals. If v is high, c indicates the channel to which the data corresponds.
m	Input	Indicates a reconfigurable filter.
<i>continued...</i>		



Signal	Direction	Description
b	Input	Indicates multibank filter. This input appears when you add a second filter definition to the <b>Coefficients</b> parameter in the parameters dialog box.
q	Output	The fixed-point filtered data output from the block. If you request more channels than can fit on a single bus, this signal is a vector. The width in bits is a function of the input width in bits and the parameterization.
v	Output	Indicates the validity of data output signals. The output data can be non-zero when v is low.
c	Output	Indicates the channel of the data output signals. The output data can be non-zero when v is low.

### Related Information

- [DSP Builder FIR and CIC Filters](#) on page 230
- [DSP Builder FIR Filters](#) on page 233

## 12.2. Dependent Delay Library

Blocks from the DSP Builder advanced blockset **Dependent Delay** library implement delays outside of scheduled models, where the delay may depend on the latency of another model.

The **ChannelDependentDelay** block provides a sample delay on the connected signals. DSP Builder processes the signals similarly to a model that includes **ChannelIn** and **ChannelOut** blocks. DSP Builder identifies a specific `valid` and `channel` signal, then an arbitrary number of data lines.

The **GPDependentDelay** block is deprecated. Do not use this block in new designs.

The value of the sample delay may depend on the latency of referenced model, refer to the **SynthesisInfo** block.

**Table 60. Parameters for the ChannelDependentDelay Block**

Parameter	Description
Number of Data Signals	Specify the number of input and output (d and q) connections for the block. DSP Builder passes each input to the corresponding output and delays it by the latency constraint.
Latency Constraint	This option allows you to select the type of constraint and to specify its value. The value can be a workspace variable or an expression but must evaluate to a positive integer. You can select the following types of constraint: <ul style="list-style-type: none"> <li>• <b>&gt;</b>: Greater than</li> <li>• <b>&gt;=</b>: Greater than or equal to</li> <li>• <b>=</b>: Equal to</li> <li>• <b>&lt;=</b>: Less than or equal to</li> <li>• <b>&lt;</b>: Less than</li> </ul> Select either <b>+</b> or <b>-</b> and type in a reference model in the text field. Specify the reference as a Simulink path string e.g. <code>'design/topLevel/model'</code> . DSP Builder then ensures the latency depends on that model, otherwise the default is that DSP Builder depends on no model.
Local Reset-Minimization	Turn on to allow DSP Builder to apply reset minimization to the delays. You must also turn on <b>Global Reset Minimization</b> . The values are: <ul style="list-style-type: none"> <li>• <b>Off</b>. Default, no reset minimization.</li> <li>• <b>On</b>. DSP Builder applies no reset to all delay stages.</li> </ul>

**Table 61. Port Interface for the ChannelDependentDelay Block**

Parameter	Direction	Description
dv	Input	The input valid signal to delay.
qv	Output	The output for the valid signal.
dc	Input	The input channel number to delay.
qc	Output	The output for the channel.
d	Input	The input data to delay.
q	Output	The output for the corresponding to the input data.

## 12.3. FFT IP Library

Use the DSP Builder advanced blockset **FFT IP** library blocks to implement full FFT IP functions. These blocks are complete primitive subsystems.

1. [Bit Reverse Core C \(BitReverseCoreC and VariableBitReverse\)](#) on page 258
2. [FFT \(FFT, FFT\\_Light, VFFT, VFFT\\_Light\)](#) on page 259

### 12.3.1. Bit Reverse Core C (BitReverseCoreC and VariableBitReverse)

The **BitReverseCoreC** block performs buffering and bit reversal of incoming FFT frames.

A single synthesis time parameter specifies the length  $N$  of the fast Fourier transform.

The bit reversal that this block applies is appropriate only for transform sizes that are an integer power of two. The block is single-buffered to support full streaming operation with minimal overhead.

The **VariableBitReverse** block performs buffering and bit-reversal of variable-sized FFT frames for designs with a **VFFT** or **VFFT\_Float** block. A single synthesis-time parameter  $N$  specifies the length  $2^N$  of the largest frame that the block handles. The **VariableBitReverse** block has an additional input: `size`, which specifies the length  $2^{\text{size of the current frame}}$ .

To reconfigure the **VariableBitReverse** block between frames, observe the following rules:

- Ensure the size input is always in the range  $0 \leq \text{size} \leq N$ .
- Keep the size input constant while the **VariableBitReverse** block is processing a frame.
- When you reconfigure the **VariableBitReverse**, you must completely flush the **VariableBitReverse** block before changing the value of the size input. You must wait at least  $2^{\text{oldSize}}$  (where *oldSize* is the previous value of the size input) cycles before providing valid input to the VFFT.

**Table 62. Parameters for the BitReverseCoreC Block**

Parameter	Description
FFT Size	Specifies the size of the FFT.

**Table 63. Parameters for the VariableBitReverse Block**

Parameter	Description
N	Logarithm of the maximum frame size.

**Table 64. Port Interface for the BitReverseCoreC Block**

Signal	Direction	Type	Description
v	Input	Boolean	Valid input signal.
c	Input	Unsigned 8-bit integer	Channel input signal.
size	Input	Unsigned integer	Logarithm of the current input frame size. <b>VariableBitReverse</b> only.
x	Input	Any complex fixed-point ( <b>BitReverseCoreC</b> ); any ( <b>VariableBitReverse</b> )	Complex data input signal.
qv	Output	Boolean	Valid output signal.
qsize	Output	Unsigned integer	Logarithm of the current output frame size. <b>VariableBitReverse</b> only.
qc	Output	Unsigned 8-bit integer	Channel output signal.
q	Output	Any	Complex data output signal.

### 12.3.2. FFT (FFT, FFT\_Light, VFFT, VFFT\_Light)

The FFT and VFFT blocks support processing multiple interleaved FFTs. The number of interleaved FFTs must be a power of 2. Each FFT is independent except that all the input for all of the FFTs must arrive as a single contiguous block of data. For example, with 8 FFTs each of size 1K each input block must contain 8K points.

The following blocks are in the **Primitives FFT Design Elements** library:

- FFT\_Light
- VFFT\_Light

For floating-point FFTs, select either correct or faithful rounding. Correct rounding corresponds to the normal IEEE semantics; faithful rounding delivers less accurate results but requires less logic to implement.

The **FFT** block provides a full radix-2<sup>2</sup> streaming FFT or IFFT. Use the FFT block for fixed-point or floating-point data. The block is a scheduled subsystem.

The **FFT\_Light** block is a light-weight variant. However, it is not a scheduled subsystem, and it doesn't implement the *c* (channel) signal. The blocks provide an output signal, *g*, which pulses high at the start of each output block.

The FFT blocks all support block-based flow control. You must supply all the input data required for a single FFT iteration (one block) on consecutive clocks cycles, but an arbitrary large (or small) gap can exist between consecutive blocks. The **BitReverseCoreC** and **Transpose** blocks produce data in blocks that respect this protocol.

You may provide the input data to any of these block in either natural or bit-reversed order; the output result is in bit-reversed or natural order, respectively.

The **VFFT** block provides a variable-size streaming FFT or IFFT. For these blocks, you statically specify the largest and smallest FFT that the block handles. You can dynamically configure the number of points processed in each FFT iteration using the size signal.

Use the **VFFT** block for fixed-point or floating-point data. The **VFFT** block is a scheduled subsystem and implements *v* (valid) and *c* (channel) signals.

The **VFFT\_light** block is a light-weight variant of the **VFFT** block. It is not a scheduled subsystem, and it doesn't implement the *c* (channel) signal. Instead, it provides an output *g* signal, which pulses high at the start of each output block.

The **VFFT** blocks all support block-based flow control. You must supply all the input data required for a single VFFT iteration (one block) on consecutive clocks cycles. If you use two successive FFT iterations that use the same FFT size, the inter-block gap can be as small (or as large) as you like.

However, if you want to reconfigure the **VFFT** block between FFT iterations, you must use the following rules:

- The size input should always be in the range  $minSize \leq size \leq maxSize$ .
- The size input must be kept constant while the **VFFT** block processes an FFT iteration.
- When you reconfigure the **VFFT**, you must completely flush **VFFT** pipeline before changing the value of the size input. You must wait at least  $2^{oldSize}$  (where *oldSize* is the previous value of the size input) cycles before providing valid input to the VFFT.

*Note:* The **VariableBitReverse** block also requires an inter-block gap of  $2^{oldSize}$  cycles when you reconfigure its size. If you use both the **VariableBitReverse** block and the VFFT block, you need to provide an interblock gap of  $2 \cdot (2^{oldSize})$  cycles to allow both blocks to reconfigure successfully.

Not all parameters are available with all blocks.

**Table 65. Parameters for the FFT and VFFT Blocks**

Parameter	Description
<b>iFFT</b>	<b>true</b> to implement an IFFT, otherwise <b>false</b> .
<b>Number of interleaved subchannels</b>	Enter how many FFTs that DSP Builder interleaves in each block.
<b>Bit-reversed input</b>	<b>true</b> if you expect bit-reversed input, otherwise <b>false</b> .
<b>N</b>	The logarithm of the FFT size. FFT and FFT_Light only
<b>maxSize</b>	The logarithm of the maximum FFT size. VFFT and VFFT_Light only.
<b>minSize</b>	The logarithm of the minimum FFT size. VFFT and VFFT_Light only.
<b>Input type</b>	Input signal type.
<b>Input scaling exponent</b>	The fixed-point scaling factor of the input.
<b>Twiddle/pruning specification</b>	Refer to <i>About Pruning and Twiddle for FFT Blocks</i> .
<b>Use faithful rounding</b>	<b>true</b> if the block uses faithful (rather than correct) rounding for floating-point operations. Fixed-point FFTs ignore this parameter.

Not all signals are available with all blocks

**Table 66. Port Interface for the FFT Blocks**

Signal	Direction	Type	Description
v	Input	Boolean.	Valid input signal.
c	Input	Unsigned 8-bit integer.	Channel input signal. <b>FFT</b> and <b>VFFT</b> , only.
size	Input	Unsigned integer.	Logarithm of the current FFT size. <b>VFFT</b> and <b>VFFT_Light</b> only.
d	Input	Any complex fixed-point.	Complex data input signal. <b>VFFT</b> and <b>VFFT_Light</b> only.
x	Input	Any complex fixed-point type ( <b>FFT</b> and <b>FFT_light</b> ). Any floating-point type ( <b>FFT_Float</b> or <b>FP_FFT_Light</b> ).	Complex data input signal.
qv	Output	Boolean.	Valid output signal.
qc	Output	Unsigned 8-bit integer.	Channel output signal. <b>FFT</b> and <b>VFFT</b> , only.
q	Output	Same as x.	Complex data output signal.
g	Output	Boolean	Start of output block. <b>VFFT_Light</b> only.

**Related Information**

[About Pruning and Twiddle for FFT Blocks on page 284](#)

## 13. Interfaces Library

---

Use the DSP Builder advanced blockset **Interfaces** library blocks to set and use Avalon interfaces. DSP Builder treats design-level ports that do not route via Avalon interface blocks as individual conduits.

1. [Memory-Mapped Library](#) on page 262
2. [Streaming Library](#) on page 273

### 13.1. Memory-Mapped Library

The DSP Builder advanced blockset **Memory Mapped** library blocks provide memories and registers that you can access both in your DSP datapath and with an external interface. You can use these blocks to configure coefficients or run-time parameters and read calculated values.

This library also provides blocks that you can use to simulate the bus interface in the Simulink environment.

*Note:* Do not turn on **Bit Accurate Simulation** when your design includes **Memory-Mapped** library blocks, otherwise the simulation is all zeros.

1. [Bus Slave \(BusSlave\)](#) on page 262
2. [Bus Stimulus \(BusStimulus\)](#) on page 262
3. [Bus Stimulus File Reader \(Bus StimulusFileReader\)](#) on page 263
4. [External Memory, Memory Read, Memory Write](#) on page 265
5. [Register Bit \(RegBit\)](#) on page 269
6. [Register Field \(RegField\)](#) on page 270
7. [Register Out \(RegOut\)](#) on page 271
8. [Shared Memory \(SharedMem\)](#) on page 271

#### 13.1.1. Bus Slave (BusSlave)

The DSP Builder **BusSlave** block is deprecated.

#### 13.1.2. Bus Stimulus (BusStimulus)

The DSP Builder **Bus StimulusFileReader** block with the **BusStimulus** block simulates accesses over the processor interface in the Simulink environment.

The **BusStimulus** block performs hidden accesses to the registers and **SharedMem** blocks in the memory hierarchy of your model. It is an interface that allows another block to read and write to any address. The `address` and `data` ports act as though an external processor reads and writes to your system.

The **BusStimulus** block transmits data from its input ports (`address`, `writedata` and `write`) over the processor interface, and thus modifies the internal state of the memory-mapped registers and memories as appropriate. Any response from the simulated processor interface is output on the `readdata` and `readvalid` output ports.

For example, to use the **BusStimulus** block connect constants to the address and data inputs. A pulse on the `write` port then writes the data to any register mapped to the specified address. Put a counter on the `address` input to provide all the data in every memory location on the `readdata` port. DSP Builder asserts the `readdatavalid` output when a valid read data is on the `readdata` port.

**Table 67. Parameters for the BusStimulus Block**

Parameter	Description
Sample Time	Specifies the Simulink sample time.
Show read enable	Turn on to show read enable port. If you use the <b>BusStimulus</b> with the <b>BusStimulusFileReader</b> blocks in a design, ensure this parameter is turned on or turned off in both blocks.

**Table 68. Port Interface for the BusStimulus Block**

Signal	Direction	Type	Description
<code>address</code>	Input	Unsigned integer	Address to access.
<code>writedata</code>	Input	16-, 32-, or 64-bit unsigned integer	Write data.
<code>write</code>	Input	Boolean	Write enable.
<code>read</code>	Input	Boolean	Read enable.
<code>readdata</code>	Output	16-, 32-, or 64-bit unsigned integer	Read data.
<code>readdatavalid</code>	Output	Boolean	Read data valid.

### 13.1.3. Bus Stimulus File Reader (Bus StimulusFileReader)

The DSP Builder **BusStimulus** block with the **BusStimulusFileReader** block simulates accesses over the processor interface in the Simulink environment.

The **BusStimulusFileReader** block reads a stimulus file (`.stm`) and generates signals that match the **BusStimulus** block.

A bus stimulus file describes a sequence of transactions to occur over the processor interface, together with expected read back values. This block reads such files and produces outputs for each entry in the file.

Bus stimulus files automatically write to any blocks that have processor mapped registers when you simulate a design. Any design with useful register files generates a bus stimulus file that you can use to bring your design out of reset (all registers 0). You can also write your own bus stimulus files with the following format:

```
MemSpace Address WriteData WE ExpReadData Mask
```

or

MemSpace Address WriteData WE RE ExpReadData Mask

where:

MemSpace specifies the memory space (the format supports multiple memory spaces).

Address is the word address.

WriteData is the data to write if any.

WE performs a write when 1.

ExpReadData is the expected read data. The value that is read from a location is checked against this value to allow self checking tests.

Mask specifies when the expected read data is checked, only the bits in this mask are checked, to allows you to read, write, or check specified bits in a register.

RE performs a read when 1. If RE is not present, assume RE is 1 when WE is 0.

The mask also masks the written data and performs a read-update-write cycle if you write to certain bits (i.e. not overwrite all of them).

During simulation, any mismatch between the expected read data (as the bus stimulus file describes) and the incoming read data (as the **BusStimulus** block provides) highlights and DSP Builder issues a warning.

**Table 69. Parameters for the BusStimulusFileReader Block**

Parameter	Description
Enabled	Turn on to enable reading of the bus stimulus file data. You must turn on <b>Has read enable</b> in the <b>BusStimulusFileReader</b> block if you turn on <b>Show read enable</b> in the <b>BusStimulus</b> block.
Stimulus File Name	Specifies the file from which to read bus stimulus data.
Log File Name	Specifies the file to store a log of all attempted bus stimulus accesses.
Space Width	Specifies the width of the memory space as described in the bus stimulus file—must be the same as the width specified in the <b>DSP Builder &gt; Avalon Interfaces &gt; Avalon memory-mapped Agent</b> menu.
Addr Width	Specifies the width of the address space as described in the bus stimulus file—must be the same as the width specified in the <b>DSP Builder &gt; Avalon Interfaces &gt; Avalon memory-mapped Agent</b> menu.
Data Width	Specifies the width of the data as described in the bus stimulus file—must be the same as the width specified in the <b>DSP Builder &gt; Avalon Interfaces &gt; Avalon memory-mapped Agent</b> menu.
Sample Time	Specifies the Simulink sample time.
Has read enable	Turn on to show read enable port. If you use the BusStimulusFileReader with the BusStimulus block in a design, ensure this parameter is turned on or turned off in both blocks.



**Table 70. Port Interface for the BusStimulusFileReader Block**

Signal	Direction	Type	Description
<b>address</b>	Output	Unsigned integer	Address from file.
<b>checkstrobe</b>	Output	Boolean	Indicates when the <b>readexpected</b> and <b>mask</b> signals should be checked against <b>readdata</b> .
<b>endofstimulus</b>	Output	Boolean	Generated signal to indicate when the end of the bus stimulus file is reached.
<b>read</b>	Output	Boolean	Read signal from file.
<b>readdatavalid</b>	Input	Boolean	Read data valid.
<b>readdata</b>	Input	16-bit or 32-bit unsigned integer	Read data.
<b>readexpected</b>	Output	16-bit or 32-bit unsigned integer	Expected read data from file.
<b>space</b>	Output	Unsigned integer	Memory space from file.
<b>write</b>	Output	Boolean	Write signal from file.
<b>writedata</b>	Output	16-bit or 32-bit unsigned integer	Data from file.
<b>mask</b>	Output	16-bit or 32-bit unsigned integer	Mask value from file.

### 13.1.4. External Memory, Memory Read, Memory Write

The DSP Builder **External Memory** block specifies characteristics of external memory and related Avalon-MM interfaces. DSP Builder uses this information to set bit widths on related interface ports in simulation, generated HDL and to build an external memory simulation model. The **Memory Read** and **Memory Write** blocks provide (read or write) access to associated external memory models in simulation. In HDL, each of these blocks is driving dedicated Avalon memory-mapped host interface. Associate read and write ports with **External Memory** blocks using identifiers. Connect these interfaces to a DDR3 SDRAM controller in your system-level design in Platform Designer.

Always add the **External Memory** block to the top-level of your DSP Builder design (similar to **Control** or **Signals** blocks).

Your design can have several instances of these blocks, but you must give them separate identifiers. DSP Builder creates a separate simulation model for each of these blocks.

**Table 71. External Memory Block Parameters**

Parameter	Values	Description
Identifier	Numeric value	A unique identifier for <b>External Memory</b> block that you should set on <b>Memory Read</b> or <b>Memory Write</b> block to associate these blocks with the <b>External Memory</b> block.
Avalon-MM Interface Data Width	A valid Avalon-MM interface data width value. Should be power of 2.	The width of the data signal in the generated Avalon memory-mapped host interfaces for associated <b>Memory Read</b> and <b>Memory Write</b> blocks. Set the data ports on these blocks to the same width.
<i>continued...</i>		

Parameter	Values	Description
Memory Data Width	Should be less than or equal to <b>Avalon-MM Interface Data Width</b> . The ratio between these two widths should be a power of 2.	The data width of the actual external memory. Only use to calculate the size of the memory which affects the width of address bus. Set this parameter to a quarter of the <b>Avalon MM Interface Data Width</b> parameter to define DDR memory operating at half rate.
Number of Rows	Numeric value. Should be power of 2.	The number of rows, columns, and banks of the actual physical memory that you connect to the DSP Builder design. Carefully chose access patterns based on these values to get the best performance of external memory.
Number of Columns		
Number of Banks		
Memory Size	Read-Only parameter	This parameter displays the size of the external memory based on the specified number of rows, columns, banks and memory data width. DSP Builder uses the following equation: $memory\_size = rows * columns * banks * memory\_data\_width$ The width of the address bus on Avalon memory-mapped host interfaces generated for associated <b>Memory Read</b> and <b>Memory Write</b> blocks, and the width of address input on these blocks, is: $address\_width = \log_2(memory\_size)$

**Table 72. Parameters that Only Affect Simulation**

Parameter	Values	Description
Signal Busy for Specified Amount of Simulation Time (%)	off 12.5, 25, 50, 75, 87.5	Any value other than <b>off</b> forces external memory simulation model into a busy state (the model refuses read or write requests) at random points during simulation. The actual value limits the overall busy time compared to design simulation time. The busy state of the memory model will be indicated with low value on ready ports for associated <b>Memory Read</b> or <b>Memory Write</b> blocks. If this feature is enabled, you may need to increase overall simulation time in order to get all requests to external memory through. Longer simulation time will be required for higher limits.
Show Diagnostic ports	Boolean switch	Turn on this option to add diagnostic ports, to <b>External Memory</b> blocks, which display the state of the simulation model.
Dump Memory Region into File	Boolean switch	Turn on so the <b>External Memory</b> block dumps its content for the specified region into a file. Each <b>Avalon MM Interface Data Width</b> value occupies a line in the file and is printed as a sequence of 8-bit decimal values. For <b>External Memory</b> blocks with <b>Avalon MM Interface Data Width</b> set to 16, the lines in the dump file have the following format <code>a[7:0] a1[15:8]</code>
Dump File Name	Valid file name	The name of the dump file with extension (DSP Builder does not add an extension) The dump file is created in the current directory.
Dump Region Start Address	Valid word address	The start address of the region in external memory that should be dumped.

*continued...*

Parameter	Values	Description
Dump Region Size	Non negative number	The number of words that should be dumped starting with the specified address.
Memory Initialization Start Address	A valid word address	—
Memory Initialization Data	A 1-D or 2-D matrix	DSP Builder writes this data to the memory from the <b>Start Address</b> . DSP Builder does not convert the data type. You must use the data type you want. The type must be: <ul style="list-style-type: none"> <li>• single</li> <li>• double</li> <li>• INT8</li> <li>• UINT8</li> <li>• INT16</li> <li>• UINT16</li> <li>• INT32</li> <li>• UINT32</li> <li>• INT64</li> <li>• UINT64</li> </ul> Where the Avalon-MM interface data width is wider than the data type, DSP Builder packs the initialization data into the words first, before moving on to the next address. You can use a 2-D matrix, where each column represents an address in the memory and each row represents a bit-slice in the word.

**Table 73. External Memory Block Diagnostic Ports**

You can enable these ports through dedicated parameters.

Name	Direction	Type	Description
reading	Output	Boolean	High when external memory model is performing reading operation; otherwise low.
writing	Output	Boolean	High when external memory model is performing writing operation; otherwise low.
busy	Output	Boolean	High when external memory model is in busy state; otherwise low.

### Memory Read Block

This block is an access point for reading from the associated **External Memory** block. It provides a simple interface with ready and valid based handshaking for reading. In generated HDL, use this block as an adapter between the provided interface and the Avalon memory-mapped host interface. You can place these blocks at any level of hierarchy under the DSP Builder device level block. The design can contain several of these blocks, with each of the blocks accessing the associated **External Memory** block.

**Table 74. Memory Read Parameters**

Parameter	Values	Description
<b>Identifier</b>	One of identifiers set for External Memory blocks in the design	Set to match an identifier on one of the <b>External Memory</b> blocks in the design.
<b>Maximum Burst Size</b>	off 2, 4, 8, 16, 32, 64, 128, 256, 512, 1024	If the value is set to <b>off</b> , DSP Builder does not allow burst requests. For other values, DSP Builder adds a new port to specify an actual size (less than or equal specified <b>Maximum Burst Size</b> ) for each burst request.

**Table 75. Memory Read Ports**

Name	Direction	Type	Description
read	Input	Boolean	Set this port to high to indicate a new read request.
address	Input	Unsigned Integer	Sets the address for the request. The width of this port is: $\log_2(\text{memory\_size})$ , <i>memory_size</i> is the size of associated <b>External Memory</b> .
burstcount	Input	Unsigned Integer	Optional. DSP Builder adds if <b>Maximum Burst Count</b> is not <b>off</b> . Sets the actual number of bursts for the read request. If you initiate a burst request, update this port and the <code>read</code> and <code>address</code> ports once at the beginning of request. The width of this port is: $\log_2(\text{max\_burst\_count}) + 1$
valid	Output	Boolean	Indicates that the valid response is available on the data port.
ready	Output	Boolean	Indicates that the block (associated memory) is ready to accept new request. Do not update input ports if this value is low.
data	Output	Unsigned Integer	Contains the read data. The width of this port is based on the <b>Avalon MM Interface Data Width</b> parameter of the associated <b>External Memory</b> .

### Memory Write Block

This block is an access point for writing to the associated external memory model. It provides a simple interface with `ready` and `valid` based handshaking for writing. In generated HDL, this block is an adapter between the provided interface and the actual Avalon memory-mapped host interface. Place these blocks at any level of hierarchy under DSP Builder device level block. The design can contain several of these blocks, with each of the blocks accessing the associated **External Memory** block.

**Table 76. Memory Write Parameters**

Parameter	Values	Description
<b>Identifier</b>	One of identifiers set for <b>External Memory</b> blocks in the design	Set to match an identifier on one of the <b>External Memory</b> blocks in the design.
<b>Byte Enables</b>	Boolean width	Activate this parameter to use byte enables for the write request. If enabled, DSP Builder adds a separate port to provide byte enable values.
<b>Maximum Burst Size</b>	off 2, 4, 8, 16, 32, 64, 128, 256, 512, 1024	If the value is set to <b>off</b> , DSP Builder does not allow burst requests. For any other values, DSP Builder adds a new port to specify an actual size (less than or equal specified <b>Maximum Burst Size</b> ) for each burst request. If you initiate a burst write, <b>External Memory</b> blocks ignore subsequent addresses until the burst is completed.

*continued...*

Parameter	Values	Description
		When a burst write is in progress, DSP Builder queues the read and write requests from associated <b>Memory Read</b> and <b>Memory Write</b> blocks until the write burst is completed.

**Table 77. Memory Write Ports**

Name	Direction	Type	Description
write	Input	Boolean	Set this port to high to indicate new write request to associated <b>External Memory</b> blocks.
address	Input	Unsigned integer	Sets the address for write request. The width of this port is the <b>Avalon MM Interface Data Width</b> parameter value on the associated <b>External Memory</b> block.
data	Input	Unsigned integer	Sets the write data. The width of this port is: $\log_2(\text{memory\_size})$ , <i>memory_size</i> is the size of associated <b>External Memory</b> .
byteenable	Input	Unsigned integer	Optional. DSP Builder adds <i>byteenable</i> when the <b>Byte Enables</b> parameter is on. Sets the byte enables for write data. The width of this port is: $\text{data\_port\_width} / 8$
burstcount	Input	Unsigned integer	Optional. DSP Builder adds <i>burstcount</i> when the <b>Maximum Burst Count</b> parameter is not <b>off</b> . Sets the actual burst count for burst write requests. When you initiate a burst request, ensure you update the <i>address</i> port and this port once at the beginning of request. Update the <i>write</i> port every time you update the <i>data</i> port to supply the next portion of burst data. For example, if you provide a new portion of data every cycle, keep the <i>write</i> port high throughout the burst. The width of this port is set as: $\log_2(\text{mac\_burst\_count})+1$
ready	Output	Boolean	Indicates whether the block is ready to accept a new write request or a continuation of ongoing burst request. Do not update input ports if this output is low.

### 13.1.5. Register Bit (RegBit)

The DSP Builder **RegBit** block provides a register bit that you can read in your model and read or write with the processor interface.

**Table 78. Parameters for the RegBit Block**

Parameter	Description
Register Offset	Specifies the address of the register. Must evaluate to an integer address.
Read/Write Mode	Specifies the mode of the memory as viewed from the processor:
<i>continued...</i>	

Parameter	Description
	<ul style="list-style-type: none"> <li>• <b>Read</b>: processor can only read over specified address range.</li> <li>• <b>Write</b>: processor can only write over specified address range.</li> <li>• <b>Read/Write</b>: processor can read or write over specified address range.</li> <li>• <b>Constant</b>: processor cannot access specified address range. This option continues to reserve space in the memory map.</li> </ul>
Bit	Specifies the bit location of the memory-mapped register in a processor word (allows different registers to share same address).
Initial Value	Specifies the initial state of the register.
Description	Text describing the register. The description is propagated to the generated memory map.
Sample Time	Specifies the Simulink sample time.

**Table 79. Port Interface for the RegBit Block**

Signal	Direction	Type	Description
q	Output	Boolean	Data.

**Related Information**

[Avalon Memory-Mapped Agent Settings \(AvalonMemoryMappedAgentSettings\)](#) on page 221

### 13.1.6. Register Field (RegField)

The DSP Builder **RegField** block provides a register field that you can read in your model and read or write with the processor interface.

**Table 80. Parameters for the RegField Block**

Parameter	Description
Register Offset	Specifies the address of the register. Must evaluate to an integer address.
Read/Write Mode	Specifies the mode of the memory as viewed from the processor: <ul style="list-style-type: none"> <li>• <b>Write</b>: processor can only write over specified address range.</li> <li>• <b>Read/Write</b>: processor can read or write over specified address range.</li> <li>• <b>Constant</b>: processor cannot access specified address range. This option continues to reserve space in the memory map.</li> </ul>
Most Significant Bit	Specifies the MSB of the memory-mapped register in a processor word (allows different registers to share same address). When multiple <b>RegBit</b> , <b>RegOut</b> , and <b>RegField</b> blocks specify the same address, they refer to the same Avalon-MM register. To avoid conflicts, ensure that the ranges that you specify do not overlap.
Least Significant Bit	Specifies the LSB of the memory-mapped register in a processor word (allows different registers to share same address). When multiple <b>RegBit</b> , <b>RegOut</b> , and <b>RegField</b> blocks specify the same address, they refer to the same Avalon-MM register. To avoid conflicts, ensure that the ranges that you specify do not overlap.
Register Output Type	Specifies the width and sign of the data type that the register stores. The size should equal (MSB – LSB + 1).
Register Output Scale	Specifies the scaling of data type that the register stores. For example, 2 <sup>-15</sup> for 15 of the above bits as fractional bits.
Initial Value	Specifies the initial state of the register.
Description	Text describing the register. The description is propagated to the generated memory map.
Sample Time	Specifies the Simulink sample time.

**Table 81. Port Interface for the RegField Block**

Signal	Direction	Type	Description
q	Output	As specified in <b>Register Output Type</b> .	Data.

**Related Information**

[Avalon Memory-Mapped Agent Settings \(AvalonMemoryMappedAgentSettings\)](#) on page 221

### 13.1.7. Register Out (RegOut)

The **RegOut** block provides a register field that you can write to your model and read from the processor interface.

**Table 82. Parameters for the RegOut Block**

Parameter	Description
Register Offset	Specifies the address of the register. Must evaluate to an integer address.
Most Significant Bit	Specifies the MSB of the memory-mapped register in a processor word (allows different registers to share same address). When multiple <b>RegBit</b> , <b>RegOut</b> , and <b>RegField</b> blocks specify the same address, they refer to the same Avalon-MM register. To avoid conflicts, ensure that the ranges that you specify do not overlap.
Least Significant Bit	Specifies the LSB of the memory-mapped register in a processor word (allows different registers to share same address). When multiple <b>RegBit</b> , <b>RegOut</b> , and <b>RegField</b> blocks specify the same address, they refer to the same Avalon-MM register. To avoid conflicts, ensure that the ranges that you specify do not overlap.
Description	Text describing the register. The description is propagated to the generated memory map.
Sample Time	Specifies the Simulink sample time.

**Table 83. Port Interface for the RegOut Block**

Signal	Direction	Type	Description
d	Input	Any fixed-point type	Write data.
w	Input	Boolean	Write enable.

### 13.1.8. Shared Memory (SharedMem)

The DSP Builder **SharedMem** block provides a memory block that you can read from or write to your model and read to or write from the processor interface.

The length of the **Initial Data** parameter, 1-D array, determines the size of the memory. You can optionally initialize the generated HDL with this data.

**Table 84. Parameters for the SharedMem Block**

Parameter	Description
Memory-Mapped Address	Specifies the address of the memory block. Must evaluate to an integer address.
Enable bit slicing	Turn on to allow multiple <b>SharedMem</b> blocks to occupy the same address range and each to take a slice of the data bus. When you turn on this parameter, enter the most and least significant bits of the bus that this <b>SharedMem</b> block connects to in the <b>MSB</b> and <b>LSB</b> parameters. When using this feature, some restrictions apply to the <b>SharedMem</b> block: <ul style="list-style-type: none"> <li>The bit-slice width must be equal to or less than the bus width (i.e. the <b>SharedMem</b> cannot be asymmetric)</li> <li>The bit-slice of one <b>SharedMem</b> block cannot overlap the bit-slice of another</li> <li>The bit-slice must match the size of the data type specified in the <b>Memory Output Type</b> parameter. If <b>SharedMem</b> blocks share address ranges, their address ranges must overlap exactly</li> <li>Only other <b>SharedMem</b> blocks can share an address with a <b>SharedMem</b> block</li> <li>The <b>SharedMem</b> block must have an auto-generated address map (i.e. the <b>Memory Mapped Address</b> parameter must be a scalar value)</li> </ul>
Read/Write Mode	Specifies the mode of the memory as viewed from the processor: <ul style="list-style-type: none"> <li><b>Read</b>: processor can only read over specified address range.</li> <li><b>Write</b>: processor can only write over specified address range.</li> <li><b>Read/Write</b>: processor can read or write over specified address range.</li> <li><b>Constant</b>: processor cannot access specified address range. This option continues to reserve space in the memory map.</li> </ul>
Initial Data	Specifies the initialization data. The size of the 1-D array determines the memory size.
Initialize Hardware Memory Blocks with Initial Data Contents	Turn on to initialize the memory with the specified initial data values. Turn off to use only the size of the initial data for the size of memory. The <b>Uninitialized means</b> parameter on the Control Block <b>Optimization</b> tab determines the size.
Description	Text describing the memory block. The description is propagated to the generated memory map.
Memory Output Type	Specifies the data type that the memory block stores.
Memory Output Scale	Specifies the scale factor to apply to the data stored in the memory block.
Sample Time	Specifies the Simulink sample time.

**Table 85. Port Interface for the SharedMem Block**

Signal	Direction	Type	Description
a	Input	Unsigned integer	Address.
wd	Input	Any fixed-point type	Write data.
we	Input	Boolean	Write enable.
rd	Output	Any fixed-point type	Read data.

### Intel Hyperflex Architecture Support for SharedMem Block

Intel Hyperflex architectures do not support all modes of memory operation and some modes are performance limited. For more information, refer to the *Intel Agilex Embedded Memory User Guide* or the *Intel Stratix 10 Embedded Memory User Guide*.

In designs on Intel Hyperflex architectures, Intel recommends you use a **SharedMem** block for one-way communication between internal and external Avalon-MM interfaces. Do not select **Read/Write** for **Read/Write Mode**; only use **Read** or **Write** for **Read/Write Mode** not both read and write. On the internal side, either do



not connect the `rd` interface or drive `we` to constant zero. Do not both dynamically drive `we` and use the `rd` output. Only use the **SharedMem** block in your design for one-way communication.

DSP Builder may duplicate your memory to provide support for up to one write with two reads on Intel Hyperflex architectures. Reads on the bus and system side are from separate copies of the memory and any writes are applied to both copies. DSP Builder offers **SharedMem** support in true dual port memory configurations depending on the constraints of the Intel Hyperflex architecture M20K block. **SharedMem** blocks have no support for dual clocks (bus clock must run at system rate) and no support for mixed widths (**SharedMem** data width must match bus width).

#### Related Information

- [Avalon Memory-Mapped Agent Settings \(AvalonMemoryMappedAgentSettings\)](#) on page 221
- [Control](#) on page 223
- [Intel Agilix Embedded Memory User Guide](#)
- [Intel Stratix 10 Embedded Memory User Guide](#)

## 13.2. Streaming Library

The **Streaming** library contains the extensible Avalon-ST interface blocks, which are masked subsystems.

1. [Avalon-ST Input \(AStInput\)](#) on page 273
2. [Avalon-ST Input FIFO Buffer \(AStInputFIFO\)](#) on page 274
3. [Avalon-ST Output \(AStOutput\)](#) on page 274

#### Related Information

- [Modifying Avalon-ST Blocks](#) on page 67
- [Restrictions for DSP Builder Designs with Avalon Streaming Interface Blocks](#) on page 67

### 13.2.1. Avalon-ST Input (AStInput)

Place this block at the front end of a system to generate the appropriate `hw.tcl` code for an Avalon Streaming interface with same name as the name of this block.

**Table 86. AStInput Block External Interface Signals**

Name	Direction	Description
<code>sink_channel</code>	input	Channel number.
<code>sink_data</code>	input	The data (which may be, or include control data).
<code>sink_eop</code>	input	Indicates end of packet.
<code>sink_ready</code>	output	Indicates to upstream components that the DSPBA component can accept <code>sink_data</code> on this rising clock edge.
<code>sink_sop</code>	input	Indicates start of packet.
<code>sink_valid</code>	input	Indicates that <code>sink_data</code> , <code>sink_channel</code> , <code>sink_sop</code> , and <code>sink_eop</code> are valid.

**Table 87. AStInput Block Internal Interface Signals**

Name	Direction	Description
input_channel	output	Channel number.
input_data	output	The data (which may be, or include control data).
input_eop	output	Indicates end of packet.
input_ready	input	indicates from the output of the DSP Builder component that it can accept sink_data on this rising clock edge.
input_sop	output	Indicates start of packet.
input_valid	output	indicates that input_data, input_channel, input_sop and input_eop are valid.

### 13.2.2. Avalon-ST Input FIFO Buffer (AStInputFIFO)

The **AStInputFIFO** block is the same as the **AStInput** block but includes FIFO buffers to capture data to implement backpressure. Specify FIFO characteristics (e.g. depth) from the parameters window of this block.

### 13.2.3. Avalon-ST Output (AStOutput)

Place this block at the back end of a system to generate the appropriate hw.tcl code for an Avalon Streaming interface with same name as the name of this block.

**Table 88. AStOutput Block External Interface Signals**

Name	Direction	Description
source_channel	Output	Channel number.
source_data	Output	The data to be output (which may be, or include control data).
source_eop	Output	Indicates end of packet,
source_ready	Input	Indicates from downstream components that they can accept source_data on this rising clock edge.
source_sop	Output	Indicates start of packet.
source_valid	Output	Indicates that source_data, source_channel, source_sop, and source_eop are valid.

**Table 89. AStOutput Block Internal Interface Signals**

Name	Direction	Description
output_channel	input	Channel number.
output_data	input	The output data (which may be, or include control data).
output_eop	input	Indicates end of packet.
output_ready	output	Indicates from the output of the DSP Builder component that it can accept sink_data on this rising clock edge.
output_sop	input	Indicates start of packet.
output_valid	input	Indicates that output_data, output_channel, output_sop, and output_eop are valid.

The downstream system component may not accept data and so may back pressure this block by forcing Avalon ST signal `source_ready = 0`. However, the design may still have valid outputs in the pipeline. You must store these outputs in memory. DSP Builder writes the output data for the design into a data FIFO buffer, with the Avalon-ST signals channel. It writes `sop` and `eop` into the respective channel, FIFO buffers.

Connect the backpressure signal (`source_ready`) from downstream components to port `ready` in this subsystem. Then DSP Builder reads the FIFO buffers when the downstream block can accept data (`read_fifo = 1`) and data in FIFO to output (`fifo_empty_n = 1`) exists.

If the downstream component is continually backpressuring this design, these FIFO buffers start to fill up. If you continue to feed data into the component, eventually the FIFO buffers overflow, which you must not allow to happen. Therefore, when the FIFO buffers reach a certain fill level, they assert signal `nearly_full = 1`. Use this signal to apply backpressure to upstream component (forcing Avalon ST signal `sink_ready = 0`). So that upstream components stop sending in more data and so that the FIFO buffer should not overflow, set the fill level at which `nearly_full = 1` to a value that depends on the latency of this design. For example, if the design contains a single **Primitive** subsystem and the `ChannelOut` block indicates a latency of `L`, assert the `nearly_full` flag at the latest point when `L` free entries are in the FIFO buffer. Setting this threshold is a manual process and the full threshold must be greater than or equal to (depth of FIFO buffer - `L`).

## 14. Primitives Library

---

Use the DSP Builder advanced blockset **Primitives** library blocks to create fast and efficient designs. DSP Builder captures these designs in the behavioral domain rather than the implementation domain by combining primitive functions. The **Primitives** library contains primitive operators such as add, multiply, and delay. It also includes functions to manipulate signal types that support building hardware functions that use MATLAB fixed-point types. You do not need to understand the details of the underlying FPGA architecture, as DSP Builder automatically maps the **Primitives** blocks into efficient FPGA constructs.

1. [Vector and Complex Type Support](#) on page 276
2. [DFT Design Elements Library](#) on page 278
3. [FFT Design Elements Library](#) on page 283
4. [Primitive Basic Blocks Library](#) on page 301
5. [Primitive Configuration Library](#) on page 359
6. [Primitive Design Elements Library](#) on page 363

### 14.1. Vector and Complex Type Support

The DSP Builder **Primitive** libraries provide automatic support for arrays and complex types.

These modes of operation engage with type propagation, and provide a convenient automatic method for generating repeated design elements to operate on all the data elements within vector and complex signals.

Blocks automatically determine whether the data they process is in scalar or vector format and operate accordingly.

Using complex data (where it is supported) automatically causes DSP Builder to generate blocks internally, which processes both real and imaginary data elements.

The hardware elements that these processes generate fully incorporate into the optimization schemes available within DSP Builder advanced blockset.

No restrictions on the combination of vector and complex modes exist.

1. [Vector Type Support](#) on page 276
2. [Complex Support](#) on page 277

#### 14.1.1. Vector Type Support

1. [Element by Element Mode](#) on page 277
2. [Mathematical Vector Mode](#) on page 277

3. [Interactions with Simulink](#) on page 277

#### 14.1.1.1. Element by Element Mode

The blocks in the primitive library exhibit an element by element mode of operation when you use them with vector types.

This mode provides a convenient way to generate a uniform array to handle each element of data in a vector signal, without having to manually instantiate multiple blocks.

Internally, DSP Builder generates identical block instantiations, one for each element in the vector signal. The vector width propagates through the Simulink system.

Change this mode of operation in one of the following two ways:

- Drive the block with a vector signal.
- Initialize the block with a vector of values. This option is only available for blocks that you can initialize with a user-specified value.

The following restrictions exist on the vectors:

- Vector signals must be of uniform type.
- Signals associated with a block must either be vectors of identical width, or scalar.

When you use a scalar value with vectors, DSP Builder uses a copy of the single scalar value with each data element in the vector signal.

This behavior is analogous to the scalar expansion that occurs with Simulink blocks.

#### 14.1.1.2. Mathematical Vector Mode

The blocks in the primitive vector library perform mathematical operations with vector data.

The outputs of these blocks are potentially a function of any or all of the inputs. Vector width does not necessarily propagate.

The **SumOfElements** block exhibits this behavior.

#### 14.1.1.3. Interactions with Simulink

You can use Simulink **Mux** and **Demux** to manipulate signals within DSP Builder advanced blockset designs.

#### 14.1.2. Complex Support

Some DSP Builder **Primitive** library blocks can automatically process complex data, which provides a convenient way to simultaneously generate data and control pathways for the real and imaginary components of such data.

For each complex value, two identical block instantiations generate internally, for the real and imaginary components.

The complex nature of the data propagates. Strictly real signals expand to provide a value for the imaginary component with complex data. The exact behavior depends on the nature of the port associated with the real signal. The real value is duplicated for

control or address signals. The real and imaginary parts of complex data are subject to identical control signals. A zero imaginary value generates real data signals in a complex data context. Real data values,  $x$ , expand, when required, to  $x + 0i$ .

Not all **Primitive** library blocks support complex data. Data signals are the only signal type permitted to be complex. DSP Builder issues an error message if an attempt is made to drive control or address signals with complex values.

1. [Interactions with Simulink](#) on page 278

### 14.1.2.1. Interactions with Simulink

You can use the complex Simulink function **complex(x,y)** to generate initialization values. Use this function to ensure DSP Builder always treats data as complex

The following elements of the Simulink environment are available for use with the primitive blocks

- Simulink **Complex to Real-Imag** and **Real-Imag to Complex** blocks may manipulate complex signals within DSP Builder advanced blockset designs.
- Simulink **Scope** blocks can display signals, but they do not directly support complex data. Attempting to view complex data generates a type propagation error.

Use a **Complex to Real-Imag** block to convert the complex signal.

Simulink automatically converts complex values of form  $(x + 0i)$  to real values, which can cause type propagation errors. The **complex()** function can resolve this problem.

Use **complex(x,0)** to ensure such data is treated as complex.

## 14.2. DFT Design Elements Library

### 14.2.1. DFT (DFT)

The **DFT** block performs a discrete Fourier transform (DFT) or an inverse DFT (IDFT) of a fixed-point complex input sequence and produces a fixed-point complex output sequence. The `demo_dft.mdl` example design demonstrates the DFT block.

You can specify the transform length (DFT size) at runtime (on a block-by-block basis) to any one of the 53 sizes specified by *3GPP TS 36.101 version 8.29.0 Release 8 (Table 1)*. In addition to the DFT block, the DSP Builder block-set provides an FFT block. However, the FFT block cannot support these 53 sizes as it is limited to powers of two sizes.

The DFT block is a back-to-back streaming DFT. This design allows the block to process DFTs continuously, with no breaks required when the DFT size changes, and no busy periods when it is unable to accept new input.

When the DFT block pipeline is empty, the input-to-output latency for a DFT operation (first input to first output) depends only on the current DFT size and the pipeline overhead. The pipeline overhead is constant for all dynamic DFT sizes.

When the pipeline is not empty, the input-to-output latency for a DFT operation is affected by the current contents of the pipeline. The DFT block implements a queuing system to prevent small DFT operations from overtaking (or colliding with) larger

ones. For example, if a 12-point DFT enters the pipeline immediately after a 3,240-point DFT, it can't produce its first output until the larger DFT has output all 3,240 points.

**Table 90. DFT Block Supported Sizes**

The DFT block supports the following 53 sizes specified by 3GPP TS 36.101 version 8.29.0 Release 8:

12	120	324	648	1,152	1,800	2,880
24	144	360	720	1,200	1,920	2,916
36	180	384	768	1,296	1,944	3,000
48	192	432	864	1,440	2,160	3,072
60	216	480	900	1,500	2,304	3,240
72	240	540	960	1,536	2,400	
96	288	576	972	1,620	2,592	
108	300	600	1,080	1,728	2,700	

**Table 91. DFT Block Signals**

Signal	Direction	Description
v	Input	Indicates if the input data is valid. Assert this Boolean signal when input data is valid. If <b>Low Latency Implementation</b> parameter is turned on, the valid input must remain high from the first input of a DFT operation to the last. If <b>Low Latency Implementation</b> parameter is turned off, any invalid (v deasserted) cycles during a DFT operation are ignored.
cmd	Input	An encoding of the size of the current DFT. In DFT/IDFT dynamic mode, an extra MSB indicates whether to calculate a DFT (0) or an IDFT (1). The size of the DFT and DFT/IDFT mode, if appropriate, must be specified at the same time (cycle) as a valid sop signal. The value of cmd is ignored at all other times. When all 53 sizes are supported, in DFT-only mode, a 6-bit integer encodes the size of the DFT currently input. In dynamic DFT/IDFT-mode, it is a 7-bit integer. The bottom 6 bits encode the size of the current DFT/IDFT, the MSB signifies whether to calculate a DFT or an IDFT. The DFT specifies the size of the DFT by specifying the zero-based index of the required size in the <b>FFT Size Table</b> parameter. For example, where <b>FFT Size Table</b> is [12 24 36], 0 specifies size 12 and 1 specifies size 24.
sop	Input	A start-of-packet Boolean signal that pulses high on the first (and only the first) valid input of a DFT/IDFT. Assert sop to indicate the first data input of a DFT. The generated RTL requires at least one cycle gap between the reset deasserting and the first sop signal.
data	Input	A complex input data signal of the user-specified type ( <b>Input Type</b> ).
v	Output	v indicates if the output data is valid. This boolean signal will be asserted for valid output data.
cmd	Output	An encoding of the size of the current DFT. In DFT/IDFT dynamic mode, an extra MSB indicates whether the current output is a DFT (0) or an IDFT (1). When all 53 sizes are supported, cmd is a 6-bit signal with the size of the current DFT. In dynamic DFT/IDFT mode cmd is a 7-bit signal with the current DFT/IDFT size and status of the DFT/IDFT. Specify the size of the DFT by specifying the zero-based index of the required size in the <b>FFT Size Table</b> vector. For example, where <b>FFT Size Table</b> is [12 24 36], 0 specifies size 12 and 1 specifies size 24.

*continued...*

Signal	Direction	Description
sop	Output	A start-of-packet Boolean signal that pulses high on the first (and only the first) valid output of a DFT/IDFT.
pos	Output	Indicates the natural order position of the current input. The data outputs are produced in radix-reversed order. Writing each sample to memory at location <code>pos</code> means the memory contains the samples in natural order.
data	Output	A full word growth complex output signal. When all 53 sizes are specified, DSP builder calculates the type of this signal by applying 13 bits of word-growth to the input data type.


**Table 92. DFT Block Parameters**

Parameter	Description
<b>FFT Size Table</b>	Specify the list of DFT sizes to support, in the form of a MATLAB vector. For example, [12 24 36]. Allowed sizes are specified in the <i>DFT Block Supported Sizes</i> table.
<b>Input Type</b>	Specify the fixed-point MATLAB type of the input data. For example, <code>fixdt(1,16,15)</code>
<b>Twiddle Type</b>	Specify the fixed-point MATLAB type of the twiddle constants, and the multiplicative constants needed for radix-3 and radix-5 FFT stages. For example, <code>fixdt(1,18,17)</code>
<b>Support dynamic DFT/IDFT</b>	When false, the block only supports DFTs; when true it also supports IDFTs.
<b>Low Latency implementation</b>	Select to choose low latency implementation. If you choose the low latency implementation the data valid input <code>v</code> must remain high from the first input of a DFT operation to the last.

### 14.2.2. Reorder (ReorderBlock)

The DFT block produces a `pos` signal to indicate where each output appears in a natural order sequence but doesn't actually reorder the outputs. The **ReorderBlock** performs the necessary physical reordering.

**Figure 96. ReorderBlock Block Diagram**



**Table 93. ReorderBlock Block Signals**

Port Name	Direction	Description
<code>v</code>	Input	Indicates if the input data is valid. Assert this Boolean signal when input data is valid. Connect <code>v</code> output of the preceding DFT to this input.
<code>cmd</code>	Input	An encoding of the size of the current DFT. In DFT/IDFT dynamic mode, an extra MSB indicates whether to calculate a DFT (0) or an IDFT (1). When all 53 sizes are supported, in DFT-only mode, a 6-bit integer encodes the size of the DFT currently input. In dynamic DFT/IDFT-mode, it is a 7-bit integer. The bottom 6 bits encode the size of the current DFT/IDFT, the MSB signifies whether to calculate a DFT or an IDFT. Specify the size of the DFT by specifying the zero-based index of the required size in the <b>FFT Size Table</b> parameter. For example, where <b>FFT Size Table</b> is [12 24 36], 0 specifies size 12 and 1 specifies size 24.

*continued...*



Port Name	Direction	Description
		Connect the <code>cmd</code> output of the preceding DFT to this input.
<code>sop</code>	Input	A start-of-packet Boolean signal that pulses high on the first (and only the first) valid input of an DFT/IDFT. The generated RTL requires at least one cycle gap between the reset deasserting and the first <code>sop</code> signal. Connect the <code>sop</code> output of the preceding DFT to this input.
<code>pos</code>	Input	Indicates the natural order position of the current input. Writing each sample to memory at location <code>pos</code> , gives the memory containing the samples in natural order. Connect the <code>pos</code> output of the preceding DFT to this input.
<code>data</code>	Input	A complex input data signal. Connect <code>data</code> output of the preceding DFT to this input.
<code>v</code>	Output	Indicates if the output data is valid. This Boolean signal is asserted for valid output data.
<code>cmd</code>	Output	An encoding of the size of the current DFT. In DFT/IDFT dynamic mode, an extra MSB indicates whether the current output is a DFT (0) or an IDFT (1). When all 53 sizes are supported, <code>cmd</code> is a 6-bit signal with the size of the current DFT. In dynamic DFT/IDFT mode <code>cmd</code> is a 7-bit signal with the current DFT/IDFT size and status of the DFT/IDFT. The DFT specifies the size of the DFT by specifying the zero-based index of the required size in the <b>Size Table</b> parameter. For example, where <b>Size Table</b> is [12 24 36], 0 specifies size 12 and 1 specifies size 24.
<code>sop</code>	Output	A start-of-packet Boolean signal that pulses high on the first (and only the first) valid output of a reordered DFT/IDFT.
<code>data</code>	Output	A complex output data signal, which is the reordered data.

**Table 94. ReorderBlock Block Parameters**

Parameter	Description
<b>Size Table</b>	Specify the of DFT sizes to support, in the form of a MATLAB vector. For example, [12 24 36]. Allowed sizes are specified in the <i>DFT Block Supported Sizes</i> table. This vector must match the <b>FFT Size Table</b> specification of the preceding DFT.

### 14.2.3. Reorder and Rescale (ReorderAndRescale)

The **ReorderAndRescale** block is a modified version of the **Reorder** block that performs dynamic rescaling in addition to reordering, so it produces its output in block floating-point format.

The scaling operation first calculates the minimum number of duplicated sign bits across all data outputs of a DFT. These duplicated MSBs can be dropped without any loss of information. Then the block drops LSBs, if required, to achieve the requested width of the output (**Output Width** parameter).

**Table 95. ReorderAndRescale Block Signals**

Port Name	Direction	Description
<code>v</code>	Input	Indicates if the input data is valid. Assert this Boolean signal when input data is valid. Connect <code>v</code> output of the preceding DFT to this input.
<code>cmd</code>	Input	An encoding of the size of the current DFT. For dynamic DFT/IDFTs, an extra MSB indicates whether to calculate a DFT (0) or an IDFT (1).

*continued...*

Port Name	Direction	Description
		<p>When all 53 sizes are supported, in DFT-only mode, a 6-bit integer encodes the size of the DFT currently input. In dynamic DFT/IDFT-mode, it is a 7-bit integer. The bottom 6 bits encode the size of the current DFT/IDFT, the MSB signifies whether to calculate a DFT or an IDFT.</p> <p>Specify the size of the DFT by specifying the zero-based index of the required size in the <b>FFT Size Table</b> vector. For example, where the <b>FFT Size Table</b> is [12 24 36], 0 specifies size 12 and 1 specifies size 24.</p> <p>Connect the <code>cmd</code> output of the preceding DFT to this input.</p>
<code>sop</code>	Input	<p>A start-of-packet Boolean signal that pulses high on the first (and only the first) valid input of an DFT/IDFT.</p> <p>The generated RTL requires at least one cycle gap between the reset deasserting and the first <code>sop</code> signal.</p> <p>Connect the <code>sop</code> output of the preceding DFT to this input.</p>
<code>pos</code>	Input	<p>Indicates the natural order position of the current input. Writing each sample to memory at location <code>pos</code> means the memory containing the samples in natural order.</p> <p>Connect the <code>pos</code> output of the preceding DFT to this input.</p>
<code>data</code>	Input	<p>A complex input data signal.</p> <p>Connect <code>data</code> output of the preceding DFT to this input.</p>
<code>v</code>	Output	<p>Indicates that the output data is valid. This Boolean signal is asserted for valid output data.</p>
<code>cmd</code>	Output	<p>An encoding of the size of the current DFT. In DFT/IDFT dynamic mode, an extra MSB indicates whether the current output is a DFT (0) or an IDFT (1).</p> <p>When all 53 sizes are supported, <code>cmd</code> is a 6-bit signal with the size of the current DFT. In dynamic DFT/IDFT mode <code>cmd</code> is a 7-bit signal with the current DFT/IDFT size and status of the DFT/IDFT.</p> <p>The DFT specifies the size of the DFT by specifying the zero-based index of the required size in the <b>Size Table</b> parameter. For example, where <b>Size Table</b> is [12 24 36], 0 specifies size 12 and 1 specifies size 24.</p>
<code>sop</code>	Output	<p>A start-of-packet Boolean signal that pulses high on the first (and only the first) valid output of a DFT/IDFT.</p>
<code>data</code>	Output	<p>A complex output data signal, which is the reordered and rescaled data.</p> <p>The output data width is specified by the <b>Output Width</b> parameter and it has the same number of fractional bits as the input type (<b>Input Type</b> parameter).</p>
<code>shift</code>	Output	<p>The rescaling factor.</p> <p>The <code>shift</code> output is constant for each DFT processed but varies between DFTs.</p> <p>For the <b>ReorderAndRescale</b> block, <math>data \times 2^{shift}</math> is approximately equal (within rounding error) to <code>data</code> for the <b>Reorder</b> block.</p>

**Table 96. ReorderAndRescale Parameters**

Parameter	Description
<b>Size Table</b>	<p>Specify the list of DFT sizes to support, in the form of a MATLAB vector. For example, [12 24 36]. Allowed sizes are specified in the <i>DFT Block Supported Sizes</i> table.</p> <p>This vector must match the <b>FFT Size Table</b> specification of the preceding DFT.</p>
<b>Input Type</b>	<p>Specify the fixed-point input type of the <b>ReorderAndRescale</b> block. The value must be the same as the output type of the preceding DFT block.</p> <p>For example, <code>fixdt(0,29,15)</code></p>
<b>Output Width</b>	<p>Specify the requested (integer and fractional) width of rescaled data output.</p> <p>The output type has the same number of fractional bits as the <b>Input Type</b>.</p>

### 14.3. FFT Design Elements Library

Use the DSP Builder advanced blockset **FFT Design Elements** library blocks to support FFT designs. The library also includes several blocks that support for a radix-2<sup>2</sup> algorithm.

The radix-2<sup>2</sup> architecture is a serial version of the radix-4 architecture. It computes a radix-4 butterfly over four (not necessarily consecutive) inputs and produces four (not necessarily consecutive) outputs.

For more information about the radix-2<sup>2</sup> algorithm, refer to *A New Approach to Pipeline FFT Processor – Shousheng He & Mats Torkleson, Department of Applied Electronics, Lund University, Sweden.*

1. [About Pruning and Twiddle for FFT Blocks](#) on page 284
2. [Bit Vector Combine \(BitVectorCombine\)](#) on page 285
3. [Butterfly Unit \(BFU\)](#) on page 285
4. [Butterfly I C \(BFIC\) \(Deprecated\)](#) on page 286
5. [Butterfly II C \(BFIIC\) \(Deprecated\)](#) on page 287
6. [Choose Bits \(ChooseBits\)](#) on page 288
7. [Crossover Switch \(XSwitch\)](#) on page 288
8. [Dual Twiddle Memory \(DualTwiddleMemoryC\)](#) on page 288
9. [Edge Detect \(EdgeDetect\)](#) on page 289
10. [Floating-Point Twiddle Generator \(TwiddleGenF\) \(Deprecated\)](#) on page 289
11. [Fully-Parallel FFTs \(FFT2P, FFT4P, FFT8P, FFT16P, FFT32P, and FFT64P\)](#) on page 290
12. [Fully-Parallel FFTs with Flexible Ordering \(FFT2X, FFT4X, FFT8X, FFT16X, FFT32X, and FFT64X\)](#) on page 290
13. [General Multitwiddle and General Twiddle \(GeneralMultiTwiddle, GeneralMultVTwiddle, GeneralTwiddle, GeneralVTwiddle\)](#) on page 291
14. [Hybrid FFT \(Hybrid\\_FFT, HybridVFFT\)](#) on page 293
15. [Multiwire Transpose \(MultiwireTranspose\)](#) on page 294
16. [Parallel Pipelined FFT \(PFFT\\_Pipe\)](#) on page 295
17. [Pulse Divider \(PulseDivider\)](#) on page 295
18. [Pulse Multiplier \(PulseMultiplier\)](#) on page 296
19. [Single-Wire Transpose \(Transpose\)](#) on page 296
20. [Split Scalar \(SplitScalar\)](#) on page 297
21. [Streaming FFTs \(FFT2, FFT4, VFFT2, and VFFT4\)](#) on page 297
22. [Stretch Pulse \(StretchPulse\)](#) on page 298
23. [Twiddle Angle \(TwiddleAngle\)](#) on page 298
24. [Twiddle Generator \(TwiddleGenC\) \(Deprecated\)](#) on page 299
25. [Twiddle and Variable Twiddle \(Twiddle and VTwiddle\)](#) on page 299
26. [Twiddle ROM \(TwiddleRom, TwiddleMultRom and TwiddleRomF \(deprecated\)\)](#) on page 300

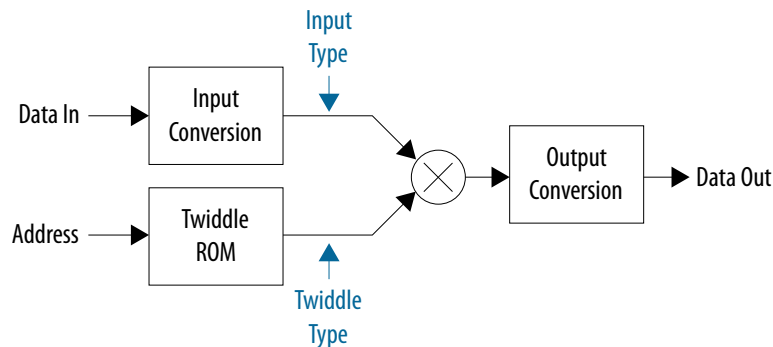
### 14.3.1. About Pruning and Twiddle for FFT Blocks

DSP Builder allows you to specify: the type of the data values before each twiddle multiplication; the type of the twiddle constants; the type of the data values after each twiddle multiplication.

For example:

```
dspba.fft.full_wordgrowth(true,false,2,fixdt(1,16,15),fixdt(1,18,17))
```

**Figure 97. Pruning and Twiddle for FFT Blocks**



An FFT with  $2N$  points has  $N$  radix-2 stages and (conceptually)  $N-1$  twiddle multipliers. In practice, DSP Builder optimizes away many of the twiddle multipliers. However, they still need entries in the twiddle specification.

The twiddle and pruning specification for this FFT consists of a  $(N-1) \times 3$  array ( $N-1$  rows with 3 entries in each row) of strings which specify these types. DSP Builder uses strings because Simulink does not pass raw types into the Simulink GUI.

DSP Builder provides three utility functions to generate twiddle and pruning specifications, each of which implements a different pruning strategy:

- `dspba.fft.full_wordgrowth(complexFFT,radix2,N,input_type,twiddle_type)`
- `dspba.fft.mild_pruning(complexFFT,radix2,N,input_type,twiddle_type)`
- `dspba.fft.prune_to_width(maxWidth,complexFFT,radix2,N,input_type,twiddle_type)`

In addition, DSP Builder provides a fourth function for floating-point FFTs (where no pruning is required)

- `dspba.fft.all_float(N, float_type)`

This function generates a pruning specification where the input, twiddle and output types are all `float_type`.

The legacy FFT interfaces use `dspba.fft.full_wordgrowth()` pruning strategy. It grows the datapath by one bit for each radix-2 FFT stage.

The `dspba.fft.mild_pruning()` grows the datapath by one bit for each two radix-2 FFT stages.

The `dspba.fft.prune_to_width(maxWidth)` grows the datapath by one bit for each radix-2 FFT stage up to the specified maximum width. At that point, it applies drastic pruning to ensure that the data input to the twiddle multiplier is never more than `maxWidth` bits wide.

Intel provides these built-in strategies only for your convenience. If you need a different pruning strategy, you can define and use your own pruning function (or just construct the pruning or twiddle array manually).

Each of these utility functions generate an array in the appropriate format ( $N-1$  rows, each containing three entries).

In each case:

- `complexFFT` is a Boolean number (usually true) that indicates whether the FFT's input is complex.
- `radix2` is a Boolean number (usually false) that indicates whether the FFT can have two consecutive twiddle stages.
- `N` is an integer indicating the number of radix-2 stages in the FFT. For example, 10 for a 1,024-point FFT.
- `input_type` is the type of the input signal.
- `twiddle_type` is the type of the twiddle constants.

### 14.3.2. Bit Vector Combine (BitVectorCombine)

The **BitVectorCombine** block concatenates a vector of bits to form a scalar. The scalar is an unsigned integer of the appropriate width. The first element of the vector becomes the least significant bit of the scalar (little-endian ordering).

Use the **BitVectorCombine** block to recombine scalars that the **SplitScalar** block splits.

**Table 97. Parameters for the BitVectorCombine Block**

Parameter	Description
Width	Width of the input vector (and the output scalar).

**Table 98. Port Interface for the BitVectorCombine Block**

Signal	Direction	Type	Description
d	Input	Boolean vector.	Data input.
q	Output	Unsigned integer.	Data output.

### 14.3.3. Butterfly Unit (BFU)

The **BFU**, **BFU\_long**, **BFU\_short**, and **BFU\_simple** blocks each implement a butterfly unit for use in floating-point streaming FFTs.

The **BFU\_long** block corresponds to a classical radix-2<sup>2</sup> butterfly I block plus its associated feedback path.

The **BFU\_short** block has exactly the same functionality, but it uses only one floating-point adders. It uses twice as many memory resources as the **BFU\_long** block, but also uses considerably less logic resources.

The **BFU** block automatically reconfigures to use either **BFU\_long** or **BFU\_short** to minimize the total (memory plus logic) resource usage.

Each BFU block performs a two-point FFT pass over a block of data of size  $2^N$  (where N is a compile-time parameter).

During the first  $2^{(N-1)}$  cycles, the control signal, **s**, is 0. During this time, the BFU block stores the first half of the input block.

During the second  $2^{(N-1)}$  cycles, **s** is 1. During this time, the **BFU** block reads the second half of the input block and produces the first result of each of  $2^{(N-1)}$  two-point FFTs on the output.

During the third  $2^{(N-1)}$  cycles, **s** is 0 again. During this time, the BFU unit produces the second result of each of the  $2^{(N-1)}$  two-point FFTs, while simultaneously storing the first half of the next input block.

**Table 99. Parameters for the BFU Block**

Parameter	Description
<b>N</b>	Specifies the input block size to be $2^N$ .

**Table 100. Port Interface for the BFU Block**

Signal	Direction	Type	Description
d	Input	Any floating-point type.	Input samples.
q	Output	Same as d.	Output results.
s	Input	Boolean.	Control pin. Drive with external logic. Ensure it is 0 for $2^{(N-1)}$ cycles and 1 for the next $2^{(N-1)}$ cycles.

### 14.3.4. Butterfly I C (BFIC) (Deprecated)

The **BFIC** block implements the butterfly I functionality associated with the radix-2<sup>2</sup> fully streaming FFT architecture.

You should parameterize this block with the incoming data type to ensure that DSP Builder maintains the necessary data precision. At the output, DSP Builder applies an additional bit of growth.

The **s** port connects to the control logic. This control logic is the extraction of the appropriate bit of a modulo N counter. The value of **s** determines the signal routing of each sample and the mathematical combination with other samples.

**Table 101. Parameters for the BFIC Block**

Parameter	Description
Input bits	Specifies the number of input bits.
Input scaling exponent	Specifies the fixed-point scaling factor of the input.

**Table 102. Port Interface for the BFIC Block**

Signal	Direction	Type	Description
s	Input	Boolean or unsigned integer <b>uint(1)</b>	Control pin.
x1	Input	Complex fixed-point data-type determined by parameterization	Complex data input from <b>ComplexSampleDelay</b> .
x2	Input	Complex fixed-point data-type determined by parameterization	Complex data input from previous stage.
z1	Output	Complex fixed-point data-type determined by parameterization	Complex data output to next stage.
z2	Output	Complex fixed-point data-type determined by parameterization	Complex data output to <b>ComplexSampleDelay</b> .

### 14.3.5. Butterfly II C (BFIIC) (Deprecated)

The **BFIIC** block implements the butterfly II functionality associated with the radix-2<sup>2</sup> fully streaming FFT or iFFT architecture.

You should parameterize this block with the incoming data type to ensure that DSP Builder maintains the necessary data precision. At the output, DSP Builder applies an additional bit of growth.

The *s* port connects to the control logic. This control logic is the extraction of the appropriate bit of a modulo N counter. The value of *s* determines the signal routing of each sample and the mathematical combination with other samples. The *t* port also connects to the control logic, but the extracted bit is different from the *s* port. The value of *t* determines whether an additional multiplication by  $-j$  occurs inside the butterfly unit.

**Table 103. Parameters for the BFIIC Block**

Parameter	Description
IFFT	Specifies that the design uses the <b>BFIIC</b> block in an IFFT.
Input bits	Specifies the number of input bits.
Input scaling exponent	Specifies the exponent part of the input scaling factor ( $2^{-\text{exponent}}$ ).
Allow output bitwidth growth	Specifies that the output is one bit wider than the input.

**Table 104. Port Interface for the BFIIC Block**

Signal	Direction	Type	Description
s	Input	Boolean	Control pin.
t	Input	Boolean	Control pin.
x1	Input	Complex fixed-point data-type determined by parameterization	Complex input from <b>ComplexSampleDelay</b> .
x2	Input	Complex fixed-point data-type determined by parameterization	Complex input from previous stage.
z1	Output	Derived complex fixed-point type	Complex output to next stage.
z2	Output	Derived complex fixed-point type	Complex output to <b>ComplexSampleDelay</b> .

### 14.3.6. Choose Bits (ChooseBits)

The **ChooseBits** block selects individual bits from its input (scalar) signal and concatenates them to form its (scalar) output signal.

You specify the bits that occur in the output signal by providing a vector of non-negative integers. Each integer specifies an input bit appears in the output. The block numbers the input bits from 0 (least significant bit) and lists the output bits starting from the least significant bit (little-endian ordering).

The block has no restriction on how many times each input bit may appear in the output. You can omit, reorder, or duplicate bits.

For example, the vector [0,1,4,4,6,5] keeps bits 0 and 1 unchanged, omits bit 3, duplicates bit 4 and swaps the positions of bits 5 and 6.

**Table 105. Parameters for the ChooseBits Block**

Parameter	Description
Selected bits	A vector of non-negative integers.

**Table 106. Port Interface for the ChooseBits Block**

Signal	Direction	Type	Description
d	Input	Any scalar.	Data input.
q	Output	Unsigned integer.	Data output.

### 14.3.7. Crossover Switch (XSwitch)

The **XSwitch** block a simple crossover switch.

**Table 107. Port Interface for the XSwitch Block**

Signal	Direction	Type	Description
swap	Input	Boolean.	Control input. When swap is 0, d0 is routed to q0 and d1 is routed to q1; when swap is 1, d0 is routed to q1 and d1 is routed to q0.
d0	Input	Any.	Data input.
d1	Input	Same as d0.	Data input.
q0	Output	Same as d0.	Data output.
q1	Output	Same as d0.	Data output.

### 14.3.8. Dual Twiddle Memory (DualTwiddleMemoryC)

The **DualTwiddleMemory** block calculates the complex twiddle factors associated with the evaluation of  $\exp(-2\pi i.k1/N)$  and  $\exp(-2\pi i.k2/N)$ .

This block uses an efficient dual-port architecture to minimize the size of the internal lookup table while supporting the generation of two complex twiddle factors per clock cycle. The block provides **k1** and **k2** at the input and they must be less than or equal to a synthesis time parameter N. Enter the width in bits and fixed-point scaling of the twiddle factors.



A cosine/sine wave has a range of  $[-1:1]$ , so you must provide at least two integer bits, and as many fractional bits as are appropriate. A good starting point is a twiddle width in bits of 16 bits (enter **16** as the **Precision**), and a scaling of  $2^{-14}$  (enter 14 as the **Scaling** exponent). The resulting fixed-point type is **sfixed16\_en14** (2.14 in fixed-point format).

**Table 108. Parameters for the DualTwiddleMemoryC Block**

Parameter	Description
Number of points (N)	Specifies the number of points on the unit circle.
Precision	Specifies the precision in bits of the twiddle factors.
Twiddle scaling exponent	Specifies the fixed-point scaling factor of the complex twiddle factor.

**Table 109. Port Interface for the DualTwiddleMemoryC Block**

Signal	Direction	Type	Description
k1	Input	Unsigned integer in range 0 to (N - 1)	Desired twiddle factor index.
k2	Input	Unsigned integer in range 0 to (N - 1)	Desired twiddle factor index.
q1	Output	Type determined by parameterization	Twiddle factor 1 (complex).
q2	Output	Type determined by parameterization	Twiddle factor 2(complex).

### 14.3.9. Edge Detect (EdgeDetect)

The **EdgeDetect** block implements a simple circuit that detects edges on its input. It outputs 0 if the current input is the same as the previous input and 1 if the inputs are different.

The **EdgeDetect** block has no parameters.

**Table 110. Port Interface for the EdgeDetect Block**

Signal	Direction	Type	Description
d	Input	Boolean or <b>ufixed(1)</b> .	Data input.
q	Output	Same as input	Data output.

#### Related Information

[About Pruning and Twiddle for FFT Blocks on page 284](#)

### 14.3.10. Floating-Point Twiddle Generator (TwiddleGenF) (Deprecated)

The **TwiddleGenF** block is the floating-point version of the fixed-point **TwiddleGenC** block. The **TwiddleGenF** block generates the appropriate complex coefficients that multiply the streaming data in a radix- $2^2$  streaming FFT or IFFT architecture.

**Table 111. Parameters for the TwiddleGenC Block**

Parameter	Description
FFT type	Specifies whether to generate twiddle factors for an FFT or an IFFT.
Twiddle type	Specifies the floating-point type used for the twiddle factors.

**Table 112. Port Interface for the TwiddleGenF Block**

Signal	Direction	Type	Description
<b>counter</b>	Input	Unsigned integer	Counter signal.
<b>w</b>	Output	Complex floating-point type	Complex data output.

### 14.3.11. Fully-Parallel FFTs (FFT2P, FFT4P, FFT8P, FFT16P, FFT32P, and FFT64P)

The **FFT2P**, **FFT4P**, **FFT8P**, **FFT16P**, **FFT32P**, and **FFT64P** blocks implement fully-parallel FFTs for 2, 4, 8, 16, 32, and 64 points respectively.

The blocks expect bit-reversed input and produce natural-order output.

Not all parameters are available with all blocks.

**Table 113. Parameters for the FFT2P Blocks**

Parameter	Description
<b>Twiddle/pruning specification</b>	
<b>Use faithful rounding</b>	True if the block uses faithful (rather than correct) rounding for floating-point operations. Fixed-point FFTs ignore this parameter.

**Table 114. Port Interface for the FFT2P Blocks**

Signal	Direction	Type	Description
d	Input	Any complex.	Complex data input signal.
q	Output	Determined by pruning specification.	Complex data output signal.

#### Related Information

[About Pruning and Twiddle for FFT Blocks on page 284](#)

### 14.3.12. Fully-Parallel FFTs with Flexible Ordering (FFT2X, FFT4X, FFT8X, FFT16X, FFT32X, and FFT64X)

The **FFT2X**, **FFT4X**, **FFT8X**, **FFT16X**, **FFT32X**, and **FFT64X** blocks implement fully-parallel FFTs (or iFFTs) for 2, 4, 8, 16, 32, and 64 points respectively.

Unlike the corresponding P blocks (**FFT2P**, **FFT4P**, etc), they implement both FFTs and iFFTs and offer flexible ordering of the input and output wires.

Each block can also be internally parallelized to process several FFTs at once. For example, if there are 16 wires, each **FFT8P** block can calculate two 8-point FFTs (by specifying the number of spatial bits to be 4). With 32 wires, the same block can calculate four 8-point FFTs (by specifying the number of spatial bits to be 5).

Not all parameters are available with all blocks.

**Table 115. Parameters for the FFT2X Blocks**

Parameter	Description
<b>iFFT</b>	<b>true</b> to implement an IFFT, otherwise <b>false</b> .
<b>Number of spatial bits</b>	M for $2^M$ wires.
<b>Bit-reversed input</b>	<b>true</b> if you expect bit-reversed input, otherwise <b>false</b> .
<b>Bit-reversed output</b>	<b>true</b> if you want bit-reversed output, otherwise <b>false</b> .
<b>Twiddle/pruning specification</b>	
<b>Use faithful rounding</b>	<b>true</b> if the block uses faithful (rather than correct) rounding for floating-point operations. Fixed-point FFTs ignore this parameter.

**Table 116. Port Interface for the FFT2X Blocks**

Signal	Direction	Type	Description
v	Input	Boolean.	Valid input signal.
d	Input	Any complex.	Complex data input signal.
qv	Output	Boolean.	Valid output signal.
q	Output	Determined by pruning specification.	Complex data output signal.

**Related Information**

[About Pruning and Twiddle for FFT Blocks](#) on page 284

**14.3.13. General Multitwiddle and General Twiddle (GeneralMultiTwiddle, GeneralMultVTwiddle, GeneralTwiddle, GeneralVTwiddle)**

Use the **GeneralTwiddle** and **GeneralMultiTwiddle** blocks to construct supersampled FFTs. The blocks have the same external interface but use different internal implementations.

The **GeneralTwiddle** block generates its twiddle factors using the **TwiddleRom** block; the **GeneralMultiTwiddle** block uses the **TwiddleMultRom** block. The **GeneralMultiTwiddle** uses approximately twice as many DSP blocks as the **GeneralTwiddle** block, but (for large FFTs) uses far fewer memory blocks.

Each data sample in the input stream has a unique address. The address consists of the timeslot in which it arrived *tbits* concatenated with the number of wires on which it arrived *sbits*. The *sbits* forms the least significant part of the address; the *tbits* forms the most significant part.

Each data sample is multiplied by a twiddle factor. For an FFT, the twiddle factor is:

$$\text{twiddle} = \exp(-2 \cdot \pi \cdot i \cdot \text{angle} / K)$$

For an IFFT, the twiddle factor is:

$$\text{twiddle} = \exp(2 \cdot \pi \cdot i \cdot \text{angle} / K)$$

where K items exist in each block of data.

For each data sample, the twiddle angle is calculated as:

$$\text{angle} = X * Y$$

where X and Y depend on the position of that data sample in the input stream.

Obtain the value of X (or Y) by extracting user-specified bits from the address of the data sample, and concatenating them.

The **GeneralVTwiddle** block (and the memory-optimized **GeneralMultVTwiddle** offer variable size and include an additional **size** input. They both use a **General Twiddle Counter** rather than a **Counter** block.

**Table 117. Parameters for the GeneralTwiddle and GeneralMultTwiddle Block**

Parameter	Description
iFFT	Generate twiddle factors for an FFT or an IFFT.
sbits	The number of spatial address bits i.e. $\log_2(N)$ where there are $N$ wires.
xbits	The vector of bit positions for X.
ybits	The vector of bit positions for Y.
Input type	The type of the input before the twiddle.
Twiddle type	The type of the twiddle factor.
Output type	The type of the output after the twiddle.
Use faithful rounding	Use faithful rather than correct rounding. Only for floating-point twiddle types.

**Table 118. Parameters for the GeneralVTwiddle and GeneralMultVTwiddle Block**

Parameter	Description
iFFT	Generate twiddle factors for an FFT or an IFFT.
Number of spatial bits	The number of spatial address bits i.e. $\log_2(N)$ where there are $N$ wires.
maxsize	Maximum FFT size is $2^{\text{maxsize}}$ .
Size of the parallel section	The number of radix-2 stages assigned to the parallel section in the surrounding <b>HybridVFFT</b> (the <b>GeneralVTwiddle</b> links between the serial and parallel sections of the <b>HybridVFFT</b> )
Input type	The type of the input before the twiddle.
Twiddle type	The type of the twiddle factor.
Output type	The type of the output after the twiddle.
Use faithful rounding	Use faithful rather than correct rounding. Only for floating-point twiddle types.

**Table 119. Port Interface for the GeneralTwiddle Block**

Signal	Direction	Type	Description
v	Input	Boolean.	Input valid signal.
d	Input	Compatible with user-specified input type.	Vector of N data inputs.
drop	Input	Unsigned integer	The number of stages to drop. <b>GeneralMultiVTwiddle</b> only.
<i>continued...</i>			

Signal	Direction	Type	Description
qv	Output	Boolean.	Output valid signal.
q	Output	User-specified output type.	Vector of N data outputs.
qdrop	Output	Unsigned integer	<b>GeneralMultiVTwiddle</b> only.

### 14.3.14. Hybrid FFT (Hybrid\_FFT, HybridVFFT)

The **Hybrid\_FFT** block implements a hybrid serial or parallel implementation of a supersampled FFT (or IFFT) that processes  $2^M$  points per cycle (with  $0 < M$ ).

The hybrid implementation consists of an optional serial section (built using single-wire streaming FFTs) associated twiddle block, and a parallel section (implemented using the **PFFT\_Pipe** block).

You control the length of the serial section by a user-supplied parameter. For an FFT with  $2^N$  points that processes  $2^M$  points per cycle, this parameter must be no greater than  $N-M$ .

In general, the serial section is more space-efficient; the parallel section is more multiplier-efficient. So changing the value of this parameter provides a trade-off between DSP usage and memory usage.

The **HybridVFFT** serial section absorbs all the variability and the size of the parallel section is fixed. The variable-size hybrid FFT includes multiple variable-size streaming FFTs, a variable-size **GeneralTwiddle** and a parallel FFT.

**Table 120. Parameters for the Hybrid\_FFT and HybridVFFT Blocks**

Parameter	Description
<b>iFFT</b>	<b>true</b> to implement an IFFT, otherwise <b>false</b> .
<b>maxsize</b>	The maximum FFT size is $2^{\text{maxsize}}$ . <b>HybridVFFT</b> only.
<b>minsize</b>	The minimum FFT size is $2^{\text{minsize}}$ and is limited by the value of <b>sbits M</b> . It cannot be smaller than $2^{\text{sbits}}$ ( $2^M$ ). <b>HybridVFFT</b> only.
<b>N</b>	Log2 of the number of points in the FFT.
<b>Bit-reversed input</b>	<b>true</b> if you expect bit-reversed input, otherwise <b>false</b> .
<b>M</b>	Log2 of the number of input wires.
<b>Number of serial stages</b>	Length of the serial section (in radix-2 stages).
<b>Twiddle/pruning specification</b>	-
<b>Optimize twiddle memory usage</b>	<b>true</b> to use <b>GeneralMultTwiddle</b> (rather than <b>GeneralTwiddle</b> ) for top-level twiddle.
<b>Use faithful rounding</b>	<b>true</b> if the block uses faithful (rather than correct) rounding for floating-point operations. Fixed-point FFTs ignore this parameter.

**Table 121. Port Interface for the Hybrid\_FFT and HybridVFFT Blocks**

Signal	Direction	Type	Description
v	Input	Boolean.	Valid input signal.
d	Input	As specified.	Complex data input signal.
<i>continued...</i>			

Signal	Direction	Type	Description
qv	Output	Boolean.	Valid output signal.
q	Output	Determined by pruning specification.	Complex data output signal.
size	Input	Unsigned integer	FFT serial section size, which must be at least equal to the difference between <b>maxsize</b> and <b>minsize</b> .

Type style	Elements used
Bold	b
Italic	i
Underlined	u

### Related Information

[About Pruning and Twiddle for FFT Blocks](#) on page 284

## 14.3.15. Multiwire Transpose (MultiwireTranspose)

The DSP Builder **MultiwireTranspose** block performs a specialized reordering of a block of data and presents it on multiple wires. The size of the block and the number of wires must both be a power of 2.

Each element in the block has a logical address, which DSP Builder forms by concatenating its spatial address (wire number) with its temporal address (slot number). The spatial address is the least-significant part of the logical address; the temporal address is the most significant part. The block specifies the reordering as an arbitrary permutation of the address bits. The block numbers the address bits from 0 (least significant). The block specifies the permutation by listing the address bits in order, starting with the least significant.

For example: specifying:

[7 6 5 4 3 2 1 0] bit-reverses a block of 256 elements

[6 7 4 5 2 3 0 1] digit reverses a block (radix 4)

[0 1 2 3 4 5 6 7] leaves the order of the data unchanged

[6 7 0 1 2 3 4 5] rotates the address bits

[2 3 4 5 6 7 0 1] is be the inverse rotation.

**Table 122. Parameters for the MultiwireTranspose Block**

Parameter	Description
<b>Address permutation</b>	A vector of integers that describes how to rearrange the block of data.
N	The number of spatial address bits. The block has $2^N$ data wires.

**Table 123. Port Interface for the MultiwireTranspose Block**

Signal	Direction	Type	Description
v	Input	Boolean.	Input valid signal.
d	Input	Any type.	Data input.
qv	Output	Boolean.	Output valid signal.
q	Output	Same as d	Data output.

### 14.3.16. Parallel Pipelined FFT (PFFT\_Pipe)

The **PFFT\_Pipe** block implements a supersampled FFT (or IFFT) that processes  $2^M$  points per cycle (with  $0 < M$ ).

The **PFFT\_Pipe** block uses a pipeline of (small) fully-parallel FFTs, twiddle, and transpose blocks. This FFT uses only a small number of DSP blocks but has a relative high latency (and associated memory usage).

Not all parameters are available with all blocks.

**Table 124. Parameters for the PFFT\_Pipe Blocks**

Parameter	Description
<b>iFFT</b>	<b>true</b> to implement an IFFT otherwise <b>false</b> .
<b>N</b>	Log2 of the number of points in the FFT.
<b>Bit-reversed input</b>	<b>true</b> if you expect bit-reversed input, otherwise <b>false</b> .
<b>Number of spatial bits</b>	$M$ for $2^M$ wires.
<b>Twiddle/pruning specification</b>	-.
<b>Use faithful rounding</b>	<b>true</b> if the block uses faithful (rather than correct) rounding for floating-point operations. Fixed-point FFTs ignore this parameter.

**Table 125. Port Interface for the PFFT\_Pipe Blocks**

Signal	Direction	Type	Description
v	Input	Boolean.	Valid input signal.
d	Input	Any.	Complex data input signal.
qv	Output	Boolean.	Valid output signal.
q	Output	Determined by pruning specification.	Complex data output signal.

#### Related Information

[About Pruning and Twiddle for FFT Blocks](#) on page 284

### 14.3.17. Pulse Divider (PulseDivider)

The **PulseDivider** block generates a single-cycle one on its output for each  $2^N$  ones on its input.

**Table 126. Parameters for the PulseDivider Block**

Parameter	Description
<b>N</b>	Specifies the input block size $2^N$ .

**Table 127. Port Interface for the PulseDivider Block**

Signal	Direction	Type	Description
v	Input	Boolean or <b>uint(1)</b> .	Data valid.
g	Output	<b>uint(1)</b> .	Block containing $2^N$ elements received.

### 14.3.18. Pulse Multiplier (PulseMultiplier)

The **PulseMultiplier** block stretches a single-cycle pulse on its input into a  $2^N$ -cycle pulse on its output. The block ignores any input pulse that arrives within  $2^N$  cycles of the previous one. If the **PulseMultiplier** block receives a second 1 on its input while it is producing an existing stream, its behavior is undefined.

The **PulseMultiplier** is a special version of the **StretchPulse** block.

**Table 128. Parameters for the PulseMultiplier Block**

Parameter	Description
<b>N</b>	Specifies the output length pulse size $2^N$ .

**Table 129. Port Interface for the PulseMultiplier Block**

Signal	Direction	Type	Description
g	Input	Boolean or <b>uint(1)</b> .	Start of $2^N$ data block.
v	Output	<b>uint(1)</b> .	Data valid.

### 14.3.19. Single-Wire Transpose (Transpose)

The DSP Builder **Transpose** block performs a specialized reordering of a block of data. The size of the block must be a power of 2.

You specify the reordering as an arbitrary permutation of the address bits. The block numbers the address bits from 0 (least significant). The block specifies the permutation by listing the address bits in order, starting with the least significant.

For example, specifying:

[5 4 3 2 1 0] bit-reverses a block of 64 elements

[4 5 2 3 0 1] digit-reverse it (radix 4)

[0 1 2 3 4 5] leaves the order of the data unchanged

[4 5 0 1 2 3] interleaves four blocks of 16 elements each

[2 3 4 5 0 1] deinterleaves four blocks of 16 elements



**Table 130. Parameters for the Transpose Block**

Parameter	Description
<b>Address permutation</b>	A vector of integers that describes how to rearrange the block of data.

**Table 131. Port Interface for the Transpose Block**

Signal	Direction	Type	Description
v	Input	Boolean.	Input valid signal.
d	Input	Any type.	Data input.
qv	Output	Boolean.	Output valid signal.
q	Output	Same as d	Data output.
g	Output	Boolean,	Start of output block.

### 14.3.20. Split Scalar (SplitScalar)

The **SplitScalar** block splits its input (typically an unsigned integer) into a vector of Booleans. The least significant bit of the scalar becomes the first entry in the vector (little-endian ordering).

FFT implementations often contain various bit-twiddling operations as part of their control structure. Use the **SplitScalar** block to make these bit-twiddling operations easier to implement.

**Table 132. Parameters for the SplitScalar Block**

Parameter	Description
<b>Width</b>	Width of the scalar (in bits), which is also the width of the output vector.

**Table 133. Port Interface for the SplitScalar Block**

Signal	Direction	Type	Description
d	Input	Any scalar.	Data input.
q	Output	Boolean vector.	Data output.

### 14.3.21. Streaming FFTs (FFT2, FFT4, VFFT2, and VFFT4)

The **FFT2**, **FFT4**, **VFFT2**, and **VFFT4** blocks are low-level blocks that implement streaming FFTs.

**Table 134. Parameters for the FFT2, FFT4, VFFT2, and VFFT4 Block**

Parameter	Description
<b>iFFT</b>	<b>true</b> for an iFFT, otherwise <b>false</b> . FFT4 and VFFT4 only.
<b>Bit reversed input</b>	<b>true</b> for bit-reversed inputs. FFT4 and VFFT4 only.
<b>Stages before this</b>	The number of stages to the left of this FFT.
<b>Stages after this</b>	The number of stages to the right of this FFT.
<b>Input type</b>	The type of the input signal. For example: fixdt(1,16,15).
<b>Use faithful rounding</b>	<b>true</b> to use faithful rather than correct rounding. Fixed-point FFTs ignore this parameter.

**Table 135. Port Interface for the FFT2, FFT4, VFFT2, and VFFT4 Blocks**

Signal	Direction	Type	Description
v	Input	Boolean	Valid input signal.
d	Input	Any complex type	Complex data input signal.
drop	Input	uint(k) for some k	Total number of FFT stages to bypass.
qv	Output	Boolean	Valid output signal.
q	Output	Any complex type	Complex data output signal.
qdrop	Output	uint(k) for some k	Total number of FFT stages to bypass.

### 14.3.22. Stretch Pulse (StretchPulse)

The DSP Builder **StretchPulse** block implements a general purpose, loadable pulse stretching circuit. A single-bit single-cycle `go` input loads a counter with a `count`. The single-bit output `q` remains high for `count` consecutive cycles then stays low until receiving another `go` signal.

### 14.3.23. Twiddle Angle (TwiddleAngle)

The **Twiddleangle** block generates FFT twiddle factors when you use it between a counter and the **TwiddleRom** (or **TwiddleRomF**) blocks.

The **TwiddleAngle** block takes the output of the counter and splits it into three parts:

- The channel field (LSBs of the counter)
- The index field
- The pivot field (MSBs of the counter)

It provides  $\text{bitreverse}(\text{pivot}) * \text{index}$  at the output. The calculation is optimized to use no multipliers and only a small amount of logic.

The **TwiddleAngle** block has an additional input: `v`, which keeps the internal state of the **TwiddleAngle** block synchronized with the counter. This input should be identical to the enable input to the counter.

**Table 136. Parameters for the TwiddleAngle Block**

Parameter	Description
K	Width of the channel field.
Pivot width	Width of the pivot field.
Index width	Width of the index field.

**Table 137. Port Interface for the TwiddleAngle Block**

Signal	Direction	Type	Description
v	Input	Boolean	Input to upcounter.
c	Input	Unsigned integer	Output of upcounter.
angle	Output	Unsigned integer	Input to <b>TwiddleRom</b> .

### 14.3.24. Twiddle Generator (TwiddleGenC) Deprecated

The **TwiddleGenC** block generates the appropriate complex coefficients that multiplies the streaming data in a radix-2<sup>2</sup> streaming FFT or iFFT architecture.

Feed at the input by a modulo N counter (where N is an integer power of two) and the appropriate complex sequence generates at the output.

To parameterize this block, set the **Counter bit width** parameter with  $\log_2(N)$  and enter the width in bits and fixed-point scaling of the twiddle factors. A cosine or sine wave has a range of [-1:1], therefore you must provide at least two integer bits, and as many fractional bits as are appropriate. Starting with a twiddle bit width of 16 bits (enter 16 as the **twiddle bit width**), and a scaling of  $2^{-14}$  (enter 14 as the **Twiddle scaling exponent**). The resulting fixed-point type is **sfix16\_en14** (2.14 fixed-point format).

**Table 138. Parameters for the TwiddleGenC Block**

Parameter	Description
FFT type	Specifies whether to generate twiddle factors for an FFT or an IFFT.
Counter bit width	Specifies the counter width in bits.
Twiddle bit width	Specifies the twiddle width in bits.
Twiddle scaling exponent value	Specifies the fixed-point scaling factor of the complex twiddle factor.

**Table 139. Port Interface for the TwiddleGenC Block**

Signal	Direction	Type	Description
counter	Input	Any fixed-point type	Counter signal.
w	Output	Derived complex fixed-point type	Complex data output.

### 14.3.25. Twiddle and Variable Twiddle (Twiddle and VTwiddle)

The Twiddle and VTwiddle blocks are low-level blocks that implement streaming FFTs.

Each twiddle block joins two FFTs (the left constituent FFT and the right constituent FFT) to form a larger FFT. For variable-size FFTs, use the VTwiddle block. Each of the constituent FFTs is either a primitive FFT (e.g. an FFT4 block) or is multiple FFT and twiddle blocks.

The Twiddle block FFT may be part of an even larger FFT. In fact, the pipeline is formed by linearizing a binary tree of FFTs (leaf nodes) and twiddle blocks (internal nodes).

Each Twiddle block requires you to specify three types:

- The data signal prior to the twiddle multiplications
- The twiddle factors
- The data signal after the twiddle multiplication.

**Table 140. Parameters for the Twiddle and VTwiddle Block**

Parameter	Description
iFFT	Generate twiddle factors for an FFT or an IFFT.
Stages before this	The number of stages to the left(*) of this composite FFT.
Left width	The number of stages in the left(*) constituent FFT.
Right width	The number of stages in the right(*) constituent FFT.
Stages after this	The number of stages to the right(*) of this composite FFT.
Input type	The type to which DSP Builder should convert the input. Refers to the type of the input after you apply an explicit type conversion. It doesn't have to exactly match the actual input type to the Twiddle block.
Twiddle type	The type of the twiddle factor.
Output type	The type of the output after the twiddle.
Use faithful rounding	Use faithful rather than correct rounding. Fixed-point FFTs ignore this parameter.

*Note:* For bit-reversed FFTs, reverse left and right, so left refers to the number of stages to the right of the current block.

**Table 141. Port Interface for the Twiddle and VTwiddle Blocks**

Signal	Direction	Type	Description
<b>v</b>	Input	Boolean	Valid input signal.
<b>d</b>	Input	Any complex type	Complex data input signal.
drop	Input	uint(k) for some k	Total number of FFT stages to bypass
qv	Output	Boolean	Valid output signal.
q	Output	Any complex type	Complex data output signal.
qdrop	Output	uint(k) for some k	Total number of FFT stages to bypass

### 14.3.26. Twiddle ROM (TwiddleRom, TwiddleMultRom and TwiddleRomF (deprecated))

The DSP Builder twiddle ROM blocks generate FFT twiddle factors, converting the input angle into a cos-sin pair. These block are memory optimized for use with wide counters.

*Note:* **TwiddleRomF** is deprecated; **TwiddleRom** has a new parameters (therefore appears in the obsolete and the common directory).

The **TwiddleRom** and **TwiddleMutlRom** block construct FFTs. They map an angle (specified as an unsigned integer) to a complex number (the twiddle factor). For an FFT, the mapping is:

$$\text{twiddle} = \exp(-2 * \pi * i * \text{angle} / N)$$

For an IFFT, the mapping is:

$$\text{twiddle} = \exp(2 * \pi * i * \text{angle} / N)$$

where  $N = 2^{\text{anglewidth}}$  and *anglewidth* is the width of the *angle* input signal.

The **TwiddleRom** and **TwiddleMultRom** blocks have the same external interface but different internal implementations. **TwiddleRom** uses a single large memory; **TwiddleMultRom** uses two smaller memories and constructs the twiddle factors using complex multiplication.

**TwiddleMultRom** consumes more DSP blocks but generally uses fewer memory blocks than **TwiddleRom**. **TwiddleMultRom** also produces slightly less accurate results than **TwiddleRom**.

**Table 142. Parameters for the TwiddleRom and TwiddleMultRom Blocks**

Parameter	Description
iFFT	True to generate twiddle factors for an IFFT.
Angle bit width	The width of the angle input signal in bits.
Twiddle type	The type of the twiddle output. For example: fixdt(1,18,17).

**Table 143. Port Interface for the TwiddleROM and TwiddleMultRom Blocks**

Signal	Direction	Type	Description
angle	Input	Unsigned	Input angle.
twiddle	Output	User specified	Output twiddle factor.

## 14.4. Primitive Basic Blocks Library

Use the DSP Builder advanced blockset **Primitive Basic Blocks** library blocks to implement low-level basic functions.

1. [Absolute Value \(Abs\)](#) on page 303
2. [Accumulator \(Acc\)](#) on page 303
3. [Add](#) on page 305
4. [Add SLoad \(AddSLoad\)](#) on page 306
5. [AddSub](#) on page 307
6. [AddSubFused](#) on page 307
7. [AND Gate \(And\)](#) on page 307
8. [Bit Combine \(BitCombine\)](#) on page 308
9. [Bit Extract \(BitExtract\)](#) on page 309
10. [Bit Reverse \(BitReverse\)](#) on page 310
11. [Compare \(CmpCtrl\)](#) on page 310
12. [Complex Conjugate \(ComplexConjugate\)](#) on page 310
13. [Compare Equality \(CmpEQ\)](#) on page 311
14. [Compare Greater Than \(CmpGE\)](#) on page 312

15. [Compare Less Than \(CmpLT\)](#) on page 312
16. [Compare Not Equal \(CmpNE\)](#) on page 312
17. [Constant \(Const\)](#) on page 313
18. [Constant Multiply \(Const Mult\)](#) on page 314
19. [Convert](#) on page 314
20. [CORDIC](#) on page 315
21. [Counter](#) on page 325
22. [Count Leading Zeros, Ones, or Sign Bits \(CLZ\)](#) on page 326
23. [Dual Memory \(DualMem\)](#) on page 327
24. [Demultiplexer \(Demux\)](#) on page 329
25. [Divide](#) on page 330
26. [Fanout](#) on page 331
27. [FIFO](#) on page 332
28. [Floating-point Classifier \(FloatClass\)](#) on page 333
29. [Floating-point Multiply Accumulate \(MultAcc\)](#) on page 333
30. [ForLoop](#) on page 334
31. [Load Exponent \(LdExp\)](#) on page 335
32. [Left Shift \(LShift\)](#) on page 336
33. [Loadable Counter \(LoadableCounter\)](#) on page 336
34. [Look-Up Table \(Lut\)](#) on page 337
35. [Loop](#) on page 339
36. [Math](#) on page 340
37. [Minimum and Maximum \(MinMax\)](#) on page 341
38. [MinMaxCtrl](#) on page 342
39. [Multiply \(Mult\)](#) on page 343
40. [Multiplexer \(Mux\)](#) on page 343
41. [NAND Gate \(Nand\)](#) on page 344
42. [Negate](#) on page 345
43. [NOR Gate \(Nor\)](#) on page 345
44. [NOT Gate \(Not\)](#) on page 346
45. [OR Gate \(Or\)](#) on page 347
46. [Polynomial](#) on page 347
47. [Ready](#) on page 348
48. [Reinterpret Cast \(ReinterpretCast\)](#) on page 348
49. [Round](#) on page 349
50. [Sample Delay \(SampleDelay\)](#) on page 349
51. [Scalar Product](#) on page 350
52. [Select](#) on page 351

- 53. [Sequence](#) on page 352
- 54. [Shift](#) on page 353
- 55. [Sqrt](#) on page 354
- 56. [Subtract \(Sub\)](#) on page 354
- 57. [Sum of Elements \(SumOfElements\)](#) on page 355
- 58. [Trig](#) on page 356
- 59. [XNOR Gate \(Xnor\)](#) on page 358
- 60. [XOR Gate \(Xor\)](#) on page 358

### 14.4.1. Absolute Value (Abs)

The **Abs** block outputs the absolute value of the input:

$$q = \text{abs}(a)$$

**Table 144. Parameters for the Abs Block**

Parameter	Description
Output data type mode	<p>Determines how the block sets its output data type:</p> <ul style="list-style-type: none"> <li>• <b>Inherit via internal rule:</b> the number of integer and fractional bits is the maximum of the number of bits in the input data types.</li> <li>• <b>Inherit via internal rule with word growth:</b> the number of fractional bits is the maximum of the number of fractional bits in the input data types. The number of integer bits is the maximum of the number of integer bits in the input data types plus one. This additional word growth allows for subtracting the most negative number from 0, which exceeds the maximum positive number that the number of bits of the input can store.</li> <li>• <b>Specify via dialog:</b> you can set the output type of the block explicitly using additional fields that are available when this option is selected. This option reinterprets the output bit pattern from the LSB up according to the specified type.</li> <li>• <b>Boolean:</b> the output type is Boolean.</li> </ul>
Output data type	Specifies the output data type. For example, <b>sfix(16)</b> , <b>uint(8)</b> .
Output scaling value	Specifies the output scaling value. For example, <b>2<sup>-15</sup></b> .

**Table 145. Port Interface for the Abs Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
a	Input	Any fixed- or floating-point type	Operand	Yes	No
q	Output	Derived fixed- or floating-point type	Result	Yes	No


### 14.4.2. Accumulator (Acc)

The **Acc** block implements an application-specific floating-point accumulator.

$$r = \text{acc}(x, n)$$

The **acc** block allows accumulating data sets of variable lengths. The block indicates a new data set by setting **n** high with the first element of the accumulation.

**Figure 98. New Data Set**



This example accumulates  $x_0 + x_1 + x_2$  and  $y_0 + y_1 + y_2$ .

The **acc** block has single and double-precision floating-point data inputs and outputs.

**Table 146. Parameters for the Add Block**

Parameter	Description
<b>LSBA</b>	This parameter defines the weight of the accumulator's LSB, and therefore the accuracy of the accumulation. This value and the maximum number of terms to be accumulated sets the accuracy of the accumulator. The maximum number of terms the design can accumulate can invalidate the $\log_2(N)$ lower bits of the accumulator. For instance, if an accuracy of $2^{(-30)}$ is enough, and you add 1k of numbers, $LSBA = -30 - \log_2(1k)$ , which is approximately -40.
<b>MSBA</b>	The weight of the MSB of the accumulation result. Adding a few guard bits to the value has little impact on the implementation size. You can set this parameter in one of the following ways: <ul style="list-style-type: none"> <li>For a stock simulation, to limit the value of any stock to \$100k before the simulation is invalid, use a value of <math>\text{ceil}(\log_2(100K)) \sim \text{ceil}(16.6) = 17</math></li> <li>For a simulation where the implemented circuit adds numbers <math>\leq 1</math>, for one year, at 400MHz, use <math>\text{ceil}(\log_2(365 * 86400 * 400 * 10^6)) \sim 54</math></li> </ul>
<b>maxMSBX</b>	The maximum weight of the inputs. When adding probabilities $\leq 1$ set this weight to 0. When adding data from sensors, set bounds on the input ranges. Alternatively, set <b>MaxMSBX</b> = <b>MSBA</b> . However, the size of the architecture may increase.

**Table 147. Port Interface for the Acc Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
x	Input	Single or double	Operand	Yes	No
n	Input	Boolean	Control	Yes	No
r	Output	Single or double	Output	Yes	No
xO	Output	Boolean	This flag goes high when the input value has a weight larger than selected value for <b>MaxMSBX</b> . The result of the accumulation is then invalid.	Yes	No
xU	Output	Boolean	If this flag goes high, an input value is completely shifted out of the accumulator. This flag warns that the value of <b>LSBA</b> is possibly too large.	Yes	No
aO	Output	Boolean	This flag goes high when the accumulated value has a weight larger than <b>MSBA</b> . The result of the accumulation is then invalid.	Yes	No



### 14.4.3. Add

The **Add** block outputs the sum of the inputs:

$$q = a + b$$

For two or more inputs, the **Add** block outputs the sum of the inputs:

$$q = a + b + \dots$$

For a single vector input, the **Add** block outputs the sum of elements:

$$q = \sum a_n$$

For a single scalar input, the **Add** block outputs the input value:

$$q = a$$

**Table 148. Parameters for the Add Block**

Parameter	Description
Output data type mode	Determines how the block sets its output data type: <ul style="list-style-type: none"> <li>• <b>Inherit via internal rule:</b> the number of integer and fractional bits is the maximum of the number of bits in the input data types.</li> <li>• <b>Inherit via internal rule with word growth:</b> the number of fractional bits is the maximum of the number of fractional bits in the input data types. The number of integer bits is the maximum of the number of integer bits in the input data types plus one. This additional word growth allows for subtracting the most negative number from 0, which exceeds the maximum positive number that you can store in the number of bits of the input.</li> <li>• <b>Specify via dialog:</b> you can set the output type of the block explicitly using additional fields that are available when this option is selected. This option reinterprets the output bit pattern from the LSB up according to the specified type.</li> <li>• <b>Boolean:</b> the output type is Boolean.</li> </ul>
Output data type	Specifies the output data type. For example, <b>sfixed(16)</b> , <b>uint(8)</b> .
Number of Inputs	Specifies the number of inputs.
Fused datapath	This option affects the floating-point architectures. Turn on this option to save hardware by omitting normalization stages between adder stages. The output deviates from that expected of IEEE compliance.
Floating point rounding	Specifies what rounding to apply to the result: <ul style="list-style-type: none"> <li>• <b>Correct.</b> IEEE compliant unbiased round to nearest output value.</li> <li>• <b>Faithful.</b> Saves hardware by sometimes rounding to the second nearest value. Error is about double that of correct rounding.</li> </ul>
Output scaling value	Specifies the output scaling value. For example, <b>2<sup>-15</sup></b> .

**Table 149. Port Interface for the Add Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
a	Input	Any fixed- or floating-point type	Operand 1	Yes	Yes
b	Input	Any fixed- or floating-point type	Operand 2	Yes	Yes
q	Output	Derived fixed- or floating-point type	Result	Yes (scalar output in one input case).	Yes

**Related Information**

[Forcing Soft Floating-point Data Types with the Advanced Options](#) on page 220

**14.4.4. Add SLoad (AddSLoad)**

The **AddSLoad** block performs the following function:

$$q = s ? v : (a + b)$$

If the **s** input is low, output the sum of the first 2 inputs, **a + b**, else if **s** is high, then output the value **v**.

**Table 150. Parameters for the AddSLoad Block**

Parameter	Description
Output data type mode	Determines how the block sets its output data type: <ul style="list-style-type: none"> <li><b>Inherit via internal rule:</b> the number of integer and fractional bits is the maximum of the number of bits in the input data types.</li> <li><b>Inherit via internal rule with word growth:</b> the number of fractional bits is the maximum of the number of fractional bits in the input data types. The number of integer bits is the maximum of the number of integer bits in the input data types plus one. This additional word growth allows for subtracting the most negative number from 0, which exceeds the maximum positive number that you can store in the number of bits of the input.</li> <li><b>Specify via dialog:</b> you can set the output type of the block explicitly using additional fields that are available when this option is selected. This option reinterprets the output bit pattern from the LSB up according to the specified type.</li> <li><b>Boolean:</b> the output type is Boolean.</li> </ul>
Output data type	Specifies the output data type. For example, <b>sfix(16), uint(8)</b> .
Output scaling value	Specifies the output scaling value. For example, <b>2<sup>-15</sup></b> .

**Table 151. Port Interface for the AddSLoad Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
a	Input	Any fixed- or floating-point type	Operand 1	Yes	Yes
b	Input	Any fixed- or floating-point type	Operand 2	Yes	Yes
<i>continued...</i>					

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
s	Output	Any fixed- or floating-point type	Synchronous load	Yes	No
v	Output	Any fixed- or floating-point type	Value to load if s is true	Yes	Yes
q	Output	Derived fixed- or floating-point type	Result	Yes	Yes

### 14.4.5. AddSub

The **AddSub** block produces either the sum ( $a + b$ ) or the difference ( $a - b$ ) depending on the input you select (1 for add; 0 for subtract).

*Note:* For single-precision inputs and designs targeting any device with a floating-point DSP block, the block uses a mixture of resources including the DSP blocks in floating-point mode.

**Table 152. Port Interface for the AddSub Block**

Signal	Direction	Type	Description	Vector	Complex
a	Input	Single or double	Operand 1	Yes	Yes
b	Input	Single or double	Operand 2	Yes	Yes
add	Input	Boolean	Select input	Yes	Yes
q	Output	Single or double	Result	Yes	Yes

### 14.4.6. AddSubFused

The **AddSubFused** block produces both the sum and the difference of the IEEE floating-point signals that arrive on the input ports.

**Table 153. Port Interface for the AddSubFused Block**

Signal	Direction	Type	Description	Vector	Complex
	Input	Single or double	Operand 1	Yes	Yes
b	Input	Single or double	Operand 2	Yes	Yes
+	Output	Single or double	Result 1	Yes	Yes
-	Output	Single or double	Result 2	Yes	Yes

### 14.4.7. AND Gate (And)

The **And** block outputs the logical **AND** of the input values.

If the number of inputs is set to 1, then the logical **and** of all the individual bits of the input word is output.

**Table 154. Parameters for the And Block**

Parameter	Description
Number of inputs	Specifies the number of inputs.
Output data type mode	Determines how the block sets its output data type: <ul style="list-style-type: none"> <li>• <b>Inherit via internal rule:</b> the number of integer and fractional bits is the maximum of the number of bits in the input data types.</li> <li>• <b>Specify via dialog:</b> you can set the output type of the block explicitly using additional fields that are available when this option is selected. This option reinterprets the output bit pattern from the LSB up according to the specified type.</li> <li>• <b>Boolean:</b> the output type is Boolean.</li> </ul>
Output data type	Specifies the output data type. For example, <b>sfix(16), uint(8)</b> .
Output scaling value	Specifies the output scaling value. For example, <b>2<sup>-15</sup></b> .

**Table 155. Port Interface for the And Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
unnamed	Input	Any fixed-point type	Operands 1 to n	Yes	No
q	Output	Derived fixed-point type	Result	Yes	No

### 14.4.8. Bit Combine (BitCombine)

The **BitCombine** block outputs the bit concatenation of the input values:

$$((h \ll \text{bitwidth}(i)) | i)$$

You can change the number of inputs on the **BitCombine** block according to your requirements. When Boolean vectors are input on multiple ports, DSP Builder combines corresponding components from each vector and outputs a vector of signals. The widths of all input vectors must match. However, the widths of the signals arriving on different inputs do not have to be equal. The one input **BitCombine** block is a special case that concatenates all the components of the input vector, so that one wide scalar signal is output. Use with logical operators to apply a 1-bit reducing operator to Boolean vectors.

**Table 156. Parameters for the BitCombine Block**

Parameter	Description
Number of inputs	Specifies the number of inputs.
Output data type mode	Determines how the block sets its output data type: <ul style="list-style-type: none"> <li>• <b>Inherit via internal rule:</b> the number of integer and fractional bits is the maximum of the number of bits in the input data types.</li> <li>• <b>Specify via dialog:</b> you can set the output type of the block explicitly using additional fields that are available when this option is selected. This option reinterprets the output bit pattern from the LSB up according to the specified type.</li> <li>• <b>Boolean:</b> the output type is Boolean.</li> </ul>
Output data type	Specifies the output data type. For example, <b>sfix(16), uint(8)</b> .
Output scaling value	Specifies the output scaling value. For example, <b>2<sup>-15</sup></b> .

**Table 157. Port Interface for the BitCombine Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
i	Input	Any fixed-point type	Operand	Yes	No
h	Input	Any fixed-point type	Operand	Yes	No
q	Output	Derived fixed-point type	Result	Yes	No

### 14.4.9. Bit Extract (BitExtract)

The **BitExtract** block outputs the bits extracted from the input, and recast as the specified data type:

$$q = (a \gg \text{LSB})$$

If bit position is a negative number, the bit position is an offset from the MSB instead of LSB.

If the **BitExtract** block initialization parameter is a vector of LSB positions, the output is a vector of matching width, even if the input is a scalar signal. Use this feature to split a wide data line into a vector of Boolean signals. The components are in the same order as you specify in the initialization parameter. If the input to the **BitCombine** block is a vector, the width of any vector initialization parameter must match, and then a different bit can be selected from each component in the vector. The output data type does not always have to be Boolean signals. For example, setting to **uint8** provides a simple way to split one wide signal into a vector of unsigned 8-bit data lines.

**Table 158. Parameters for the BitExtract Block**

Parameter	Description
Output data type mode	Determines how the block sets its output data type: <ul style="list-style-type: none"> <li>• <b>Inherit via internal rule:</b> the number of integer and fractional bits is the maximum of the number of bits in the input data types.</li> <li>• <b>Specify via dialog:</b> you can set the output type of the block explicitly using additional fields that are available when this option is selected. This option reinterprets the output bit pattern from the LSB up according to the specified type.</li> <li>• <b>Boolean:</b> the output type is Boolean.</li> </ul>
Output data type	Specifies the output data type. For example, <b>sfix(16)</b> , <b>uint(8)</b> .
Output scaling value	Specifies the output scaling value. For example, <b>2<sup>-15</sup></b> .
Least Significant Bit Position from Input Word	Specifies the bit position from the input word as the LSB in the output word.

**Table 159. Port Interface for the BitExtract Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
a	Input	Any fixed-point type	Operand	Yes	No
q	Output	Derived fixed-point type	Result	Yes	No

### 14.4.10. Bit Reverse (BitReverse)

The **BitReverse** primitive block reverses the bits at the input. The MSB is output as the LSB.

**Table 160. Port Interface for the BitReverse Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
a	Input	Any fixed-point type	Operand	Yes	No
q	Output	Derived fixed-point type	Result	Yes	No

### 14.4.11. Compare (CmpCtrl)

The **CmpCtrl** block produces the Boolean result of comparing two IEEE floating-point input signals. A select line controls the comparison. The select line is at least three-bits wide to select from five different comparison operators.

**Table 161. Comparison Operators**

s	Comparison Operator
0	Less than.
1	Less than or equal to.
2	Equal.
3	Greater than or equal to.
4	Greater than.
5	Not equal.

**Table 162. Port Interface for the CmpCtrl Block**

Signal	Direction	Type	Description	Vector	Complex
	Input	Single or double	Operand 1	Yes	No
b	Input	Single or double	Operand 2	Yes	No
s	Input	Fixed-point (unsigned)	Select input	Yes	No
q	Output	Boolean	Result	Yes	No

### 14.4.12. Complex Conjugate (ComplexConjugate)

The **ComplexConjugate** block outputs the complex conjugate of its input value.

If  $a = x + yi$ ,

then  $q = x - yi$

If the input value is real, an unchanged real value is output

$q = a$

**Table 163. Parameters for the ComplexConjugate Block**

Parameter	Description
Output data type mode	Determines how the block sets its output data type: <ul style="list-style-type: none"> <li>• <b>Inherit via internal rule:</b> the number of integer and fractional bits is the maximum of the number of bits in the input data types.</li> <li>• <b>Inherit via internal rule with word growth:</b> the number of fractional bits is the maximum of the number of fractional bits in the input data types. The number of integer bits is the maximum of the number of integer bits in the input data types plus one. This additional word growth allows for subtracting the most negative number from 0, which exceeds the maximum positive number that you can store in the number of bits of the input.</li> <li>• <b>Specify via dialog:</b> you can set the output type of the block explicitly using additional fields that are available when this option is selected. This option reinterprets the output bit pattern from the LSB up according to the specified type.</li> <li>• <b>Boolean:</b> the output type is Boolean.</li> </ul>
Output data type	Specifies the output data type. For example, <b>sfix(16)</b> , <b>uint(8)</b> .
Output scaling value	Specifies the output scaling value. For example, <b>2<sup>-15</sup></b> .
Number of Inputs	Specifies the number of inputs.

**Table 164. Port Interface for the ComplexConjugate Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
a	Input	Any fixed- or floating-point type	Operand	Yes	Yes
q	Output	Derived fixed- or floating-point type	Result	Yes	Yes

#### 14.4.13. Compare Equality (CmpEQ)

The **CmpEQ** block outputs true if and only if the two inputs have the same value:

$a == b$

**Table 165. Port Interface for the CmpEQ Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
a	Input	Any fixed- or floating-point type	Operand 1	Yes	Yes
b	Input	Any fixed- or floating-point type	Operand 2	Yes	Yes
q	Output	Boolean	Result	Yes	Yes

### 14.4.14. Compare Greater Than (CmpGE)

The **CmpGE** block outputs true if and only if the first input is greater than or equal to the second input:

$$a \geq b$$

**Table 166. Port Interface for the CmpGE Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
a	Input	Any fixed- or floating-point type	Operand 1	Yes	No
b	Input	Any fixed- or floating-point type	Operand 2	Yes	No
q	Output	Boolean	Result	Yes	No

### 14.4.15. Compare Less Than (CmpLT)

The **CmpLT** block outputs true if and only if the first input is less than the second input:

$$a < b$$

**Table 167. Port Interface for the CmpLT Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
a	Input	Any fixed- or floating-point type	Operand 1	Yes	No
b	Input	Any fixed- or floating-point type	Operand 2	Yes	No
q	Output	Boolean	Result	Yes	No

### 14.4.16. Compare Not Equal (CmpNE)

The **CmpNE** block outputs true if the two inputs do not have the same value:

$$a \neq b$$

**Table 168. Port Interface for the CmpNE Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
a	Input	Any fixed- or floating-point type	Operand 1	Yes	Yes
b	Input	Any fixed- or floating-point type	Operand 2	Yes	Yes
q	Output	Boolean	Result	Yes	Yes



### 14.4.17. Constant (Const)

The **Const** block outputs a specified constant value.

**Table 169. Parameters for the Const Block**

Parameter	Description
Output data type mode	Determines how the block sets its output data type: <ul style="list-style-type: none"> <li>• <b>Inherit via internal rule:</b> the number of integer and fractional bits is the maximum of the number of bits in the input data types.</li> <li>• <b>Inherit via back projection:</b> a downstream block that this block drives determines the output data type. If the driven block does not propagate a data type to the driver, you must use a Simulink SameDT block to copy the required data type to the output wire.</li> <li>• <b>Specify via dialog:</b> you can set the output type of the block explicitly using additional fields that are available when this option is selected. This option reinterprets the output bit pattern from the LSB up according to the specified type.</li> <li>• <b>Boolean:</b> the output type is Boolean.</li> <li>• <b>Single:</b> single-precision floating-point data.</li> <li>• <b>Double:</b> double-precision floating-point data.</li> <li>• <b>Variable precision floating point:</b> variable precision floating-point output type</li> </ul>
Output data type	Specifies the output data type. For example, <b>sfix(16), uint(8)</b> .
Output scaling value	Specifies the output scaling value. For example, <b>2^-15</b> .
Value	Specifies the constant value. This parameter may also be a fi object when specifying data of arbitrarily high precision.
Floating point precision	Specifies the floating-point precision. For example, <b>float32_m23</b> .
Warn when value is saturated	Turn off if the constant if you design the constant to be saturated. The <b>Constant</b> block generates a warning in the Simulink Diagnostic Viewer if the bit-width is not sufficient to represent the value. For example: Warning: Constant block 'constant_saturation_UUT/Const1' has saturated due to insufficient bit-width. SUGGESTION: Increase the bit-width or disable this warning in the block parameters.

Every **Primitive** library block accepts double-precision floating-point values when specifying mask parameters. This format limits precision to no more than 53 bits, which is more than sufficient for most of the blocks. For higher precision, the **Const**, **DualMem**, or **LUT** blocks optionally accept values using Simulink's Fixed Point data type. For example:

```
constValue = fi(0.142, 1, 16, 15)
```

```
vectorValue = fi(sin([0:10]'), 1, 18, 15)
```

To configure a **Const**, **DualMem**, or **LUT** with data of precision higher than IEEE double precision, create a MATLAB fi object of the required precision that contains the high precision data. Avoid truncation when creating this object. Use the fi object to specify the **Value** of the **Const**, the **Initial Contents** of the **DualMem** block, or the **Output value map** of the **LUT** block.

**Table 170. Port Interface for the Const Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
q	Output	Any fixed- or floating-point type	Constant value.	Yes	Yes

### 14.4.18. Constant Multiply (Const Mult)

The **Const Mult** block scales the input by a user configurable coefficient and outputs the result.

The **Value** parameter is a floating-point scaling factor that is multiplied by the input signal. If this parameter is a vector, the output is a vector. If both the input and the **Value** parameter are vectors, they must have the same length. If the **Value** parameter is complex, the block performs a complex multiply and the output is complex. When used in fixed-point designs, the Simulink simulation of the **Const Mult** block is only an approximation of the generated hardware. Turn on **Bit-accurate simulation** to simulate the correct behavior.

**Table 171. Port Interface for the Const Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
x	Input	Any fixed- or floating-point type	Operand 1	Yes	Yes
q	Output	Any fixed- or floating-point type	Result	Yes	Yes

### 14.4.19. Convert

The **Convert** block performs a type conversion of the input, and outputs the new data type.

You can optionally perform truncation, biased, or unbiased rounding if the output data type is smaller than the input. The LSB must be a value in the width in bits of the input type.

**Table 172. Parameters for the Convert Block**

Parameter	Description
Output data type mode	Determines how the block sets its output data type: <ul style="list-style-type: none"> <li><b>Inherit via internal rule:</b> the number of integer and fractional bits is the maximum of the number of bits in the input data types.</li> <li><b>Inherit via back projection:</b> a downstream block that this block drives determines the output data type. If the driven block does not propagate a data type to the driver, you must use a Simulink SameDT block to copy the required data type to the output wire.</li> <li><b>Specify via dialog:</b> you can set the output type of the block explicitly using additional fields that are available when this option is selected.</li> <li><b>Boolean:</b> the output type is Boolean.</li> <li><b>Single:</b> single-precision floating-point data.</li> <li><b>Double:</b> double-precision floating-point data.</li> <li><b>Variable precision floating point:</b> variable precision floating-point output type</li> </ul>
Output data type	Specifies the output data type. For example, <b>sfix(16)</b> , <b>uint(8)</b> .
Output scaling value	Specifies the output scaling value. For example, <b>2<sup>-15</sup></b> .
<i>continued...</i>	

Parameter	Description
Rounding method	Determines the rounding mode: <ul style="list-style-type: none"> <li>• <b>Truncate</b>: Discard any bits that fall below the new least significant bit.</li> <li>• <b>Biased</b>: Add 0.5 LSB and then truncate. This rounds towards infinity.</li> <li>• <b>Unbiased</b>: If the discarded bits equal 0.5 LSB of the new value then round towards the even integer, otherwise perform add 0.5 LSB and then truncate. This prevents the rounding operation introducing a DC bias where 0.5 always rounds towards positive infinity.</li> </ul>
Saturation	The <b>Convert</b> block allows saturation, which has an optional clip detect output that outputs 1 if any clipping has occurred. Saturation choices are <b>none</b> , <b>symmetric</b> , or <b>asymmetric</b> .
Floating point precision	Specifies the floating-point precision. For example, <b>float32_m23</b> .

For example, for an **Add** or **Mult** block, you can select the output word-length and fractional part using dialog.

Specifying the output type is a casting operation, which does not preserve the numerical value, only the underlying bits. This method never adds hardware to a block — just changes the interpretation of the output bits.

For example, for a multiplier with both input data-types, **sfix16\_En15** has output type **sfix32\_En30**. If you select output format **sfix32\_En28**, the output numerical value multiplies by four. For example, **1\*1** input gives an output value of 4.

If the you select output format **sfix32\_En31**, the output numerical value is divided by two. For example **1\*1** input gives an output value of 0.5.

If you want to change data-type format in a way that preserves the numerical value, use a **convert** block, which adds the corresponding hardware. Adding a **convert** block directly after a **primitive** block lets you specify the data-type to preserve the numerical value.

For example, a **Mult** block followed by a **Convert** block, with input values **1\*1** always give output value 1.

**Table 173. Port Interface for the Convert Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
a	Input	Any fixed-point type	Data	Yes	Yes
q	Output	Specified fixed-point type	Data	Yes	Yes

### 14.4.20. CORDIC

The **CORDIC** block performs a coordinate rotation using the coordinate rotation digital computer algorithm.

The CORDIC algorithm is a simple and efficient algorithm that calculates:

- Vector rotations
- Trigonometric functions such as sine and cosine
- Inverse trigonometric functions such as the arctangent
- Vector magnitudes

The **CORDIC** block can calculate all these functions. However, it does not calculate hyperbolic functions, square roots, logarithms and exponentials.

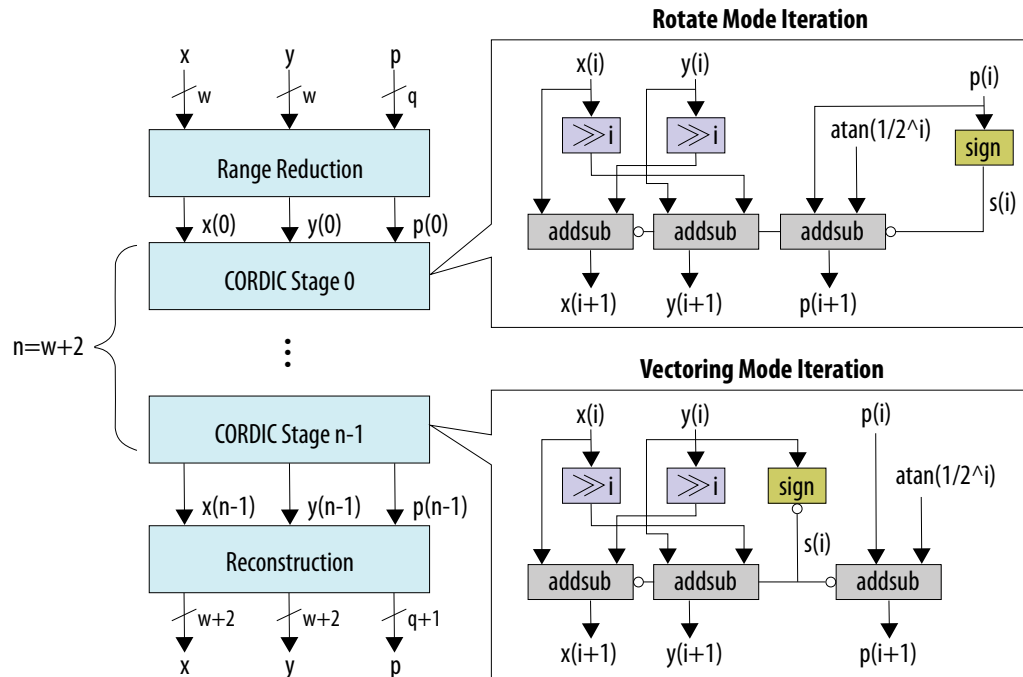
Generally, the **CORDIC** block is faster than other approaches in the absence of hardware multipliers or when multipliers become a scarce resource in your design. For low precision designs, you can also consider implementations based on direct table lookups.

The **CORDIC** block takes four inputs  $x$ ,  $y$ ,  $p$  and  $v$ . The  $x$  and  $y$  inputs provide the cartesian coordinates  $(x, y)$  of an input vector. The  $p$  input provides an input angle. Input  $v$  allows run-time switching between the two modes of the **CORDIC** block. The block has three outputs denoted by the same  $x$ ,  $y$ , and  $p$ .

The CORDIC algorithm converges to a result by performing a fixed number of iterations, where each iteration corresponds to a microrotation. If  $w$  represents the number of bits of the  $x$  and  $y$  inputs, and the block uses default type propagation, the number of iterations it performs is equal to  $w + 2$ . If the output type is **Specify via dialog** (refer to [Table 174](#) on page 318), the number of iterations equals the specified width. The **CORDIC** block implements the iterations in an unrolled fashion, allowing the block to provide a new valid output at every clock cycle.

**Figure 99. Unrolled CORDIC Architecture**

The figure shows the CORDIC architecture on the left and the two modes on the right



You can configure the **CORDIC** block at runtime to change between rotation and vector modes. Rotation mode allows rotating a vector defined by its cartesian coordinates by an angle provided in radians. By providing specific input vector coordinates, rotation mode also allows computing the sine and cosine of the input angle. The vector mode inputs a vector and rotates the vector until its  $y$  component is zero, and its  $x$  component is positive. The mode allows computing the arctangent of the input angle in radians and the magnitude of the input vector.

The generic CORDIC iterations are:

### Equation 2. Generic Cordic Iterations

$$\begin{cases} x[i+1] = x[i] - s[i]2^{-i}y[i] \\ y[i+1] = y[i] + s[i]2^{-i}x[i] \\ p[i+1] = p[i] - s[i]atan(2^{-i}) \end{cases}$$

The computation of  $s[i]$  depends on the mode:

Rotation Mode	$s[i] = +1$ if $p[i] \geq 0$ , otherwise $-1$
Vector Mode	$s[i] = +1$ if $y[i] < 0$ , otherwise $-1$

In rotation mode, active when input  $v$  is 0, the **CORDIC** rotates the input vector by a specified angle. The input vector is defined by its cartesian coordinates provided on the  $x$  and  $y$  inputs. You provide the angle by which this vector is rotated on input  $p$ . The angle must have a value in the  $[-\pi, \pi]$  interval. You read the output vector coordinates from the  $x$  and  $y$  outputs of the block. The block does not use output  $p$  in this mode. The output equations for this mode are:

### Equation 3. Rotation Mode Output Equations

$$\begin{cases} x_{out} = K(n)(x_{in}\cos(p_{in}) - y_{in}\sin(p_{in})) \\ y_{out} = K(n)(x_{in}\sin(p_{in}) + y_{in}\cos(p_{in})) \end{cases}$$

The equation also shows how you can obtain the sine and cosine of the input angle  $p$ .

Setting  $x_{in} = \frac{1}{K(n)}$  and  $y_{in} = 0$  gives:

### Equation 4. Sine and Cosine Equations using the Rotation Mode of the CORDIC

$$\begin{cases} x_{out} = \cos(p_{in}) \\ y_{out} = \sin(p_{in}) \end{cases}$$

In vector mode, active when the input  $v$  is 1, the **CORDIC** block rotates the input vector to the  $x$ -axis and records the angle required to make that rotation. The same as for rotation mode, you provide the input vector coordinates on the  $x$  and  $y$  inputs. Initialize the input  $p$  with the value 0. The output  $p$  contains the angle that the input vector  $v$  makes with the positive  $x$ -axis. This angle is in the interval  $[-\pi, \pi]$ . The first output  $x$  saves the magnitude of the input vector  $v$ .

### Equation 5. Magnitude of Input Vector

$$\begin{cases} x_{out} = K(n)\text{sqrt}(x_{in}^2 + y_{in}^2) \\ y_{out} = 0 \\ p_{out} = \text{atan2}(y_{in}, x_{in}) \end{cases}$$

The output equations of both the rotation and vector modes show that  $K(n)$  multiplies either the coordinates of the resulting output vector, or the magnitude of the vector.  $K(n)$  is a CORDIC-specific scale factor (also referred to as gain) that amplifies the

outputs. You can statically compute the gain for a known number of iterations of the algorithm. You can compensate for the gain effects, external to the **CORDIC** block, by using a constant multiplication by  $1/K(n)$ .

You can calculate the gain  $K(n)$  incrementally as the product of individual gains per step. The step gain denoted by  $K_i$  is :

**Equation 6. Step Gain**

$$K_i = \text{sqrt}(1 + 2^{-2i})$$

where  $i$  represents the iteration index, starting from 0.

The total gain is the product of the step gains  $K_i$ . Therefore, for  $n$  iterations of the CORDIC algorithm the gain is equal to:

$$K(n) = \prod_{i=0}^{n-1} K_i = \prod_{i=0}^{n-1} \text{sqrt}(1 + 2^{-2i})$$

The CORDIC gain is not automatically compensated in the **CORDIC** block. Using  $w$  as the width in bits of the  $x$  and  $y$  inputs (their widths and formats needs to match), the width of the  $x$  and  $y$  outputs is by default set to  $w+2$ , where the 2 additional bits are added in the MSB position. The additional bits allow accounting for the gain factor without overflow.

The  $p$  input is the angular value and has a range between  $-\pi$  and  $+\pi$ , which requires three integer bits to fully represent the range. The CORDIC block requires that the format of the signal feeding into the  $p$  input has exactly three integer bits, and that it is signed. Denoting by  $q$  the width in bits of the input  $p$ , the width of the output  $p$  is  $q + 1$ . The additional bit is added in the MSB position. It allows for subtracting an angle provided on the input  $p$  from the computed arctangent of the input vector (when the block is in vectoring mode) without overflowing.

**Table 174. Parameters for the CORDIC Block**

Parameter	Description
Output data type mode	Determines how the block sets its output data type:

*continued...*

Parameter	Description
	<ul style="list-style-type: none"> <li>• <b>Inherit via internal rule:</b> the number of integer and fractional bits of outputs x and y is two bits wider than the inputs x and y to accommodate for the gain. Output type for port p is one bit wider than the input port p. The outputs x, y share the same number of fractional bits as the inputs x, y and the output p is one bit wider, but has the same alignment as input p.</li> <li>• <b>Specify via dialog:</b> you can set the output type of the block explicitly using additional fields that are available when this option is selected. When you select <b>Specify via dialog</b>, the number of iterations (stages) performed by the CORDIC block equals the specified output width (including the potential sign). The following scenarios or a combination of scenarios can occur: <ul style="list-style-type: none"> <li>– LSB truncation. For example, if the internal rule output is fixdt(1, 19, 16) and you specify the output type fixdt(1, 17, 14), the 2 LSBs that are computed by the algorithm are discarded.</li> <li>– LSB extension: Example: if the internal rule output is fixdt(1, 19, 16) and you specify the output type fixdt(1, 20, 17), the LSB of the output is set to 0.</li> <li>– MSB truncation. Intel does not recommend this option, but you can also specify an output type which is narrower than the types produced by the default type propagation. In such a case a no-saturation bit selection is performed. Example: if the internal rule output is fixdt(1, 19, 16) and you specify the output type fixdt(1, 16, 16), then the 3 MSBs computed by the algorithm are dropped.</li> <li>– MSB extension. For example, if the internal rule output is fixdt(1, 19, 16) and you specify the output type fixdt(1, 20, 16), the MSB is populated via sign extension. Also, the sign extension occurs if your output type is set to fixdt(0, 20, 16). However, the type propagated on the signal on block exit is fixdt(0, 20, 16).</li> </ul> </li> <li>• <b>Boolean:</b> unsupported – the block behavior when this mode is set is unspecified.</li> </ul>
Output data type	Specifies the output data type. For example, <b>sfix(16)</b> , <b>uint(8)</b> .
Output scaling value	Specifies the output scaling value. For example, <b>2<sup>-15</sup></b> .

**Table 175. Port Interface for the CORDIC Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
x	Input	Signed fixed-point type	x coordinate of the input vector.	Yes	Yes
y	Input	Signed fixed-point type	y coordinate of the input vector.	Yes	Yes
p	Input	Signed type with three integer bits.	Required angle of rotation in the range between $-\pi$ and $+\pi$ .	Yes	Yes
v	Input	Unsigned 1-bit wide integer or Boolean.	Selects the mode of operation 0 = rotation mode, 1 = vector mode..	Yes	Yes
x	Output	When the output data type mode is set to <b>inherit via internal rule</b> , the type of the output port x is the same as the input port x, with the same alignment but two additional bits to cope with the gain.	x coordinate of the output vector.	Yes	Yes
y	Output	When the output data type mode is set to <b>inherit via internal rule</b> , the type of the output port x is the same as the input port x,	y coordinate of the output vector.	Yes	Yes

*continued...*

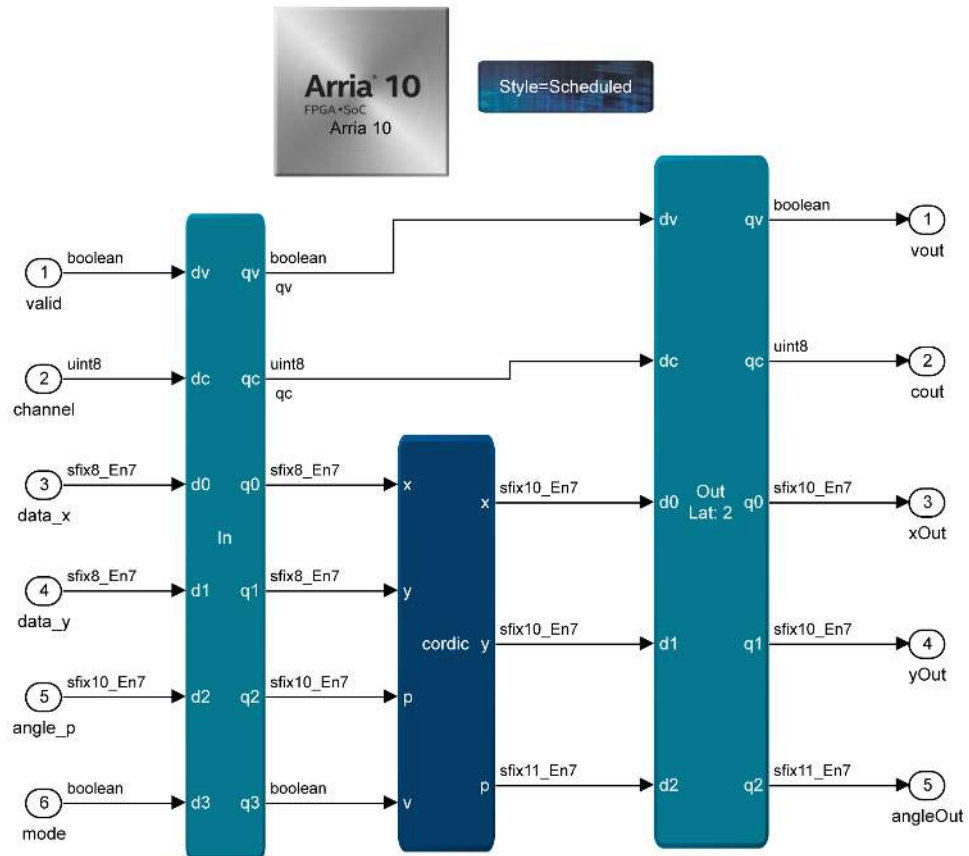
Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
		with the same alignment but two additional bits to cope with the gain.			
p	Output	Same alignment as the input port p, with one additional MSB.	In vector mode (v=1), it returns the angle of input vector defined by its cartesian coordinates x and y. The returned angle is in the interval [-pi, pi].	Yes	Yes

### Numeric Example

This complete numeric example for both modes of the **CORDIC** block uses this configuration:

- Input data types for the x and y inputs are `sfix8_en7` (corresponding to a w = 8-bit wide signed format having 7 fractional bits)
- The input type for input p is `sfix10_en7` (corresponding to a q=10-bit wide signed format having 7 fractional bits).

Figure 100. Input and Output Configuration for the Numeric Example





Whenever the example uses binary representation of values, it suffixes the values by  $_2$  to distinguish them from base 10 values.

### Rotation Mode Example

To depict the functioning of the block in rotation mode, set the inputs of the block to:

$$x = 0.75 = 0.1100000_2$$

$$y = 0.75 = 0.1100000_2$$

$$p = \frac{\pi}{6} = 000.1000011_2$$

$$v = 0$$

Table 176 on page 321 shows the operation of the **CORDIC** block in rotation mode, in which the  $(x,y)$  vector is rotated by the angle  $p$ . The first column shows the iteration index. The **CORDIC** block inputs  $w=8$ -bit signals and performs 10 iterations internally. At each iteration, the signals  $x[i]$ ,  $y[i]$  and  $p[i]$  are updated based on the values  $x[i-1]$ ,  $y[i-1]$  and  $p[i-1]$  computed by the previous iteration.

For each iteration, the first row shows the equations for computing the values, for instance:

$$y[4] = y[3] + s[3] * 2^{(-3)} x[3]$$

Whereas the second row shows the actual value computed, in both binary and decimal representation (001.1010010 $_2$  (1.640625)). The binary representation uses the 2's complement notation. The rightmost column shows the formula to compute the corresponding step-gain and the gain-step value, in decimal.

**Table 176. CORDIC Block Operating in Rotation Mode**

Iter.	x	y	p	s	$k_i$
init.	$x[0] = 0.1100000_2$ (0.75)	$y[0] = 0.1100000_2$ [0.75]	$p[0] = 0.1000011_2$ [0.5234375]	$s[0] = +1$	
1	$x[1] = x[0] - s[0]$ $(2^{-0})y[0]$	$y[1] = y[0] + s[0]$ $(2^{-0})x[0]$	$p[1] = p[0] -$ $s[0] \text{atan}(2^{-0})$		$k_0 = \sqrt{1 + 2^{-2 \times 0}}$
	000.0000000 $_2$ (0.0)	001.1000000 $_2$ ; (1.5)	111.1011110 $_2$ ; (-0.265625)	$s[1] = -1$	$k_0 \approx 1.4142$
2	$x[2] = x[1] -$ $s[1]2^{(-1)}y[1]$	$y[2] = y[1] +$ $s[1]2^{(-1)}x[1]$	$p[2] = p[1] -$ $s[1] \text{atan}(2^{-1})$		$k_1 = \sqrt{1 + 2^{-2 \times 1}}$
	000.1100000 $_2$ (0.75)	001.1000000 $_2$ (1.5)	000.0011001 $_2$ ; (0.1953125)	$s[2] = +1$	$k_1 \approx 1.11803$
3	$x[3] = x[2] -$ $s[2]2^{(-2)}y[2]$	$y[3] = y[2] +$ $s[2]2^{(-2)}x[2]$	$p[3] = p[2] -$ $s[2] \text{atan}(2^{-2})$		$k_2 = \sqrt{1 + 2^{-2 \times 2}}$
	000.0110000 $_2$ (0.375)	001.1011000 $_2$ (1.6875)	111.1111010 $_2$ ; (-4.6875e-2)	$s[3] = -1$	$k_2 \approx 1.030776$
4	$x[4] = x[3] -$ $s[3]2^{(-3)}y[3]$	$y[4] = y[3] +$ $s[3]2^{(-3)}x[3]$	$p[4] = p[3] -$ $s[3] \text{atan}(2^{-3})$		$k_3 = \sqrt{1 + 2^{-2 \times 3}}$
	000.1001011 $_2$ (0.5859375)	001.1010010 $_2$ (1.640625)	000.0001010 $_2$ (7.8125e-2)	$s[4] = +1$	$k_3 \approx 1.0077822$
5	$x[5] = x[4] -$ $s[4]2^{(-4)}y[4]$	$y[5] = y[4] +$ $s[4]2^{(-4)}x[4]$	$p[5] = p[4] -$ $s[4] \text{atan}(2^{-4})$		$k_4 = \sqrt{1 + 2^{-2 \times 4}}$

*continued...*

Iter.	x	y	p	s	$k_i$
	000.0111110_2 (0.484375)	001.1010110_2 (1.671875)	000.0000010_2 (1.5625e-2)	s[5]=+1	$k_4 \approx 1.001951$
6	$x[6] = x[5] - s[5]2^{(-5)}y[5]$	$y[6] = y[5] + s[5]2^{(-5)}x[5]$	$p[6] = p[5] - s[5]atan(2^{(-5)})$		$k_5 = \text{sqrt}(1 + 2^{-2 \times 5})$
	000.0111000_2 (0.4375)	001.1010111_2 (1.6796875)	111.1111110_2 (-1.5625e-2)	s[6]=-1	$k_5 \approx 1.000488$
7	$x[7] = x[6] - s[6]2^{(-6)}y[6]$	$y[7] = y[6] + s[6]2^{(-6)}x[6]$	$p[7] = p[6] - s[6]atan(2^{(-6)})$		$k_6 = \text{sqrt}(1 + 2^{-2 \times 6})$
	000.0111011_2 (0.4609375)	1.1010111_2 (1.6796875)	000.0000000_2; (0.0)	s[7]=+1	$k_6 \approx 1.000122$
8	$x[8] = x[7] - s[7]2^{(-7)}y[7]$	$y[8] = y[7] + s[7]2^{(-7)}x[7]$	$p[8] = p[7] - s[7]atan(2^{(-7)})$		$k_7 = \text{sqrt}(1 + 2^{-2 \times 7})$
	000.0111010_2 (0.453125)	1.1010111_2 (1.6796875)	111.1111111_2; (-7.8125e-3)	s[8]=-1	$k_7 \approx 1.000030$
9	$x[9] = x[8] - s[8]2^{(-8)}y[8]$	$y[9] = y[8] + s[8]2^{(-8)}x[8]$	$p[9] = p[8] - s[8]atan(2^{(-8)})$		$k_8 = \text{sqrt}(1 + 2^{-2 \times 8})$
	000.0111010_2; (0.453125)	1.1010111_2 (1.6796875)	111.1111111_2; (-7.8125e-3)	s[9]=-1	$k_8 \approx 1.000007$
10	$x[10] = x[9] - s[9]2^{(-9)}y[9]$	$y[10] = y[9] + s[9]2^{(-9)}x[9]$	$p[10] = p[9] - s[9]atan(2^{(-9)})$		$k_9 = \text{sqrt}(1 + 2^{-2 \times 9})$
	000.0111010_2; (0.453125)	1.1010111_2 (1.6796875)	111.1111111_2; (-7.8125e-3)		$k_9 \approx 1.000001$

The CORDIC gain after 10 iterations is  $K(n)=K(w+2) = K(10)$

**Equation 7. CORDIC Gain After 10 Iterations**

$$K(10) = K_0K_1K_2K_3K_4K_5K_6K_7K_8K_9 = \prod_{j=0}^9 K_j = \prod_{j=0}^9 \text{sqrt}(1 + 2^{-2j})$$

**Accuracy of the Rotate Mode Example**

The example achieves compensation for the gain by multiplying the output vector coordinate values by  $\frac{1}{K(10)}$ . Post compensation and after rounding to the 7-bit fraction, the output vector coordinates are:

$$x = 000.0100011_2 ; (0.2734375)$$

$$y = 001.0000011_2 ; (1.0234375)$$


The absolute error for each coordinate is:

$$|x_{\text{gold}} - x_{\text{compute}}| \approx 1.010001_2 \times 2^{-10}$$

$$|y_{\text{gold}} - y_{\text{compute}}| \approx 1.010011_2 \times 2^{-10}$$

The unit in the last place (ULP) of the output format is  $2^{-7}$ , which makes the  $2^{-10}$  output error particularly small for this set of inputs.

Figure 101. Visualization of the Execution of the CORDIC Algorithm in Rotate Mode



1. The input vector  $(x(0), y(0))$  is rotated to  $(x(1), y(1))$  by  $\text{atan}(1)=\pi/4$ .
2. The vector is rotated back by  $\text{atan}(1/2)$  to  $(x(2), y(2))$ .
3. This vector is rotated forward by  $\text{atan}(1/4)$  to  $(x(3), y(3))$ .
4. This vector is rotated back by  $\text{atan}(1/8)$  to  $(x(4), y(4))$ .

This process continues for ten iterations (for clarity the figure does not show all intermediary vectors). The figure shows in blue the resulting vector after ten iterations,  $(x(10), y(10))$ . It comes as a result of successive microrotations by angles corresponding to  $\text{atan}(2^{-i})$ . The sign of the values in the "s" column in [Table 176](#) on page 321 gives the direction of each microrotation.

The figure also shows how the magnitude of the input vector increases with the iterations. The largest increase happens in the first three iterations (refer to the step gains  $K_0, K_1$  and  $K_2$  in [Table 176](#) on page 321).

**Equation 8. Step Gains**

$$K_0 K_1 K_2 = 1.62977$$

**Vector Mode Example**

For this example, the inputs are:

- $x = 0.75 = 0.1100000_2;$
- $y = 0.75 = 0.1100000_2;$
- $p = 0.00 = 000.0000000_2;$
- $v = 1$

**Table 177. Execution of the Ten Iterations**

Iter.	x	y	p	s	$k_i$
Init	$x[0]=0.1100000_2$ (0.75)	$0.1100000_2$ (0.75)	$p[0] = 000.0000000_2$ (0.0)	$s[0]=-1$	
1	$x[1] = x[0] - s[0]$ $(2^{-0})y[0]$	$y[1] = y[0] + s[0]$ $(2^{-0})x[0]$	$p[0] - s[0]atan(2^{-0})$		$k_0 = \text{sqrt}(1 + 2^{-2 \times 0})$
	$001.1000000_2;$ (1.5)	$000.0000000_2;$ (0.0)	$000.1100101_2$ (+0.7890625)	$s[1]=-1$	$k_0 \approx 1.4142$
2	$x[2] = x[1] - s[1]$ $(2^{-1})y[1]$	$y[2] = y[1] + s[1]$ $(2^{-1})x[1]$	$p[1] - s[1]atan(2^{-1})$		$k_1 = \text{sqrt}(1 + 2^{-2 \times 1})$
	$001.1000000_2;$ (1.5)	$111.0100000_2;$ (-0.75)	$001.0100000_2$ (1.25)	$s[2]=+1$	$k_1 \approx 1.11803$
3	$x[3] = x[2] - s[2]$ $(2^{-2})y[2]$	$y[3] = y[2] + s[2]$ $(2^{-2})x[2]$	$p[3] = p[2] -$ $s[2]atan(2^{-2})$		$k_2 = \text{sqrt}(1 + 2^{-2 \times 2})$
	$001.1011000_2;$ (1.6875)	$111.1010000_2;$ (-0.375)	$001.0000001_2;$ (1.0078125)	$s[3]=+1$	$k_2 \approx 1.030776$
4	$x[4] = x[3] - s[3]$ $(2^{-3})y[3]$	$y[4] = y[3] + s[3]$ $(2^{-3})x[3]$	$p[4] = p[3] -$ $s[3]atan(2^{-3})$		$k_3 = \text{sqrt}(1 + 2^{-2 \times 3})$
	$001.1011110_2$ (1.734375)	$111.1101011_2$ (-0.1640625)	$000.1110001_2$ (0.8828125)	$s[4]=+1$	$k_3 \approx 1.0077822$
5	$x[5] = x[4] - s[4]$ $(2^{-4})y[4]$	$y[5] = y[4] + s[4]$ $(2^{-4})x[4]$	$p[5] = p[4] -$ $s[4]atan(2^{-4})$		$k_4 = \text{sqrt}(1 + 2^{-2 \times 4})$
	$001.1100000_2$ (1.75)	$111.1111000_2$ (-6.25e-2)	$000.1101001_2$ (0.8203125)	$s[5]=+1$	$k_4 \approx 1.001951$
6	$x[6] = x[5] - s[5]$ $(2^{-5})y[5]$	$y[6] = y[5] + s[5]$ $(2^{-5})x[5]$	$p[6] = p[5] -$ $s[5]atan(2^{-5})$		$k_5 = \text{sqrt}(1 + 2^{-2 \times 5})$
	$001.1100001_2$ (1.7578125)	$111.1111111_2$ (-7.8125e-3)	$000.1100101_2$ (0.7890625)	$s[6]=+1$	$k_5 \approx 1.000488$
7	$x[7] = x[6] - s[6]$ $(2^{-6})y[6]$	$y[7] = y[6] + s[6]$ $(2^{-6})x[6]$	$p[7] = p[6] -$ $s[6]atan(2^{-6})$		$k_6 = \text{sqrt}(1 + 2^{-2 \times 6})$
	$001.1100010_2$ (1.765625)	$000.0000010_2$ (1.5625e-2)	$000.1100011_2$ (0.7734375)	$s[7]=-1$	$k_6 \approx 1.000122$
8	$x[8] = x[7] - s[7]$ $(2^{-7})y[7]$	$y[8] = y[7] + s[7]$ $(2^{-7})x[7]$	$p[8] = p[7] -$ $s[7]atan(2^{-7})$		$k_7 = \text{sqrt}(1 + 2^{-2 \times 7})$
	$001.1100010_2$ (1.765625)	$000.0000001_2$ (7.8125e-3)	$000.1100100_2$ (0.78125)	$s[8]=-1$	$k_7 \approx 1.000030$
<i>continued...</i>					

Iter.	x	y	p	s	$k_i$
9	$x[9] = x[8] - s[8]$ $(2^{-8})y[8]$	$y[9] = y[8] + s[8]$ $(2^{-8})x[8]$	$p[9] = p[8] -$ $s[8]\text{atan}(2^{-8})$		$k_8 = \text{sqrt}(1 + 2^{-2 \times 8})$
	001.1100010_2 (1.765625)	000.0000001_2 (7.8125e-3)	000.1100100_2 (0.78125)	$s[9] = -1$	$k_8 \approx 1.000007$
10	$x[10] = x[9] - s[9]$ $(2^{-9})y[9]$	$y[10] = y[9] + s[9]$ $(2^{-9})x[9]$	$p[10] = p[9] -$ $s[9]\text{atan}(2^{-9})$		$k_9 = \text{sqrt}(1 + 2^{-2 \times 9})$
	001.1100010_2 (1.765625)	000.0000001_2 (7.8125e-3)	000.1100100_2 (0.78125)		$k_9 \approx 1.000001$

Equation 9. CORDIC Gain After 10 Iterations

$$K(10) = K_0K_1K_2K_3K_4K_5K_6K_7K_8K_9 = \prod_{j=0}^9 K_j = \prod_{j=0}^9 \text{sqrt}(1 + 2^{-2j})$$

To obtain the magnitude of the input vector, divide the output on  $X[10]$  by  $K(10)$  to give:

$$\begin{aligned} \frac{X(10)}{K(10)} &= \frac{001.1100010_2}{1.64675921113982\dots} \\ &= 1.07218164504931\dots \\ &= 1.000100100111101001111\dots_2 \end{aligned}$$

When rounding to a 7-bit fraction format, the computed value, post scale factor compensation is:

$$mag_{\text{compute}} = 1.0001001_2 = 1.0703125$$

The golden value corresponding to the vector magnitude is:

$$mag_{\text{gold}} = \text{sqrt}(0.75^2 + 0.75^2) = 1.0000111110000111\dots_2 = 1.06066017177982\dots$$

The absolute error on the vector magnitude is:

$$|mag_{\text{gold}} - mag_{\text{compute}}| = 1.0011\dots_2 \times 2^{(-7)}$$

The vector magnitude error for this set of inputs is the same order of magnitude as the ULP value.

Compute the accuracy of the returned angle by subtracting it from the golden value, or  $\pi/4$ :

$$\left| \frac{\pi}{4} - 000.1100100_2 \right| = 1.00001111101101\dots_2 \times 2^{(-8)}$$

Thus the relative error on the angle component for this set of inputs is roughly half ULP.

### 14.4.21. Counter

The **Counter** block maintains a counter and outputs the counter value each cycle.

The input is a counter enable and allows you to implement irregular counters. The counter initializes to the value that you provide, and counts with the modulo, with the step size you provide:

```
count = _pre_initialization_value;
while (1) { if (en) count = (count + _step_size) % _modulo }
```

**Note:** If you create a counter with a preinitialization value of 0 and with a step of 1, it outputs the value 1 (not 0) on its first enabled cycle. If you want the counter to output 0 on its first valid output, initialize with:

```
[<(modulo - step size)> <modulo> <step size>]
```

**Note:** Modulo and step size cannot be coprime—the step size must exactly divide into the modulo value.

**Table 178. Parameters for the Counter Block**

Parameter	Description
Output data type mode	Determines how the block sets its output data type: <ul style="list-style-type: none"> <li><b>Inherit via internal rule:</b> the number of integer and fractional bits is the maximum of the number of bits in the input data types.</li> <li><b>Specify via dialog:</b> you can set the output type of the block explicitly using additional fields that are available when this option is selected. This option reinterprets the output bit pattern from the LSB up according to the specified type.</li> <li><b>Boolean:</b> the output type is Boolean.</li> </ul>
Output data type	Specifies the output data type. For example, <b>sfix(16), uint(8)</b> .
Output scaling value	Specifies the output scaling value. For example, <b>2<sup>-15</sup></b> .
Counter setup	A vector that specifies the counter in the format: <b>[&lt;pre_initialization_value&gt; &lt;modulo&gt; &lt;step size&gt;]</b> For example, <b>[0 32 1]</b>

**Table 179. Port Interface for the Counter Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
en	Input	Boolean	Count enable	Yes	No
q	Output	Specified fixed-point type	Result	Yes	No

### 14.4.22. Count Leading Zeros, Ones, or Sign Bits (CLZ)

The **CLZ** block counts the leading zeros, ones, or sign bits of the input, and outputs that count.

**Table 180. Parameters for the CLZ Block**

Parameter	Description
Mode	Determines how the block sets its output data type: <ul style="list-style-type: none"> <li><b>Count Leading Zeros:</b> returns the count of the leading zeros in the input</li> <li><b>Count Leading Ones:</b> returns the count of the leading ones in the input</li> <li><b>Count Leading Digits:</b> returns the count of the leading sign digits in the input</li> </ul>

**Table 181. Port Interface for the CLZ Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
a	Input	Any fixed-point type	Operand	Yes	No
q	Output	Derived fixed-point type	Number of consecutive zero bits in input word starting from the MSB	Yes	No

### 14.4.23. Dual Memory (DualMem)

The DSP Builder **DualMem** block models a dual interface memory structure. You can read or write the first data interface (inputs *d*, *a*, and *w*). The second data interface (specified on the second address input) is read only by default, but can be read and write if you turn on **Allow write on both ports**. The memory size is inferred from the size of the initialization array.

The behavior of read during write cycles of the memories depends on the interface to which you read:

- Reading from *q*<sub>1</sub> while writing to interface 1 outputs the new data on *q*<sub>1</sub> (write first behavior).
- Reading from *q*<sub>2</sub> while writing to interface 1 outputs the old data on *q*<sub>2</sub> (read first behavior).

Turning on **DONT\_CARE** may give a higher *f*<sub>MAX</sub> for your design, especially if you implement the memory as a MLAB. When this option is on, the output is not double-registered (and therefore, in the case of MLAB implementation, uses fewer external registers), and you gain an extra half-cycle on the output. The word don't care overlaid on the block symbol indicates the current setting is **DON'T CARE**. The default is off, which outputs old data for read-during-write.

**Table 182. Parameters for the DualMem Block**

Parameter	Description
Output data type mode	Determines how the block sets its output data type: <ul style="list-style-type: none"> <li>• <b>Inherit via internal rule:</b> the number of integer and fractional bits is the maximum of the number of bits in the input data types.</li> <li>• <b>Specify via dialog:</b> you can set the output type of the block explicitly using additional fields that are available when this option is selected. This option reinterprets the output bit pattern from the LSB up according to the specified type.</li> <li>• <b>Boolean:</b> the output type is Boolean.</li> </ul>
Output data type	Specifies the output data type. For example, <b>sfix(16), uint(8)</b> .
Output scaling value	Specifies the output scaling value. For example, <b>2<sup>-15</sup></b> .
Initial contents	Specifies the initialization data. The size of the 1-D array determines the memory size. This parameter may also be a fi object when specifying data of arbitrarily high precision.
Use DONT_CARE when reading from and writing to the same address	Turn this option on to produce faster hardware ( a higher <i>f</i> <sub>MAX</sub> ) but with uncertain read data in hardware if you are simultaneously reading from and writing to the same address. Ensure that you do not read from or write to the same address at the same time to guarantee valid read data.  Intel Hyperflex architectures restrict permissible configurations of the <b>DualMem</b> block. You might need to turn on this option for a valid configuration, if DSP Builder gives a warning. To avoid this restriction, implement a simple dual-port RAM, with port 1 as write-only (e.g. connect <i>q</i> <sub>1</sub> read on port 1 to a Simulink terminator block), and port 2 as read-only (with separate addressing).

*continued...*

Parameter	Description
	<p>When you turn on this option and you have a write on one port, a read on another port, and both have the same address, the read data is undefined. Simulink simulations represent these undefined values as zeros; the ModelSim simulation shows Xs. This difference in representation may cause simulation mismatches if you allow such undefined values to be generated. To prevent simulation mismatches:</p> <ul style="list-style-type: none"> <li>• Either avoid generating accesses that cause undefined values or detect the conditions of address equality and a write access at the input</li> <li>• Do not propagate that output.</li> </ul>
Allow write on both ports	Turn on to read and write on both ports.
Initialize Hardware Memory Blocks with Initial Data Contents	Turn on to initialize the memory with the specified initial data values. Turn off to use only the size of the initial data for the size of memory. The <b>Uninitialized means</b> parameter on the Control Block <b>Optimization</b> tab determines the content when the memory is uninitialized..
Provide read-access inputs	<p>Turn on to expose the memory read-access input control ('r1', 'r2') for each output port ('q1', 'q2'). You can use a read-access input to indicate when the corresponding output is not used. Disabling the output when the read-access input is '0', offers potential power-saving benefits and avoids write-read contention. When a read-access input is '0', you cannot assume the value of the corresponding output. The hardware response to read-access input '0' depends on the underlying physical memory's capabilities and design configuration.</p> <p>Turn off (default) to expose no read-access input controls. This setting is equivalent to permanently enabling read access to the outputs by driving the inputs with constant '1'.</p>

You can specify the contents of the **DualMem** block in one of the following ways:

- Use a single row or column vector to specify table contents. The length of the 1D row or column vector determines the number of addressable entries in the table. If DSP Builder reads vector data from the table, all components of a given vector share the same value.
- When a look-up table contains vector data, you can provide a matrix to specify the table contents. The number of rows in the matrix determines the number of addressable entries in the table. Each row specifies the vector contents of the corresponding table entry. The number of columns must match the vector length, otherwise DSP Builder issues an error.

Every **Primitive** library block accepts double-precision floating-point values when specifying mask parameters. This format limits precision to no more than 53 bits, which is more than sufficient for most of the blocks. For higher precision, the **Const**, **DualMem**, or **LUT** blocks optionally accept values using Simulink's Fixed Point data type. For example:

```
constValue = fi(0.142, 1, 16, 15)
```

```
vectorValue = fi(sin([0:10]'), 1, 18, 15)
```

To configure a **Const**, **DualMem**, or **LUT** with data of precision higher than IEEE double precision, create a MATLAB `fi` object of the required precision that contains the high precision data. Avoid truncation when creating this object. Use the `fi` object to specify the **Value** of the **Const**, the **Initial Contents** of the **DualMem** block, or the **Output value map** of the **LUT** block.



**Table 183. Port Interface for the DualMem Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
d1	Input	Any	Data to write for interface 1	Yes	Yes
a1	Input	Unsigned integer	Address to read or write from for interface 1	Yes	No
w1	Input	Boolean	Write is enabled for interface 1 when 1	Yes	No
r1	Input	Boolean	Read-access control for q output of interface 1. Only present when you turn on <b>Provide Read-Access Inputs</b>	Yes	No
d2	Input	Any	Data to write for interface 2. Only present when you turn on <b>Allow write on both ports</b>	Yes	Yes
a2	Input	Unsigned integer	Address to read or write from for interface 2	Yes	No
w2	Input	Boolean	Write is enabled for interface 2 when 1. Only present when you turn on <b>Allow write on both ports</b>	Yes	No
r2	Input	Boolean	Read-Access control for q output of interface 1. Only present when you turn on <b>Provide Read-Access Inputs</b>	Yes	No
q1	Output	Any	Data out from interface 1. (1)	Yes	Yes
q2	Output	Any	Data out from interface 2. (1)	Yes	Yes

Note:  
1. If the address for interface 1 exceeds the memory size, q1 is not defined. If the address for interface 2 exceeds the memory size, q2 is not defined.

**Related Information**

- [Control](#) on page 223
- [RAM Megafunction User Guide.](#)  
For more information about this option

**14.4.24. Demultiplexer (Demux)**

The **Demux** block deserializes the DSP Builder protocol bus on its inputs to produce a configurable number of output signals without TDM.

The **Demux** block is a primitive version of the **ChannelViewer** block

**Table 184. Parameters for the Demux Block**

Parameter	Description
Number of output channels	A vector of the channel number you want to see for example [0 1 3].
Number of input channels	The number of input channels. The block takes valid, channel, and (vector) data inputs. The channel is the normal channel count, which varies across 0 to <b>NumberOfChannelsPerWire</b> .

### 14.4.25. Divide

The **Divide** block outputs the first input, a, divided by the second input, b.

$$q = a/b$$

**Table 185. Parameters for the Divide Block**

Parameter	Description
Output data type mode	Determines how the block sets its output data type: <ul style="list-style-type: none"> <li><b>Inherit via internal rule:</b> if the input data types are floating-point, the output data type is the same floating-point data type. Mixing different precisions is not allowed. If the input data types are fixed-point, the output data type is fixed point with bitwidth equal the to the sum of the bitwidths of the input data types. The fraction width is equal to the sum of the fraction width of the a-input data type, and the integer bitwidth of the b-input data type.</li> <li><b>Specify via dialog:</b> you can set the output type of the block explicitly using additional fields that are available when this option is selected. This option type casts the output to the chosen fixed-point type. Attempting to type cast floating-point input is disallowed. You can only use this option to trim bits off the least significant end of the output data type that is otherwise inherited.</li> </ul>
Output data type	Specifies the output data type. For example, fixdt(1,16,15)
Rounding mode	Specifies what rounding to apply to the result: <ul style="list-style-type: none"> <li>Correct. IEEE compliant unbiased round to nearest output value.</li> <li>Faithful. Saves hardware by sometimes rounding to the second nearest value. Error is about double that of correct rounding.</li> </ul>
Output scaling value	Specifies the output scaling value. For example, 2 <sup>-15</sup> .
Float point rounding	This option only has an effect for floating-point inputs.: <ul style="list-style-type: none"> <li><b>Correct:</b> the result is correctly rounded IEEE</li> <li><b>Faithful:</b> the result is may be rounded up or may be rounded down</li> </ul>

Table 17–46 shows the data-type inheritance for fixed-point inputs.

**Table 186. Data-Type Inheritance for Fixed-Point Inputs**

Port	Fixed-Point Data Type	Integer Bits	Fraction Bits
a	sfix16_en10	6	10
b	sfix12_en7	5	7
q (Inherit via internal rule)	sfix28_en15	6 + 7 = 13	10 + 5 = 15

If you specify **Specify via dialog** for the **Output data type mode**, the block restricts the allowed data types to: **sfix28\_en15**, **sfix27\_en14**, **sfix26\_en13**, etc.

**Table 187. Port Interface for the Divide Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
a	Input	Any fixed- or floating-point type.	Numerator	Yes	No
b	Input	Any fixed- or floating-point type.	Denominator	Yes	No
q	Output	Any fixed- or floating-point type.	Result	Yes	No

### 14.4.26. Fanout

The **Fanout** block behaves like a wire, connecting its single input to one or more outputs. The **Fanout** and **VectorFanout** are similar blocks.

The number of outputs is one of the parameters of the **Fanout** block. Use a **Fanout** block instead of a simple wire to provide a hint to DSP Builder that the wire is expected to be long. DSP Builder might ignore the hint (which amounts to implementing the **Fanout** block using a simple wire), or might insert one or more additional registers on the wire to improve the physical routing of the design. The number of registers it inserts depends on the target device, target  $f_{MAX}$  and other properties of your design. Inserting a **Fanout** block does not change the behavior of your design. If DSP Builder chooses to insert extra registers, it automatically adjusts the latency of any parallel paths to preserve the original wire-like behaviour. By default, DSP Builder implements all **Fanout** blocks as simple wires on non-HyperFlex devices. FFTs and FIRs (which both contain embedded **Fanout** blocks) retain the same QoR characteristics as in DSP Builder v15.0 and earlier (which has no **Fanout** blocks). To enable DSP Builder to choose different implementations for the **Fanout** blocks in your design, specify `DSPBA_Features.EnableFanoutBlocks = true;` at the MATLAB command line. This command increases the number of registers your design uses, but potentially increases its  $f_{MAX}$ . You can specify that DSP Builder doesn't need to initialize any registers that it chooses to insert. Then DSP Builder inserts hyper-registers (instead of ordinary, ALM registers) on devices that support the HyperFlex architecture. You should use this option for datapaths where the initial value is unimportant, but you should avoid using it for control paths.

**Table 188. Parameters for the Fanout Block**

Parameter	Type	Description
Number of outputs	Integer > 1	Number of output ports
Uninitialized	Check box	Specifies whether DSP Builder can use hyper registers. When you apply automatic reset minimization, turn off <b>Uninitialized</b> , which allows reset minimization to choose the correct behavior automatically. Turning on <b>Uninitialized</b> forces no reset.

**Table 189. Port Interface for the Fanout Block**

Signal	Direction	Type	Description
d	Input	Any	Input
q0, q1, q2, etc	Output	Same as d.	Copy of d.

### 14.4.27. FIFO

The **FIFO** block models a FIFO memory. DSP Builder writes data through the  $d$  input when the write-enable input  $w$  is high. After some implementation-specific number of cycles, DSP Builder presents data at output  $q$  and the valid output  $v$  goes high. DSP Builder holds this data at output  $q$  until the read acknowledge input  $r$  is set high.

The **FIFO** block wraps the Intel single clock FIFO (SCFIFO) megafunction operating in show-ahead mode. That is, the read input,  $r$ , is a read acknowledgement which means the DSP Builder has read the output data,  $q$ , from the FIFO buffer, so you can delete it and show the next data output on  $q$ . The data you present on  $q$  is only valid if the output valid signal,  $v$ , is high

**Table 190. Parameters for the FIFO Block**

Parameter	Description
FIFO Setup	<p>A vector of three non-zero integers in the format: [<math>\langle \text{depth} \rangle</math> <math>\langle \text{fill\_threshold} \rangle</math> <math>\langle \text{full\_period} \rangle</math>]</p> <ul style="list-style-type: none"> <li><math>\text{depth}</math> specifies the maximum number of data values that the FIFO can store.</li> <li><math>\text{fill\_threshold}</math> specifies a low-threshold for empty-detection. If the number of data items in the memory is greater than the low-threshold, the <math>t</math> output is 1 (otherwise it is 0).</li> <li><math>\text{full\_period}</math> specifies a high-threshold for full-detection. If the number of data items is greater than the high-threshold, output <math>f</math> is 1 (otherwise it is 0).</li> </ul>

If the inputs  $w$  or  $r$  is a vector, the FIFO setup parameter must be a three column matrix with the number of rows equal to the number of components in the vector. Each row in the matrix independently configures the depth, fill\_threshold, and full\_period of the FIFO buffer for the corresponding vector component.

**Table 191. Port Interface for the FIFO Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
$w$	Input	Boolean	Write enable.	Yes	No
$d$	Input	Fixed-point	Data.	Yes	Yes
$r$	Input	Boolean	Read acknowledge.	Yes	No
$v$	Output	Boolean	Valid.	Yes	No
$q$	Output	Fixed-point	Data.	Yes	Yes
$t$	Output	Boolean	Fill threshold.	Yes	No
$f$	Output	Boolean	Fullness.	Yes	No

You can to set  $\text{fill\_threshold}$  to a low number ( $<3$ ) and arrive at a state such that output  $t$  is high and output  $v$  is low, because of differences in latency across different pairs of ports—from  $w$  to  $v$  is 3 cycles, from  $r$  to  $t$  is 1 cycle, from  $w$  to  $t$  is 1 cycle. If this situation arises, do not send a read acknowledgement to the FIFO buffer. Ensure that when the  $v$  output is low, the  $r$  input is also low, otherwise a warning appears in the MATLAB command window. If the read acknowledgement is derived from a feedback from the  $t$  output, ensure that the  $\text{fill\_threshold}$  is set to a sufficiently high number (3 or above). Likewise for the  $f$  output and the  $\text{full\_period}$ .

You may supply vector data to the  $d$  input, and vector data on the  $q$  output is the result. DSP Builder does not support vector signals on the  $w$  or  $r$  inputs, and the behavior is unspecified. The  $v$ ,  $t$ , and  $f$  outputs are always scalar.

### 14.4.28. Floating-point Classifier (FloatClass)

The **FloatClass** block indicates whether a floating-point input is equal to zero, is signed (negative), is infinity, or is equal to not a number.

**Table 192. Port Interface for the FloatClass Block**

Signal	Direction	Vector Data Support	Complex Data Support	Description
x	Input	Yes	—	—
inf	Output	Yes	—	Infinity output.
nan	Output	Yes	—	Not a number output.
sig	Output	Yes	—	Signed negative output.
zer	Output	Yes	—	Zero output.

### 14.4.29. Floating-point Multiply Accumulate (MultAcc)

The **MultAcc** block instantiates a DSP block in multiply-accumulate mode. This block only works on any device with a floating-point DSP block and supports a hardware single-precision multiply accumulate structure. The block latency is 4 cycles.

**Table 193. Parameters for the MultAcc Block**

Parameter	Description
Output data type mode	<p>Determines how the block sets its output data type:</p> <ul style="list-style-type: none"> <li><b>Inherit via internal rule:</b> the number of integer and fractional bits is the maximum of the number of bits in the input data types.</li> <li><b>Specify via dialog:</b> you can set the output type of the block explicitly using additional fields that are available when this option is selected. This option reinterprets the output bit pattern from the LSB up according to the specified type.</li> <li><b>Boolean:</b> the output type is Boolean.</li> </ul>
Function	Accumulate (fpAcc) or multiply accumulate (fpMultAcc). When you select <b>fpMultAcc</b> , the block flushes denormalized numbers to zero on inputs and output; when you select <b>fpAcc</b> , the block flushes subnormal numbers to zero on inputs and output.

The fpMultAcc function implements the following equation:

$$q_n = (\text{acc} \& q_{(n-1)}) + x * y$$

- When acc is high (1) the result is equal to the sum between the previous accumulated result and the product  $x*y$ .
- When acc is low (0) the output value is the product value  $x*y$

The fpAcc function implements the following equation:

$$q_n = (\text{acc} \& q_{(n-1)}) + x$$

- When acc is high the result consists of the sum between the previous accumulated result and the input x.
- When acc is low the x input is forwarded to the output q.
- Subnormal numbers are flushed to zero on inputs and output

**Table 194. Port Interface for the MultAcc Block**

Signal	Direction	Vector Data Support	Complex Data Support	Description
x	Input	—	—	—
y	Input	—	—	Multiply accumulate only. Tied to 1 for the accumulate parameter.
acc	Input	—	—	—
q	Output	—	—	—

### 14.4.30. ForLoop

The **ForLoop** block extends the basic loop, providing a more flexible structure that implements all common loop structures—for example, triangular loops, parallel loops, and sequential loops.

Each **ForLoop** block manages a single counter with a token-passing scheme that allows you to link these counters in a variety of ways.

Each **ForLoop** block has a static loop test parameter, which may be  $\leq$ ,  $<$ ,  $>$  or  $\geq$ . Loops that count up should use  $\leq$  or  $<$ , depending on whether you consider the limit value, supplied by the limit signal, is within the range of the loop. Loops that count down should use  $\geq$  or  $>$ .

The latency of the **ForLoop** block is non-zero. At loop end detection there are some cycles that may be invalid overhead required to build nested loop structures. The second activation of an inner loop does not necessarily begin immediately after the end of the first activation.

**Table 195. Port Interface for the ForLoop Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
bs	Output	Boolean	Token-passing inputs and outputs. The four signals <i>ls</i> (loop start), <i>bs</i> (body start), <i>bd</i> (body done) and <i>ld</i> (loop done) pass a control token between different <b>ForLoop</b> blocks, to create a variety of different control structures.  When the <i>ls</i> port receives a token, the <b>ForLoop</b> block initializes. The loop counter is set to its initial value (that the <i>i</i> signal specifies). When the <i>bd</i> port receives a token, the step value ( <i>s</i> ) increments the loop counter. In either case, the new value of the counter is compared with the limit value ( <i>l</i> ) with the statically-configured loop test.	Yes	No
bd	Input	Boolean		Yes	No
ld	Output	Boolean		Yes	No
ls	Input	Boolean		Yes	No

*continued...*

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
			<p>If the loop test passes, the <b>ForLoop</b> block outputs the control token on the <b>bs</b> port to initiate the loop body and the valid signal, <b>v</b>, becomes active. If the loop test fails, the <b>ForLoop</b> block outputs the control token on <b>ld</b> port to indicate that the loop is complete and <b>v</b> becomes inactive.</p> <p>The <b>ForLoop</b> block becomes active when it receives a token on its <b>ls</b> port, and remains active until it finally outputs a token on its <b>ld</b> port. Changing any of the loop parameterization inputs (<b>i</b>, <b>s</b>, or <b>l</b>) while the loop is active, is not supported and produces unpredictable results.</p>		
<b>c</b>	Output	Derived unsigned fixed-point type	The signal <b>c</b> is the count output from the loop. Its value is reliable only when the valid signal, <b>v</b> , is active	Yes	No
<b>e</b>	Input	Boolean	Use the enable input, <b>e</b> , to suspend and resume operation of the <b>ForLoop</b> block. When you disable the loop, the valid signal, <b>v</b> , goes low but DSP Builder makes no changes to the internal state of the block. When you re-enable the block, it resumes counting from the state at which you suspended it.	Yes	No
<b>i</b>	Input	Any unsigned fixed-point type	<p>Loop parameterization inputs. The signals <b>i</b>, <b>s</b>, and <b>l</b> set the initial value, step and limit value (respectively) of the loop. Use with the loop test parameter, to control the operation of the loop. The loop parameter signals must be held constant while the loop is active, but you may them when the loop is inactive. Different activations of a <b>ForLoop</b> block can have different start or end points, which is useful for creating nested triangular loops, for example.</p>	Yes	No
<b>s</b>	Input	Any unsigned fixed-point type		Yes	No
<b>l</b>	Input	Any unsigned fixed-point type		Yes	No
<b>e1</b>	Output	Boolean	<p>Auxiliary loop outputs: the signals <b>f1</b> and <b>l1</b> are active on the first loop iteration and last loop iteration, respectively. The signal <b>e1</b> is active when the <b>ForLoop</b> block is processing an empty loop.</p>	Yes	No
<b>f1</b>	Output	Boolean		Yes	No
<b>l1</b>	Output	Boolean		Yes	No

### 14.4.31. Load Exponent (LdExp)

Table 196. Functions for the LdExp Block

Function	Description
<b>ilogb(a)</b>	Outputs the integer logarithm of input, <b>a</b> , to the base 2. $q = \text{floor}(\log_2 a )$ .
<b>ldexp(a,b)</b>	outputs the first input, <b>a</b> , scaled by 2 raised to the power of the second input, <b>b</b> . $q = a \cdot 2^b$ .

The **Function mask** parameter selects either **ldexp** or **ilogb**. The number of input ports on the block change according to the number of operands.

**Table 197. Port Interface for the LdExp Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
a	Input	Single or double	Operand 1	Yes	No
b	Input	Integer	Operand 2	Yes	No
q	Output	floating-point for ldexp; Integer for ilogb	Result	Yes	No

### 14.4.32. Left Shift (LShift)

The **LShift** block outputs the left shifted version of the input value.

The shift is specified by the input b:

$$q = (a \ll b)$$

The width of the data type **a** determines the maximum size of the shift. Shifts of more than the input word width result in an output of 0.

**Table 198. Parameters for the LShift Block**

Parameter	Description
Output data type mode	Determines how the block sets its output data type: <ul style="list-style-type: none"> <li><b>Inherit via internal rule:</b> the number of integer and fractional bits is the maximum of the number of bits in the input data types.</li> <li><b>Specify via dialog:</b> you can set the output type of the block explicitly using additional fields that are available when this option is selected. This option reinterprets the output bit pattern from the LSB up according to the specified type.</li> <li><b>Boolean:</b> the output type is Boolean.</li> </ul>
Output data type	Specifies the output data type. For example, <b>sfix(16)</b> , <b>uint(8)</b> .
Output scaling value	Specifies the output scaling value. For example, <b>2<sup>-15</sup></b> .

**Table 199. Port Interface for the LShift Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
a	Input	Any fixed-point type	Operand 1	Yes	No
b	Input	Any fixed-point type	Operand 2	Yes	No
q	Output	Derived fixed-point type	Result	Yes	No

### 14.4.33. Loadable Counter (LoadableCounter)

The **LoadableCounter** block maintains a counter that you can reload with new parameters as needed in-circuit. The value of the counter increments by the step value every cycle for which the enable input is high. If the counter exceeds or equals the modulo value, or underflows in the case of a negative step value, it wraps around to zero or the value minus the modulo value as applicable. The current counter value is always available from the block's only output.



Internal registers hold the value, modulo, and step size of the counter. The values of these registers on reset are parameters that you can set on the block. Additionally, you can reload these registers with new values in-circuit by raising the **ld** signal high. While **ld** is high, DSP Builder writes the values of the **i**, **s**, and **m** input signals into the value, step, and modulo registers, respectively. The value of **i** passes through to the counter output. When **ld** falls low again, the counter resumes its normal operation starting from these new values.

If the initial or step values exceed the modulo value, the behavior is undefined. Using signed step values increases logic usage in hardware.

**Table 200. Parameters**

Parameter	Description
Counter setup	A vector that specifies the counter settings on reset in the following format: [<initial value> <modulo> <step size>]

**Table 201. Signals**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
en	Input	Boolean	Enable the counter.	Yes	No
ld	Input	Boolean	Load the counter.	Yes	No
i	Input	Any unsigned integer	New initial value to load.	Yes	No
s	Input	Any integer	New step value to load.	Yes	No
m	Input	Any non-zero unsigned integer	New modulo value to load.	Yes	No
q	Output	Unsigned integer	Counter value.	Yes	No

### 14.4.34. Look-Up Table (Lut)

The DSP Builder **Lut** block outputs the contents of a look-up table, indexed by the input.

$$q = \text{LUT}[a]$$

The size of the table determines the size of the initialization arrays.

**Table 202. Parameters for the Lut Block**

Parameter	Description
Output data type mode	Determines how the block sets its output data type: <ul style="list-style-type: none"> <li>• <b>Inherit via internal rule:</b> the number of integer and fractional bits is the maximum of the number of bits in the input data types.</li> <li>• <b>Inherit via back projection:</b> a downstream block that this block drives determines the output data type. If the driven block does not propagate a data type to the driver, you must use a Simulink SameDT block to copy the required data type to the output wire.</li> <li>• <b>Specify via dialog:</b> you can set the output type of the block explicitly using additional fields that are available when this option is selected. This option reinterprets the output bit pattern from the LSB up according to the specified type.</li> <li>• <b>Boolean:</b> the output type is Boolean.</li> </ul>
<i>continued...</i>	

Parameter	Description
	<ul style="list-style-type: none"> <li>• <b>Single</b>: single-precision floating-point data.</li> <li>• <b>Double</b>: double-precision floating-point data.</li> <li>• <b>Variable precision floating point</b>: variable precision floating-point output type</li> </ul>
Output data type	Specifies the output data type. For example, <b>sfix(16), uint(8)</b> .
Output scaling value	Specifies the output scaling value. For example, <b>2^-15</b> .
Output value map	Specifies the location of the output values. For example, <b>round([0;254]/17)</b> . This parameter may also be a fi object when specifying data of arbitrarily high precision.
Floating point precision	Specifies the floating-point precision. For example, <b>float32_m23</b> .

You can specify the contents of the **Lut** block in one of the following ways:

- Specify table contents with a single row or column vector. The length of the 1D row or column vector determines the number of addressable entries in the table. If DSP Builder reads vector data from the table, all components of a given vector share the same value.
- When a look-up table contains vector data, you can provide a matrix to specify the table contents. The number of rows in the matrix determines the number of addressable entries in the table. Each row specifies the vector contents of the corresponding table entry. The number of columns must match the vector length, otherwise DSP Builder issues an error.

*Note:* The default initialization of the LUT is a row vector **round([0:255]/17)**. This vector is inconsistent with the default for the **DualMem** block, which is a column vector **[zeros(16, 1)]**. The latter form is consistent with the new matrix initialization form in which the number of rows determines the addressable size.

Every **Primitive** library block accepts double-precision floating-point values when specifying mask parameters. This format limits precision to no more than 53 bits, which is more than sufficient for most of the blocks. For higher precision, the **Const**, **DualMem**, or **LUT** blocks optionally accept values using Simulink's Fixed Point data type. For example:

```
constValue = fi(0.142, 1, 16, 15)
vectorValue = fi(sin([0:10]'), 1, 18, 15)
```

To configure a **Const**, **DualMem**, or **LUT** with data of precision higher than IEEE double precision, create a MATLAB fi object of the required precision that contains the high precision data. Avoid truncation when creating this object. Use the fi object to specify the **Value** of the **Const**, the **Initial Contents** of the **DualMem** block, or the **Output value map** of the **LUT** block.

**Table 203. Port Interface for the Lut Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
a	Input	Any fixed--point type	Operand	Yes	No
q	Output	Derived fixed- or floating-point type	Result	Yes	Yes (with output value map)

### 14.4.35. Loop

The **Loop** block maintains a set of counters that implement the equivalent of a nested for loop in software. The counted values range from 0 to limit values provided with an input signal.

When the `go` signal is asserted on the `g` input, limit-values are read into the block with the `c` input. The dimension of the vector determines the number of counters (nested loops). When DSP Builder enables the block with the `e` input, it presents the counter values as a vector value at the `q` output each cycle. The valid output is set to 1 to indicate that a valid output is present.

There are vectors of flags indicating when first values (output `f`) and last values (output `l`) occur.

A particular element in these vector outputs is set to 1 when the corresponding loop counter is set at 0 or at count-1 respectively.

Use the **Loop** block to drive datapaths that operate on regular data either from an input port or data stored in a memory. The enable input, and corresponding valid output, facilitate forward flow control.

For a two dimensional loop the equivalent C++ code to describe the general loop is:

```
for (int i = 0; i < c[0]; i++)
for (int j = 0; j < c[1]; j++) {
q[0] = i;
q[1] = j;
f[0] = (i==0);
f[1] = (j==0);
l[0] = (i==(c[0]-1));
l[1] = (j==(c[1]-1));
}
```

**Table 204. Port Interface for the Loop Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
<code>g</code>	Input	Boolean	Go.	Yes	No
<code>c</code>	Input	Unsigned integer	Counter limit values.	Yes	No
<code>e</code>	Input	Boolean	Enable.	Yes	No
<code>v</code>	Output	Boolean	Valid.	Yes	No
<code>q</code>	Output	Unsigned integer	Counter output values.	Yes	No
<code>f</code>	Output	Boolean	First value flags.	Yes	No
<code>l</code>	Output	Boolean	Last value flags.	Yes	No

### 14.4.36. Math

The **Math** block applies a mathematical operation to its floating-point inputs and outputs the floating-point result. A mask parameter popup menu selects the required elementary mathematical function that DSP Builder applies.

**Table 205. Functions for the Math Block**

Function	Description
exp(x) (1)	e raised to the exponent x.
exp2(x)	2 raised to the exponent x.
exp10(x)	10 raised to the exponent x.
expm1(x)	exp(x) - 1.
inverse(x)	The reciprocal of x. For <b>Floating-point rounding</b> , select either <b>correct</b> or <b>faithful</b> .
hypot(x,y)	Hypotenuse of right-angled triangle with other two sides length x and y.
hypot3d	The Euclidean norm of a vector in a three-dimensional space.
log(x) (1)	Natural logarithm.
log2(x)	Logarithm of x to the base 2.
log10(x)	Logarithm of x to the base 10.
log1p(x)	log(x+1).
pow(x,y)	x raised to the power of y.
powr(x,y)	x raised to the power of y, where I is non-negative.
mod(x,y)	(x - n × y) where n = y/x rounded toward zero.
Note: 1. For single-precision input and designs targeting any device with a floating-point DSP block, the block uses a mixture of resources including the DSP blocks in floating-point mode. This implementation uses fewer ALMs at the expense of more DSP blocks.	

Table 17-66 shows the functions for the **Math** block.

The **Function** mask parameter selects one of five elementary functions. The number of input ports on the block change as required by the semantics of the function that you select:

- One-input function:  $\exp(x)$ ,  $\exp2(x)$ ,  $\exp10(x)$ ,  $\expm1(x)$ ,  $\log(x)$ ,  $\log2(x)$ ,  $\log10(x)$ ,  $\log1p(x)$ ,  $\text{inverse}(x)$
- Two-input functions:  $\text{hypot}(x,y)$ ,  $\text{mod}(x,y)$ ,  $\text{pow}(x,y)$ ,  $\text{powr}(x,y)$
- Three input function:  $\text{hypot3d}(x,y,z)$

**Table 206. Port Interface for the Math Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
x	Input	Single or double	Operand 1	Yes	No
y	Input	Single or double	Operand 2 (hypot, mod, pow, and powr only)	Yes	No
z	Input	Single or double	Hypot3 only.	Yes	No
q	Output	Single or double	Result	Yes	No

### 14.4.37. Minimum and Maximum (MinMax)

The **MinMax** block allows you to select a bounding function to apply to the inputs.

**Table 207. Functions for the MinMax Block**

Function	Data types	Description
max	Fixed- or floating-point	Outputs a if $a > b$ , otherwise outputs b.
min	Fixed or floating-point	Outputs a if $a < b$ , otherwise outputs b.
maxmag	Floating-point	Outputs a if $ a  >  b $ , b if $ b  >  a $ , otherwise $\max(a,b)$ .
minmag	Floating-point	Outputs a if $ a  <  b $ , b if $ b  <  a $ , otherwise $\min(a,b)$ .
dim	Floating-point	Outputs $(a - b)$ if $a > b$ , otherwise outputs 0.
sat	Floating-point	Saturate input a to interval $[c,b]$ .

The **Function** mask parameter selects one of six bounding functions. The number of input ports on the block change as required by the semantics of the function that you select:

- Two-input functions: max, min, maxmag, minmag, dim
- Three-input functions: sat

The **Output data type mode** mask parameter applies only if the input is fixed-point format.

**Table 208. Port Interface for the MinMax Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
	Input	Fixed-point, single or double	Operand 1	Yes	No
b	Input	Fixed-point, single or double	Operand 2	Yes	No
c	Input	Single or double	Operand 3 (saturate only)	Yes	No
q	Output	Single or double	Result	Yes	No

### 14.4.38. MinMaxCtrl

The **MinMaxCtrl** block applies a minimum or maximum operator to the inputs depending on the Boolean signal it receives on the **control** port.

**Table 209. Functions of the MinMaxCtrl Block**

s	Function	Description
0	Minimum	Output a if a < b, otherwise output b.
1	Maximum	Output a if a > b, otherwise output b.

The **Output data type mode** mask parameter applies only if the inputs **a** and **b** are fixed-point format.

**Table 210. Parameters for the MinMaxCtrl Block**

Parameter	Description
Output data type mode	Determines how the block sets its output data type: <ul style="list-style-type: none"> <li><b>Inherit via internal rule:</b> the number of integer and fractional bits is the maximum of the number of bits in the input data types.</li> <li><b>Specify via dialog:</b> you can set the output type of the block explicitly using additional fields that are available. This option reinterprets the output bit pattern from the LSB up according to the specified type.</li> <li><b>Boolean:</b> the output type is Boolean.</li> </ul>
Output data type	Specifies the output data type. For example, sfix(16), uint(8).
Output scaling value	Specifies the output scaling value. For example, 2 <sup>-15</sup> .

**Table 211. Port Interface for the MinMaxCtrl Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
	Input	Fixed-point, single or double	Operand 1	Yes	No
b	Input	Fixed-point, single or double	Operand 2	Yes	No
s	Input	Boolean	Select input	Yes	No
q	Output	Fixed-point, single or double	Result	Yes	No

### 14.4.39. Multiply (Mult)

The **Mult** block outputs the product of the inputs:

$$q = a \times b$$

*Note:* For single-precision inputs and designs targeting any device with a floating-point DSP block, the block uses a mixture of resources including the DSP blocks in floating-point mode.

**Table 212. Parameters for the Mult Block**

Parameter	Description
Output data type mode	Determines how the block sets its output data type: <ul style="list-style-type: none"> <li>• <b>Inherit via internal rule:</b> the number of integer and fractional bits is the maximum of the number of bits in the input data types.</li> <li>• <b>Specify via dialog:</b> you can set the output type of the block explicitly using additional fields that are available when this option is selected. This option reinterprets the output bit pattern from the LSB up according to the specified type.</li> <li>• <b>Boolean:</b> the output type is Boolean.</li> <li>• <b>Variable precision floating point:</b> variable precision floating-point output type</li> </ul>
Output data type	Specifies the output data type. For example, <b>sfixed(16)</b> , <b>uint(8)</b> .
Floating point rounding	Specifies what rounding to apply to the result: <ul style="list-style-type: none"> <li>• Correct. IEEE compliant unbiased round to nearest output value.</li> <li>• Faithful. Saves hardware by sometimes rounding to the second nearest value. Error is about double that of correct rounding.</li> </ul>
Output scaling value	Specifies the output scaling value. For example, <b>2<sup>-15</sup></b> .
Floating point precision	Specifies the floating-point precision. For example, <b>float32_m23</b> .

**Table 213. Port Interface for the Mult Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
a	Input	Any fixed- or floating-point type	Operand 1	Yes	Yes
b	Input	Any fixed- or floating-point type	Operand 2	Yes	Yes
q	Output	Derived fixed- or floating-point type	Result	Yes	Yes

#### Related Information

[Forcing Soft Floating-point Data Types with the Advanced Options](#) on page 220

### 14.4.40. Multiplexer (Mux)

The **Mux** block allows a variable number of inputs and outputs the selected input, or zero if the select value is invalid (outside the number of data signals).

**Note:** You can make a multiple input multiplexer by combining more than one **mux2** blocks in a tree or by using a **Select** block.

**Table 214. Parameters for the Mux Block**

Parameter	Description
Number of data signals	The input type for <i>s</i> is an unsigned integer of width $\log_2(\text{number of data signals})$ . Boolean is also allowed in the case of two data inputs.
Output data type mode	Determines how the block sets its output data type: <ul style="list-style-type: none"> <li><b>Inherit via internal rule:</b> the number of integer and fractional bits is the maximum of the number of bits in the input data types.</li> <li><b>Specify via dialog:</b> you can set the output type of the block explicitly using additional fields that are available when this option is selected. This option reinterprets the output bit pattern from the LSB up according to the specified type.</li> <li><b>Boolean:</b> the output type is Boolean.</li> </ul>
Output data type	Specifies the output data type. For example, <b>sfixed(16)</b> , <b>uint(8)</b> .
Output scaling value	Specifies the output scaling value. For example, $2^{-15}$ .

**Table 215. Port Interface for the Mux Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
<i>s</i>	Input	Any fixed-point type	Select input	Yes	No
0	Input	Any fixed- or floating-point type	Input 0	Yes	Yes
1	Input	Any fixed- or floating-point type	Input 1	Yes	Yes
<i>q</i>	Output	Derived fixed- or floating-point type	Result	Yes	Yes

### 14.4.41. NAND Gate (Nand)

The **Nand** block outputs the logical **NAND** of the input values:

$$q = \sim(a \& b)$$

If the number of inputs is set to 1, then output the logical **NAND** of all the individual bits of the input word.

**Table 216. Parameters for the Nand Block**

Parameter	Description
Number of inputs	Specifies the number of inputs.
Output data type mode	Determines how the block sets its output data type:

*continued...*



Parameter	Description
	<ul style="list-style-type: none"> <li>• <b>Inherit via internal rule:</b> the number of integer and fractional bits is the maximum of the number of bits in the input data types.</li> <li>• <b>Specify via dialog:</b> you can set the output type of the block explicitly using additional fields that are available when this option is selected. This option reinterprets the output bit pattern from the LSB up according to the specified type.</li> <li>• <b>Boolean:</b> the output type is Boolean.</li> </ul>
Output data type	Specifies the output data type. For example, <b>sfixed(16)</b> , <b>uint(8)</b> .
Output scaling value	Specifies the output scaling value. For example, <b>2<sup>-15</sup></b> .

**Table 217. Port Interface for the Nand Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
a	Input	Any fixed-point type	Operand 1	Yes	No
b	Input	Any fixed-point type	Operand 2	Yes	No
q	Output	Derived fixed-point type	Result	Yes	No

#### 14.4.42. Negate

The **Negate** block outputs the negation of the input value.

The **Output datatype mode** determines how the block infers its output data type:

- **Inherit via internal rule.** The output data type is the same as the input data type.
- **Inherit via internal rule with word growth.** The output data type is the same as the input data type. If the input data type is fixed-point, word growth is applied to the output data type.

**Table 218. Port Interface for the Not Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
a	Input	Any fixed- or floating-point type	Operand	Yes	Yes
q	Output	Any fixed- or floating-point type	Result	Yes	No

#### 14.4.43. NOR Gate (Nor)

The **Nor** block outputs the logical **NOR** of the input values:

$$q = \sim(a | b)$$

Set the number of inputs to 1, to output the logical **NOR** of all the individual bits of the input word.

**Table 219. Parameters for the Nor Block**

Parameter	Description
Number of inputs	Specifies the number of inputs.
Output data type mode	Determines how the block sets its output data type: <ul style="list-style-type: none"> <li>• <b>Inherit via internal rule:</b> the number of integer and fractional bits is the maximum of the number of bits in the input data types.</li> <li>• <b>Specify via dialog:</b> you can set the output type of the block explicitly using additional fields that are available when this option is selected. This option reinterprets the output bit pattern from the LSB up according to the specified type.</li> <li>• <b>Boolean:</b> the output type is Boolean.</li> </ul>
Output data type	Specifies the output data type. For example, <b>sfix(16), uint(8)</b> .
Output scaling value	Specifies the output scaling value. For example, <b>2<sup>-15</sup></b> .

**Table 220. Port Interface for the Nor Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
a	Input	Any fixed-point type	Operand 1	Yes	No
b	Input	Any fixed-point type	Operand 2	Yes	No
q	Output	Derived fixed-point type	Result	Yes	No

#### 14.4.44. NOT Gate (Not)

The **Not** block outputs the logical **NOT** of the input value:

$$q = \sim a$$

**Table 221. Parameters for the Not Block**

Parameter	Description
Output data type mode	Determines how the block sets its output data type: <ul style="list-style-type: none"> <li>• <b>Inherit via internal rule:</b> the number of integer and fractional bits is the maximum of the number of bits in the input data types.</li> <li>• <b>Specify via dialog:</b> you can set the output type of the block explicitly using additional fields that are available when this option is selected. This option reinterprets the output bit pattern from the LSB up according to the specified type.</li> <li>• <b>Boolean:</b> the output type is Boolean.</li> </ul>
Output data type	Specifies the output data type. For example, <b>sfix(16), uint(8)</b> .
Output scaling value	Specifies the output scaling value. For example, <b>2<sup>-15</sup></b> .

**Table 222. Port Interface for the Not Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
a	Input	Any fixed-point type	Operand	Yes	No
q	Output	Derived fixed-point type	Result	Yes	No

#### 14.4.45. OR Gate (Or)

The **Or** block outputs the logical **OR** of the input values:

$$q = a | b$$

Set the number of inputs to 1, to output the logical **OR** of all the individual bits of the input word.

**Table 223. Parameters for the Or Block**

Parameter	Description
Number of inputs	Specifies the number of inputs.
Output data type mode	Determines how the block sets its output data type: <ul style="list-style-type: none"> <li>• <b>Inherit via internal rule:</b> the number of integer and fractional bits is the maximum of the number of bits in the input data types.</li> <li>• <b>Specify via dialog:</b> you can set the output type of the block explicitly using additional fields that are available when this option is selected. This option reinterprets the output bit pattern from the LSB up according to the specified type.</li> <li>• <b>Boolean:</b> the output type is Boolean.</li> </ul>
Output data type	Specifies the output data type. For example, <b>sfix(16)</b> , <b>uint(8)</b> .
Output scaling value	Specifies the output scaling value. For example, <b>2<sup>-15</sup></b> .

**Table 224. Port Interface for the Or Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
a	Input	Any fixed-point type	Operand 1	Yes	No
b	Input	Any fixed-point type	Operand 2	Yes	No
q	Output	Derived fixed-point type	Result	Yes	No

#### 14.4.46. Polynomial

The **Polynomial** block takes input **x**, and provides the result of evaluating a polynomial of degree, **n**:

$$f(x) = a_0 + a_1x + a_2x^2 + \dots + a_nx^n$$

**Table 225. Parameters for the Polynomial Block**

Parameter	Description
Coefficient banks	A vector of (n + 1) components. Specify the coefficients in the order a0, a1, a2, ..., an. If input x is driven by a vector signal, then a matrix with (n+1) columns, and one row per vector component can be specified. Each output component will be the result of evaluating an independently defined polynomial of degree, n. If there is more than one coefficient bank, the number of rows in the matrix should be v*u, for v vector components, and u banks. The coefficients for a given bank are ordered contiguously in the matrix.
Number of coefficient banks	<ul style="list-style-type: none"> <li>Set to the default value of 1, for only one input, x.</li> <li>Set to greater than 1, for a second input, b, to specify which bank of coefficients DSP Builder uses to evaluate the polynomial.</li> </ul>

**Table 226. Port Interface for the Polynomial Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
x	Input	Floating-point	Data	Yes	No
b	Input	Integer	Bank selector	No	No
q	Output	Floating-point	Data	Yes	No

#### 14.4.47. Ready

Use the **Ready** block in designs with ALU folding. The **Ready** block adds a ready signal to your design.

#### 14.4.48. Reinterpret Cast (ReinterpretCast)

The **ReinterpretCast** block outputs the same bit pattern that it reads on its input port, but casts it to a data type that you specify with the block parameters. This data type should use the same number of bits as the bit width of the input signal.

**Table 227. Parameters for the ReinterpretCast Block**

Parameter	Description
Output data type mode	Determines how the block sets its output data type: <ul style="list-style-type: none"> <li><b>Inherit via internal rule:</b> the number of integer and fractional bits is the maximum of the number of bits in the input data types.</li> <li><b>Inherit via back projection:</b> a downstream block that this block drives determines the output data type. If the driven block does not propagate a data type to the driver, you must use a Simulink SameDT block to copy the required data type to the output wire.</li> <li><b>Specify via dialog:</b> you can set the output type of the block explicitly using additional fields that are available when this option is selected. This option reinterprets the output bit pattern from the LSB up according to the specified type.</li> <li><b>Boolean:</b> the output type is Boolean.</li> </ul>

*continued...*

Parameter	Description
	<ul style="list-style-type: none"> <li>• <b>Single</b>: single-precision floating-point data.</li> <li>• <b>Double</b>: double-precision floating-point data.</li> <li>• <b>Variable precision floating point</b>: variable precision floating-point output type</li> </ul>
Output data type	Specifies the output data type. For example, <b>sfix(16)</b> , <b>uint(8)</b> .
Output scaling value	Specifies the output scaling value. For example, <b>2<sup>-15</sup></b> .
Floating point precision	Specifies the floating-point precision. For example, <b>float32_m23</b> .

**Table 228. Port Interface for the ReinterpretCast Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
a	Input	Any fixed- or floating-point type	Operand	Yes	Yes
q	Output	User specified	Result	Yes	Yes

#### 14.4.49. Round

The **Round** block applies a rounding operation to the floating-point input. A mask parameter popup menu selects the required rounding function that you apply.

**Table 229. Functions for the Round Block**

Function	Description
ceil(x)	Lowest integer not less than input x.
floor(x)	Highest integer not exceeding input x.
rint(x)	Round to nearest integer; halfway cases rounded to even number.
round(x)	Round to nearest integer; halfway cases rounded away from zero.

The **Function** mask parameter selects one of four rounding functions.

**Table 230. Port Interface for the Round Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
x	Input	Single or double	Operand	Yes	Yes
q	Output	Single or double	Result	Yes	Yes

#### 14.4.50. Sample Delay (SampleDelay)

The **SampleDelay** block outputs a delayed version of the input.

*Note:* **SampleDelay** blocks might not reset to zero. Do not use designs that rely on **SampleDelays** output of zero after reset. Use the `valid` signal to indicate valid data and its propagation through the design.

**Table 231. Parameters for the SampleDelay Block**

Parameter	Description
Output data type mode	Determines how the block sets its output data type: <ul style="list-style-type: none"> <li>• <b>Inherit via internal rule:</b> the number of integer and fractional bits is the maximum of the number of bits in the input data types.</li> <li>• <b>Specify via dialog:</b> you can set the output type of the block explicitly using additional fields that are available when this option is selected. This option reinterprets the output bit pattern from the LSB up according to the specified type.</li> <li>• <b>Boolean:</b> the output type is Boolean.</li> <li>• <b>Single:</b> single floating-point data.</li> <li>• <b>Double:</b> double floating-point data.</li> </ul>
Output data type	Specifies the output data type. For example, <b>sfix(16)</b> , <b>uint(8)</b> .
Output scaling value	Specifies the output scaling value. For example, <b>2<sup>-15</sup></b> .
Number of delays	Specifies the number of samples to delay.
Minimum delay	Checks if the delay can grow as needed, so that the specified length becomes the lower bound.
Equivalence group	Sample delays that share the same equivalence group string grow by the same increment.

**Table 232. Port Interface for the SampleDelay Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
a	Input	Any fixed- or floating-point type	Data input	Yes	Yes
q	Output	Derived fixed- or floating-point type	Data output	Yes	Yes

### 14.4.51. Scalar Product

The **Scalar Product** block accepts two vector inputs of the same dimension and produces the inner product on the output. If one or more inputs are complex, the output is complex. If one of the inputs is a scalar signal, the same factor scales all vector components of the other input port.

*Note:* For single-precision inputs and designs targeting devices with floating-point DSP blocks, the block uses a mixture of resources including the device DSP blocks in floating-point mode.

**Table 233. Parameters for the Scalar Block**

Parameter	Description
Output data type mode	Determines how the block sets its output data type:
<i>continued...</i>	

Parameter	Description
	<ul style="list-style-type: none"> <li>• <b>Inherit via internal rule:</b> the number of integer and fractional bits is the maximum of the number of bits in the input data types.</li> <li>• <b>Specify via dialog:</b> you can set the output type of the block explicitly using additional fields that are available when this option is selected. This option reinterprets the output bit pattern from the LSB up according to the specified type.</li> <li>• <b>Inherit via internal rule with word growth:</b> the number of fractional bits is the maximum of the number of fractional bits in the input data types. The number of integer bits is the maximum of the number of integer bits in the input data types plus one. This additional word growth allows for subtracting the most negative number from 0, which exceeds the maximum positive number that you can store in the number of bits of the input.</li> <li>• <b>Boolean:</b> the output type is Boolean.</li> <li>• <b>Variable precision floating point:</b> variable precision floating-point output type.</li> </ul>
Output data type	Specifies the output data type. For example, <b>fixdt(1, 16, 15)</b> . Only available for <b>Specify via dialog</b>
Output scaling value	Specifies the output scaling value. For example, <b>2<sup>-15</sup></b> . Only available for <b>Specify via dialog</b>
Floating-point precision	Specifies a predefined floating-point type. Only available for <b>Variable precision floating point:</b>
Fused datapath	This option affects the floating-point architectures. Turn on this option to save hardware by omitting normalization stages between adder stages. The output deviates from that expected of IEEE compliance.
Floating Point Rounding	Specifies what rounding to apply to the result: <ul style="list-style-type: none"> <li>• Correct. IEEE compliant unbiased round to nearest output value.</li> <li>• Faithful. Saves hardware by sometimes rounding to the second nearest value. Error is about double that of correct rounding.</li> </ul>

**Table 234. Port Interface for the Scalar Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
a	Input	Any fixed- or floating-point type	Operand 1	Yes	Yes
b	Input	Any fixed- or floating-point type	Operand 2	Yes	Yes
q	Output	Derived fixed- or floating-point type	Result	No	Yes

### 14.4.52. Select

The **Select** block outputs one of the data signals (a, b, ...) if its paired select input (0, 1, ...) has a non-zero value.

$$q = 0 ? a : (1 ? b : d)$$

If all select inputs are 0, the **Select** block outputs the default value d. At most one select input should be high at a time.

**Table 235. Parameters for the Select Block**

Parameter	Description
Output data type mode	Determines how the block sets its output data type: <ul style="list-style-type: none"> <li>• <b>Inherit via internal rule:</b> the number of integer and fractional bits is the maximum of the number of bits in the input data types.</li> <li>• <b>Specify via dialog:</b> you can set the output type of the block explicitly using additional fields that are available when this option is selected. This option reinterprets the output bit pattern from the LSB up according to the specified type.</li> <li>• <b>Boolean:</b> the output type is Boolean.</li> </ul>
Output data type	Specifies the output data type. For example, <b>sfix(16)</b> , <b>uint(8)</b> .
Output scaling value	Specifies the output scaling value. For example, <b>2<sup>-15</sup></b> .
Number of cases	Specifies the number of non-default data inputs.

**Table 236. Port Interface for the Select Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
d	Input	Any fixed- or floating-point type	Default input	Yes	Yes
0, 1, 2, ...	Input	Boolean	One-hot select inputs	Yes	No
a, b, c, ...	Input	Any fixed- or floating-point type	Data input	Yes	Yes
q	Output	Derived fixed- or floating-point type	Result	Yes	Yes

### 14.4.53. Sequence

The **Sequence** block outputs a Boolean pulse of configurable duration and phase.

The input acts as an enable for this sequence. Usually, this block initializes with an array of Boolean pulses of length period. The first step\_value entries are zero, and the remaining values are one.

A counter steps along this array, one entry at a time, and indexes the array. The output value is the contents of the array. The counter is initialized to initial\_value. The counter wraps at step period, back to zero, to index the beginning of the array.

**Table 237. Parameters for the Sequence Block**

Parameter	Description
Output data type mode	Determines how the block sets its output data type:
<i>continued...</i>	



Parameter	Description
	<ul style="list-style-type: none"> <li>• <b>Inherit via internal rule:</b> the number of integer and fractional bits is the maximum of the number of bits in the input data types.</li> <li>• <b>Specify via dialog:</b> you can set the output type of the block explicitly using additional fields that are available when this option is selected. This option reinterprets the output bit pattern from the LSB up according to the specified type.</li> <li>• <b>Boolean:</b> the output type is Boolean.</li> </ul>
Output data type	Specifies the output data type. For example, <b>sfix(16), uint(8)</b> .
Output scaling value	Specifies the output scaling value. For example, <b>2<sup>-15</sup></b> .
Sequence setup	A vector that specifies the counter in the format: [ <b>&lt;initial_value&gt; &lt;step_value&gt; &lt;period&gt;</b> ] For example, [0 50 100]

**Table 238. Port Interface for the Sequence Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
a	Input	Boolean	Sequence enable	Yes	No
q	Output	Boolean	Result	Yes	No

#### 14.4.54. Shift

The **Shift** block outputs the logical right shifted version of the input value if unsigned, or outputs the arithmetic right shifted version of the input value if signed. The shift is specified by the input b:

$$q = (a \gg b)$$

The width of the data type b determines the maximum size of the shift.

Shifts of more than the input word width result in an output of 0 for non-negative numbers and (0 - 2<sup>-F</sup>) for negative numbers (where F is the fraction length).

**Table 239. Parameters for the Shift Block**

Parameter	Description
Output data type mode	Determines how the block sets its output data type: <ul style="list-style-type: none"> <li>• <b>Inherit via internal rule:</b> the number of integer and fractional bits is the maximum of the number of bits in the input data types.</li> <li>• <b>Specify via dialog:</b> you can set the output type of the block explicitly using additional fields that are available when this option is selected. This option reinterprets the output bit pattern from the LSB up according to the specified type.</li> <li>• <b>Boolean:</b> the output type is Boolean.</li> </ul>
Output data type	Specifies the output data type. For example, <b>sfix(16), uint(8)</b> .
Output scaling value	Specifies the output scaling value. For example, <b>2<sup>-15</sup></b> .

**Table 240. Port Interface for the Shift Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
a	Input	Any fixed-point type	Operand 1	Yes	No
b	Input	Unsigned integer	Operand 2	Yes	No
q	Output	Derived fixed-point type	Result	Yes	No

### 14.4.55. Sqrt

The **Sqrt** block applies a numerical root operation to its input and produces the result. The mask parameter pop-up menu selects the required root function that you apply.

**Table 241. Functions for the Sqrt Block**

Function	Description
cbrt(x)	Cube root.
recipsqrt(x)	Reciprocal square root.
sqrt(x)	Square root.

**Table 242. Parameters for Sqrt Block**

Parameter	Values	Description
<b>Function</b>	<b>cbrt(x)</b> , <b>recipsqrt(x)</b> , or <b>sqrt(x)</b>	Selects the numerical root function.
<b>Floating-point rounding</b>	<b>Correct</b> or <b>Faithful</b>	Only for sqrt(x) function.
<b>Advanced Options</b>	Blank or <code>struct('method', 256)</code>	The sqrt(x) function with integer input and output has two semantics: floor semantics, <code>floor(sqrt(x))</code> , for a logic reduction on wider data types or the default round-to-nearest semantics. To select floor semantics, type <code>struct('method', 256)</code> .

**Table 243. Port Interface for the Sqrt Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
x	Input	For, <b>recipsqrt(x)</b> , or <b>sqrt(x)</b> , any fixed- or floating-point 32-bit integer. For <b>cbrt(x)</b> floating-point input only.	Operand	Yes	No
q	Output	Output type depends on input type.	Result	Yes	No

### 14.4.56. Subtract (Sub)

The **Sub** block outputs the difference between the inputs:

$$q = a - b.$$

*Note:* For single-precision inputs and designs targeting any device with a floating-point DSP block, the block uses a mixture of resources including the DSP blocks in floating-point mode.

**Table 244. Parameters for the Sub Block**

Parameter	Description
Output data type mode	Determines how the block sets its output data type: <ul style="list-style-type: none"> <li>• <b>Inherit via internal rule:</b> the number of integer and fractional bits is the maximum of the number of bits in the input data types.</li> <li>• <b>Inherit via internal rule with word growth:</b> the number of fractional bits is the maximum of the number of fractional bits in the input data types. The number of integer bits is the maximum of the number of integer bits in the input data types plus one. This additional word growth allows for subtracting the most negative number from 0, which exceeds the maximum positive number that you can store in the number of bits of the input.</li> <li>• <b>Specify via dialog:</b> you can set the output type of the block explicitly using additional fields that are available when this option is selected. This option reinterprets the output bit pattern from the LSB up according to the specified type.</li> <li>• <b>Boolean:</b> the output type is Boolean.</li> </ul>
Output data type	Specifies the output data type. For example, <b>sfix(16), uint(8)</b> .
Rounding mode	Specifies what rounding to apply to the result: <ul style="list-style-type: none"> <li>• Correct. IEEE compliant unbiased round to nearest output value.</li> <li>• Faithful. Saves hardware by sometimes rounding to the second nearest value. Error is about double that of correct rounding.</li> </ul>
Output scaling value	Specifies the output scaling value. For example, <b>2<sup>-15</sup></b> .

**Table 245. Port Interface for the Sub Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
a	Input	Any fixed- or floating-point type	Operand 1	Yes	Yes
b	Input	Any fixed- or floating-point type	Operand 2	Yes	Yes
q	Output	Derived fixed- or floating-point type	Result	Yes	Yes

**Related Information**

[Forcing Soft Floating-point Data Types with the Advanced Options on page 220](#)

**14.4.57. Sum of Elements (SumOfElements)**

The **SumOfElements** block outputs the sum of the elements within its single data input.

$$q = \sum a_n$$

If the input is a scalar, the **SumOfElements** block outputs an unchanged value.

$$q = a$$

**Table 246. Parameters for the SumOfElements Block**

Parameter	Description
Output data type mode	Determines how the block sets its output data type: <ul style="list-style-type: none"> <li>• <b>Inherit via internal rule:</b> the number of integer and fractional bits is the maximum of the number of bits in the input data types.</li> <li>• <b>Inherit via internal rule with word growth:</b> the number of fractional bits is the maximum of the number of fractional bits in the input data types. The number of integer bits is the maximum of the number of integer bits in the input data types plus one. This additional word growth allows for subtracting the most negative number from 0, which exceeds the maximum positive number that you can store in the number of bits of the input.</li> <li>• <b>Specify via dialog:</b> you can set the output type of the block explicitly using additional fields that are available when this option is selected.</li> <li>• <b>Boolean:</b> the output type is Boolean.</li> </ul>
Output data type	Specifies the output data type. For example, <b>sfix(16), uint(8)</b> .
Output scaling value	Specifies the output scaling value. For example, <b>2<sup>-15</sup></b> .
Number of Inputs	Specifies the number of inputs.

**Table 247. Port Interface for the SumOfElements Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
a	Input	Any fixed-point type	Operand	Yes	Yes
q	Output	Derived fixed-point type	Result	No (scalar output only).	Yes

**Related Information**

[Forcing Soft Floating-point Data Types with the Advanced Options](#) on page 220

**14.4.58. Trig**

The **Trig** block applies a trigonometric operation to its floating-point inputs and produces the floating-point result.

*Note:* Your design may use up to 50% less resources if you use the pi functions.

**Table 248. Functions for the Trig Block**

Function	Description
acos(a)	Arc cosine (inverse cosine) output in radians.
asin(a)	Arc sine (inverse sine) output in radians.
atan(a)	Arc tangent (inverse tangent) output in radians.
acos(x)/pi	Arc cosine (inverse cosine) output as fraction of half circle.
asin(x)/pi	Arc sine (inverse sine) output as fraction of half circle.
atan(x)/pi	Arc tangent (inverse tangent) output as fraction of half circle.
atan2(y,x)	Four quadrant inverse tangent, output angle in interval [-n,+n] radians.
cos(a)	Cosine of input in radians.
cos(pi*x)	Cosine of input angle specified as fraction of half circle.
cot(a)	Cotangent of input in radians.
cot(pi*x)	Cotangent of input angle specified as fraction of half circle.
sin(a)	Sine of input in radians.
sin(pi*x)	Sine of input angle specified as fraction of half circle.
sincos(a)	Outputs both sine and cosine of input a in radians.
tan(a)	Tangent of input in radians.
tan(pi*x)	Tangent of input angle specified as fraction of half circle.

The **Function** parameter selects one of the 16 trigonometric functions. The number of input ports and output ports on the block change as required by the semantics of the function that you select:

- One-input and one-output: sin, cos, tan, cot, asin, acos, atan
- Two-inputs and one-output: atan2
- One-input and two-outputs: sincos

If you reduce the input range for the sin(x) and cos(x) functions to the interval  $[-2\pi, 2\pi]$ , and you target devices with floating-point DSP blocks, in **Advanced Options** set **struct('rangeReduction',0)**. The design then uses the floating-point mode of the DSP blocks to build more efficient architectures.

**Table 249. Port Interface for the Trig Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
x	Input	Single or double	Operand 1 (Operand 2 of atan2)	Yes	No
y	Input	Single or double	Operand 1 of atan2	Yes	No
q	Output	Single or double	Result 1	Yes	No
r	Output	Single or double	Result 2 (Cosine output of sincos)	Yes	No

### 14.4.59. XNOR Gate (Xnor)

The **Xnor** block outputs the logical **XNOR** of the input values:

$$q = \sim(a \text{ XOR } b)$$

Set the number of inputs to 1, to output the logical **XNOR** of all the individual bits of the input word.

**Table 250. Parameters for the Xnor Block**

Parameter	Description
Number of inputs	Specifies the number of inputs.
Output data type mode	Determines how the block sets its output data type: <ul style="list-style-type: none"> <li><b>Inherit via internal rule:</b> the number of integer and fractional bits is the maximum of the number of bits in the input data types.</li> <li><b>Specify via dialog:</b> you can set the output type of the block explicitly using additional fields that are available when this option is selected. This option reinterprets the output bit pattern from the LSB up according to the specified type.</li> <li><b>Boolean:</b> the output type is Boolean.</li> </ul>
Output data type	Specifies the output data type. For example, <b>sfix(16)</b> , <b>uint(8)</b> .
Output scaling value	Specifies the output scaling value. For example, <b>2<sup>-15</sup></b> .

**Table 251. Port Interface for the Xnor Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
a	Input	Any fixed-point type	Operand 1	Yes	No
b	Input	Any fixed-point type	Operand 2	Yes	No
q	Output	Derived fixed-point type	Result	Yes	No

### 14.4.60. XOR Gate (Xor)

The **Xor** block outputs the logical **XOR** of the input values:

$$q = (a \text{ XOR } b)$$

Set the number of inputs to 1, to output the logical **XOR** of all the individual bits of the input word.

**Table 252. Parameters for the Xor Block**

Parameter	Description
Number of inputs	Specifies the number of inputs.
Output data type mode	Determines how the block sets its output data type:
<i>continued...</i>	

Parameter	Description
	<ul style="list-style-type: none"> <li>• <b>Inherit via internal rule:</b> the number of integer and fractional bits is the maximum of the number of bits in the input data types.</li> <li>• <b>Specify via dialog:</b> you can set the output type of the block explicitly using additional fields that are available when this option is selected. This option reinterprets the output bit pattern from the LSB up according to the specified type.</li> <li>• <b>Boolean:</b> the output type is Boolean.</li> </ul>
Output data type	Specifies the output data type. For example, <b>sfix(16)</b> , <b>uint(8)</b> .
Output scaling value	Specifies the output scaling value. For example, <b>2<sup>-15</sup></b> .

**Table 253. Port Interface for the Xor Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
a	Input	Any fixed-point type	Operand 1	Yes	No
b	Input	Any fixed-point type	Operand 2	Yes	No
q	Output	Derived fixed-point type	Result	Yes	No

## 14.5. Primitive Configuration Library

Use the DSP Builder advanced blockset **Primitive Configuration** library blocks to implement blocks that change how DSP Builder synthesizes primitive subsystems, including boundary delimiters.

1. [Channel In \(ChannelIn\)](#) on page 359
2. [Channel Out \(ChannelOut\)](#) on page 360
3. [General Purpose Input \(GPIn\)](#) on page 361
4. [General Purpose Output \(GPOut\)](#) on page 361
5. [Synthesis Information \(SynthesisInfo\)](#) on page 362

### 14.5.1. Channel In (ChannelIn)

The **ChannelIn** block delineates the input boundary of a DSP Builder synthesizable **Primitive** subsystem.

The **ChannelIn** block passes its input through to the outputs unchanged, with types preserved. This block indicates to DSP Builder that these signals arrive synchronized from their source, so that the synthesis tool can interpret them.

**Table 254. Parameters for the ChannelIn Block**

Parameter	Description
Number of data signals	Specifies the number of data signals on this block.

**Table 255. Port Interface for the ChannelIn Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
v	Input	Boolean	Valid input signal	No	No
c	Input	uint(8)	Channel input signal	No	No
d0, 1, 2, ...	Input	Any fixed-or floating-point type	A number of input data signals	Yes	Yes
v	Output	Boolean	Valid signal	No	No
c	Output	uint(8)	Channel signal	No	No
q0, 1, 2, ...	Output	Any fixed- or floating-point type	A number of data signals	Yes	Yes

### 14.5.2. Channel Out (ChannelOut)

The **ChannelOut** block delineates the output boundary of a DSP Builder synthesizable **Primitive** subsystem.

The **ChannelOut** block passes its input through to the outputs unchanged, with types preserved. This block indicates to DSP Builder that these signals must synchronize, which the synthesis tool can ensure.

When you run a simulation in Simulink, DSP Builder adds additional latency from the balanced pipelining stages to meet the specified timing constraints for your model. The block accounts for this additional latency. This latency does not include any delay explicitly added to your model, by for example a **SampleDelay** block, just added pipelining for timing closure.

*Note:* You can also access the value of the latency parameter by typing a command of the following form on the MATLAB command line:

```
get_param(gcf,'latency')
```

**Table 256. Parameters for the ChannelOut Block**

Parameter	Description
Number of data signals	Specifies the number of data signals on this block.

**Table 257. Port Interface for the ChannelOut Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
v	Input	Boolean	Valid output signal	No	No
c	Input	8-bit unsigned integer	Channel output signal	No	No
d0, d1, d2, ...	Input	Any fixed-or floating-point type	A number of output data signals	Yes	Yes
<i>continued...</i>					



Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
v	Output	Boolean	Valid signal	No	No
c	Output	8-bit unsigned integer	Channel signal	No	No
q0, q1, q2, ...	Output	Any fixed-or floating-point type	A number of data signals	Yes	Yes

### 14.5.3. General Purpose Input (GPIn)

The **GPIn** block models a general purpose input to a synthesizable subsystem. It is similar to the ChannelIn block but has no `valid` or `channel` inputs.

If the signal width is greater than one, you can assume the multiple inputs are synchronized.

**Table 258. Parameters for the GPIn Block**

Parameter	Description
Number of data signals	Specifies the number of input and output signals.

**Table 259. Port Interface for the GPIn Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
a, b, ...	Input	Any fixed-point type	Operands 1 to n	Yes	Yes
a, b, ...	Output	Same type as input	Data is passed through unchanged.	Yes	Yes

### 14.5.4. General Purpose Output (GPOut)

The **GPOut** block models a general purpose output to a synthesizable subsystem. It is similar to the **ChannelOut\_help** block but has no `valid` or `channel` inputs.

If the width is greater than one, the multiple outputs generate and are synchronized.

**Table 260. Parameters for the GPOut Block**

Parameter	Description
Number of data signals	Specifies the number of input and output signals.

**Table 261. Port Interface for the GPOut Block**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
a, b, ...	Input	Any fixed- or floating-point type	Operands 1 to n	Yes	Yes
a, b, ...	Output	Same type as input	Data is passed through unchanged.	Yes	Yes

### 14.5.5. Synthesis Information (SynthesisInfo)

Use the **SynthesisInfo** block to set the synthesis mode and label a primitive subsystem as the top-level synthesizable subsystem. DSP Builder flattens and synthesizes the subsystem, and all those subsystems below as a unit. Primitive subsystems must have a **SynthesisInfo** block. DSP Builder creates pipelines and redistributes memories optimally to achieve the desired clock frequency. The **SynthesisInfo** block controls the synthesis flow for the current model.

The inputs and outputs to this subsystem become the primary inputs and outputs of the RTL entity that DSP Builder creates.

The **SynthesisInfo** block can be at the same level as the **Device** block (if the synthesizable subsystem is the same as the generated hardware subsystem). However, it is often convenient to create a separate subsystem level that contains the **Device** block. Refer to the design examples for some examples of design hierarchy.

**Table 262. Parameters for the SynthesisInfo Block**

Parameter	Description
Constrain Latency	<p>This option allows you to select the type of constraint and to specify its value. The value can be a workspace variable or an expression but must evaluate to a positive integer.</p> <p>You can select the following types of constraint:</p> <ul style="list-style-type: none"> <li>• &gt;: Greater than</li> <li>• &gt;=: Greater than or equal to</li> <li>• =: Equal to</li> <li>• &lt;=: Less than or equal to</li> <li>• &lt;: Less than</li> </ul> <p>Select either + or - and type in a reference model in the text field. Specify the reference as a Simulink path string e.g. 'design/topLevel/model'. DSP Builder then ensures the latency depends on that model, otherwise the default is that DSP Builder depends on no model.</p> <p>Constrain Latency only applies to subsystems which use the <b>ChannelIn</b> or <b>ChannelOut</b> blocks and not to subsystems that use the <b>GPIn</b> or <b>GPOut</b> blocks.</p>
Bit accurate simulation	<p>Turn on in floating-point designs to give bit accurate rather than mathematical simulations. Fixed point designs always use bit accurate.</p>
Local reset minimization	<p>Select the reset minimization for the associated synthesizable subsystem. Valid only if <b>Control</b> block <b>Global Enable</b> is <b>On</b>.</p> <p>The default is <b>Conditional – On for ChannelIn/Out only</b>.</p> <p>Select <b>Off</b> to disable reset minimization on this synthesizable subsystem.</p> <p>Select <b>On – Always (for ChannelIn/Out or GPIn/Out)</b> to apply reset minimization to a synthesizable subsystem that uses GPIn/Out blocks. In a GPIn/Out subsystem with reset minimization, the whole subsystem is data flow and has no valid signal to be control flow.</p>

The **SynthesisInfo** block has no inputs or outputs.

1. [Scheduled Synthesis](#) on page 362
2. [Updated Help](#) on page 363

#### Related Information

[Reset Minimization](#) on page 197

#### 14.5.5.1. Scheduled Synthesis

The **Scheduled** style of operation uses a pipelining and delay distribution algorithm that creates fast hardware implementations from an easily described untimed block diagram. This style takes full advantage of the automatic pipelining capability.

The algorithm performs the following operations:

1. Reads in and flattens your design example for any subsystem that contains a **SynthesisInfo** block.
2. Builds an internal graph to represent the logic.
3. Based on the absolute clock frequency requested, adds enough pipeline stages to meet that clock frequency. For example, you may pipeline long adders into several shorter adders. This additional pipelining helps reach high clock frequencies.

#### 14.5.5.2. Updated Help

After you run a simulation, DSP Builder updates the help pages with specific information about each instance of a block. This updated help overrides the default help link. To find the updated help, click on the help link on the block after simulation. This updated help includes a link back to the help for the general block and the following information about the generated instance:

- Date and time of generation
- The latency introduced by this block.
- Port interface table.

### 14.6. Primitive Design Elements Library

Use the DSP Builder advanced blockset **Primitive Design Elements** library blocks to implement configurable blocks and common design patterns built from primitive blocks.

1. [Anchored Delay](#) on page 364
2. [Complex to Real-Imag](#) on page 364
3. [Enabled Delay Line](#) on page 364
4. [Enabled Feedback Delay](#) on page 364
5. [Expand Scalar \(ExpandScalar\)](#) on page 364
6. [Finite State Machine](#) on page 364
7. [Nested Loops \(NestedLoop1, NestedLoop2, NestedLoop3\)](#) on page 370
8. [Pause](#) on page 371
9. [Reset-Priority Latch \(SRlatch\\_PS\)](#) on page 372
10. [Same Data Type \(SameDT\)](#) on page 372
11. [Set-Priority Latch \(SRlatch\)](#) on page 372
12. [Single-Cycle Latency Latch \(latch\\_1L\)](#) on page 373
13. [Tapped Line Delay \(TappedLineDelay\)](#) on page 373
14. [Variable Super-Sample Delay \(VariableDelay\)](#) on page 373
15. [Vector Fanout \(VectorFanout\)](#) on page 374
16. [Vector Multiplexer \(VectorMux\)](#) on page 374
17. [Zero-Latency Latch \(latch\\_0L\)](#) on page 374

### 14.6.1. Anchored Delay

DSP Builder **SampleDelay** blocks are often not suitable in FSMs, which are common in control unit designs. To ensure that DSP Builder's simulation of FSMs matches the synthesized hardware, use the **Anchored Delay** block not the **SampleDelay** block.

The **Anchored Delay** block has a data input and a valid input port. Connect the valid input port to the valid in of the enclosing **Primitive** subsystem to allow DSP Builder to correctly schedule the starting state of your control unit design.

### 14.6.2. Complex to Real-Imag

The DSP Builder **Complex to Real-Imag** block handles custom data types that the enhanced precision floating-point types support.

### 14.6.3. Enabled Delay Line

The DSP Builder **Enabled Delay Line** block takes a single data signal  $a$  and an enable signal  $e$  and implements an enabled delay line, with  $q$  as the delayed data output. Internally, the block is either a Latch\_0L or Latch\_1L (depending on the **Zero or one initial delay** parameter) followed by a series of Latch\_1Ls. The final output connects to the output port  $q$ . When you use the block in a feedback loop, DSP Builder cannot redistribute the enabled sample delays around the feedback path. In these instances, use the **Enabled Feedback Delay** block."

### 14.6.4. Enabled Feedback Delay

The DSP Builder **Enabled Feedback Delay** block takes a single data signal  $a$  and an enable signal  $e$  and implements an enabled delay, with  $q$  as the delayed data output. The block is a non-enabled **SampleDelay** block followed by a FIFO buffer. When you use the block in a feedback loop, DSP Builder can distribute the non-enabled sample delay around the feedback path, while retaining the right enabled feedback behavior.

### 14.6.5. Expand Scalar (ExpandScalar)

The DSP Builder **ExpandScalar** block takes a single connection and replicates it  $N$  times to form a width  $N$  vector of duplicate signals. The block passes on the width parameter to a Simulink multiplexer under the mask, and uses some standard Simulink commands to add the connections lines.

### 14.6.6. Finite State Machine

The **Finite State Machine** block adds state machines to a DSP Builder design. You can describe your **Finite State Machine** using the FSM specification language, which you enter into a text file. Then you load that text file using the **Finite State Machine** block and wire up the input and output ports of the **Finite State Machine** block to the rest of your design. DSP Builder then generates appropriate loop (**ForLoop**) and lookup-table (**LUT**) structures to implement the state-machine described in the text file you provide. However, those **ForLoop** and **LUT** blocks are not directly visible in the DSP Builder design. DSP Builder translates the **ForLoop** and **LUT** blocks to RTL with automatic device mapping and latency balancing with the rest of your DSP Builder design. DSP Builder provides a finite state machine example design, `demo_fsm.mdl`, which shows how to use the **Finite State Machine** block to filter out specific numerical values from an input stream.

## FSM Specification Language

The FSM specification language supports the following features:

- Boolean expressions to form state transition style FSMs
- Nested, dependent, loops which allow you to express, for example, rectangular and triangular counters
- Automatic dead-cycle hiding at loop completion

The configuration text file specifies the **Finite State Machine** block to generate.

The file has two sections:

- Options that are optional settings to override certain defaults such as the name of the XML file and external port names
- Netlist components that define the **ForLoop** blocks and state transitions, and how they fit together (i.e. nesting or sequencing structure) that defines the behaviour of the state machine

The keyword `netlist` separates these two sections:

```
option settings ...
netlist...
netlist components ...
```

### Options

Table 263. Options

Option	Description
<code>inputs &lt;port-name&gt;</code>	Specify one or more boolean input ports that DSP Builder uses in conditional expressions for state transitions.
<code>enable &lt;port-name&gt;</code>	Add an optional enable with the specified port name. Default is to tie enable to $V_{CC}$ .
<code>start &lt;port-name&gt;</code>	Specify a different name for the <code>go</code> port. A one-cycle pulse on this port starts the <b>Finite State Machine</b> block.
<code>finish &lt;port-name&gt;</code>	Add a finish port with the name as specified. The port emits a one-cycle pulse from this port when the <b>Finite State Machine</b> completes. This port allows two or more <b>Finite State Machine</b> blocks to daisy chain, or to nest an <b>Finite State Machine</b> in a <b>ForLoop</b> block.

Figure 102. Options Example

```
# this is a comment
inputs d0 d1
enable en
start begin
finish done
netlist
...
```

## Netlist Components

**Figure 103. Netlist Components Example**

Instantiates a **ForLoop** block *<name>* and inequality operator set to *<*, *<=*, *>*, or *>=*. The values that you specify for *<init>*, *<step>*, and *<limit>* drive the input ports of the **ForLoop** block.

```
for <name> <init> <inequality-op> <limit> step <step>: <port-name> ...
... nested components
end
```

If *<step>* is 1, you can omit it. For example:

```
for x 0 < 8
end
```

The port-names are optional. If specified, they must correspond to the names of **ForLoop** output ports:

Name	Output Port
bs	Body Start
ld	Loop Done
el	Empty Loop
fl	First Loop
ll	Last Loop
v	Valid
c	Counter

For example:

```
for x 8 >= 0 step -1 : ld v c
end
```

You can nest **ForLoop** blocks:

```
for x 8 >= 3 step -1
  for y 0 < x
  end
end
```

The inner **ForLoop** *<limit>* is specified to be the counter value of the outer **ForLoop**

You can select state transitions:

```
transitions <name> : <port-name> ...
... state transitions
end
```

You specify state transitions by the following inputs:

- state *<current-state>*
- if (*<condition>*) *<next-state>* *<output-value>* ...
- Each state definition can be followed by one or more if transition statements. The last transition of a state definition can be a catch-all default transition:
- default *<next-state>* *<output-value>* ...

The output-values are a list of boolean values (0 or 1). There must be the same number of output-values as port-names specified on the `transitions` line. For example:

```
transitions Sticky : q r
state Start
if (a) Next 1 0
default Start 1 1
state Next
if (~a) Start 0 1
end
```

If you omit `default`, the **Finite State Machine** assumes an implicit catch-all default transition where the next-state is the same as the current-state and all output-values are 0.

You can specify a `finish` port for a state transition style **Finite State Machine** that produces a pulse for one cycle when the **Finite State Machine**, after it is activated on receipt of the token (or go signal), returns back to its initial state on completion. This corresponds to passing on the token to the next **Finite State Machine** component in the sequence.

```
transitions Wait : q r
finish waitDone
... state transitions
end
```

You can freely mix `transitions...end` **Finite State Machine** blocks with `for...end` **Finite State Machine** blocks to form a **Finite State Machine** sequence. Also you can nest these **Finite State Machine** sequences inside a `for...end` **Finite State Machine** that may be a member of another **Finite State Machine** sequence. The `transitions` **Finite State Machine** does not support you nesting a **Finite State Machine** sequence inside it.

### Nested ForLoops

Figure 104. Nested ForLoops

```
for row 0 < 8 : c
  for col 0 < 8 : c v

  end

end
```

### Triangular Nested Iteration

Figure 105. Triangular nested iteration

```
for row 1 < 8 : c
  for col 0 < row : c v

  end

end
```

## Strategies for Hiding Dead Cycles

Figure 106. `deadcycle endearly`

```
for row 0 < 8 : c
    deadcycle endearly
    for col 0 < 8 : c v
        end
    end
end
```

The `deadcycle endearly` strategy is a simple transformation of the netlist that reduces the number of iterations of the innermost loop by one. The example is functionally equivalent to:

```
for row 0 < 8 : c
    for col 0 < 7 : c v
        end
    end
end
```

This strategy only makes sense if the innermost **ForLoop** body is empty or has zero latency, so that each iteration takes one cycle to compensate for the one cycle spent in the outer loop iteration. The control signals also reflect the shortened iteration space such that the last-loop and loop-done signals go high one cycle earlier. You need to compensate for this in your design.

Figure 107. `deadcycle shiftrange`

```
for x 0 < 4
    deadcycle shiftrange

    for y 8 >= 1 step -1
        deadcycle shiftrange

        for z 0 < y : v c
            end
        end
    end
end
```

The `deadcycle shiftrange` strategy overlaps each iteration of the outer **ForLoop** with the first iteration of the inner **ForLoop**. However, in triangular nested iteration, the inner **ForLoop** needs to access the outer **ForLoop** block's counter value before it completes its own iteration cycle. To resolve this dependency, the outer **ForLoop** iterates over the range from  $(init+step)$  to  $(limit+step)$ , i.e. its range is shifted



forward by one iteration step. This transformation is only applicable if the outer **ForLoop** block's parameters are all constants. The following construction does not work:

```
for x 0 < 4
  for y 8 >= x step -1
    deadcycle shiftrange
  for z 1 < 5
```

The y-**ForLoop** block's limit depends on the x-**ForLoop** block's counter. It is not trivial to shift the range of the y-**ForLoop** without affecting other **ForLoop** blocks.

**Figure 108. deadcycle oneahead**

```
for x 0 < 4
  for y 8 >= x step -1
    deadcycle oneahead
    for z 0 <= y : v c
  end
end
end
```

The `deadcycle_oneahead` strategy is similar to the `shiftrange` strategy but the FSM does not require the outer **ForLoop** block's parameters to be constants. In the example, the y-**ForLoop** block's limit is the x-**ForLoop** block's counter. The netlist transformation uses more elaborate control logic that pre-advances the outer **ForLoop** by one iteration before initialising the inner **ForLoop**. This approach doesn't work for deeper nesting structures such as:

```
for x 0 < 4
  deadcycle oneahead
  for y 8 >= 1 step -1
    deadcycle oneahead
    for z 0 < y : v c
```

Although the FSM hides the dead-cycle in the x-**ForLoop**, the FSM inserts an extra cycle to pre-advance the y-**ForLoop** when applying its dead-cycle hiding strategy.

You can mix and match the different dead-cycle hiding strategies

### Related Information

[DSP Builder Finite State Machine Design Example](#) on page 108

### 14.6.7. Nested Loops (NestedLoop1, NestedLoop2, NestedLoop3)

The DSP Builder **NestedLoop1**, **NestedLoop2**, and **NestedLoop3** blocks maintain a set of counters that implement the equivalent of a nested for loop in software. They provide more flexibility than the **Loop** block and greater readability and lower latency than **ForLoop** blocks.

DSP Builder implements the **NestedLoop** blocks as masked subsystems and use existing DSP Builder **Primitive** library blocks. They do not have fixed implementations. DSP Builder generates a new implementation at runtime whenever you change any of the loop specifications.

For each loop in a **NestedLoop** block, you can specify start, increment, and end expressions. Each of these expressions may have one of the following three forms:

- A constant expression that evaluates (in the MATLAB base environment) to an integer. For example, if the MATLAB variable N has the value 256,  $(\log_2(N)+1)$  is a legal expression (and evaluates to 9).
- An instance of the loop variable controlling an enclosing loop. For example, you can use "i" (the outer loop variable) as the start expression of the "j" or "k" loops.
- A port name, optionally accompanied by a width specification in angle brackets. For example "p" or "q<4>". If no width is specified, it defaults to 8. This option generates a new input port (with the user-defined name and width) on the **NestedLoop** block.

For a **NestedLoop2** block, with user-supplied start, increment, and end expressions of S1, I1 and E1 (for the outer loop) and S2, I2 and E2 (for the inner loop), the equivalent C++ code is:

```
int i = S1;
do {
    int j = S2;
    do {
        j += I2;
    } while (j != E2);
    i += I1;
} while (i != E1);
```

Each **NestedLoop** block has two fixed input ports (*go* and *en*) and a variable number of additional user-defined input ports. DSP Builder regards each user-defined port as a signed input.

Each block also has two fixed output ports (*qv* and *q1*) and one (*NestedLoop1*), two (*NestedLoop2*) or three (*NestedLoop3*) output ports for the counter values.

When the input *en* signal is low (inactive), the output *qv* (valid) signal is also set low. The state of the **NestedLoop** block does not change, even if it receives a *go* signal.

Normal operation occurs when the *en* signal is high. The **NestedLoop** block can be in the waiting or counting state.

The **NestedLoop** block resets into the waiting state and remains there until it receives a *go* signal. While in the waiting state, the *qv* signal is low and the value of the other outputs are undefined.

When the block receives a `go` signal, the **NestedLoop** block transitions into the counting state. The counters start running and the `qv` output signal is set high. When all the counters eventually reach their final values, the `ql` (last cycle) output becomes high. On the following cycle, the **NestedLoop** block returns to the waiting state until it receives another `go` signal.

If the block receives a `go` signal while the **NestedLoop** block is already in the counting state, it remains in the counting state but all its counters are reset to their start values.

Observe the following points:

- All counters in the **NestedLoop** block are signed. To effectively use unsigned counters, zero-extend any unsigned inputs (and correspondingly increase their width specifications) by one bit.
- The end test is an equality test. Each loop continues until the current value of the counter is equal to the end value. However, if the loop counter overflows, the subsequent behavior of the **NestedLoop** block is undefined.
- The end values are inclusive. So setting the start value to 0, the increment to 1 and the end value to 10 actually produces 11 iterations of the loop.
- The previous two factors means that every loop iterates at least once. **NestedLoop** blocks (unlike **ForLoop** blocks) do not support empty loops.
- When you use user-defined ports to supply loop control values, the values on these ports must be held constant while the **NestedLoop** block is in its counting state. Otherwise the block produces undefined behavior.

**Table 264. NestedLoop Block Port interface**

Signal	Direction	Type	Description	Vector Data Support	Complex Data Support
<code>go</code>	Input	Boolean	Go.	No	No
<code>en</code>	Input	Boolean	Enable.	No	No
<code>?</code>	Input	Signed integer	Loop control values.	No	No
<code>qv</code>	Output	Boolean	Valid.	No	No
<code>ql</code>	Output	Boolean	Last iteration flag.	No	No
<code>i</code>	Output	Signed integer	Outer loop count.	No	No
<code>j</code>	Output	Signed integer	Middle loop count. ( <code>NestedLoop2</code> only).	No	No
<code>k</code>	Output	Signed integer	Inner loop count ( <code>NestedLoop3</code> only).	No	No

### 14.6.8. Pause

The **Pause** block implements a breakpoint with trigger count to break and single step through Simulink simulations.

Input is a Boolean signal enabling a counter, which counts up to the parameterized count value, then pauses the simulation. Press play to resume simulation from that point. If you do not increase the trigger count, the block causes another pause. This single stepping mode pauses the simulation at the next cycle. The block is red when enabled.

Use the **Pause** block for debugging designs. For example; run to breakpoint, then turn on **Show port values when hovering** at this point. This option permanently causes slow simulation, so only turn on when stepping through. Using display blocks allows you to see variables displayed at the paused time (similar to watch variables in a software debugger). You can change the trigger count, for example by adding 100, to simulate the next 100 cycles. The block color changes to red when you turn on the **Pause** block. Use the `valid` signal as the input, so that it counts valid steps only. Alternatively, use a different control signal, such as a `writenable`, as an input to get the system to break then. You can easily add other logic blocks to generate a break signal that you can use just for debugging.

### 14.6.9. Reset-Priority Latch (SRIatch\_PS)

DSP Builder offers two single-cycle latency latch subsystems for common operations for the valid signal, latching with set and reset. The **SRIatch** block gives priority to the reset input signal; the **SRIatch\_PS** block gives priority to the set input signal. In both blocks, if set and reset inputs are both zero the current output state is maintained.

**Table 265. Truth Table for SRIatch\_PS**

S	R	Q
0	0	Q
1	0	1
0	1	0
1	1	0

### 14.6.10. Same Data Type (SameDT)

The DSP Builder **SameDT** block duplicates the data type of the signal on port a onto the data type of signal on port b. This block is equivalent to the Simulink Data Type Duplicate block, but it also handles custom data types, such as the DSP Builder enhanced precision and wide integer data types.

### 14.6.11. Set-Priority Latch (SRIatch)

DSP Builder offers two single-cycle latency latch subsystems for common operations for the valid signal, latching with set and reset. The **SRIatch** block gives priority to the reset input signal. The **SRIatch\_PS** block gives priority to the set input signal. In both blocks if set and reset inputs are both zero the current output state is maintained.

**Table 266. Truth Table for SRLatch**

S	R	q
0	0	q
1	0	1
0	1	0
1	1	1

### 14.6.12. Single-Cycle Latency Latch (latch\_1L)

The DSP Builder **latch\_1** block enable signal affects the output on the following clock cycle.

These latches work for any data type, and for vector and complex numbers.

Right-click on the block and select **Look Under Mask**, for the structure.

The *e* signal is a `ufix(1)` enable signal. When *e* is high, the **latch\_1** block delays data from input *d* by one cycle and feeds through to output *q*. When *e* is low, the **latch\_1** block holds the last output.

A switch in *e* means the **latch\_1** block holds the output one cycle later.

### 14.6.13. Tapped Line Delay (TappedLineDelay)

The DSP Builder **TappedDelayLine** block takes a single scalar signal *a* and an enable signal *e* and produces an enabled tapped delay line of signals as a vector, *q*, and a corresponding valid signal, *qv*. The block is either a `Latch_0L` or `Latch_1L` (depending on the **Zero or One Initial Delay** parameter) followed by a series of `Latch_1L`s. The block forms a vector from all the latch outputs. By default, the vector has the least delayed signal at the top (unless you turn on **Reverse order** parameter). This block uses latches from the Control library. DSP Builder does not support vector signal input for this block.

### 14.6.14. Variable Super-Sample Delay (VariableDelay)

The DSP Builder **VariableDelay** blocks provides a sample delay for super-sample data where multiple data samples arrive per clock cycle. For multiple data per clock cycle, a delay of one sample shifts the signals across the wires, and uses a register delay for just the one signal that wraps round to the beginning on the next cycle.

For example:

[1 2 3 4]' [5 6 7 8]' [9 10 11 12]'...

when delayed by one sample becomes

\n[X 1 2 3]' [4 5 6 7]' [8 9 10 11]' [12 ...

when delayed by two samples becomes

[X X 1 2]' [3 4 5 6]' [7 8 9 10]' [11 12 ...

where [1 2 3 4]<sup>1</sup> means 1,2,3 and 4 arrive in parallel on 4 separate wires. The **Phases** parameter specifies the number of parallel data samples. The delay input must be a unsigned integer less than the number of parallel data samples.

### 14.6.15. Vector Fanout (VectorFanout)

The **VectorFanout** block behaves like a wire, connecting its single input to one vector output that contains several copies of the single input. The **Fanout** and **VectorFanout** are similar blocks. For a description of the **VectorFanout** block, refer to the **Fanout** block.

**Table 267. Parameters for the VectorFanout Block**

Parameter	Type	Description
Vector width	Integer > 0	Vector width of output port.
Allow use of uninitialized registers	Check box	Turn on to allow DSP Builder to use hyper registers. DSP Builder does not initialize the inserted routing registers on reset

**Table 268. Port Interface for the VectorFanout Block**

Signal	Direction	Type	Description
d	Input	Any scalar	Input
q	Output	Vector of d's input type	(Vector width) copies of d.

### 14.6.16. Vector Multiplexer (VectorMux)

The DSP Builder **VectorMux** block dynamically selects a single scalar signal from an input vector of signals. If *n* is a vector of width *N*, *sel* takes the range [0:*N*-1] and the block produces the (*sel*)th signal from the vector.

This block is an autogenerating masked subsystem that **Primitive** library blocks build. Internally, it is a demultiplexer and multiplexer, but parameterizable such that you do not have to manually draw and reconnect the connections between the demultiplexer and multiplexer if the vector width parameter changes.

### 14.6.17. Zero-Latency Latch (latch\_0L)

The DSP Builder **latch\_0** block enable signal has an immediate effect on the output. While the enable is high, the data passes straight through. When the enable goes low, the **latch\_0** block outputs and holds the data input from the previous cycle.

The *e* signal is a `ufix(1)` enable signal. When *e* is high, the **latch\_0** block feeds data from input *d* through to output *q*. When *e* is low, the **latch\_0** block holds the last output.

A switch in *e* is effective immediately.

## 15. Utilities Library

---

The DSP Builder advanced blockset **Utilities** library contains miscellaneous blocks that support building and refining designs.

1. [Analyze and Test Library](#) on page 375

### 15.1. Analyze and Test Library

1. [Capture Values](#) on page 375
2. [HDL Import](#) on page 375
3. [HDL Import Config](#) on page 377
4. [Pause](#) on page 371

#### 15.1.1. Capture Values

The **Capture Values** block can capture a variable number of signal inputs and supports vector and complex types.

You can add the block anywhere in the Simulink design. The block only supports the `.vcd` file format. DSP Builder writes this file in the RTL directory and it derives its name from the name given to the block. The specific arrangement of `.vcd` is based on what ModelSim writes out - i.e. only Boolean wires are used. You can import it into ModelSim using the `vcd2wlf` tool. The waveforms should match with those generated by the HDL simulation, although you might see an offset because of the Simulink model latency correction.

#### 15.1.2. HDL Import

You can import VHDL, Verilog HDL, and System Verilog into DSP Builder designs when you add a HDL Import block to your design. You can only configure **HDL Import** blocks after you configure the **HDL Import Config** block.

Figure 109. HDL Import Block Parameters

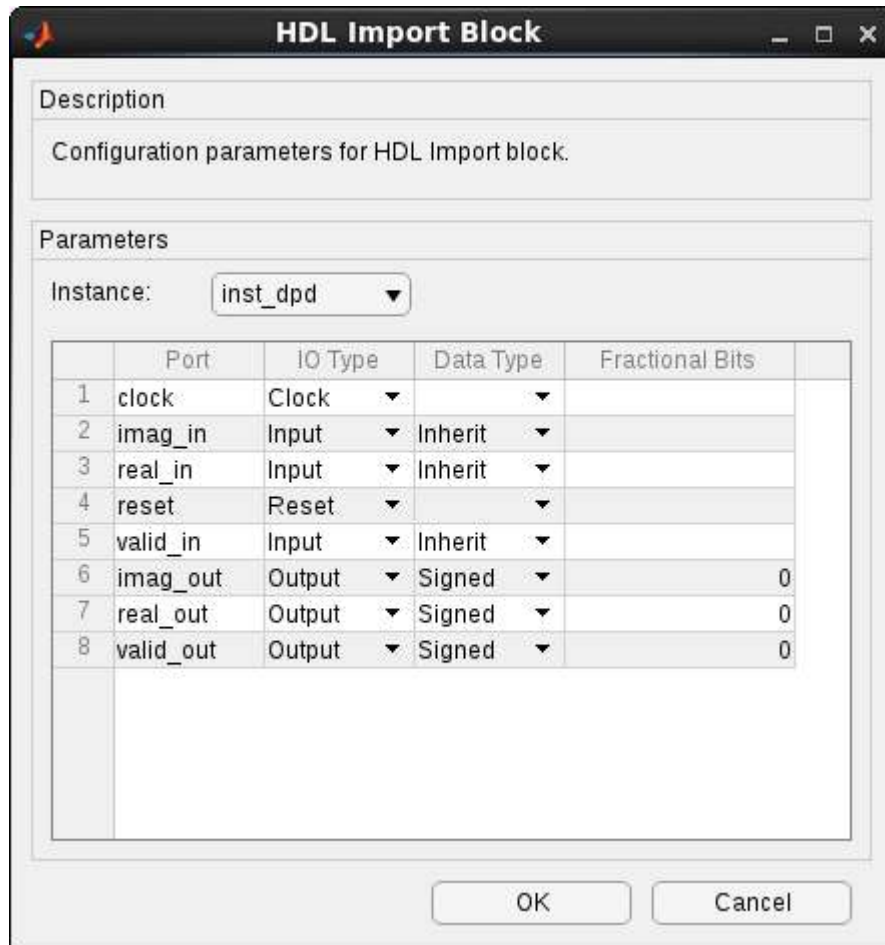


Table 269. HDL Import Parameters

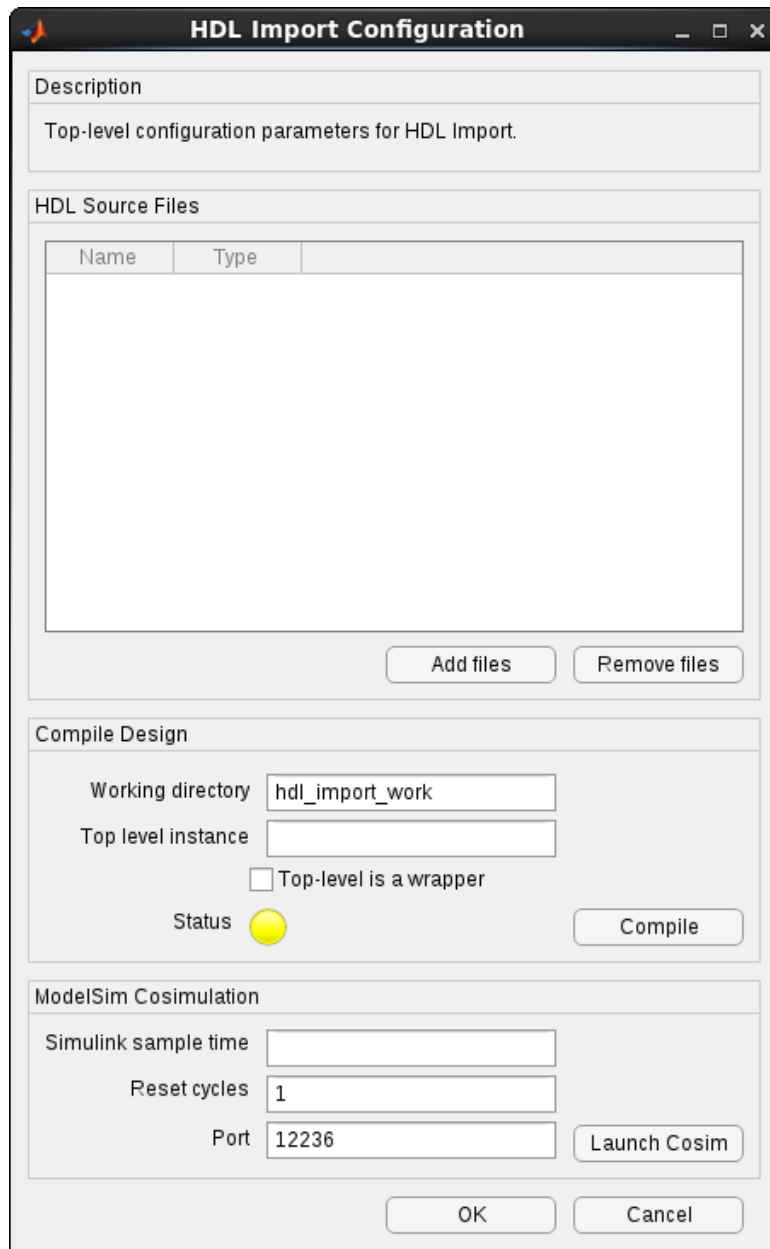
Parameter	Description
<b>Instance</b>	Select from any instance in your imported HDL. Each <b>HDL Import</b> block must represent a unique instance.
Port	DSP Builder automatically populates this column.
I/O Type	DSP Builder determines the IO type based on the name of the port. You can change any entry to <b>Input</b> , <b>Output</b> , <b>Clock</b> , or <b>Reset</b> . <b>HDL Import</b> only allows one clock and one reset.
Data Type	Informs Simulink and DSP Builder how they should interpret the ModelSim data. Set the <b>Data Type</b> of inputs to <b>Inherit</b> ; the <b>Data Type</b> of outputs defaults to <b>Signed</b> . For Boolean or std_logic data type, select <b>Unsigned with 0 fractional bits</b> .



### 15.1.3. HDL Import Config

The **HDL Import Config** block contains the top-level information to implement the HDL import feature. To use the HDL import feature, you must have one **HDL Import Config** block be at the top level of your design.

Figure 110. HDL Import Config



The HDL import feature needs the time relationship between ModelSim and Simulink. ModelSim uses the **Control** block-defined clock rate.

**Table 270. HDL Import Parameters**

Parameter	Description
Working Directory	DSP Builder creates this working directory for the ModelSim library and other intermediate files.
Top-level instance	Enter the name of the top-level instance. If that instance is not the HDL you want to import but a wrapper for multiple instances, turn on <b>Top-level is a wrapper</b> .
Compile	Click to compile imported RTL. Reclick if the imported RTL changes. For the cosimulation, DSP Builder creates a ModelSim library and performs ModelSim compilation followed by a series of Quartus synthesis compilations, one for each imported instance. The output from the compilations, including any errors, is printed to the MATLAB Command Window The status light is yellow when the compile status is unknown, red when an error has occurred, and green on success. DSP Builder prints the compilation output to the MATLAB Command Window.
Simulink sample time	Specify the sample time of the DSP Builder part of your Simulink model.
Reset cycles	Allows you to hold your imported HDL in reset for an arbitrary number of cycles before the cosimulation begins.
Port	The TCP/IP port number that the cosimulation uses for communication.

### 15.1.4. Pause

The **Pause** block implements a breakpoint with trigger count to break and single step through Simulink simulations.

Input is a Boolean signal enabling a counter, which counts up to the parameterized count value, then pauses the simulation. Press play to resume simulation from that point. If you do not increase the trigger count, the block causes another pause. This single stepping mode pauses the simulation at the next cycle. The block is red when enabled.

Use the **Pause** block for debugging designs. For example; run to breakpoint, then turn on **Show port values when hovering** at this point. This option permanently causes slow simulation, so only turn on when stepping through. Using display blocks allows you to see variables displayed at the paused time (similar to watch variables in a software debugger). You can change the trigger count, for example by adding 100, to simulate the next 100 cycles. The block color changes to red when you turn on the **Pause** block. Use the `valid` signal as the input, so that it counts valid steps only. Alternatively, use a different control signal, such as a `writenable`, as an input to get the system to break then. You can easily add other logic blocks to generate a break signal that you can use just for debugging.

## 16. Simulink Supported Blocks

---

DSP Builder supports some Simulink blocks within or at the interface of synthesizable DSP Builder subsystems. If you use unsupported blocks, DSP Builder generates errors or the Simulink simulation may mismatch with the generated RTL

**Table 271. Simulink Supported Blocks**

Simulink Block	Notes
Bus Creator	Cannot connect the output of a <b>Bus Creator</b> to the DSP Builder IO blocks ( <b>ChannelIn</b> , <b>ChannelOut</b> , <b>GPIIn</b> , <b>GPOut</b> )
Bus Selector	
Mux	
Demux	
Complex to Real-Imag	
Real-Imag to Complex	
Const	Does not support DSP Builder variable precision floating-point blocks. Use DSP Builder <b>Const</b> block instead.
Terminator	
Goto	
From	
Selector	Can only be used with vector inputs. Cannot be used with bus inputs.
Inport	
Outport	
Vector Concatenate	
Subsystem	

## 17. Document Revision History for DSP Builder for Intel FPGAs (Advanced Blockset) Handbook

Version	Software Version	Changes
2022.03.23	21.3	Added two numerical examples to <i>CORDIC</i>
2021.12.07	21.3	Updated <i>CORDIC</i>
2021.09.30	21.3	<ul style="list-style-type: none"> <li>Added support for Cyclone V devices</li> <li>Removed <b>Bus Slave</b> block references</li> <li>Renamed Avalon-MM slave settings block: to <b>Avalon memory-mapped Agent Settings</b>.</li> <li>Updated <b>DualMem</b> block</li> <li>Added DFT library blocks</li> </ul>
2021.03.29	21.1	<ul style="list-style-type: none"> <li>Added <b>Finite State Machine</b> block and design example.</li> <li>Added <i>Simulink Supported Blocks</i>.</li> </ul>
2020.10.05	20.3	<ul style="list-style-type: none"> <li>Removed Intel Quartus Prime Standard Edition from <i>Device Support</i></li> <li>Removed Arria V, Arria V, Arria V GZ, Cyclone IV, Cyclone V, Intel MAX<sup>®</sup> 10, Stratix IV, Stratix V from <i>Device Support</i>.</li> <li>Added Intel Agilex devices to <i>Device Support</i></li> <li>Changed Intel Cyclone 10 device support to Intel Cyclone 10 GX in <i>Device Support</i></li> <li>Added new parameter to <b>DualMem</b> block.</li> <li>Added <b>Uninitialized means</b> parameter to <b>Control</b> block.</li> <li>Removed ... device selector option from <b>Device</b> block.</li> <li>Updated <b>SynthesisInfo</b> block description.</li> <li>Removed <b>GPDependentDelay</b> block.</li> </ul>
2020.04.22	20.1	Corrected input type for <b>Lut</b> and <b>Mux</b> blocks
2020.04.15	20.1	<ul style="list-style-type: none"> <li>Updated <i>Using Bit-Accurate Simulation</i>.</li> <li>Updated <i>DSP Builder Testbench Verification</i>.</li> <li>Removed <i>Cholesky Solver Single Channel</i>.</li> <li>Added <b>Report FIFO fill level</b> parameter to <i>Control</i>.</li> <li>Added more parameters to <i>Control Block Testbenches Tab Parameters</i></li> </ul>
2020.01.06	19.3	Removed "You can also use this option to implement efficient phase-shift keying (PSK) modulators in which the input to the phase modulator varies according to a data stream." from <i>NCO Block</i>
2019.10.31	19.3	Corrected <b>Fanout</b> and <b>VectorFanout</b> descriptions
2019.10.10	19.3	<ul style="list-style-type: none"> <li>Removed all references to standard blockset.</li> <li>Removed DSP Builder Standard and Advanced Blockset Interoperability</li> <li>Added QRD solver reference design</li> </ul>
2019.06.10	19.1	Added new parameters to <b>External Memory</b> block.
<i>continued...</i>		

Version	Software Version	Changes
2019.03.01	19.1	<ul style="list-style-type: none"> <li>Updated supported floating-point data types.</li> <li>Removed "except for Forloop blocks" from <i>Verifying your DSP Builder Design with C++ Software Models</i></li> <li>Added new parameters to: <ul style="list-style-type: none"> <li>Sqrt block</li> <li>Control block</li> </ul> </li> <li>Removed <i>Using Latency Constraints in DSP Builder Designs</i></li> </ul>
2018.09.17	18.1	<ul style="list-style-type: none"> <li>Updated <b>SharedMem</b> block description.</li> <li>Updated: <ul style="list-style-type: none"> <li>HDL Import feature description.</li> </ul> </li> <li>Removed: <ul style="list-style-type: none"> <li>Running the Simple Complex Multiplication Design Example</li> <li>About the Complex Multilication Design Example</li> <li>Cosimulation Block Parameters</li> <li>Configuring the HDL Import Block Parameters</li> <li>Setting up a Modelsim Cosimulation</li> <li>Adding Ports</li> <li>Adding a HDL Import Block</li> <li>Running a Cosimulation</li> <li>Verifying with HDL Import in the ModelSim</li> <li>Adding HDL Import Design to Intel Quartus Prime</li> </ul> </li> </ul>
2018.06.27	18.0	Updated <i>Arria 10</i> to <i>any device with a floating-point block</i> in floating-point designs
2018.06.08	18.0	Added HDL import.
2018.05.09	18.0	<ul style="list-style-type: none"> <li>Added new parameter to <b>Constant</b> block.</li> <li>Added extra description to <b>Fanout</b> block <b>Uninitialized</b> parameter.</li> <li>Added extra info to <b>DualMem</b> block parameters.</li> <li>Added reset minimization feature</li> <li>Added new parameter to <b>SharedMem</b> block.</li> </ul>
2017.11.06	17.1	<ul style="list-style-type: none"> <li>Improved description on <b>NCO</b> block <b>Accumulator Bit Width</b> parameter.</li> <li>Corrected parameters on <b>Scalar Product</b> block.</li> <li>Added <i>Forcing Soft Floating-point Data Types with the Advanced Options</i> topic</li> <li>Added super-sample NCO design example.</li> <li>Added support for Intel Cyclone 10 and Intel Stratix 10 devices.</li> <li>Removed instances of <b>Signals</b> block.</li> <li>Changed input type on <b>GPIIn</b> block; changed output type on <b>GPOut</b> block.</li> <li>Deleted WYSIWYG option on <b>SynthesisInfo</b> block.</li> </ul>
2017.05.02	17.0	<ul style="list-style-type: none"> <li>Rebranded as Intel</li> <li>Deprecated <b>Signals</b> block</li> <li>Corrected Interpolating FIR Filter design clock to say 240 MHz</li> <li>Added description on how to get output b on decimating, fractional rate, interpolating, and single-rate FIR filters</li> <li>Added Gaussian and Random Number Generator design examples</li> <li>Added variable-size supersampled FFT design example</li> <li>Added <b>HybridVFFT</b> block</li> <li>Added <b>GeneralVTwiddle</b> and <b>GeneralMultVTwiddle</b> blocks</li> <li>Corrected device support and removed Stratix 10 devices.</li> </ul>
2016.11.01	16.1	<ul style="list-style-type: none"> <li>Added device support</li> <li>Added 4-channel 2-antenna DUC and DDC for LTE reference design</li> <li>Added <b>BFU_simple</b> block</li> </ul>

*continued...*

Version	Software Version	Changes
2016.05.01	16.0	<ul style="list-style-type: none"> <li>• Revised getting started</li> <li>• Revised design rules</li> <li>• Revised setting up Simulink</li> <li>• Revised Primitive library description</li> <li>• Revised DSP Builder design structure</li> <li>• Added a library list</li> <li>• Removed <b>Run Quartus</b> and <b>Run ModelSim</b> blocks</li> <li>• Moved primitive subsystem designs to avoid from <i>Troubleshooting</i> chapter to <i>Design Rules and Recommendations</i></li> <li>• Moved <i>Troubleshooting</i> to <i>Design Rules, Recommendations, and Troubleshooting</i>.</li> <li>• Revised <b>EditParams</b> block description</li> <li>• Moved hardware verification from <i>Techniques for Advanced Users</i> to <i>Design Flow</i> chapter</li> <li>• Changed threshold names and descriptions on <b>Localthreshold</b> and <b>Control</b> blocks</li> </ul>
2016.05.01 cont	16.0 cont	<ul style="list-style-type: none"> <li>• Reorganised libraries:               <ul style="list-style-type: none"> <li>– Deleted <b>ModelVectorPrim</b> library. Moved <b>SumofElements</b> block to <b>Primitives &gt; Primitives Basic Blocks</b></li> <li>– Renamed <b>Channel</b> library to <b>Channel Filter and Waveform</b> library</li> <li>– Renamed <i>Filter</i> chapter to <i>IP</i> chapter</li> <li>– Created <b>FFT IP</b> library and added <b>BitReversecoreC</b>, <b>FFT</b>, <b>FFT_Float</b>, <b>VariablebitReverse</b>, <b>VFFT</b>, and <b>VFFT_Float</b> blocks to it.</li> <li>– Deleted <b>Waveform Synthesis</b> library and moved blocks into <b>Channel Filter and Waveform</b> library.</li> <li>– Renamed <b>Base</b> library to <b>Design Configuration</b></li> <li>– Moved <b>ChanView</b> block from <b>Base</b> to <b>Added transmit and receive modem referenceChannel Filter and Waveform</b> library</li> <li>– Moved <b>Scale</b> block from <b>Base</b> to <b>Channel Filter and Waveform</b> library</li> <li>– Renamed <b>FFT</b> to <b>FFT Design Elements</b> library and moved to <b>Primitives</b> library</li> <li>– Created <b>Primitives Basic Blocks</b> library and move all blocks from <b>Primitive</b> library to it.</li> <li>– Created <b>Primitive Configuration</b> library and moved <b>ChannelIn</b>, <b>ChannelOut</b>, <b>GPIIn</b>, <b>GPOut</b>, and <b>SynthesisInfo</b> blocks into it.</li> <li>– Renamed <b>ModelIP</b> to <b>IP</b> library</li> <li>– Renamed <b>ModelBus</b> to <b>Memory Mapped</b> library</li> <li>– Renamed <i>ModelBus</i> chapter to <i>Interfaces</i></li> <li>– Created <b>Streaming</b> library and moved the <b>AStInput</b>, <b>AStOutput</b>, and <b>AStInputFIFO</b> blocks into it</li> <li>– Renamed <b>Additional</b> to <b>Primitive Design Elements</b> library</li> <li>– Renamed <i>Additional</i> chapter to <i>Utilities</i></li> <li>– Created <b>Analyze and Test</b> library and moved <b>Capture Values</b> and <b>Pause</b> blocks to it.</li> <li>– Deleted <b>External Memories</b> library and moved <b>External Memory</b> block to <b>Interfaces &gt; Memory Mapped</b>.</li> <li>– Moved <i>DDC Design Example</i> to <i>Design Examples and Reference Design</i> chapter.</li> </ul> </li> </ul>
2015.11.01	15.1	<ul style="list-style-type: none"> <li>• Changed Quartus II to Quartus Prime software</li> <li>• Changed <b>Run Quartus II</b> block to <b>Run Quartus Prime</b> block</li> <li>• Removed <b>Turn on coverage in testbenches</b> and <b>Signal view depth</b> options from <b>Control</b> block</li> <li>• Improved FIR Filter Avalon-MM port descriptions</li> <li>• Changed some <b>Avalon-MM Slave Settings</b> block descriptions</li> <li>• Added design rules for Modelbus blocks.</li> </ul>

*continued...*

Version	Software Version	Changes
		<ul style="list-style-type: none"> <li>• Added reconfigurable FIR filter information.</li> <li>• Removed <b>Enhanced Precision Support</b> block</li> <li>• Removed <b>ScalarProduct</b> graphs</li> <li>• Added new blocks:               <ul style="list-style-type: none"> <li>– <b>Capture Values</b></li> <li>– <b>Fanout</b></li> <li>– <b>Pause</b></li> <li>– <b>Vectorfanout</b></li> </ul> </li> <li>• Added IIR: full-rate fixed-point and IIR: full-rate floating-point demos</li> <li>•</li> <li>• Changed <code>set_param</code> to <code>dspba.set_param</code></li> </ul>
2015.05.01	15.0	<ul style="list-style-type: none"> <li>• Added external memories library</li> <li>• Added <b>External Memory</b> block</li> <li>• Added new <b>Allow write on both ports</b> parameter to <b>DualMem</b> block</li> <li>• Removed read/write note from <b>SharedMem</b> block</li> <li>• Changed parameters on <b>AvalonMMSlaveSettings</b> block</li> <li>• Added note to latency constraints topic: latency constraints only apply between <b>ChannelIn</b> and <b>ChannelOut</b> blocks.</li> <li>• Added support for Verilog HDL implementation</li> <li>• Removed architecture versus implementation information</li> <li>• Removed <b>SynthesisInfo</b> block <b>WYSIWYG</b> option</li> <li>• Removed the following design examples:               <ul style="list-style-type: none"> <li>– 1K floating-point FFT</li> <li>– Radix-2 streaming FFT</li> <li>– Radix-4 streaming FFT</li> </ul> </li> </ul>
December 2014	14.1	<ul style="list-style-type: none"> <li>• Added step about disabling virtual pins in the Quartus Prime software when using HIL with advanced blockset designs</li> <li>• Added information on <code>_mmap.h</code> file, which contains register information on your design</li> <li>• Corrected Has read enable and Show read enable descriptions in <b>BusStimulus</b> and <b>BusStimulusFileReader</b> blocks</li> <li>• Added <b>BusStimulus</b> and <b>BusStimulusFileReader</b> blocks to memory-mapped registers design example.</li> <li>• Added <b>AvalonMMSlaveSettings</b> block and DSP Builder &gt; Avalon Interfaces &gt; Avalon-MM slave menu option</li> <li>• Removed bus parameters from <b>Control</b> and <b>Signal</b> blocks</li> <li>• Removed the following design examples:               <ul style="list-style-type: none"> <li>– Color Space Converter (Resource Sharing Folding)</li> <li>– Interpolating FIR Filter with Updating Coefficients</li> <li>– Primitive FIR Filter (Resource Sharing Folding)</li> <li>– Single-Stage IIR Filter (Resource Sharing Folding)</li> <li>– Three-stage IIR Filter (Resource Sharing Folding)</li> <li>– Added system-in-the-loop support</li> </ul> </li> <li>• Added new blocks:               <ul style="list-style-type: none"> <li>– Floating-point classifier</li> <li>– Floating-point multiply accumulate</li> <li>– Added hypotenuse function to <b>math</b> block</li> </ul> </li> </ul>

*continued...*

Version	Software Version	Changes
		<ul style="list-style-type: none"> <li>• Added design examples:               <ul style="list-style-type: none"> <li>– Color space converter</li> <li>– Complex FIR design example</li> <li>– CORDIC from Primitive Blocks</li> <li>– Crest factor reduction</li> <li>– Folding FIR</li> <li>– Variable Integer Rate Decimation Filter</li> <li>– Vector sort - sequential and iterative</li> </ul> </li> <li>• Added reference designs:               <ul style="list-style-type: none"> <li>– Crest factor reduction</li> <li>– Direct RF with Synthesizable Testbench</li> <li>– Dynamic Decimation Filter</li> <li>– Reconfigurable Decimation Filter</li> <li>– Variable Integer Rate Decimation Filter</li> </ul> </li> <li>• Changed directory structure</li> <li>• Added correct floating-point rounding for reciprocal and square root blocks.</li> <li>• Corrected signal descriptions for LoadableCounter block</li> <li>• Removed resource sharing folder</li> <li>• Added new ALU folder information:               <ul style="list-style-type: none"> <li>– Start of packet signal</li> <li>– Clock-rate mode</li> </ul> </li> </ul>
June 2014	14.0	<ul style="list-style-type: none"> <li>• Added new blocks:               <ul style="list-style-type: none"> <li>– Enabled Delay Line</li> <li>– Enabled Feedback Delay</li> <li>– FFT2P, FFT4P, FFT8P, FFT16P, FFT32P, and FFT64PFFT2X, FFT4X, FFT8X, FFT16X, FFT32X, and FFT64XFFT2, FFT4, VFFT2, and VFFT4</li> <li>– General Multitwiddle and General Twiddle (GeneralMultiTwiddle, GeneralTwiddle</li> <li>– )Hybrid FFT (Hybrid_FFT)</li> <li>– Parallel Pipelined FFT (PFFT_Pipe)</li> <li>– Ready</li> </ul> </li> <li>• Updated port descriptions for               <ul style="list-style-type: none"> <li>– Convert</li> <li>– Const</li> <li>– LUT</li> <li>– Reinterpretcast</li> </ul> </li> </ul>



Version	Software Version	Changes
		<ul style="list-style-type: none"> <li>• New parameter interfaces for:                             <ul style="list-style-type: none"> <li>– FFT</li> <li>– FFT_Light</li> <li>– VFFT</li> <li>– VFFT_Light</li> </ul> </li> <li>• Added new design examples:                             <ul style="list-style-type: none"> <li>– Avalon-ST Interface (Input and Output FIFO Buffer) with Backpressure</li> <li>– Avalon-ST Interface (Output FIFO Buffer) with Backpressure</li> <li>– Fixed-point maths functions</li> <li>– Fractional square root using CORDIC</li> <li>– Normalizer</li> <li>– Square root using CORDIC</li> <li>– Switchable FFT/iFFT</li> <li>– Variable-Size Fixed-Point FFT</li> <li>– Variable-Size Fixed-Point FFT without BitReverseCoreC Block</li> <li>– Variable-Size Fixed-Point iFFTVariable-Size Fixed-Point iFFT without BitReverseCoreC Block</li> <li>– Variable-Size Floating-Point FFTVariable-Size Floating-Point FFT without BitReverseCoreC Block</li> <li>– Variable-Size Floating-Point iFFTVariable-Size Floating-Point iFFT without BitReverseCoreC Block</li> </ul> </li> <li>• Added new ready signal for ALU folding.</li> </ul>