

Synthesis and Simulation Design Guide

Getting Started

HDL Coding Hints

***Understanding High-Density
Design Flow***

Designing FPGAs with HDL

Simulating Your Design

***Accelerate FPGA Macros
with One-Hot Approach***

Report Files



The Xilinx logo shown above is a registered trademark of Xilinx, Inc.

FPGA Architect, FPGA Foundry, NeoCAD, NeoCAD EPIC, NeoCAD PRISM, NeoROUTE, Timing Wizard, TRACE, XACT, XILINX, XC2064, XC3090, XC4005, XC5210, and XC-DS501 are registered trademarks of Xilinx, Inc.



The shadow X shown above is a trademark of Xilinx, Inc.

All XC-prefix product designations, Alliance Series, AllianceCORE, BITA, CLC, Configurable Logic Cell, Dual Block, EZTag, FastCLK, FastCONNECT, FastFLASH, FastMap, Foundation, HardWire, LCA, LogiBLOX, Logic Cell, LogiCORE, LogicProfessor, MicroVia, Plus Logic, PLUSASM, Plustran, P+, PowerGuide, PowerMaze, SelectI/O, Select-RAM, Select-RAM+, Smartguide, SmartSearch, Smartspec, Spartan, TrueMap, UIM, VectorMaze, VersaBlock, VersaRing, Virtex, WebLINX, XABEL, XACTstep, XACTstep Advanced, XACTstep Foundry, XACT-Floorplanner, XACT-Performance, XAM, XAPP, X-BLOX, X-BLOX plus, XChecker, XDM, XDS, XEPLD, Xilinx Foundation Series, XPP, XSI, and ZERO+ are trademarks of Xilinx, Inc. The Programmable Logic Company and The Programmable Gate Array Company are service marks of Xilinx, Inc.

All other trademarks are the property of their respective owners.

Xilinx, Inc. does not assume any liability arising out of the application or use of any product described or shown herein; nor does it convey any license under its patents, copyrights, or maskwork rights or any rights of others. Xilinx, Inc. reserves the right to make changes, at any time, in order to improve reliability, function or design and to supply the best product possible. Xilinx, Inc. will not assume responsibility for the use of any circuitry described herein other than circuitry entirely embodied in its products. Xilinx, Inc. devices and products are protected under one or more of the following U.S. Patents: 4,642,487; 4,695,740; 4,706,216; 4,713,557; 4,746,822; 4,750,155; 4,758,985; 4,820,937; 4,821,233; 4,835,418; 4,855,619; 4,855,669; 4,902,910; 4,940,909; 4,967,107; 5,012,135; 5,023,606; 5,028,821; 5,047,710; 5,068,603; 5,140,193; 5,148,390; 5,155,432; 5,166,858; 5,224,056; 5,243,238; 5,245,277; 5,267,187; 5,291,079; 5,295,090; 5,302,866; 5,319,252; 5,319,254; 5,321,704; 5,329,174; 5,329,181; 5,331,220; 5,331,226; 5,332,929; 5,337,255; 5,343,406; 5,349,248; 5,349,249; 5,349,250; 5,349,691; 5,357,153; 5,360,747; 5,361,229; 5,362,999; 5,365,125; 5,367,207; 5,386,154; 5,394,104; 5,399,924; 5,399,925; 5,410,189; 5,410,194; 5,414,377; 5,422,833; 5,426,378; 5,426,379; 5,430,687; 5,432,719; 5,448,181; 5,448,493; 5,450,021; 5,450,022; 5,453,706; 5,455,525; 5,466,117; 5,469,003; 5,475,253; 5,477,414; 5,481,206; 5,483,478; 5,486,707; 5,486,776; 5,488,316; 5,489,858; 5,489,866; 5,491,353; 5,495,196; 5,498,979; 5,498,989; 5,499,192; 5,500,608; 5,500,609; 5,502,000; 5,502,440; 5,504,439; 5,506,518; 5,506,523; 5,506,878; 5,513,124; 5,517,135; 5,521,835; 5,521,837; 5,523,963; 5,523,971; 5,524,097; 5,526,322; 5,528,169; 5,528,176; 5,530,378; 5,530,384; 5,546,018; 5,550,839; 5,550,843; 5,552,722; 5,553,001; 5,559,751; 5,561,367; 5,561,629; 5,561,631; 5,563,527; 5,563,528; 5,563,529; 5,563,827; 5,565,792; 5,566,123; 5,570,051; 5,574,634; 5,574,655; 5,578,946; 5,581,198; 5,581,199; 5,581,738; 5,583,450; 5,583,452; 5,592,105; 5,594,367; 5,598,424; 5,600,263; 5,600,264; 5,600,271; 5,600,597; 5,608,342; 5,610,536; 5,610,790; 5,610,829; 5,612,633; 5,617,021; 5,617,041; 5,617,327; 5,617,573; 5,623,387; 5,627,480; 5,629,637; 5,629,886; 5,631,577; 5,631,583; 5,635,851; 5,636,368; 5,640,106; 5,642,058; 5,646,545; 5,646,547; 5,646,564; 5,646,903; 5,648,732; 5,648,913; 5,650,672; 5,650,946; 5,652,904; 5,654,631; 5,656,950; 5,657,290; 5,659,484; 5,661,660; 5,661,685; 5,670,896; 5,670,897; 5,672,966; 5,673,198; 5,675,262; 5,675,270; 5,675,589; 5,677,638; 5,682,107; 5,689,133; 5,689,516; 5,691,907; 5,691,912; 5,694,047; 5,694,056; 5,724,276; 5,694,399; 5,696,454; 5,701,091; 5,701,441; 5,703,759; 5,705,932; 5,705,938; 5,708,597; 5,712,579; 5,715,197; 5,717,340; 5,719,506; 5,719,507; 5,724,276; 5,726,484; 5,726,584; Re. 34,363, Re. 34,444, and Re. 34,808. Other U.S. and foreign patents pending. Xilinx, Inc. does not represent that devices shown or products described herein are free from patent infringement or from any other third party right. Xilinx, Inc. assumes no obligation to correct

any errors contained herein or to advise any user of this text of any correction if such be made. Xilinx, Inc. will not assume any liability for the accuracy or correctness of any engineering or software support or assistance provided to a user.

Xilinx products are not intended for use in life support appliances, devices, or systems. Use of a Xilinx product in such applications without the written consent of the appropriate Xilinx officer is prohibited.

Copyright 1991-1998 Xilinx, Inc. All Rights Reserved.

Preface

About This Manual

This manual provides a general overview of designing Field Programmable Gate Arrays (FPGAs) with Hardware Description Languages (HDLs). It includes design hints for the novice HDL user, as well as for the experienced user who is designing FPGAs for the first time.

The design examples in this manual were created with Verilog and VHSIC Hardware Description Language (VHDL); compiled with the various synthesis tools; and targeted for XC4000, Spartan, and XC5200 devices. Xilinx equally endorses both Verilog and VHDL. VHDL may be more difficult to learn than Verilog and usually requires more explanation.

This manual does not address certain topics that are important when creating HDL designs, such as the design environment; verification techniques; constraining in the synthesis tool; test considerations; and system verification. Refer to your synthesis tool's reference manuals and design methodology notes for additional information.

Before using this manual, you should be familiar with the operations that are common to all Xilinx software tools, such as starting the software, selecting a tool for use, specifying operations, and managing design data. These topics are covered in the *Quick Start Guide*.

You must consult *The Programmable Logic Data Book* for device-specific information on Xilinx device characteristics, including readback, boundary scan, configuration, length count, and debugging. *The Programmable Logic Data Book* is available in hard copy and on the Xilinx web site (<http://www.xilinx.com>). See <http://www.xilinx.com/partinfo/databook.htm> for the current version of this book.

For specific design issues or problems, use the Answers Search function on the Web (<http://www.xilinx.com/support/searchtd.htm>) to access the following.

- Answers Database: current listing of solution records for the Xilinx software tools
- Applications Notes: descriptions of device-specific design techniques and approaches
- Data Sheets: pages from *The Programmable Logic Data Book*
- XCELL Journal: quarterly journals for Xilinx programmable logic users
- Expert Journals: the latest news, design tips, and patch information on the Xilinx design environment

If you cannot access the Web, you can install and access the Answers book with the DynaText online browser in the same manner as the Xilinx online book collection. The Answers book includes information in the Answers Database at the time of this release.

Manual Contents

- Chapter 1, “Getting Started,” provides a general overview of designing Field Programmable Gate Arrays (FPGAs) with HDLs. This chapter also includes installation requirements and instructions.
- Chapter 2, “HDL Coding Hints,” includes HDL coding hints and design examples to help you develop an efficient coding style.
- Chapter 3, “Understanding High-Density Design Flow,” provides synthesis and Xilinx implementation techniques to increase design performance and utilization.
- Chapter 4, “Designing FPGAs with HDL,” includes coding techniques to help you improve synthesis results.
- Chapter 5, “Simulating Your Design,” describes simulation methods for verifying the function and timing of your designs.
- Appendix A, “Accelerate FPGA Macros with One-Hot Approach,” reprints an article describing one-hot encoding in detail.

-
- Appendix B, “Report Files,” includes area and timing report files from various synthesis vendors.

Conventions

Typographical

This manual uses the following conventions. An example illustrates each convention.

- `Courier font` indicates messages, prompts, and program files that the system displays.

```
speed grade: -100
```

- **Courier bold** indicates literal commands that you enter in a syntactical statement. However, braces “{ }” in Courier bold are not literal and square brackets “[]” in Courier bold are literal only in the case of bus specifications, such as bus [7:0].

```
rpt_del_net=
```

Courier bold also indicates commands that you select from a menu.

```
File → Open
```

- *Italic font* denotes the following items.
 - Variables in a syntax statement for which you must supply values

```
edif2ngd design_name
```
 - References to other manuals

See the *Development System Reference Guide* for more information.
 - Emphasis in text

If a wire is drawn so that it overlaps the pin of a symbol, the two nets are *not* connected.

- Square brackets “[]” indicate an optional entry or parameter. However, in bus specifications, such as bus [7:0], they are required.

```
edif2ngd [option_name] design_name
```

Square brackets also enclose footnotes in tables that are printed out as hardcopy in DynaText[®].

- Braces “{ }” enclose a list of items from which you must choose one or more.

```
lowpwr = {on | off}
```

- A vertical bar “|” separates items in a list of choices.

```
lowpwr = {on | off}
```

- A vertical ellipsis indicates repetitive material that has been omitted.

```
IOB #1: Name = QOUT'  
IOB #2: Name = CLKIN'  
.  
.  
.
```

- A horizontal ellipsis “...” indicates that an item can be repeated one or more times.

```
allow block block_name loc1 loc2 . . . locn;
```


Contents

Preface

| | |
|-------------------------|----|
| About This Manual | i |
| Manual Contents | ii |

Conventions

| | |
|--------------------|---|
| Typographical..... | v |
|--------------------|---|

Chapter 1 Getting Started

| | |
|--|------|
| Advantages of Using HDLs to Design FPGAs | 1-2 |
| Designing FPGAs with HDLs | 1-3 |
| Using Verilog..... | 1-3 |
| Using VHDL | 1-3 |
| Comparing ASICs and FPGAs..... | 1-3 |
| Using Synthesis Tools | 1-4 |
| Using FPGA System Features..... | 1-4 |
| Designing Hierarchy..... | 1-4 |
| Specifying Speed Requirements..... | 1-5 |
| Installing Design Examples and Tactical Software | 1-5 |
| Software Requirements | 1-5 |
| Memory Requirements | 1-5 |
| Disk Space Requirements | 1-6 |
| Xilinx Internet Site | 1-7 |
| Retrieving Tactical Software and Design Examples | 1-8 |
| From Xilinx Internet FTP Site..... | 1-8 |
| Extracting the Files | 1-10 |
| Extracting .tar.Z File in UNIX | 1-10 |
| Extracting .zip File in UNIX | 1-10 |
| Extracting .zip File in MS-DOS™ | 1-10 |
| Directory Tree Structure..... | 1-11 |
| Technical Support | 1-14 |
| Xilinx World Wide Web Site | 1-14 |

| | |
|--|------|
| Technical and Applications Support Hotlines..... | 1-14 |
| Xilinx FTP Site | 1-15 |
| XDOCS E-mail Server | 1-15 |

Chapter 2 HDL Coding Hints

| | |
|---|------|
| Comparing Synthesis and Simulation Results | 2-1 |
| Omit the Wait for XX ns Statement | 2-2 |
| Omit the ...After XX ns or Delay Statement | 2-2 |
| Use Case and If-Else Statements | 2-3 |
| Order and Group Arithmetic Functions | 2-3 |
| Omit Initial Values | 2-4 |
| Selecting HDL Formatting Styles | 2-4 |
| Selecting a Capitalization Style | 2-4 |
| Verilog | 2-5 |
| VHDL | 2-5 |
| Using Xilinx Naming Conventions | 2-5 |
| Matching File Names to Entity and Module Names | 2-6 |
| Naming Identifiers, Types, and Packages | 2-6 |
| Using Labels | 2-7 |
| Labeling Flow Control Constructs | 2-7 |
| Using Variables for Constants (VHDL Only) | 2-8 |
| Using Constants to Specify OPCODE Functions (VHDL) | 2-8 |
| Using Parameters for Constants (Verilog Only) | 2-9 |
| Using Parameters to Specify OPCODE Functions (Verilog) | 2-9 |
| Using Parameters to Specify Bus Size (Verilog) | 2-9 |
| Using Named and Positional Association | 2-9 |
| Managing Your Design | 2-10 |
| Creating Readable Code | 2-10 |
| Indenting Your Code | 2-10 |
| Using Empty Lines | 2-12 |
| Using Spaces | 2-12 |
| Breaking Long Lines of Code | 2-12 |
| Adding Comments | 2-13 |
| Using Std_logic Data Type (VHDL only) | 2-13 |
| Declaring Ports | 2-14 |
| Minimizing the Use of Ports Declared as Buffers | 2-14 |
| Comparing Signals and Variables (VHDL only) | 2-15 |
| Using Signals (VHDL) | 2-16 |
| Using Variables (VHDL) | 2-17 |
| Using Schematic Design Hints with HDL Designs | 2-17 |

| | |
|--|------|
| Barrel Shifter Design | 2-18 |
| 16-bit Barrel Shifter (VHDL) | 2-18 |
| 16-bit Barrel Shifter (Verilog) | 2-20 |
| 16-bit Barrel Shifter with Two Levels of Multiplexers (VHDL) | 2-23 |
| 16-bit Barrel Shifter with Two Levels of Multiplexers (Verilog) | 2-25 |
| Implementing Latches and Registers..... | 2-27 |
| D Latch Inference | 2-27 |
| Converting a D Latch to a D Register | 2-29 |
| Resource Sharing | 2-31 |
| Gate Reduction | 2-37 |
| Preset Pin or Clear Pin | 2-38 |
| Register Inference | 2-39 |
| Using Clock Enable Pin Instead of Gated Clocks..... | 2-42 |
| Using If Statements..... | 2-45 |
| Using Case Statements | 2-47 |
| Using Nested If Statements | 2-48 |
| Inefficient Use of Nested If Statement | 2-48 |
| Nested If Example Modified to Use If-Case..... | 2-51 |
| Comparing If Statement and Case Statement | 2-55 |

Chapter 3 Understanding High-Density Design Flow

| | |
|--|-----|
| Design Flow | 3-1 |
| Entering your Design and Selecting Hierarchy | 3-3 |
| Design Entry Recommendations | 3-3 |
| Use RTL Code | 3-3 |
| Carefully Select Design Hierarchy | 3-3 |
| Functional Simulation of your Design..... | 3-4 |
| Synthesizing and Optimizing your Design..... | 3-4 |
| Creating an Initialization File..... | 3-4 |
| Creating a Compile Run Script | 3-4 |
| Compiling Your Design | 3-5 |
| Modifying your Design | 3-5 |
| Compiling Large Designs..... | 3-5 |
| Saving Compiled Design as XNF or EDIF | 3-5 |
| Setting Timing Constraints | 3-5 |
| Naming and Grouping Signals Together..... | 3-6 |
| TNMs | 3-6 |
| TIMEGRPs..... | 3-6 |
| TPSYNC Specification..... | 3-7 |
| Specifying Timing Constraints | 3-7 |

| | |
|--|------|
| Period Constraint | 3-8 |
| FROM:TO Style Constraint | 3-8 |
| Offset Constraint | 3-9 |
| Ignoring Timing Paths | 3-9 |
| Controlling Signal Skew | 3-10 |
| Timing Constraint Priority..... | 3-10 |
| Evaluating Design Size and Performance..... | 3-12 |
| Using your Synthesis Tool to Estimate Device Utilization and Performance..... | 3-12 |
| Using the Timing Report Command | 3-13 |
| Determining Actual Device Utilization and Pre-routed Performance | 3-13 |
| Using the Design Manager to Map your Design | 3-13 |
| Using the Command Line to Map your Design | 3-15 |
| Evaluating your Design for Coding Style and System Features | 3-17 |
| Tips for Improving Design Performance | 3-17 |
| Modifying your Code..... | 3-17 |
| Using FPGA System Features..... | 3-18 |
| Using Xilinx-specific Features of your Synthesis Tool | 3-18 |
| Placing and Routing your Design..... | 3-18 |
| Decreasing Implementation Time | 3-19 |
| Improving Implementation Results..... | 3-21 |
| Multi-Pass Place and Route Option | 3-21 |
| Turns Engine Option | 3-21 |
| Re-entrant Routing Option..... | 3-21 |
| Cost-Based Clean-up Option | 3-22 |
| Delay-Based Clean-up Option | 3-23 |
| Guide Option (not recommended) | 3-23 |
| Timing Simulation of your Design | 3-23 |
| Downloading to the Device and In-system Debugging | 3-23 |
| Creating a PROM File for Stand-Alone Operation | 3-24 |

Chapter 4 Designing FPGAs with HDL

| | |
|---|------|
| Using Global Low-skew Clock Buffers | 4-2 |
| Inserting Clock Buffers..... | 4-6 |
| Instantiating Global Clock Buffers..... | 4-6 |
| Instantiating Buffers Driven from a Port..... | 4-6 |
| Instantiating Buffers Driven from Internal Logic..... | 4-6 |
| Using Dedicated Global Set/Reset Resource | 4-9 |
| Startup State | 4-10 |
| Preset vs. Clear (XC4000, Spartan) | 4-10 |

| | |
|---|------|
| Increasing Performance with the GSR/GR Net..... | 4-10 |
| Design Example without Dedicated GSR/GR Resource | 4-11 |
| Design Example with Dedicated GSR/GR Resource | 4-16 |
| Design Example with Active Low GSR/GR Signal..... | 4-23 |
| Encoding State Machines | 4-26 |
| Using Binary Encoding..... | 4-26 |
| Binary Encoded State Machine VHDL Example..... | 4-27 |
| Binary Encoded State Machine Verilog Example | 4-30 |
| Using Enumerated Type Encoding | 4-32 |
| Enumerated Type Encoded State Machine VHDL | |
| Example..... | 4-33 |
| Enumerated Type Encoded State Machine Verilog | |
| Example..... | 4-34 |
| Using One-Hot Encoding | 4-35 |
| One-hot Encoded State Machine VHDL Example | 4-36 |
| One-hot Encoded State Machine Verilog Example | 4-37 |
| Summary of Encoding Styles..... | 4-38 |
| Comparing Synthesis Results for Encoding Styles..... | 4-39 |
| Initializing the State Machine | 4-39 |
| Using Dedicated I/O Decoders..... | 4-40 |
| Instantiating LogiBLOX Modules..... | 4-45 |
| Using LogiBLOX in HDL Designs | 4-47 |
| Instantiating a LogiBLOX “Black Box” Component | 4-51 |
| VHDL Example | 4-51 |
| Implementing Memory..... | 4-52 |
| Implementing XC4000 and Spartan ROMs..... | 4-53 |
| RTL Description of a ROM VHDL Example | 4-53 |
| RTL Description of a ROM Verilog Example | 4-54 |
| Implementing XC4000 Family RAMs | 4-55 |
| Instantiating RAM VHDL Example..... | 4-56 |
| Instantiating RAM Verilog Example | 4-57 |
| Inferring RAM VHDL Example | 4-58 |
| Using LogiBLOX to Implement Memory..... | 4-59 |
| Implementing Boundary Scan (JTAG 1149.1) | 4-63 |
| Instantiating the Boundary Scan Symbol | 4-63 |
| Boundary Scan VHDL Example..... | 4-64 |
| Boundary Scan Verilog Example | 4-65 |
| Implementing Logic with IOBs..... | 4-67 |
| XC4000E/EX/XL/XV and Spartan IOBs | 4-67 |
| Inputs | 4-67 |
| Outputs | 4-68 |
| Slew Rate | 4-68 |

| | |
|--|------|
| Pull-ups and Pull-downs | 4-69 |
| XC4000EX/XL/XV Output Multiplexer/2-Input Function Generator | 4-69 |
| XC5200 IOBs | 4-71 |
| Bi-directional I/O | 4-71 |
| Inferring Bi-directional I/O | 4-72 |
| Instantiating Bi-directional I/O | 4-74 |
| Using LogiBLOX to Create Bi-directional I/O | 4-76 |
| Specifying Pad Locations..... | 4-79 |
| Moving Registers into the IOB | 4-79 |
| Use -pr Option with Map | 4-80 |
| Using Unbonded IOBs (XC4000E/EX/XL/XV and Spartan Only)..... | 4-81 |
| 4-bit Shift Register Using Unbonded I/O VHDL Example | 4-81 |
| 4-bit Shift Register Using Unbonded I/O Verilog Example | 4-83 |
| Implementing Multiplexers with Tristate Buffers..... | 4-84 |
| Mux Implemented with Gates VHDL Example | 4-85 |
| Mux Implemented with Gates Verilog Example | 4-85 |
| Mux Implemented with BUFTs VHDL Example | 4-86 |
| Mux Implemented with BUFTs Verilog Example..... | 4-87 |
| Using Pipelining | 4-89 |
| Before Pipelining..... | 4-89 |
| After Pipelining..... | 4-90 |
| Design Hierarchy..... | 4-90 |
| Using Synthesis Tools with Hierarchical Designs | 4-91 |
| Restrict Shared Resources to Same Hierarchy Level | 4-91 |
| Compile Multiple Instances Together | 4-91 |
| Restrict Related Combinatorial Logic to Same Hierarchy Level..... | 4-92 |
| Separate Speed Critical Paths from Non-critical Paths | 4-92 |
| Restrict Combinatorial Logic that Drives a Register to Same Hierarchy Level | 4-92 |
| Restrict Module Size | 4-92 |
| Register All Outputs | 4-93 |
| Restrict One Clock to Each Module or to Entire Design | 4-93 |

Chapter 5 Simulating Your Design

| | |
|---|-----|
| Functional Simulation..... | 5-4 |
| Register Transfer Level (RTL) | 5-4 |
| Post-Synthesis (Pre-NGDBuild) Gate-Level Simulation | 5-5 |

| | |
|---|------|
| Post-NGDBuild (Pre-Map) Gate-Level Simulation | 5-6 |
| Post-Map Partial Timing (CLB and IOB Block Delays) | 5-6 |
| Creating a Simulation Test Bench | 5-7 |
| Timing Simulation..... | 5-8 |
| Post-Route Full Timing (Block and Net Delays)..... | 5-8 |
| Creating a Timing Simulation Netlist..... | 5-9 |
| From the Design Manager | 5-9 |
| From the Command Line | 5-10 |
| Using VHDL/Verilog Libraries and Models..... | 5-11 |
| Adhering to Industry Standards | 5-11 |
| Locating Library Source Files | 5-12 |
| Case Sensitivity | 5-13 |
| Using the UNISIM Library | 5-14 |
| UNISIM Library Structure..... | 5-14 |
| Compiling the UNISIM Library | 5-15 |
| Instantiating UNISIM Library Components | 5-16 |
| Using the LogiBLOX Library | 5-16 |
| Compiling the LogiBLOX Library | 5-17 |
| Instantiating LogiBLOX Modules | 5-17 |
| Using the LogiCORE Library..... | 5-18 |
| Simulating Global Signals | 5-18 |
| Adapting Schematic Global Signal Methodology for VHDL | 5-20 |
| Setting VHDL Global Set/Reset Emulation in Functional Simulation | 5-21 |
| Global Signal Considerations (VHDL)..... | 5-22 |
| GSR Network Design Cases..... | 5-23 |
| Using VHDL Reset-On-Configuration (ROC) Cell (Case 1A)..... | 5-24 |
| Using ROC Cell Implementation Model (Case 1A)..... | 5-26 |
| ROC Test Bench (Case 1A) | 5-26 |
| ROC Model in Four Design Phases (Case 1A) | 5-28 |
| Using VHDL ROCBUF Cell (Case 1B) | 5-29 |
| ROCBUF Model in Four Design Phases (Case 1B) | 5-31 |
| Using VHDL STARTBUF Block (Case 2A and 2B)..... | 5-32 |
| GTS Network Design Cases | 5-35 |
| Using VHDL Tristate-On-Configuration (TOC) | 5-37 |
| VHDL TOC Cell (Case A1) | 5-37 |
| TOC Cell Instantiation (Case A1) | 5-37 |
| TOC Test Bench (Case A1)..... | 5-39 |
| TOC Model in Four Design Phases (Case A1)..... | 5-40 |
| Using VHDL TOCBUF (Case B1)..... | 5-42 |
| TOCBUF Model Example (Case B1) | 5-42 |

| | |
|---|------|
| TOCBUF Model in Four Design Phases (Case B1)..... | 5-43 |
| Using Oscillators (VHDL) | 5-46 |
| Oscillator VHDL Example | 5-46 |
| Oscillator Test Bench..... | 5-47 |
| Compiling Verilog Libraries | 5-49 |
| Compiling Libraries for Modelsim..... | 5-49 |
| SIMPRIM Libraries..... | 5-49 |
| LogiBLOX Libraries..... | 5-49 |
| UNISIM Libraries | 5-50 |
| Verilog Simulation | 5-51 |
| Setting Verilog Global Set/Reset..... | 5-51 |
| Defining GSR in a Test Bench | 5-52 |
| Designs without a STARTUP Block | 5-55 |
| Example 1: XC4000E/L/X/XL/XV and Spartan RTL Functional Simulation (No STARTUP Block) | 5-55 |
| Example 2: XC5200 RTL Functional Simulation (No STARTUP Block) | 5-57 |
| Example 3: XC3000A/L and XC3100A/L RTL and Post-synthesis Functional Simulation (No STARTUP Block) | 5-58 |
| Designs with a STARTUP block (XC4000E/L/X/XL/XV, Spartan, and XC5200 Devices Only) | 5-59 |
| Example 1a: XC4000E/L/X/XL/XV and Spartan: RTL and Post-synthesis Simulation (With STARTUP) | 5-60 |
| Example 1b: Post-NGDBuild Functional, Post-Map Timing, and Post-Route Timing Simulation (With STARTUP) | 5-61 |
| Example 2a: XC5200: RTL or Post-synthesis Functional Simulation Designs with STARTUP Block | 5-62 |
| Example 2b: Post-NGDBuild Functional, Post-Map Timing, and Post-Route Timing Simulation (With STARTUP Block)..... | 5-63 |
| Example 3: XC3000A/L and XC3100A/L designs..... | 5-63 |
| Setting Verilog Global Tristate (XC4000, Spartan, and XC5200 Outputs Only) | 5-63 |
| Specifying GTS | 5-64 |
| Designs without a STARTUP Block | 5-65 |
| XC4000E/L/X/XL/XV, Spartan, and XC5200 RTL Functional Simulation (No STARTUP Block) | 5-65 |
| Designs with a STARTUP block (XC4000E/L/X/XL/XV, Spartan, and XC5200 Devices Only) | 5-66 |

| | |
|--|------|
| Example 1a: XC4000E/L/X/XL/XV, Spartan, and XC5200: RTL or Post-Synthesis Functional Simulation (With STARTUP, GTS Pin Connected)..... | 5-66 |
| Example 1b: Post-NGDBuild Simulation of GTS (With STARTUP, GTS Pin connected) | 5-67 |
| Example 2a: XC4000E/L/X/XL/XV, Spartan, and XC5200: Unified Library Simulation (With STARTUP, GTS Pin not connected) | 5-68 |
| Example 2b: Post-NGDBuild Simulation of GTS (With STARTUP, GTS Pin not connected)..... | 5-68 |
| Using Verilog Oscillators (OSC, OSC4, OSC5) | 5-69 |

Appendix A Accelerate FPGA Macros with One-Hot Approach

Appendix B Report Files

| | |
|------------------------------|------|
| Synplicity® | B-1 |
| Exemplar Logic | B-4 |
| Summary Report..... | B-5 |
| Log Report | B-7 |
| Synopsys FPGA Express™ | B-10 |

Index

Getting Started

Hardware Description Languages (HDLs) are used to describe the behavior and structure of system and circuit designs. This chapter includes a general overview of designing FPGAs with HDLs. System requirements and installation instructions are also provided.

To learn more about designing FPGAs with HDLs, Xilinx recommends that you enroll in the appropriate training classes offered by Xilinx and by the vendors of synthesis software. An understanding of FPGA architecture allows you to create HDL code that effectively uses FPGA system features.

Before you start to create your FPGA designs, refer to the current version of the Quick Start Guide for Xilinx Alliance Series for a description of the design flow; installation information; and general information on the Xilinx tools.

For the latest information on Xilinx parts and software, visit the Xilinx Web site at <http://www.xilinx.com>. On the Xilinx home page, click on Service and Support, and use the Customer Service and Support page to get answers to your technical questions. You can also use the File Download option to download the latest software patches, tutorials, design files, and documentation.

This chapter includes the following sections.

- “Advantages of Using HDLs to Design FPGAs”
- “Designing FPGAs with HDLs”
- “Installing Design Examples and Tactical Software”
- “Technical Support”

Advantages of Using HDLs to Design FPGAs

Using HDLs to design high-density FPGAs is advantageous for the following reasons.

- **Top-Down Approach for Large Projects**

HDLs are used to create complex designs. The top-down approach to system design supported by HDLs is advantageous for large projects that require many designers working together. After the overall design plan is determined, designers can work independently on separate sections of the code.

- **Functional Simulation Early in the Design Flow**

You can verify the functionality of your design early in the design flow by simulating the HDL description. Testing your design decisions before the design is implemented at the RTL or gate level allows you to make any necessary changes early in the design process.

- **Synthesis of HDL Code to Gates**

You can synthesize your hardware description to a design implemented with gates. This step decreases design time by eliminating the traditional gate-level bottleneck. Synthesis to gates also reduces the number of errors that can occur during a manual translation of a hardware description to a schematic design. Additionally, you can apply the techniques used by the synthesis tool (such as machine encoding styles or automatic I/O insertion) during the optimization of your design to the original HDL code, resulting in greater efficiency.

- **Early Testing of Various Design Implementations**

HDLs allow you to test different implementations of your design early in the design flow. You can then use the synthesis tool to perform the logic synthesis and optimization into gates. Additionally, Xilinx FPGAs allow you to implement your design at your computer. Since the synthesis time is short, you have more time to explore different architectural possibilities at the Register Transfer Level (RTL). You can reprogram Xilinx FPGAs to test several implementations of your design.

Designing FPGAs with HDLs

If you are more familiar with schematic design entry, you may find it difficult at first to create HDL designs. You must make the transition from graphical concepts, such as block diagrams, state machines, flow diagrams, and truth tables, to abstract representations of design components. You can ease this transition by not losing sight of your overall design plan as you code in HDL. To effectively use an HDL, you must understand the syntax of the language; the synthesis and simulator software; the architecture of your target device; and the implementation tools. This section gives you some design hints to help you create FPGAs with HDLs.

Using Verilog

Verilog[®] is popular for synthesis designs because it is less verbose than traditional VHDL, and it is standardized as IEEE-STD-1364-95. It was not originally intended as an input to synthesis, and many Verilog constructs are not supported by synthesis software. The Verilog examples in this manual were tested and synthesized with current, commonly-used FPGA synthesis software. The coding strategies presented in the remaining chapters of this manual can help you create HDL descriptions that can be synthesized.

Using VHDL

VHSIC Hardware Description Language (VHDL) is a hardware description language for designing Integrated Circuits (ICs). It was not originally intended as an input to synthesis, and many VHDL constructs are not supported by synthesis software. However, the high level of abstraction of VHDL makes it easy to describe the system-level components and test benches that are not synthesized. In addition, the various synthesis tools use different subsets of the VHDL language. The examples in this manual will work with most commonly used FPGA synthesis software. The coding strategies presented in the remaining chapters of this manual can help you create HDL descriptions that can be synthesized.

Comparing ASICs and FPGAs

Methods used to design ASICs do not always apply to FPGA designs. ASICs have more gate and routing resources than FPGAs. Because

ASICs have a large number of available resources, you can easily create inefficient code that results in a large number of gates. When designing FPGAs, you *must* create efficient code.

Using Synthesis Tools

Most of the commonly-used FPGA synthesis tools have special optimization algorithms for Xilinx FPGAs. Constraints and compiling options perform differently depending on the target device. There are some commands and constraints that do not apply to FPGAs and, if used, may adversely impact your results. You should understand how your synthesis tool processes designs before creating FPGA designs. Most synthesis vendors include information in their manuals specifically for Xilinx FPGAs.

Using FPGA System Features

You can improve device performance and area utilization by creating HDL code that uses FPGA system features, such as global reset, wide I/O decoders, and memory. FPGA system features are described in this manual.

Designing Hierarchy

Current HDL design methods are specifically written for ASIC designs. You can use some of these ASIC design methods when designing FPGAs; however, certain techniques may unnecessarily increase the number of gates or CLB levels.

Design hierarchy is important in the implementation of an FPGA and also during incremental or interactive changes. Some synthesizers maintain the hierarchical boundaries unless you group modules together. Modules should have registered outputs so their boundaries are not an impediment to optimization. Otherwise, modules should be as large as possible within the limitations of your synthesis tool. The “5,000 gates per module” rule is no longer valid, and can interfere with optimization. Check with your synthesis vendor for the current recommendations for preferred module size. As a last resort, use the grouping commands of your synthesizer, if available. The size and content of the modules influence synthesis results and design implementation. This manual describes how to create effective design hierarchy.

Specifying Speed Requirements

To meet timing requirements, you should understand how to set timing constraints in both the synthesis and placement/routing tools.

Installing Design Examples and Tactical Software

The information in this manual supplements information in your synthesis and HDL simulator manuals. Before you start designing Xilinx FPGAs, read the Xilinx-specific information in your HDL manuals. Also, read and follow the instructions in the latest version of the *Quick Start Guide for Xilinx Alliance Series*, as well as the current version of the *Alliance Series Install and Release Document*.

This manual includes numerous HDL design examples created with VHDL and Verilog. VHDL is more comprehensive than Verilog, and you may need to spend more time learning how to apply VHDL constructs to synthesis.

Software Requirements

To synthesize, simulate, and implement the design examples in this manual, you should have the current versions of your synthesis and simulation software, as well as the Alliance Series 1.4 or later version of the Xilinx Development System installed on your system.

Memory Requirements

The values provided in the following table are for typical designs, and include loading the operating system. Additional memory may be required for certain “boundary-case” or unusual designs, as well as for the concurrent operation of other applications (for example, synthesis or HDL simulation). Xilinx recommends compiling XC4000EX/ XL designs on the Ultra Sparc, HP715, or equivalent workstations. Although 64 MB of RAM and 64 MB of swap space are required to compile XC4000EX designs, Xilinx recommends that you use 128 MB of RAM and 128 MB of swap space for more efficient processing of your XC4000EX designs.

Table 1-1 Memory Requirements for Workstations and PCs

| Xilinx Device | RAM | Swap Space |
|--|------------|-------------------|
| XC3000A/L XC3100A/L XC4000E/L XC4028EX through XC4036EX XC4002XL through XC4028XL XCS (Spartan) XC5200 XC9500 (small devices) | 64 MB | 64 MB – 128 MB |
| XC4036XL through XC4062XL XC9500 (large devices) | 128 MB | 128 MB – 256 MB |
| XC4085XL XC40125XV | 256 MB | 256 MB – 512 MB |

Disk Space Requirements

Before you install the programs and files, verify that your system meets the requirements listed in the following table for the applicable options. The disk space requirements listed are approximations and may not exactly match the actual numbers. To significantly reduce the amount of disk space needed, install only the components and documentation that you will actually use. In the following table, the Data column represents files that are common to all three workstation platforms. For example, for a Solaris machine, you need ~ 110 (12 plus 98) MB of disk space.

Note: Refer to the *Alliance Series Install and Release Document* for more information on disk space requirements.

Table 1-2 Disk Space Requirements

| Software Component | Data | Sun | Sol | HP |
|--|--|--------|--------|---------|
| Xilinx Core Technology ¹ | ~12 MB | ~96 MB | ~98 MB | ~108 MB |
| Xilinx Device Data Files (All devices) ² | ~195 MB | ~26 MB | ~26 MB | ~26 MB |
| Documentation: Online Help Documentation Browser Xilinx Tutorial Files Xilinx Userware | ~30 MB total ~17 MB ~1 MB ~4 MB | ~10 MB | ~10 MB | ~10 MB |

¹ The DynaText[®] browser and its online books are not included. If you install more than one platform, the data directory is only installed once. For example, if you install both SunOS and Solaris, the total disk space for the core technology is:
data (12 MB) + SunOS (96 MB) + Solaris (98 MB) = ~206 MB

² The memory requirements specified are for the installation of all Xilinx devices. You can significantly reduce the amount of disk space required by installing only the files for the devices you want to target.

Xilinx Internet Site

To download the programs and files from the Xilinx Internet Site, you must meet the disk requirements listed in the following table.

Table 1-3 Internet Files

| Directory/Location | Description | Compressed File | Directory Size |
|--------------------------------|---|--|----------------|
| M1_VHDL_source ¹ | All VHDL source code only (no scripts, compilation, or implementation files) | m1_vhdl_src.tar.Z (size: 60 KB) or m1_vhdl_src.zip (size: 68 KB) | 271 KB |
| M1_Verilog_source ¹ | All Verilog source code only (no scripts, compilation, or implementation files) | m1_verilog_src.tar.Z (size: 57 KB) or m1_verilog.src.zip (size: 64 KB) | 256 KB |

Table 1-3 Internet Files

| Directory/Location | Description | Compressed File | Directory Size |
|----------------------------|--|--|----------------|
| M1_HDL_source ¹ | All VHDL and Verilog source code only (no scripts, compilation, or implementation files) | m1_hdl_src.tar.Z (size: 110 KB) or m1_hdl_src.zip (size: 129 KB) | 497 KB |

¹ These files are located at <ftp://ftp.xilinx.com/pub/applications/3rdparty>

Retrieving Tactical Software and Design Examples

You can retrieve the HDL design examples from the Xilinx Internet Site. If you need assistance retrieving the files, use the information listed in the “Technical Support” section of this chapter to contact the Xilinx Hotline.

You must install the retrieved files on the same system as the Xilinx software and the synthesis and simulation tools. However, do not install the files into the directory with the current release of the software since they may get overwritten during the installation of the next version of the software.

From Xilinx Internet FTP Site

You can retrieve the programs and files from the Xilinx Internet FTP (File Transfer Protocol) site. Alternatively, if you are not familiar with FTP, you can retrieve the files by going to the Xilinx Web site (<http://www.xilinx.com>), clicking on Service and Support, and using the File Download option. To access the Xilinx FTP Site, you must either have an internet-capable FTP utility available on your machine or a Web browser that has FTP. To use FTP, your machine must be connected to the Internet and you must have permission to use FTP on remote sites. If you need more information on this procedure, contact your system administrator.

To retrieve the programs and files from the Xilinx FTP site, use the following procedure.

1. Go to the directory on your local machine where you want to download the files, as follows.

```
cd directory
```

2. Invoke the FTP utility or your Web browser that provides FTP.
3. Connect to the Xilinx FTP site, `ftp.xilinx.com` as follows.

```
ftp> open ftp.xilinx.com
```

or

Enter the following URL.

```
ftp://ftp.xilinx.com
```

4. Log into a guest account if the FTP utility or Web browser does not perform this automatically. This account gives you download privileges.

```
Name (machine:user-name): anonymous
```

```
Guest login ok, send your complete e-mail  
address as the password.
```

```
Password: your_email_address
```

5. Go to the following directory.

```
ftp> cd pub/applications/3rdparty
```

6. If you are using an FTP utility, make sure you are in binary mode.

```
ftp> bin
```

7. Retrieve the appropriate design files as follows.

```
ftp> get design_files.tar.Z
```

or

```
ftp> get design_files.zip
```

or

Select the appropriate file and select a destination directory on your local machine.

8. Extract the files as described in the next section.

Extracting the Files

You must install the retrieved files on the same system as the current release of the Xilinx software and the synthesis and simulation tools. However, do not install the files in the directory with the current software because they may get overwritten during the installation of the next version of the software. The files are stored in the UNIX™ standard tar and compress form, as well as in the PC™ standard zip form. To extract the files, use one of the following procedures.

Note: If the following procedures do not work on your system, consult your system administrator for help on extracting the files.

Extracting .tar.Z File in UNIX

1. Go to the directory where you downloaded the files.

```
cd downloaded_files
```

2. Uncompress the files.

```
uncompress design.tar.Z
```

3. Extract the files.

```
tar xvf design.tar
```

Extracting .zip File in UNIX

1. Go to the directory where you downloaded the files.

```
cd downloaded_files
```

2. Uncompress the files.

```
unzip design.zip
```

Extracting .zip File in MS-DOS™

1. Go to the directory where you downloaded the files:

```
cd downloaded_files
```

2. Uncompress the files:

```
pkunzip -d design.zip
```

Directory Tree Structure

After you have completed the installation, you should have the following directory tree structure.

```
5k_preset
  /VHDL
  /Verilog
/Async_RAM_as_latch
  /VHDL
  /Verilog
/Barrel_SR
  /VHDL
    /Barrel
    /Barrel_Org
  /Verilog
    /Barrel
    /Barrel_Org
/Bidir_LogiBLOX
  /VHDL
  /Verilog
/Bidir_infer
  /VHDL
  /Verilog
/Bidir_instantiate
  /VHDL
  /Verilog
/Bnd_scan_4k
  /VHDL
  /Verilog
/Bnd_scan_5k
  /VHDL
  /Verilog
/Case_vs_if
  /VHDL
    /Case_ex
    /If_ex
  /Verilog
    /Case_ex
    /If_ex
/Clock_enable
  /VHDL
  /Verilog
/Clock_mux
  /VHDL
  /Verilog
```

```
/Constants
  /VHDL
  /Verilog
    /Parameter1
    /Parameter2
/D_latch
  /VHDL
  /Verilog
/D_register
  /VHDL
  /Verilog
/FF_example
  /VHDL
  /Verilog
/GR_5K
  /VHDL
    /Active_low_GR
    /No_GR
    /Use_GR
  /Verilog
    /Active_low_GR
    /No_GR
    /Use_GR
/GSR
  /VHDL
    /Active_low_GSR
    /No_GSR
    /Use_GSR
  /Verilog
    /Active_low_GSR
    /No_GSR
    /Use_GSR
/Gate_clock
  /VHDL
    /Gate_clock
    /Gate_clock2
  /Verilog
    /Gate_clock
    /Gate_clock2
/IO_Decoder
  /VHDL
  /Verilog
/LogiBLOX_DP_RAM
  /VHDL
  /Verilog
```

```
/LogiBLOX_SR
  /VHDL
  /Verilog
/Mux_vs_3state
  /VHDL
    /Mux_gate
    /Mux_gate16
    /Mux_tbuf
    /Mux_tbuf16
  /Verilog
    /Mux_gate
    /Mux_gate16
    /Mux_tbuf
    /Mux_tbuf16
/Nested_if
  /VHDL
    /If_case
    /Nested_if
  /Verilog
    /If_case
    /Nested_if
/OMUX_example
  /VHDL
  /Verilog
/RAM_primitive
  /VHDL
  /Verilog
/ROM_RTL
  /VHDL
  /Verilog
/Res_sharing
  /VHDL
    /Res_no_share
    /Res_sharing
  /Verilog
    /Res_no_share
    /Res_sharing
/Set_and_Reset
  /VHDL
  /Verilog
/Sig_vs_Var
  /VHDL
    /Xor_Sig
    /Xor_Var
/State_Machine
```

```
/VHDL
  /Binary
  /Enum
  /One_Hot
/Verilog
  /Binary
  /Enum
  /One_Hot
/Unbonded_IO
  /VHDL
  /Verilog
```

Technical Support

You can contact Xilinx for additional information and assistance in the following ways.

Xilinx World Wide Web Site

Enter <http://www.xilinx.com>. Click on the Service and Support option on the Xilinx Home Page. Use the Customer Service and Support page to get answers to your technical questions. You can use the Answers Search option to search the Answers database, file download area, application notes, XCELL journals, data sheets, and expert journals.

Technical and Applications Support Hotlines

The telephone hotlines give you direct access to Xilinx Application Engineers worldwide. You can also e-mail or fax your technical questions to the same locations.

Table 1-4 Technical Support

| Location | Telephone | Electronic Mail | Facsimile (Fax) |
|-----------------------|-----------------|---------------------|-----------------|
| North America | 1-800-255-7778 | hotline@xilinx.com | 1-408-879-4442 |
| Japan | 81-3-3297-9163 | jhotline@xilinx.com | 81-3-3297-0067 |
| France | 33-1-3463-0100 | frhelp@xilinx.com | 33-1-3463-0959 |
| Germany | 49-89-9915-4930 | dlhelp@xilinx.com | 49-89-904-4748 |
| United Kingdom | 44-1932-820821 | ukhelp@xilinx.com | 44-1932-828522 |
| Corporate Switchboard | 1-408-559-7778 | | |

Note: When e-mailing or faxing inquiries, provide your complete name, company name, and phone number. Also, provide a complete problem description including your design entry software and design stage.

Xilinx FTP Site

`ftp://ftp.xilinx.com`

The FTP site provides online access to automated tutorials, design examples, online documents, utilities, and published patches.

XDOCS E-mail Server

`xdocs@xilinx.com`

Include the word “help” in the subject header. This e-mail service provides access to the Customer Service and Support page from the Xilinx World Wide Web Site. On the Xilinx home page, click on Service and Support, and use the Customer Service and Support page to get answers to your technical questions.

HDL Coding Hints

HDLs contain many complex constructs that are difficult to understand at first. Also, the methods and examples included in HDL manuals do not always apply to the design of FPGAs. If you currently use HDLs to design ASICs, your established coding style may unnecessarily increase the number of gates or CLB levels in FPGA designs.

HDL synthesis tools implement logic based on the coding style of your design. To learn how to efficiently code with HDLs, you can attend training classes, read reference and methodology notes, and refer to synthesis guidelines and templates available from Xilinx and the synthesis vendors. When coding your designs, remember that HDLs are mainly hardware description languages. You should try to find a balance between the quality of the end hardware results and the speed of simulation.

The coding hints and examples included in this chapter are not intended to teach you every aspect of VHDL or Verilog, but they should help you develop an efficient coding style.

The following sections are included in this chapter.

- “Comparing Synthesis and Simulation Results”
- “Selecting HDL Formatting Styles”
- “Using Schematic Design Hints with HDL Designs”

Comparing Synthesis and Simulation Results

VHDL and Verilog are hardware description and simulation languages that were not originally intended as input to synthesis. Therefore, many hardware description and simulation constructs are not supported by synthesis tools. In addition, the various synthesis

tools use different subsets of VHDL and Verilog. VHDL and Verilog semantics are well defined for design simulation. The synthesis tools must adhere to these semantics to ensure that designs simulate the same way before and after synthesis. Follow the guidelines presented below to create code that simulates the same way before and after synthesis.

Omit the Wait for XX ns Statement

Do not use the Wait for XX ns statement in your code. XX specifies the number of nanoseconds that must pass before a condition is executed. This statement does not synthesize to a component. In designs that include this statement, the functionality of the simulated design does not match the functionality of the synthesized design. VHDL and Verilog examples of the Wait for XX ns statement are as follows.

- VHDL

```
wait for XX ns;
```
- Verilog

```
#XX;
```

Omit the ...After XX ns or Delay Statement

Do not use the ...After XX ns statement in your VHDL code or the Delay assignment in your Verilog code. Examples of these statements are as follows.

- VHDL

```
(Q <=0 after XX ns)
```
- Verilog

```
assign #XX Q=0;
```

XX specifies the number of nanoseconds that must pass before a condition is executed. This statement is usually ignored by the synthesis tool. In this case, the functionality of the simulated design does not match the functionality of the synthesized design.

Use Case and If-Else Statements

You can use If-Else statements, Case statements, or other conditional code to create state machines or other conditional logic. These statements implement the functions differently, however, the simulated designs are identical. The If-Else statement generally specifies priority-encoded logic and the Case statement generally specifies balanced behavior. The If-Else statement can, in some cases, result in a slower circuit overall. These statements vary with the synthesis tool. Refer to the “Comparing If Statement and Case Statement” section of this chapter for more information.

Order and Group Arithmetic Functions

The ordering and grouping of arithmetic functions can influence design performance. For example, the following two VHDL statements are not necessarily equivalent.

```
ADD <= A1 + A2 + A3 + A4;  
ADD <= (A1 + A2) + (A3 + A4);
```

For Verilog, the following two statements are not necessarily equivalent.

```
ADD = A1 + A2 + A3 + A4;  
ADD = (A1 + A2) + (A3 + A4);
```

The first statement cascades three adders in series. The second statement creates two adders in parallel: $A1 + A2$ and $A3 + A4$. In the second statement, the two additions are evaluated in parallel and the results are combined with a third adder. RTL simulation results are the same for both statements, however, the second statement results in a faster circuit after synthesis (depending on the bit width of the input signals).

Although the second statement generally results in a faster circuit, in some cases, you may want to use the first statement. For example, if the $A4$ signal reaches the adder later than the other signals, the first statement produces a faster implementation because the cascaded structure creates fewer logic levels for $A4$. This structure allows $A4$ to catch up to the other signals. In this case, $A1$ is the fastest signal followed by $A2$ and $A3$; $A4$ is the slowest signal.

Most synthesis tools can balance or restructure the arithmetic operator tree if timing constraints require it. However, Xilinx recommends that you code your design for your selected structure.

Omit Initial Values

Do not assign signals and variables initial values because initial values are ignored by most synthesis tools. The functionality of the simulated design may not match the functionality of the synthesized design.

For example, do not use initialization statements like the following VHDL and Verilog statements.

- VHDL

```
variable SUM:INTEGER:=0;
```

- Verilog

```
wire SUM=1'b0;
```

Selecting HDL Formatting Styles

Because HDL designs are often created by design teams, Xilinx recommends that you agree on a style for your code at the beginning of your project. An established coding style allows you to read and understand code written by your fellow team members. Also, inefficient coding styles can adversely impact synthesis and simulation, which can result in slow circuits. Additionally, because portions of existing HDL designs are often used in new designs, you should follow coding standards that are understood by the majority of HDL designers. This section of the manual provides a list of suggested coding styles that you should establish before you begin your designs.

Selecting a Capitalization Style

Select a capitalization style for your code. Xilinx recommends using a consistent style (lower or upper case) for entity or module names in FPGA designs.

Verilog

For Verilog, the following style is recommended.

- Use lower case letters for the following.
 - Module names
 - Verilog language keywords
- Use upper case letters for the following.
 - Labels
 - Reg, wire, instance, and instantiated cell names

Note: Cell names must be upper case to use the UNISIM simulation library and certain synthesis libraries. Check with your synthesis vendor.

VHDL

Note: VHDL is case-insensitive.

For VHDL, use lower case for all language constructs from the IEEE-STD 1076. Any inputs defined by you should be upper case. For example, use upper case for the names of signals, instances, components, architectures, processes, entities, variables, configurations, libraries, functions, packages, data types, and sub-types. For the names of standard or vendor packages, the style used by the vendor or uppercase letters are used, as shown for IEEE in the following example:

```
library IEEE;  
use IEEE.std_logic_1164.all;  
signal SIG: UNSIGNED (5 downto 0);
```

Using Xilinx Naming Conventions

Use the Xilinx naming conventions listed in this section for naming signals, variables, and instances that are translated into nets, buses, and symbols.

Note: Most synthesis tools convert illegal characters to legal ones.

- User-defined names can contain A-Z, a-z, \$, _, -, <, and >. A “/” is also valid, however, it is not recommended because it is used as a hierarchy separator

- Names must contain at least one non-numeric character
- Names cannot be more than 256 characters long

The following FPGA resource names are reserved and should not be used to name nets or components.

- Components (Comps), Configurable Logic Blocks (CLBs), Input/Output Blocks (IOBs), basic elements (bels), clock buffers (BUFGs), tristate buffers (BUFTs), oscillators (OSC), CCLK, DP, GND, VCC, and RST
- CLB names such as AA, AB, and R1C2
- Primitive names such as TD0, BSCAN, M0, M1, M2, or STARTUP
- Do not use pin names such as P1 and A4 for component names
- Do not use pad names such as PAD1 for component names

Refer to the language reference manual for Verilog or VHDL for language-specific naming restrictions. Xilinx does not recommend using escape sequences for illegal characters. Also, if you plan on importing schematics into your design, use the most restrictive character set.

Matching File Names to Entity and Module Names

The VHDL or Verilog source code file name should match the designated name of the entity (VHDL) or module (Verilog) specified in your design file. This is less confusing and generally makes it easier to create a script file for the compilation of your design. Xilinx also recommends that if your design contains more than one entity or module, each should be contained in a separate file with the appropriate file name. It is also a good idea to use the same name as your top-level design file for your synthesis script file with either a .do, .scr, .script, or the appropriate default script file extension for your synthesis tool.

Naming Identifiers, Types, and Packages

You can use long (256 characters maximum) identifier names with underscores and embedded punctuation in your code. Use meaningful names for signals and variables, such as CONTROL_REGISTER. Use meaningful names when defining VHDL types and packages as shown in the following examples.


```
type LOCATION_TYPE is ...;
package STRING_IO_PKG is
```

Using Labels

Use labels to group logic. Label all processes, functions, and procedures as shown in the following examples. Labeling makes it easier to debug your code.

- VHDL

```
ASYNC_FF: process (CLK,RST)
```

- Verilog

```
always @ (posedge CLK or posedge RST)
begin: ASYNC_FF
```

Labeling Flow Control Constructs

You can use optional labels on flow control constructs to make the code structure more obvious, as shown in the following VHDL and Verilog examples. However, you should note that these labels are not translated to gate or register names in your implemented design. Flow control constructs can slow down simulations in some Verilog simulators.

- VHDL Example

```
-- D_REGISTER.VHD
-- May 1997

-- Changing Latch into a D-Register

library IEEE;
use IEEE.std_logic_1164.all;

entity d_register is
    port (CLK, DATA: in STD_LOGIC;
          Q: out STD_LOGIC);
end d_register;

architecture BEHAV of d_register is
begin
My_D_Reg: process (CLK, DATA)
begin
    if (CLK'event and CLK='1') then
```

```
        Q <= DATA;
    end if;
    end process; --End My_D_Reg
end BEHAV;
```

- Verilog Example

```
/* Changing Latch into a D-Register
 * D_REGISTER.V
 * May 1997 */

module d_register (CLK, DATA, Q);

input CLK;
input DATA;
output Q;

reg Q;

    always @ (posedge CLK)
begin: My_D_Reg
    Q <= DATA;
end

endmodule
```

Using Variables for Constants (VHDL Only)

Do not use variables for constants in your code. Define constant numeric values in your code as constants and use them by name. This coding convention allows you to easily determine if several occurrences of the same literal value have the same meaning. In some simulators, using constants allows greater optimization. In the following code example, the OPCODE values are declared as constants, and the constant names refer to their function. This method produces readable code that may be easier to modify.

Using Constants to Specify OPCODE Functions (VHDL)

```
constant ZERO      : STD_LOGIC_VECTOR (1 downto 0):="00";
constant A_AND_B   : STD_LOGIC_VECTOR (1 downto 0):="01";
constant A_OR_B    : STD_LOGIC_VECTOR (1 downto 0):="10";
constant ONE       : STD_LOGIC_VECTOR (1 downto 0):="11";
```

```

process (OPCODE, A, B)
begin
    if      (OPCODE = A_AND_B) then OP_OUT <= A and B;
    elsif  (OPCODE = A_OR_B)  then OP_OUT <= A or B;
    elsif  (OPCODE = ONE)   then OP_OUT <= '1';
    else                                OP_OUT <= '0';
    end if;
end process;

```

Using Parameters for Constants (Verilog Only)

You can specify a constant value in Verilog using the *parameter* special data type, as shown in the following examples. The first example includes a definition of OPCODE constants as shown in the previous VHDL example. The second example shows how to use a parameter statement to define module bus widths.

Using Parameters to Specify OPCODE Functions (Verilog)

```

parameter ZERO = 2'b00;
parameter A_AND_B = 2'b01;
parameter A_OR_B = 2'b10;
parameter ONE = 2'b11;

always @ (OPCODE or A or B)
begin
    if      (OPCODE==ZERO)      OP_OUT=1'b0;
    else if(OPCODE==A_AND_B)   OP_OUT=A&B;
    else if(OPCODE==A_OR_B)   OP_OUT=A|B;
    else                        OP_OUT=1'b1;
end

```

Using Parameters to Specify Bus Size (Verilog)

```

parameter BUS_SIZE = 8;

output [BUS_SIZE-1:0] OUT;
input  [BUS_SIZE-1:0] X,Y;

```

Using Named and Positional Association

Use positional association in function and procedure calls, and in port lists only when you assign all items in the list. Use named association when you assign only some of the items in the list. Also, Xilinx

suggests that you use named association to prevent incorrect connections for the ports of instantiated components. Do not combine positional and named association in the same statement as illustrated in the following examples.

- VHDL

| | |
|-----------|--|
| Incorrect | <code>CLK_1: BUFGS port map (I=>CLOCK_IN,CLOCK_OUT);</code> |
| Correct | <code>CLK_1: BUFGS port map (I=>CLOCK_IN,O=>CLOCK_OUT);</code> |

- Verilog

| | |
|-----------|---|
| Incorrect | <code>BUFGS CLK_1 (.I(CLOCK_IN), CLOCK_OUT);</code> |
| Correct | <code>BUFGS CLK_1 (.I(CLOCK_IN), .O(CLOCK_OUT));</code> |

Managing Your Design

As part of your coding specifications, you should include rules for naming, organizing, and distributing your files. In VHDL designs, use explicit configurations to control the selection of components and architectures that you want to compile, simulate, or synthesize. In some synthesis tools, configuration information is ignored. In this case, you only need to compile the architecture that you want to synthesize.

Creating Readable Code

Use the recommendations in this section to create code that is easy to read.

Indenting Your Code

Indent blocks of code to align related statements. You should define the number of spaces for each indentation level and specify whether the Begin statement is placed on a line by itself. In the examples in this manual, each level of indentation is four spaces and the Begin statement is on a separate line that is not indented from the previous line of code. The examples below illustrate the indentation style used in this manual.

- VHDL Example

```
-- D_LATCH.VHD
-- May 1997

library IEEE;
use IEEE.std_logic_1164.all;

entity d_latch is
    port (GATE, DATA: in STD_LOGIC;
          Q: out STD_LOGIC);
end d_latch;

architecture BEHAV of d_latch is
begin
    LATCH: process (GATE, DATA)
    begin
        if (GATE = '1') then
            Q <= DATA;
        end if;
    end process; -- end LATCH

end BEHAV;
```

- Verilog Example

```
/* Transparent High Latch
 * D_LATCH.V
 * May 1997 */

module d_latch (GATE, DATA, Q);

input GATE;
input DATA;
output Q;

reg Q;

    always @ (GATE or DATA)
begin: LATCH
    if (GATE == 1'b1)
        Q <= DATA;
    end // End Latch

endmodule
```

Using Empty Lines

Use empty lines to separate top-level constructs, designs, architectures, configurations, processes, subprograms, and packages.

Using Spaces

Use spaces to make your code easier to read. You can omit or use spaces between signal names as shown in the following examples.

- VHDL Example

```
process (RST,CLOCK,LOAD,CE)
process (RST, CLOCK, LOAD, CE)
```

- Verilog Example

```
module test (A,B,C)
module test (A, B, C)
```

Use a space after colons as shown in the following examples.

- VHDL Example

```
signal QOUT: STD_LOGIC_VECTOR (3 downto 0);
CLK_1: BUFGS port map (I=>CLOCK_IN,O=>CLOCK_OUT);
```

- Verilog Example

```
begin: CPU_DATA
```

Breaking Long Lines of Code

Break long lines of code at an appropriate point, such as at a comma, a colon, or a parenthesis to make your code easier to read, as illustrated in the following code fragments.

- VHDL Example

```
U1: load_reg port map
(INX=>A,LOAD=>LD,CLK=>SCLK,OUTX=>B);
```

- Verilog Example

```
load_reg U1
(.INX(A), .LOAD(LD), .CLK(SCLK), .OUTX(B));
```

Adding Comments

Add comments to your code to improve readability, reduce debugging time, and make it easier to maintain your code.

- VHDL Example

```
-- Read Counter (16-bit)
-- Updated 1-25-98 to add Clock Enable, John Doe
-- Updated 1-28-98 to add Terminal Count, Joe Cool

process (RST, CLOCK, CE)
begin
.
.
.
```

- Verilog Example

```
// Read Counter (16-bit)
// Updated 1-25-98 to add Clock Enable, John Doe
// Updated 1-28-98 to add Terminal Count, Joe Cool

always @ (posedge RST or posedge CLOCK)
begin
.
.
.
```

Using Std_logic Data Type (VHDL only)

The Std_logic (IEEE 1164) type is recommended for hardware descriptions for the following reasons.

- It has nine different values that represent most of the states found in digital circuits.
- Automatically initialized to an unknown value. This automatic initialization is important for HDL designs because it forces you to initialize your design to a known state, which is similar to what is required in a schematic design. Do not override this feature by initializing signals and variables to a known value when they are declared because the result may be a gate-level circuit that cannot be initialized to a known value.
- Easy to perform a board-level simulation. For example, if you use an integer type for ports for one circuit and standard logic for ports for another circuit, your design can be synthesized;

however, you will need to perform time-consuming type conversions for a board-level simulation.

The back-annotated netlist from Xilinx implementation is in `Std_logic`. If you do not use `Std_logic` type to drive your top-level entity in the testbench, you cannot reuse your functional testbench for timing simulation. Some synthesis tools can create a wrapper for type conversion between the two top-level entities; however, this is not recommended by Xilinx.

Declaring Ports

Xilinx recommends that you use the `Std_logic` package for all entity port declarations. This package makes it easier to integrate the synthesized netlist back into the design hierarchy without requiring conversion functions for the ports. A VHDL example using the `Std_logic` package for port declarations is shown below.

```
Entity alu is
    port( A : in STD_LOGIC_VECTOR(3 downto 0);
          B : in STD_LOGIC_VECTOR(3 downto 0);
          CLK : in STD_LOGIC;
          C : out STD_LOGIC_VECTOR(3 downto 0) );
end alu;
```

Since the `downto` convention for vectors is supported in a back-annotated netlist, the RTL and synthesized netlists should use the same convention if you are using the same test bench. This is necessary because of the loss of directionality when your design is synthesized to an EDIF or XNF netlist.

Minimizing the Use of Ports Declared as Buffers

Do not use buffers when a signal is used internally and as an output port. In the following VHDL example, signal `C` is used internally and as an output port.

```
Entity alu is
    port( A : in STD_LOGIC_VECTOR(3 downto 0);
          B : in STD_LOGIC_VECTOR(3 downto 0);
          CLK : in STD_LOGIC;
          C : buffer STD_LOGIC_VECTOR(3 downto 0) );
end alu;

architecture BEHAVIORAL of alu is
begin
```



```
process begin
    if (CLK'event and CLK='1') then
        C <= UNSIGNED(A) + UNSIGNED(B) + UNSIGNED(C);
    end if;
end process;
end BEHAVIORAL;
```

Because signal C is used both internally and as an output port, every level of hierarchy in your design that connects to port C must be declared as a buffer. However, buffer types are not commonly used in VHDL designs because they can cause problems during synthesis. To reduce the amount of buffer coding in hierarchical designs, you can insert a dummy signal and declare port C as an output, as shown in the following VHDL example.

```
Entity alu is
    port( A : in STD_LOGIC_VECTOR(3 downto 0);
          B : in STD_LOGIC_VECTOR(3 downto 0);
          CLK : in STD_LOGIC;
          C : out STD_LOGIC_VECTOR(3 downto 0));
end alu;

architecture BEHAVIORAL of alu is
    -- dummy signal
    signal C_INT : STD_LOGIC_VECTOR(3 downto 0);
begin
    C <= C_INT;
    process begin
        if (CLK'event and CLK='1') then
            C_INT <= UNSIGNED(A) + UNSIGNED(B) +
                UNSIGNED(C_INT);
        end if;
    end process;
end BEHAVIORAL;
```

Comparing Signals and Variables (VHDL only)

You can use signals and variables in your designs. Signals are similar to hardware and are not updated until the end of a process. Variables are immediately updated and, as a result, can effect the functioning of your design. Xilinx recommends using signals for hardware descriptions; however, variables allow quick simulation.

The following VHDL examples show a synthesized design that uses signals and variables, respectively. These examples are shown implemented with gates in the “Gate implementation of XOR_SIG” figure and the “Gate Implementation of XOR_VAR” figure.

Note: If you assign several values to a signal in one process, only the final value is used. When you assign a value to a variable, the assignment takes place immediately. A variable maintains its value until you specify a new value.

Using Signals (VHDL)

```
-- XOR_SIG.VHD
-- May 1997
Library IEEE;
use IEEE.std_logic_1164.all;

entity xor_sig is
  port (A, B, C: in  STD_LOGIC;
        X, Y: out STD_LOGIC);
end xor_sig;

architecture SIG_ARCH of xor_sig is
  signal D: STD_LOGIC;
begin
  SIG:process (A,B,C)
  begin
    D <= A; -- ignored !!
    X <= C xor D;
    D <= B; -- overrides !!
    Y <= C xor D;
  end process;
end SIG_ARCH;
```




Figure 2-1 Gate implementation of XOR_SIG

Using Variables (VHDL)

```

-- XOR_VAR.VHD
-- May 1997


Library IEEE;
use IEEE.std_logic_1164.all;
use IEEE.std_logic_unsigned.all;

entity xor_var is
    port (A, B, C: in  STD_LOGIC;
          X, Y:      out STD_LOGIC);
end xor_var;

architecture VAR_ARCH of xor_var is
begin

    VAR:process (A,B,C)
        variable D: STD_LOGIC;
    begin
        D := A;
        X <= C xor D;
        D := B;
        Y <= C xor D;
    end process;
end VAR_ARCH;

```



X8543

Figure 2-2 Gate Implementation of XOR_VAR

Using Schematic Design Hints with HDL Designs

This section describes how to apply schematic entry design strategies to HDL designs.

Barrel Shifter Design

The schematic version of the barrel shifter design is included in the “Multiplexers and Barrel Shifters in XC3000/XC3100” application note (XAPP 026.001) available on the Xilinx web site at <http://www.xilinx.com>. In this example, two levels of multiplexers are used to increase the speed of a 16-bit barrel shifter. This design is for XC3000 and XC3100 device families; however, it can also be used for other Xilinx devices.

The following VHDL and Verilog examples show a 16-bit barrel shifter implemented using sixteen 16-to-1 multiplexers, one for each output. A 16-to-1 multiplexer is a 20-input function with 16 data inputs and four select inputs. When targeting an FPGA device based on 4-input lookup tables (such as XC4000 and XC3000 family of devices), a 20-input function requires at least five logic blocks. Therefore, the minimum design size is 80 (16 x 5) logic blocks.

16-bit Barrel Shifter (VHDL)

```
-----  
-- VHDL Model for a 16-bit Barrel Shifter      --  
-- barrel_org.vhd                             --  
-- !!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!    --  
-- THIS EXAMPLE IS FOR COMPARISON ONLY       --  
-- May 1997                                  --  
-- USE barrel.vhd                            --  
-----  
  
library IEEE;  
use IEEE.std_logic_1164.all;  
use IEEE.std_logic_arith.all;  
  
entity barrel_org is  
    port (S:in      STD_LOGIC_VECTOR (3 downto 0);  
          A_P:in    STD_LOGIC_VECTOR (15 downto 0);  
          B_P:out   STD_LOGIC_VECTOR (15 downto 0));  
end barrel_org;  
  
architecture RTL of barrel_org is  
  
begin  
    SHIFT: process (S, A_P)  
    begin  
        case S is
```

```
when "0000" =>
    B_P <= A_P;

when "0001" =>
    B_P(14 downto 0) <= A_P(15 downto 1);
    B_P(15) <= A_P(0);

when "0010" =>
    B_P(13 downto 0) <= A_P(15 downto 2);
    B_P(15 downto 14) <= A_P(1 downto 0);

when "0011" =>
    B_P(12 downto 0) <= A_P(15 downto 3);
    B_P(15 downto 13) <= A_P(2 downto 0);

when "0100" =>
    B_P(11 downto 0) <= A_P(15 downto 4);
    B_P(15 downto 12) <= A_P(3 downto 0);

when "0101" =>
    B_P(10 downto 0) <= A_P(15 downto 5);
    B_P(15 downto 11) <= A_P(4 downto 0);

when "0110" =>
    B_P(9 downto 0) <= A_P(15 downto 6);
    B_P(15 downto 10) <= A_P(5 downto 0);

when "0111" =>
    B_P(8 downto 0) <= A_P(15 downto 7);
    B_P(15 downto 9) <= A_P(6 downto 0);

when "1000" =>
    B_P(7 downto 0) <= A_P(15 downto 8);
    B_P(15 downto 8) <= A_P(7 downto 0);

when "1001" =>
    B_P(6 downto 0) <= A_P(15 downto 9);
    B_P(15 downto 7) <= A_P(8 downto 0);

when "1010" =>
    B_P(5 downto 0) <= A_P(15 downto 10);
    B_P(15 downto 6) <= A_P(9 downto 0);

when "1011" =>
    B_P(4 downto 0) <= A_P(15 downto 11);
    B_P(15 downto 5) <= A_P(10 downto 0);
```

```
        when "1100" =>
            B_P(3 downto 0)  <= A_P(15 downto 12);
            B_P(15 downto 4) <= A_P(11 downto 0);

        when "1101" =>
            B_P(2 downto 0)  <= A_P(15 downto 13);
            B_P(15 downto 3) <= A_P(12 downto 0);

        when "1110" =>
            B_P(1 downto 0)  <= A_P(15 downto 14);
            B_P(15 downto 2) <= A_P(13 downto 0);

        when "1111" =>
            B_P(0) <= A_P(15);
            B_P(15 downto 1) <= A_P(14 downto 0);

        when others =>
            B_P <= A_P;
        end case;
    end process; -- End SHIFT

end RTL;
```

16-bit Barrel Shifter (Verilog)

```
////////////////////////////////////
// BARREL_ORG.V Version 1.0 //
// Xilinx HDL Synthesis Design Guide //
// Unoptimized model for a 16-bit Barrel Shifter //
// THIS EXAMPLE IS FOR COMPARISON ONLY //
// Use BARREL.V //
// January 1998 //
////////////////////////////////////

module barrel_org (S, A_P, B_P);

    input  [3:0] S;
    input  [15:0] A_P;
    output [15:0] B_P;

    reg    [15:0] B_P;

    always @ (A_P or S)
        begin
            case (S)

```

```
4'b0000 : // Shift by 0
begin
  B_P <= A_P;
end

4'b0001 : // Shift by 1
begin
  B_P[15] <= A_P[0];
  B_P[14:0] <= A_P[15:1];
end

4'b0010 : // Shift by 2
begin
  B_P[15:14] <= A_P[1:0];
  B_P[13:0] <= A_P[15:2];
end

4'b0011 : // Shift by 3
begin
  B_P[15:13] <= A_P[2:0];
  B_P[12:0] <= A_P[15:3];
end

4'b0100 : // Shift by 4
begin
  B_P[15:12] <= A_P[3:0];
  B_P[11:0] <= A_P[15:4];
end

4'b0101 : // Shift by 5
begin
  B_P[15:11] <= A_P[4:0];
  B_P[10:0] <= A_P[15:5];
end

4'b0110 : // Shift by 6
begin
  B_P[15:10] <= A_P[5:0];
  B_P[9:0] <= A_P[15:6];
end

4'b0111 : // Shift by 7
begin
  B_P[15:9] <= A_P[6:0];
  B_P[8:0] <= A_P[15:7];
end
```

```
4'b1000 : // Shift by 8
begin
  B_P[15:8]  <= A_P[7:0];
  B_P[7:0]   <= A_P[15:8];
end

4'b1001 : // Shift by 9
begin
  B_P[15:7]  <= A_P[8:0];
  B_P[6:0]   <= A_P[15:9];
end

4'b1010 : // Shift by 10
begin
  B_P[15:6]  <= A_P[9:0];
  B_P[5:0]   <= A_P[15:10];
end

4'b1011 : // Shift by 11
begin
  B_P[15:5]  <= A_P[10:0];
  B_P[4:0]   <= A_P[15:11];
end

4'b1100 : // Shift by 12
begin
  B_P[15:4]  <= A_P[11:0];
  B_P[3:0]   <= A_P[15:12];
end

4'b1101 : // Shift by 13
begin
  B_P[15:3]  <= A_P[12:0];
  B_P[2:0]   <= A_P[15:13];
end

4'b1110 : // Shift by 14
begin
  B_P[15:2]  <= A_P[13:0];
  B_P[1:0]   <= A_P[15:14];
end

4'b1111 : // Shift by 15
begin
  B_P[15:1]  <= A_P[14:0];
  B_P[0]     <= A_P[15];
```



```

        end
    default :
        B_P      <= A_P;
    endcase
end
endmodule

```

The following modified VHDL and Verilog designs use two levels of multiplexers and are twice as fast as the previous designs. These designs are implemented using 32 4-to-1 multiplexers arranged in two levels of sixteen. The first level rotates the input data by 0, 1, 2, or 3 bits and the second level rotates the data by 0, 4, 8, or 12 bits. Since you can build a 4-to-1 multiplexer with a single CLB, the minimum size of this version of the design is 32 (32 x 1) CLBs.

16-bit Barrel Shifter with Two Levels of Multiplexers (VHDL)

```

-- BARREL.VHD
-- Based on XAPP 26 (see http://www.xilinx.com)
-- 16-bit barrel shifter (shift right)
-- May 1997

library IEEE;
use IEEE.std_logic_1164.all;
use IEEE.std_logic_arith.all;

entity barrel is
port (S:   in STD_LOGIC_VECTOR(3 downto 0);
      A_P: in STD_LOGIC_VECTOR(15 downto 0);
      B_P: out STD_LOGIC_VECTOR(15 downto 0));
end barrel;

architecture RTL of barrel is

signal SEL1,SEL2: STD_LOGIC_VECTOR(1 downto 0);
signal C:         STD_LOGIC_VECTOR(15 downto 0);

begin
    FIRST_LVL: process (A_P, SEL1)
    begin
        case SEL1 is
            when "00" => -- Shift by 0

```

```
        C <= A_P;

    when "01" => -- Shift by 1
        C(15) <= A_P(0);
        C(14 downto 0) <= A_P(15 downto 1);

    when "10" => -- Shift by 2
        C(15 downto 14) <= A_P(1 downto 0);
        C(13 downto 0) <= A_P(15 downto 2);

    when "11" => -- Shift by 3
        C(15 downto 13) <= A_P(2 downto 0);
        C(12 downto 0) <= A_P(15 downto 3);

    when others =>
        C <= A_P;
    end case;
end process; --End FIRST_LVL

SECND_LVL: process (C, SEL2)
begin
    case SEL2 is
        when "00" => --Shift by 0
            B_P <= C;

        when "01" => --Shift by 4
            B_P(15 downto 12) <= C(3 downto 0);
            B_P(11 downto 0) <= C(15 downto 4);

        when "10" => --Shift by 8
            B_P(7 downto 0) <= C(15 downto 8);
            B_P(15 downto 8) <= C(7 downto 0);

        when "11" => --Shift by 12
            B_P(3 downto 0) <= C(15 downto 12);
            B_P(15 downto 4) <= C(11 downto 0);

        when others =>
            B_P <= C;
    end case;
end process; -- End SECOND_LVL

SEL1 <= S(1 downto 0);
SEL2 <= S(3 downto 2);

end rtl;
```

16-bit Barrel Shifter with Two Levels of Multiplexers (Verilog)

```
/*
 * BARREL.V
 * XAPP 26 http://www.xilinx.com
 * 16-bit barrel shifter [shift right]
 * May 1997
 */

module barrel (S, A_P, B_P);

input [3:0] S;
input [15:0] A_P;
output [15:0] B_P;

reg [15:0] B_P;

wire [1:0] SEL1, SEL2;
reg [15:0] C;

assign SEL1 = S[1:0];
assign SEL2 = S[3:2];

always @ (A_P or SEL1)
begin
    case (SEL1)
        2'b00 : // Shift by 0
            begin
                C <= A_P;
            end

        2'b01 : // Shift by 1
            begin
                C[15] <= A_P[0];
                C[14:0] <= A_P[15:1];
            end

        2'b10 : // Shift by 2
            begin
                C[15:14] <= A_P[1:0];
                C[13:0] <= A_P[15:2];
            end

        2'b11 : // Shift by 3
    end
end
```

```
begin
C[15:13] <= A_P[2:0];
C[12:0] <= A_P[15:3];
end

default :
    C <= A_P;
endcase
end

always @ (C or SEL2)
begin
    case (SEL2)
        2'b00 : // Shift by 0
            begin
                B_P <= C;
            end

        2'b01 : // Shift by 4
            begin
                B_P[15:12] <= C[3:0];
                B_P[11:0] <= C[15:4];
            end

        2'b10 : // Shift by 8
            begin
                B_P[7:0] <= C[15:8];
                B_P[15:8] <= C[7:0];
            end

        2'b11 : // Shift by 12
            begin
                B_P[3:0] <= C[15:12];
                B_P[15:4] <= C[11:0];
            end

        default :
            B_P <= C;
    endcase
end

endmodule
```

When these two designs are implemented in an XC4005E-2 device with a popular synthesis tool, there is a 64% improvement in the gate count (88 occupied CLBs reduced to 32 occupied CLBs) in the

barrel.vhd design as compared to the barrel_org.vhd design. Additionally, there is a 19% improvement in speed from 35.58 ns (5 logic levels) to 28.85 ns (4 logic levels).

Implementing Latches and Registers

Synthesizers infer latches from incomplete conditional expressions, such as an If statement without an Else clause. This can be problematic for FPGA designs because not all FPGA devices have latches available in the CLBs. In addition, you may think that a register is created, and the synthesis tool actually created a latch. The XC4000EX/XL and XC5200 FPGAs do have registers that can be configured to act as latches. For these devices, synthesizers infer a dedicated latch from incomplete conditional expressions. XC4000E, XC3100A, and XC3000A devices do not have latches in their CLBs. For these devices, latches described in RTL code are implemented with gates in the CLB function generators. For XC4000E devices, if the latch is directly connected to an input port, it is implemented in an IOB as a dedicated input latch. For example, the D latch described in the following VHDL and Verilog designs is implemented with one function generator as shown in the “D Latch Implemented with Gates” figure.

D Latch Inference

- VHDL Example

```
-- D_LATCH.VHD
-- May 1997

library IEEE;
use IEEE.std_logic_1164.all;

entity d_latch is
    port (GATE, DATA: in STD_LOGIC;
          Q: out STD_LOGIC);
end d_latch;

architecture BEHAV of d_latch is
begin
    LATCH: process (GATE, DATA)
    begin
        if (GATE = '1') then
            Q <= DATA;
        end if;
    end process;
end architecture;
```

```

        end if;
    end process; -- end LATCH

end BEHAV;

```

- Verilog Example

```

/* Transparent High Latch
 * D_LATCH.V
 * May 1997
 */

module d_latch (GATE, DATA, Q);

input GATE;
input DATA;
output Q;

reg Q;

always @ (GATE or DATA)
begin
    if (GATE == 1'b1)
        Q <= DATA;
    end // End Latch

endmodule

```

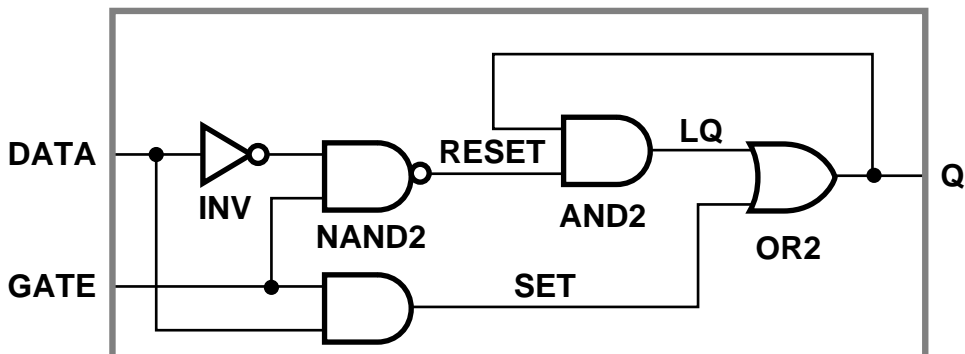


Figure 2-3 D Latch Implemented with Gates

In this example, a combinational loop results in a hold-time requirement on DATA with respect to GATE. Since most synthesis tools do not process hold-time requirements because of the uncertainty of routing delays, Xilinx does not recommend implementing latches

with combinatorial feedback loops. A recommended method for implementing latches is described in this section.

To eliminate this possible problem, use D registers instead of latches. For example, to convert the D latch to a D register, use an Else statement or modify the code to resemble the following example.

Converting a D Latch to a D Register

- VHDL Example

```
-- D_REGISTER.VHD
-- May 1997

-- Changing Latch into a D-Register

library IEEE;
use IEEE.std_logic_1164.all;

entity d_register is
    port (CLK, DATA: in STD_LOGIC;
          Q: out STD_LOGIC);
end d_register;

architecture BEHAV of d_register is
begin
MY_D_REG: process (CLK, DATA)
    begin
        if (CLK'event and CLK='1') then
            Q <= DATA;
        end if;
    end process; --End MY_D_REG
end BEHAV;
```

- Verilog Example

```
/* Changing Latch into a D-Register
 * D_REGISTER.V
 * May 1997 */

module d_register (CLK, DATA, Q);

input CLK;
input DATA;
output Q;
```

```

reg Q;

always @ (posedge CLK)
begin: My_D_Reg
  Q <= DATA;
end

endmodule

```

With some synthesis tools you can determine the number of latches that are implemented in your design. Check the manuals that came with your software for information on determining the number of latches in your design.

You should convert all If statements without corresponding Else statements and without a clock edge to registers. Use the recommended register coding styles in the synthesis tool documentation to complete this conversion.

In XC4000E devices, you can implement a D latch by instantiating a RAM 16x1 primitive, as illustrated in the following figure.

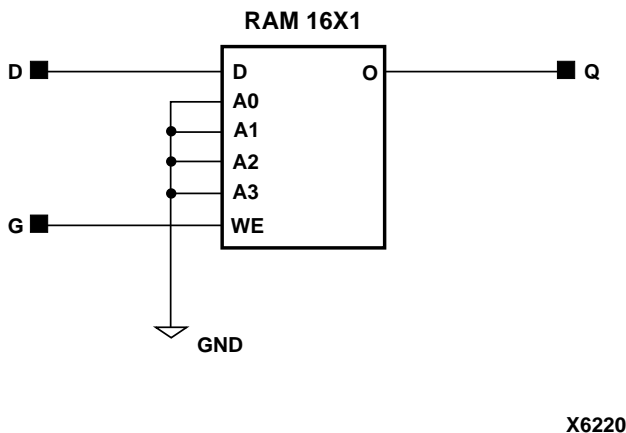


Figure 2-4 D Latch Implemented by Instantiating a RAM

In all other cases (such as latches with reset/set or enable), use a D flip-flop instead of a latch. This rule also applies to JK and SR flip-flops.

The following table provides a comparison of area and speed for a D latch implemented with gates, a 16x1 RAM primitive, and a D flip-flop.

Table 2-1 D Latch Implementation Comparison

| Comparison | Spartan, XC4000E CLB Latch Implemented with Gates | XC4000EX/XL/XV, XC5200 CLB Latch | All Spartan and XC4000 Input Latch | XC4000 E/EX/XL/XV Instantiated RAM Latch | All Families D Flip Flop |
|-------------------|---|-------------------------------------|--|--|--|
| Advantages | RTL HDL infers latch | RTL HDL infers latch, no hold times | RTL HDL infers latch, no hold times (if not specifying NODELAY, saves CLB resources) | No hold time or combinatorial loops, best for XC4000E when latch needed in CLB | No hold time or combinatorial loop. FPGAs are register abundant. |
| Disadvantages | Feedback loop results in hold time requirement, not suggested | Not available in XC4000E or Spartan | Not available in XC5200, input to latch must directly connect to port | Must be instantiated, uses logic resources | Requires change in code to convert latch to register |
| Area ^a | 1 Function Generator | 1 CLB Register/Latch | 1 IOB Register/Latch | 1 Function Generator | 1 CLB Register/Latch |

a. Area is the number of function generators and registers required. XC4000 CLBs have two function generators and two registers; XC5200 CLBs have four function generators and four register/latches.

Resource Sharing

Resource sharing is an optimization technique that uses a single functional block (such as an adder or comparator) to implement several operators in the HDL code. Use resource sharing to improve design performance by reducing the gate count and the routing congestion. If you do not use resource sharing, each HDL operation is built with

separate circuitry. However, you may want to disable resource sharing for speed critical paths in your design.

The following operators can be shared either with instances of the same operator or with an operator on the same line.

*
+ -
> >= < <=

For example, a + operator can be shared with instances of other + operators or with - operators. A * operator can be shared only with other * operators.

You can implement arithmetic functions (+, -, magnitude comparators) with gates or with your synthesis tool's module library. The library functions use modules that take advantage of the carry logic in XC4000 family, XC5200 family, and Spartan CLBs. Carry logic and its dedicated routing increase the speed of arithmetic functions that are larger than 4-bits. To increase speed, use the module library if your design contains arithmetic functions that are larger than 4-bits or if your design contains only one arithmetic function. Resource sharing of the module library automatically occurs in most synthesis tools if the arithmetic functions are in the same process.

Resource sharing adds additional logic levels to multiplex the inputs to implement more than one function. Therefore, you may not want to use it for arithmetic functions that are part of your design's time critical path.

Since resource sharing allows you to reduce the number of design resources, the device area required for your design is also decreased. The area that is used for a shared resource depends on the type and bit width of the shared operation. You should create a shared resource to accommodate the largest bit width and to perform all operations.

If you use resource sharing in your designs, you may want to use multiplexers to transfer values from different sources to a common resource input. In designs that have shared operations with the same output target, the number of multiplexers is reduced as illustrated in the following VHDL and Verilog examples. The HDL example is shown implemented with gates in the "Implementation of Resource Sharing" figure.

- VHDL Example

```
-- RES_SHARING.VHD
-- May 1997

library IEEE;
use IEEE.std_logic_1164.all;
use IEEE.std_logic_unsigned.all;
use IEEE.std_logic_arith.all;

entity res_sharing is
    port (A1,B1,C1,D1: in STD_LOGIC_VECTOR (7 downto 0);
          COND_1: in STD_LOGIC;
          Z1: out STD_LOGIC_VECTOR (7 downto 0));
end res_sharing;

architecture BEHAV of res_sharing is
begin
P1: process (A1,B1,C1,D1,COND_1)
    begin
        if (COND_1='1') then
            Z1 <= A1 + B1;
        else
            Z1 <= C1 + D1;
        end if;
    end process; -- end P1

end BEHAV;
```

- Verilog Example

```
/* Resource Sharing Example
 * RES_SHARING.V
 * May 1997 */

module res_sharing (A1, B1, C1, D1, COND_1, Z1);

input      COND_1;
input  [7:0] A1, B1, C1, D1;
output [7:0] Z1;

reg [7:0] Z1;

    always @(A1 or B1 or C1 or D1 or COND_1)
    begin
        if (COND_1)
            Z1 <= A1 + B1;
        else
            Z1 <= C1 + D1;
    end

endmodule
```

If you disable resource sharing or if you code the design with the adders in separate processes, the design is implemented using two separate modules as shown in the “Implementation without Resource Sharing” figure.

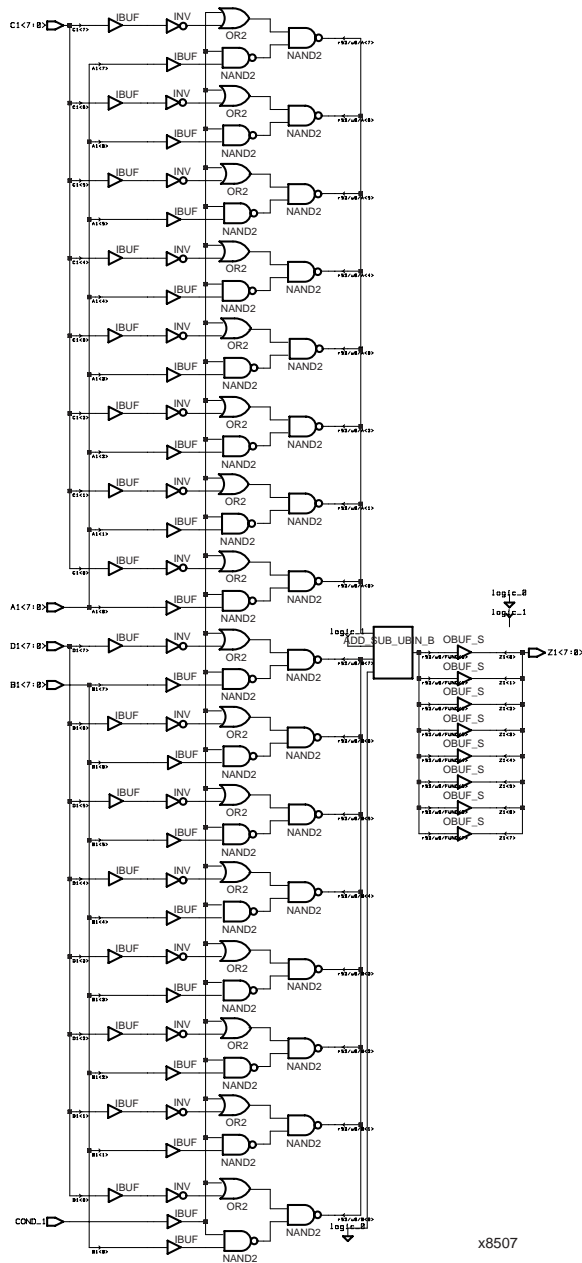


Figure 2-5 Implementation of Resource Sharing

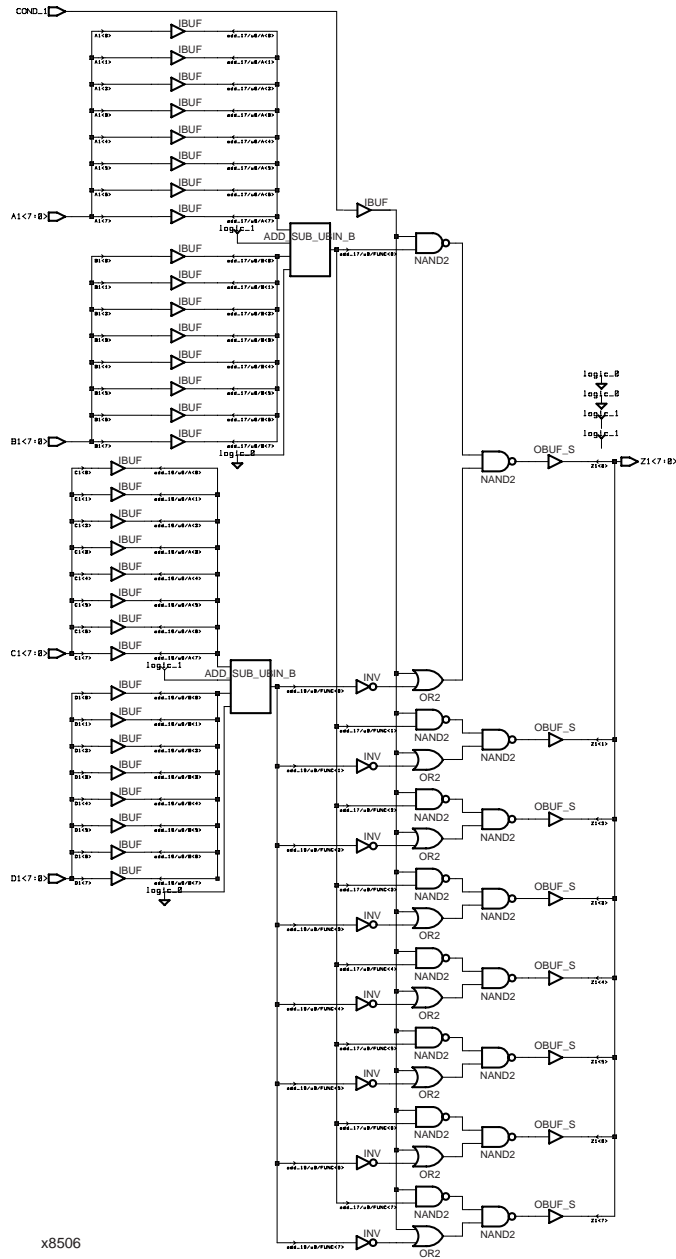


Figure 2-6 Implementation without Resource Sharing

Some synthesis tools generate modules from special Xilinx module generation algorithms. Generally, this module generation is used for operators such as adders, subtracters, incrementers, decrementers, and comparators. The following table provides a comparison of the number of CLBs used and the delay for the VHDL and Verilog designs with and without resource sharing.

Table 2-2 Resource Sharing/No Resource Sharing Comparison for XC4005EPC84-2

| Comparison | Resource Sharing with Xilinx Module Generation | No Resource Sharing with Xilinx Module Generation | Resource Sharing without Xilinx Module Generation | No Resource Sharing without Xilinx Module Generation |
|--------------------------|--|---|---|--|
| F/G Functions | 24 | 24 | 19 | 28 |
| H Function Generators | 0 | 0 | 11 | 8 |
| Fast Carry Logic CLBs | 5 | 10 | 0 | 0 |
| Longest Delay | 27.878 ns | 23.761 ns | 47.010 ns | 33.386 ns |
| Advantages/Disadvantages | Potential for area reduction | Potential for decreased critical path delay | No carry logic increases path delays | No carry logic increases CLB count |

Note: Refer to the appropriate reference manual for more information on resource sharing.

Gate Reduction

Use the generated module components to reduce the number of gates in your designs. The module generation algorithms use Xilinx carry logic to reduce function generator logic and improve routing and speed performance. Further gate reduction can occur with synthesis tools that recognize the use of constants with the modules.

Preset Pin or Clear Pin

Xilinx FPGAs consist of CLBs that contain function generators and flip-flops. The XC4000 family flip-flops have a dedicated clock enable pin and either a clear (asynchronous reset) pin or a preset (asynchronous set) pin. All synchronous preset or clear functions can be implemented with combinatorial logic in the function generators.

The XC3000 family and XC5200 family FPGAs have an asynchronous reset pin on the CLB registers. An asynchronous preset can be inferred, but is built by connecting one inverter to the D input and connecting a second inverter to the Q output of a register. In this case, an asynchronous preset is created when the asynchronous reset is activated. This may require additional logic and increase delays. If possible, the inverters are merged with existing logic connected to the register input or output.

You can configure FPGA CLB registers to have either a preset pin or a clear pin. You cannot configure the CLB register for both pins. You must modify any process that requires both pins to use only one pin or you must use two registers to implement the process. If a register is described with an asynchronous set and reset, your synthesis tool may issue an error message similar to the following during the compilation of your design.

```
Warning: Target library contains no replacement for
register 'Q_reg' (**FFGEN**) . (TRANS-4)
Warning: Cell 'Q_reg' (**FFGEN**) not translated.
(TRANS-1)
```

During the implementation of the synthesized netlist, NGDBuild issues the following error message.

```
ERROR:basnu - logical block "Q_reg" of type "_FFGEN_"
is unexpanded.
```

An XC4000 CLB is shown in the following figure.

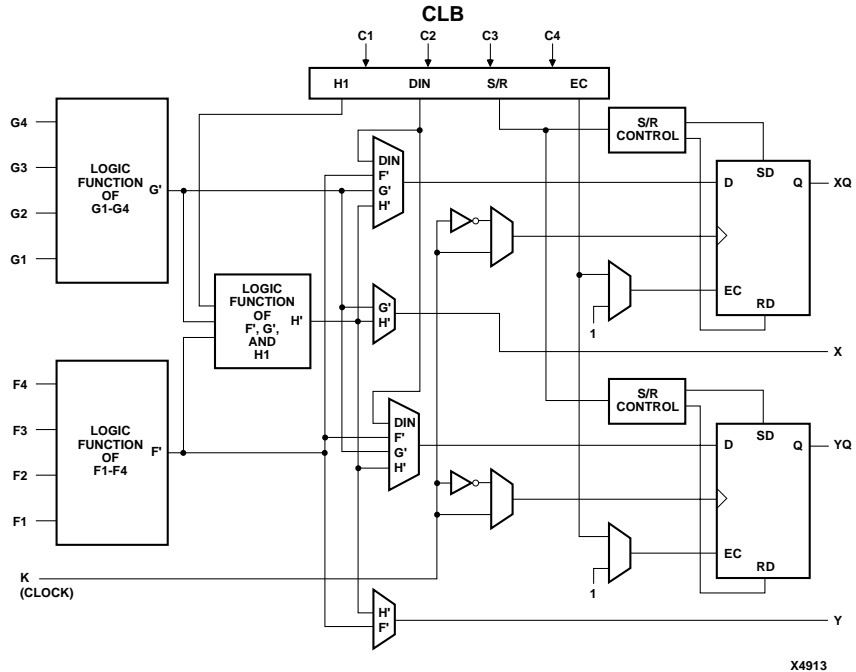


Figure 2-7 XC4000 Configurable Logic Block

The following VHDL and Verilog designs show how to describe a register with a clock enable and either an asynchronous preset or a clear.

Register Inference

- VHDL Example

```
-- FF_EXAMPLE.VHD
-- May 1997
-- Example of Implementing Registers
```

```
library IEEE;
use IEEE.std_logic_1164.all;
use IEEE.std_logic_unsigned.all;

entity ff_example is
    port ( RESET, CLOCK, ENABLE: in STD_LOGIC;
          D_IN: in STD_LOGIC_VECTOR (7 downto 0);
```

```
        A_Q_OUT: out STD_LOGIC_VECTOR (7 downto 0);
        B_Q_OUT: out STD_LOGIC_VECTOR (7 downto 0);
        C_Q_OUT: out STD_LOGIC_VECTOR (7 downto 0);
        D_Q_OUT: out STD_LOGIC_VECTOR (7 downto 0));
end ff_example;

architecture BEHAV of ff_example is
begin

    -- D flip-flop
    FF: process (CLOCK)
    begin
        if (CLOCK'event and CLOCK='1') then
            A_Q_OUT <= D_IN;
        end if;
    end process; -- End FF

    -- Flip-flop with asynchronous reset
    FF_ASYNC_RESET: process (RESET, CLOCK)
    begin
        if (RESET = '1') then
            B_Q_OUT <= "00000000";
        elsif (CLOCK'event and CLOCK='1') then
            B_Q_OUT <= D_IN;
        end if;
    end process; -- End FF_ASYNC_RESET

    -- Flip-flop with asynchronous set
    FF_ASYNC_SET: process (RESET, CLOCK)
    begin
        if (RESET = '1') then
            C_Q_OUT <= "11111111";
        elsif (CLOCK'event and CLOCK = '1') then
            C_Q_OUT <= D_IN;
        end if;
    end process; -- End FF_ASYNC_SET

    -- Flip-flop with asynchronous reset and clock
    enable
    FF_CLOCK_ENABLE: process (ENABLE, RESET, CLOCK)
    begin
        if (RESET = '1') then
            D_Q_OUT <= "00000000";
        elsif (CLOCK'event and CLOCK='1') then
            if (ENABLE='1') then
                D_Q_OUT <= D_IN;
            end if;
        end if;
    end process; -- End FF_CLOCK_ENABLE
```

```

        end if;
    end if;
end process; -- End FF_CLOCK_ENABLE

end BEHAV;

```

- **Verilog Example**

```

/* Example of Implementing Registers
 * FF_EXAMPLE.V
 * May 1997
 */

module ff_example (RESET, CLOCK, ENABLE, D_IN,
                  A_Q_OUT, B_Q_OUT, C_Q_OUT, D_Q_OUT);

input RESET, CLOCK, ENABLE;
input      [7:0] D_IN;
output     [7:0] A_Q_OUT;
output     [7:0] B_Q_OUT;
output     [7:0] C_Q_OUT;
output     [7:0] D_Q_OUT;

reg        [7:0] A_Q_OUT;
reg        [7:0] B_Q_OUT;
reg        [7:0] C_Q_OUT;
reg        [7:0] D_Q_OUT;

// D flip-flop
always @(posedge CLOCK)
begin
    A_Q_OUT <= D_IN;
end

// Flip-flop with asynchronous reset
always @(posedge RESET or posedge CLOCK)
begin
    if (RESET)
        B_Q_OUT <= 8'b00000000;
    else
        B_Q_OUT <= D_IN;
end

// Flip-flop with asynchronous set
always @(posedge RESET or posedge CLOCK)
begin
    if (RESET)

```

```
        C_Q_OUT <= 8'b11111111;
    else
        C_Q_OUT <= D_IN;
end

//Flip-flop with asynchronous reset & clock enable
always @(posedge RESET or posedge CLOCK)
begin
    if (RESET)
        D_Q_OUT <= 8'b00000000;
    else if (ENABLE)
        D_Q_OUT <= D_IN;
end

endmodule
```

Using Clock Enable Pin Instead of Gated Clocks

Use the CLB clock enable pin instead of gated clocks in your designs. Gated clocks can introduce glitches, increased clock delay, clock skew, and other undesirable effects. The first two examples in this section (VHDL and Verilog) illustrate a design that uses a gated clock. The “Implementation of Gated Clock” figure shows this design implemented with gates. Following these examples are VHDL and Verilog designs that show how you can modify the gated clock design to use the clock enable pin of the CLB. The “Implementation of Clock Enable” figure shows this design implemented with gates.

- VHDL Example

```
-----
-- GATE_CLOCK.VHD Version 1.1           --
-- Illustrates clock buffer control     --
-- Better implementation is to use      --
-- clock enable rather than gated clock --
-- May 1997                             --
-----

library IEEE;
use IEEE.std_logic_1164.all;
use IEEE.std_logic_unsigned.all;

entity gate_clock is
    port (IN1,IN2,DATA,CLK,LOAD: in STD_LOGIC;
          OUT1: out STD_LOGIC);
end gate_clock;
```

```

architecture BEHAVIORAL of gate_clock is

signal GATECLK: STD_LOGIC;

begin

GATECLK <= (IN1 and IN2 and CLK);

    GATE_PR: process (GATECLK,DATA,LOAD)
    begin
        if (GATECLK'event and GATECLK='1') then
            if (LOAD='1') then
                OUT1 <= DATA;
            end if;
        end if;
    end process; --End GATE_PR

end BEHAVIORAL;

```

- **Verilog Example**

```

////////////////////////////////////
// GATE_CLOCK.V Version 1.1          //
// Gated Clock Example               //
// Better implementation to use clock //
// enables than gating the clock     //
// May 1997                          //
////////////////////////////////////

module gate_clock(IN1, IN2, DATA, CLK, LOAD, OUT1) ;
input      IN1 ;
input      IN2 ;
input      DATA ;
input      CLK ;
input      LOAD ;
output     OUT1 ;

reg        OUT1 ;

wire GATECLK ;

assign GATECLK = (IN1 & IN2 & CLK);

always @(posedge GATECLK)
begin


```

```

        if (LOAD == 1'b1)
            OUT1 = DATA;
        end

    endmodule

```



X8628

Figure 2-8 Implementation of Gated Clock

- VHDL Example

```

-- CLOCK_ENABLE.VHD
-- May 1997

library IEEE;
use IEEE.std_logic_1164.all;
use IEEE.std_logic_unsigned.all;

entity clock_enable is
    port (IN1,IN2,DATA,CLOCK,LOAD: in STD_LOGIC;
          DOUT: out STD_LOGIC);
end clock_enable;

architecture BEHAV of clock_enable is
    signal ENABLE: STD_LOGIC;
begin

    ENABLE <= IN1 and IN2 and LOAD;

    EN_PR: process (ENABLE,DATA,CLOCK)
    begin
        if (CLOCK'event and CLOCK='1') then
            if (ENABLE='1') then
                DOUT <= DATA;
            end if;
        end if;
    end if;
end if;

```

```

        end process; -- End EN_PR

    end BEHAV;

    • Verilog Example

    /* Clock enable example
     * CLOCK_ENABLE.V
     * May 1997
     */

    module clock_enable (IN1, IN2, DATA, CLK, LOAD, DOUT);

    input IN1, IN2, DATA;
    input CLK, LOAD;
    output DOUT;

    wire ENABLE;
    reg DOUT;

    assign ENABLE = IN1 & IN2 & LOAD;

    always @(posedge CLK)
    begin
        if (ENABLE)
            DOUT <= DATA;
    end

    endmodule

```




Figure 2-9 Implementation of Clock Enable

Using If Statements

The VHDL syntax for If statements is as follows:

```

if condition then
    sequence_of_statements;

```

```
{elsif condition then
    sequence_of_statements;}
else
    sequence_of_statements;
end if;
```

The Verilog syntax for If statements is as follows:

```
if (condition)
    begin
        sequence of statements;
    end
{else if (condition)
    begin
        sequence of statements;
    end}
else
    begin
        sequence of statements;
    end
```

Use If statements to execute a sequence of statements based on the value of a condition. The If statement checks each condition in order until the first true condition is found and then executes the statements associated with that condition. After a true condition is found and the statements associated with that condition are executed, the rest of the If statement is ignored. If none of the conditions are true, and an Else clause is present, the statements associated with the Else are executed. If none of the conditions are true, and an Else clause is not present, none of the statements are executed.

If the conditions are not completely specified (as shown below), a latch is inferred to hold the value of the target signal.

- VHDL Example

```
if (L = '1') then
    Q <= D;
end if;
```
- Verilog Example

```
if (L==1'b1)
    Q=D;
```


To avoid a latch inference, specify all conditions, as shown here.

- VHDL Example

```
if (L = '1') then
    Q <= D;
else
    Q <= '0';
end if;
```

- Verilog Example

```
if (L==1'b1)
    Q=D;
else
    Q=0;
```

Using Case Statements

The VHDL syntax for Case statements is as follows.

```
case expression is
    when choices =>
        {sequence_of_statements;}
    {when choices =>
        {sequence_of_statements;}}
    when others =>
        {sequence_of_statements;}
end case;
```

The Verilog syntax for Case statements is as follows.

```
case (expression)
    choices: statement;
    {choices: statement;}
    default: statement;
endcase
```

Use Case statements to execute one of several sequences of statements, depending on the value of the expression. When the Case statement is executed, the given expression is compared to each choice until a match is found. The statements associated with the matching choice are executed. The statements associated with the Others (VHDL) or Default (Verilog) clause are executed when the given expression does not match any of the choices. The Others or Default clause is optional, however, if you do not use it, you must

include all possible values for expression. For clarity and for synthesis, each Choices statement must have a unique value for the expression. If possible, put the most likely Cases first to improve simulation speed.

Using Nested If Statements

Improper use of the Nested If statement can result in an increase in area and longer delays in your designs. Each If keyword specifies priority-encoded logic. To avoid long path delays, do not use extremely long Nested If constructs as shown in the following VHDL/Verilog examples. These designs are shown implemented in gates in the “Implementation of Nested If” figure. Following these examples are VHDL and Verilog designs that use the Case construct with the Nested If to more effectively describe the same function. The Case construct reduces the delay by approximately 3 ns (using an XC4005E-2 part). The implementation of this design is shown in the “Implementation of If-Case” figure.

Inefficient Use of Nested If Statement

- VHDL Example

```
-- NESTED_IF.VHD
-- May 1997

Library IEEE;
use IEEE.STD_LOGIC_1164.all;
use IEEE.STD_LOGIC_UNSIGNED.all;
use IEEE.STD_LOGIC_ARITH.all;

entity nested_if is
    port (ADDR_A:    in std_logic_vector (1 downto 0); -- ADDRESS Code
          ADDR_B:    in std_logic_vector (1 downto 0); -- ADDRESS Code
          ADDR_C:    in std_logic_vector (1 downto 0); -- ADDRESS Code
          ADDR_D:    in std_logic_vector (1 downto 0); -- ADDRESS Code
          RESET:     in std_logic;
          CLK :      in std_logic;
          DEC_Q:     out std_logic_vector (5 downto 0)); -- Decode OUTPUT
end nested_if;

architecture xilinx of nested_if is
begin

----- NESTED_IF PROCESS -----
```

```

NESTED_IF: process (CLK)
begin
  if (CLK'event and CLK = '1') then
    if (RESET = '0') then
      if (ADDR_A = "00") then
        DEC_Q(5 downto 4) <= ADDR_D;
        DEC_Q(3 downto 2) <= "01";
        DEC_Q(1 downto 0) <= "00";
        if (ADDR_B = "01") then
          DEC_Q(3 downto 2) <= unsigned(ADDR_A) + '1';
          DEC_Q(1 downto 0) <= unsigned(ADDR_B) + '1';
          if (ADDR_C = "10") then
            DEC_Q(5 downto 4) <= unsigned(ADDR_D) + '1';
            if (ADDR_D = "11") then
              DEC_Q(5 downto 4) <= "00";
            end if;
          else
            DEC_Q(5 downto 4) <= ADDR_D;
          end if;
        end if;
      else
        DEC_Q(5 downto 4) <= ADDR_D;
        DEC_Q(3 downto 2) <= ADDR_A;
        DEC_Q(1 downto 0) <= unsigned(ADDR_B) + '1';
      end if;
    else
      DEC_Q <= "000000";
    end if;
  end if;
end process;
end xilinx;

```

- Verilog Example

```

////////////////////////////////////
// NESTED_IF.V //
// Nested If vs. Case Design Example //
// August 1997 //
////////////////////////////////////

module nested_if (ADDR_A, ADDR_B, ADDR_C, ADDR_D, RESET, CLK, DEC_Q);

  input [1:0] ADDR_A ;
  input [1:0] ADDR_B ;
  input [1:0] ADDR_C ;
  input [1:0] ADDR_D ;

```

```
input      RESET, CLK ;
output [5:0] DEC_Q ;

reg  [5:0] DEC_Q ;

// Nested If Process //
always @ (posedge CLK)
begin
  if (RESET == 1'b1)
    begin
      if (ADDR_A == 2'b00)
        begin
          DEC_Q[5:4] <= ADDR_D;
          DEC_Q[3:2] <= 2'b01;
          DEC_Q[1:0] <= 2'b00;
          if (ADDR_B == 2'b01)
            begin
              DEC_Q[3:2] <= ADDR_A + 1'b1;
              DEC_Q[1:0] <= ADDR_B + 1'b1;
              if (ADDR_C == 2'b10)
                begin
                  DEC_Q[5:4] <= ADDR_D + 1'b1;
                  if (ADDR_D == 2'b11)
                    DEC_Q[5:4] <= 2'b00;
                end
            end
          else
            DEC_Q[5:4] <= ADDR_D;
          end
        end
      else
        DEC_Q[5:4] <= ADDR_D;
        DEC_Q[3:2] <= ADDR_A;
        DEC_Q[1:0] <= ADDR_B + 1'b1;
      end
    end
  else
    DEC_Q <= 6'b000000;
  end
endmodule
```

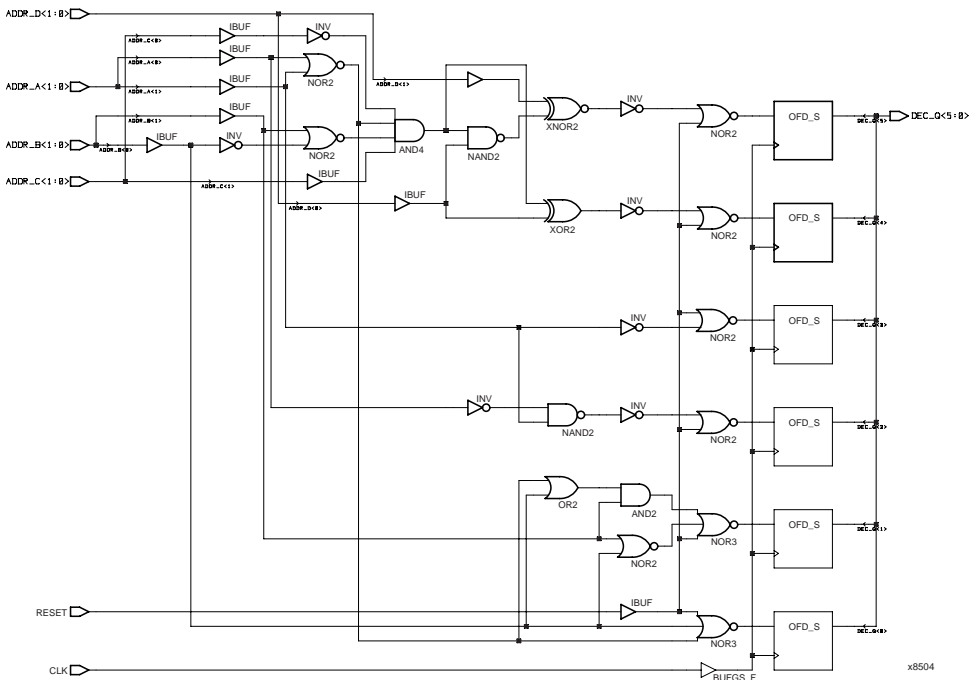


Figure 2-10 Implementation of Nested If

Nested If Example Modified to Use If-Case

Note: In the following example, the hyphens (“don’t cares”) used for bits in the Case statement may evaluate incorrectly to false for some synthesis tools.

- VHDL Example

```
-- IF_CASE.VHD
-- May 1997
```

```
Library IEEE;
use IEEE.STD_LOGIC_1164.all;
use IEEE.STD_LOGIC_UNSIGNED.all;
use IEEE.STD_LOGIC_ARITH.all;
```

```
entity if_case is
    port (ADDR_A:    in std_logic_vector (1 downto 0); -- ADDRESS Code
          ADDR_B:    in std_logic_vector (1 downto 0); -- ADDRESS Code
```

```
        ADDR_C:   in std_logic_vector (1 downto 0); -- ADDRESS Code
        ADDR_D:   in std_logic_vector (1 downto 0); -- ADDRESS Code
        RESET:    in std_logic;
        CLK :     in std_logic;
        DEC_Q:    out std_logic_vector (5 downto 0)); -- Decode OUTPUT
end if_case;
```

architecture xilinx of if_case is

```
signal ADDR_ALL : std_logic_vector (7 downto 0);
begin
```

----concatenate all address lines -----

```
ADDR_ALL <= (ADDR_A & ADDR_B & ADDR_C & ADDR_D) ;
```

-----Use 'case' instead of 'nested_if' for efficient gate netlist-----

```
IF_CASE: process (CLK)
```

```
begin
```

```
    if (CLK'event and CLK = '1') then
```

```
        if (RESET = '0') then
```

```
            case ADDR_ALL is
```

```
                when "00011011" =>
```

```
                    DEC_Q(5 downto 4) <= "00";
```

```
                    DEC_Q(3 downto 2) <= unsigned(ADDR_A) + '1';
```

```
                    DEC_Q(1 downto 0) <= unsigned(ADDR_B) + '1';
```

```
                when "000110--" =>
```

```
                    DEC_Q(5 downto 4) <= unsigned(ADDR_D) + '1';
```

```
                    DEC_Q(3 downto 2) <= unsigned(ADDR_A) + '1';
```

```
                    DEC_Q(1 downto 0) <= unsigned(ADDR_B) + '1';
```

```
                when "0001----" =>
```

```
                    DEC_Q(5 downto 4) <= ADDR_D;
```

```
                    DEC_Q(3 downto 2) <= unsigned(ADDR_A) + '1';
```

```
                    DEC_Q(1 downto 0) <= unsigned(ADDR_B) + '1';
```

```
                when "00-----" =>
```

```
                    DEC_Q(5 downto 4) <= ADDR_D;
```

```
                    DEC_Q(3 downto 2) <= "01";
```

```
                    DEC_Q(1 downto 0) <= "00";
```

```
                when others =>
```

```
                    DEC_Q(5 downto 4) <= ADDR_D;
```

```
                    DEC_Q(3 downto 2) <= ADDR_A;
```

```
                    DEC_Q(1 downto 0) <= unsigned(ADDR_B) + '1';
```

```
            end case;
```

```
        else
```

```
            DEC_Q <= "000000";
```

```
        end if;
```

```
    end if;
```

```

    end process;
end xilinx;

```

- Verilog Example

```

////////////////////////////////////
// IF_CASE.V                               //
// Nested If vs. Case Design Example       //
// August 1997                             //
////////////////////////////////////

module if_case (ADDR_A, ADDR_B, ADDR_C, ADDR_D, RESET, CLK, DEC_Q);

    input  [1:0]  ADDR_A ;
    input  [1:0]  ADDR_B ;
    input  [1:0]  ADDR_C ;
    input  [1:0]  ADDR_D ;
    input                RESET, CLK ;
    output [5:0]  DEC_Q ;

    wire  [7:0]  ADDR_ALL ;
    reg   [5:0]  DEC_Q ;

    // Concatenate all address lines //
    assign ADDR_ALL = {ADDR_A, ADDR_B, ADDR_C, ADDR_D} ;

    // Use 'case' instead of 'nested_if' for efficient gate netlist //
    always @ (posedge CLK)
    begin
        begin
            if (RESET == 1'b1)
                begin
                    casex (ADDR_ALL)
                        8'b00011011: begin
                                    DEC_Q[5:4] <= 2'b00;
                                    DEC_Q[3:2] <= ADDR_A + 1;
                                    DEC_Q[1:0] <= ADDR_B + 1'b1;
                                end
                        8'b000110xx: begin
                                    DEC_Q[5:4] <= ADDR_D + 1'b1;
                                    DEC_Q[3:2] <= ADDR_A + 1'b1;
                                    DEC_Q[1:0] <= ADDR_B + 1'b1;
                                end
                        8'b0001xxxx: begin
                                    DEC_Q[5:4] <= ADDR_D;
                                    DEC_Q[3:2] <= ADDR_A + 1'b1;
                                    DEC_Q[1:0] <= ADDR_B + 1'b1;
                                end
                    endcase
                end
            else
                DEC_Q <= 0;
            end
        end
    end

```

```

                                end
        8'b00xxxxxx: begin
            DEC_Q[5:4] <= ADDR_D;
            DEC_Q[3:2] <= 2'b01;
            DEC_Q[1:0] <= 2'b00;
        end
        default: begin
            DEC_Q[5:4] <= ADDR_D;
            DEC_Q[3:2] <= ADDR_A;
            DEC_Q[1:0] <= ADDR_B + 1'b1;
        end
    endcase
end
else
    DEC_Q <= 6'b000000;
end
endmodule
```

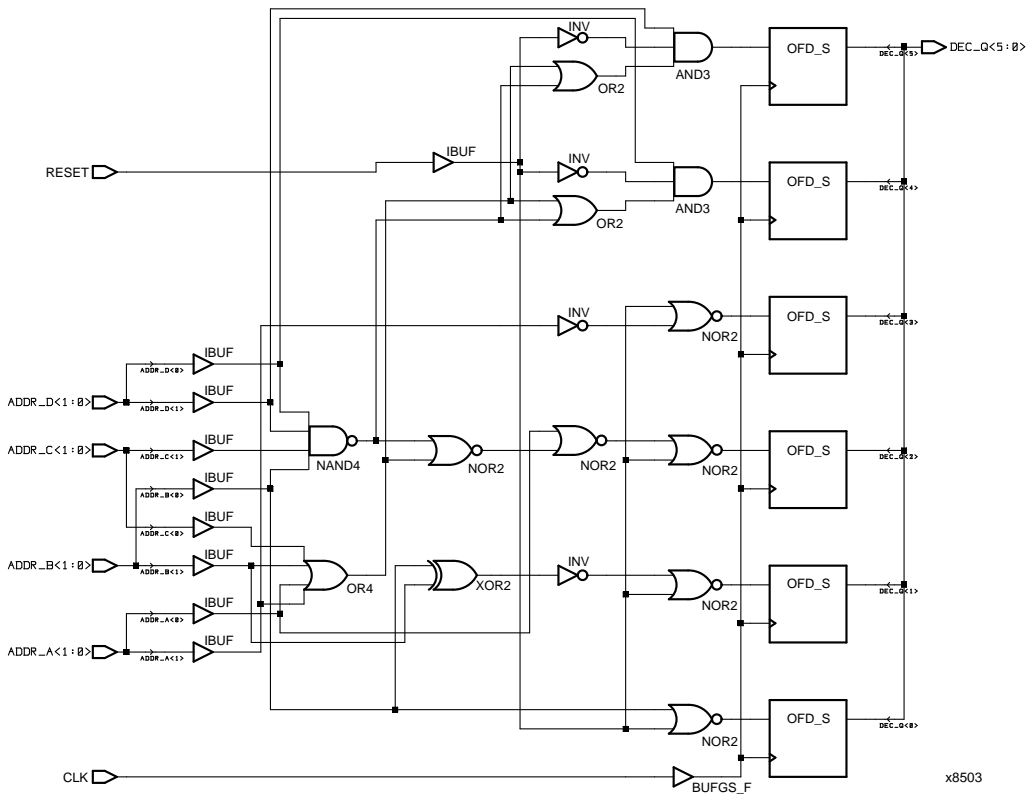



Figure 2-11 Implementation of If-Case

Comparing If Statement and Case Statement

The If statement generally produces priority-encoded logic and the Case statement generally creates balanced logic. An If statement can contain a set of different expressions while a Case statement is evaluated against a common controlling expression. In general, use the Case statement for complex decoding and use the If statement for speed critical paths.

Most current synthesis tools can determine if the if-elsif conditions are mutually exclusive, and will not create extra logic to build the priority tree.

The following examples use an If construct in a 4-to-1 multiplexer design. The “If_Ex Implementation” figure shows the implementation of these designs.

4-to-1 Multiplexer Design with If Construct

- VHDL Example

```
-- IF_EX.VHD
-- May 1997

library IEEE;
use IEEE.std_logic_1164.all;
use IEEE.std_logic_unsigned.all;

entity if_ex is
    port (SEL: in STD_LOGIC_VECTOR(1 downto 0);
          A,B,C,D: in STD_LOGIC;
          MUX_OUT: out STD_LOGIC);
end if_ex;

architecture BEHAV of if_ex is
begin

    IF_PRO: process (SEL,A,B,C,D)
    begin
        if      (SEL="00") then MUX_OUT <= A;
        elsif  (SEL="01") then MUX_OUT <= B;
        elsif  (SEL="10") then MUX_OUT <= C;
        elsif  (SEL="11") then MUX_OUT <= D;
        else
            MUX_OUT <= '0';
        end if;
    end process; --END IF_PRO

end BEHAV;
```

- Verilog Example

```
// IF_EX.V //
// Example of a If statement showing a //
// mux created using priority encoded logic //
// HDL Synthesis Design Guide for FPGAs //
// November 1997 //
////////////////////////////////////

module if_ex (A, B, C, D, SEL, MUX_OUT);

    input      A, B, C, D;
    input  [1:0] SEL;
    output     MUX_OUT;

    reg        MUX_OUT;

    always @ (A or B or C or D or SEL)
    begin
        if (SEL == 2'b00)
            MUX_OUT = A;
        else if (SEL == 2'b01)
            MUX_OUT = B;
        else if (SEL == 2'b10)
            MUX_OUT = C;
        else if (SEL == 2'b11)
            MUX_OUT = D;
        else
            MUX_OUT = 0;
    end
endmodule
```

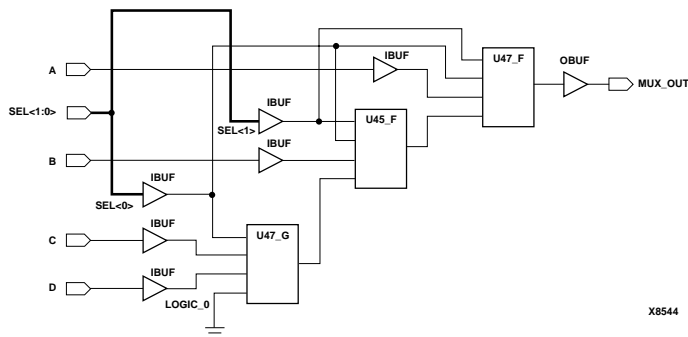


Figure 2-12 If_Ex Implementation

The following VHDL and Verilog examples use a Case construct for the same multiplexer. The “Case_Ex Implementation” figure shows the implementation of these designs. In these examples, the Case implementation requires only one XC4000 CLB while the If construct requires two CLBs in some synthesis tools. In this case, design the multiplexer using the Case construct because fewer resources are used and the delay path is shorter.

4-to-1 Multiplexer Design with Case Construct

- VHDL Example

```
-- CASE_EX.VHD
-- May 1997

library IEEE;
use IEEE.std_logic_1164.all;
use IEEE.std_logic_unsigned.all;

entity case_ex is
    port (SEL: in STD_LOGIC_VECTOR(1 downto 0);
          A,B,C,D: in STD_LOGIC;
          MUX_OUT: out STD_LOGIC);
end case_ex;

architecture BEHAV of case_ex is
begin

    CASE_PRO: process (SEL,A,B,C,D)
```

```

begin
  case SEL is
    when "00" => MUX_OUT <= A;
    when "01" => MUX_OUT <= B;
    when "10" => MUX_OUT <= C;
    when "11" => MUX_OUT <= D;
    when others => MUX_OUT <= '0';
  end case;
end process; --End CASE_PRO

```

```
end BEHAV;
```

- **Verilog Example**

```

////////////////////////////////////
// CASE_EX.V //
// Example of a Case statement showing //
// A mux created using parallel logic //
// HDL Synthesis Design Guide for FPGAs //
// November 1997 //
////////////////////////////////////

module case_ex (A, B, C, D, SEL, MUX_OUT);

input      A, B, C, D;
input [1:0] SEL;
output     MUX_OUT;

reg        MUX_OUT;

always @ (A or B or C or D or SEL)
begin
  case (SEL)
    2'b00:
      MUX_OUT = A;
    2'b01:
      MUX_OUT = B;
    2'b10:
      MUX_OUT = C;
    2'b11:
      MUX_OUT = D;
    default:
      MUX_OUT = 0;
  endcase
end

```

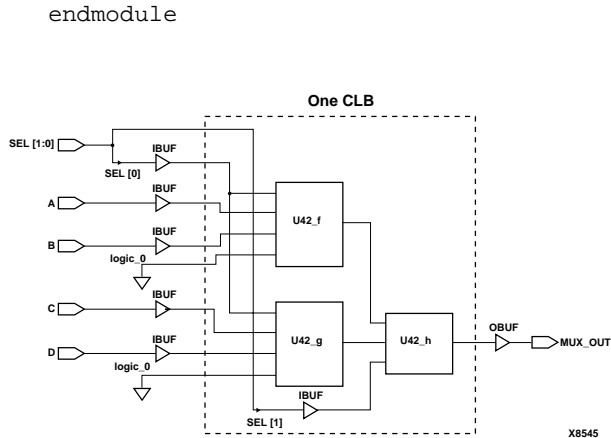


Figure 2-13 Case_Ex Implementation

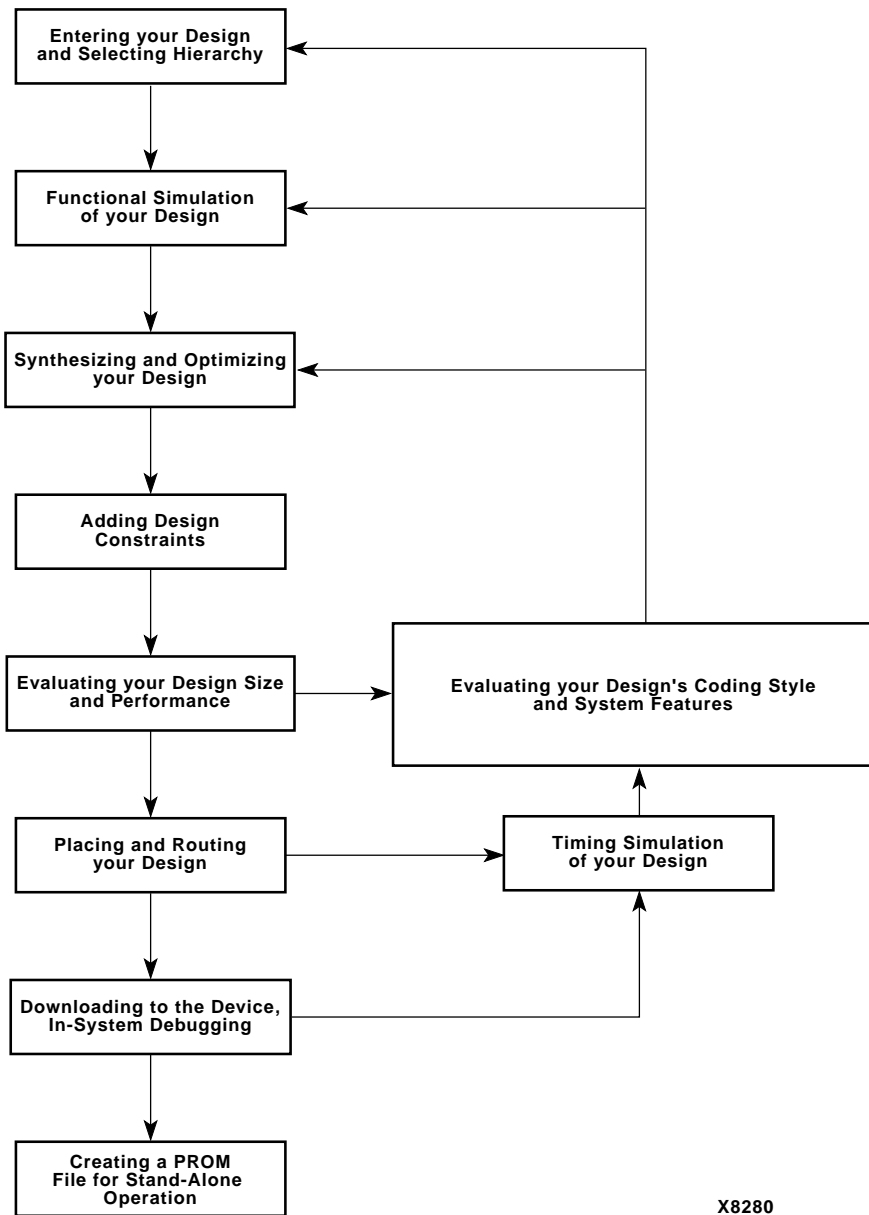
Understanding High-Density Design Flow

This chapter describes the steps in a typical HDL design flow. Although these steps may vary with each design, the information in this chapter is a good starting point for any design. If necessary, refer to the current version of the Quick Start Guide for the Xilinx Alliance Series to familiarize yourself with the Xilinx and interface tools. This chapter includes the following sections.

- “Design Flow”
- “Entering your Design and Selecting Hierarchy”
- “Functional Simulation of your Design”
- “Synthesizing and Optimizing your Design”
- “Setting Timing Constraints”
- “Evaluating Design Size and Performance”
- “Evaluating your Design for Coding Style and System Features”
- “Placing and Routing your Design”
- “Timing Simulation of your Design”
- “Downloading to the Device and In-system Debugging”
- “Creating a PROM File for Stand-Alone Operation”

Design Flow

An overview of the steps in the design flow are shown in the “Design Flow Overview” figure.



X8280

Figure 3-1 Design Flow Overview

Entering your Design and Selecting Hierarchy

The first step in implementing your design is creating the HDL code based on your design criteria.

Design Entry Recommendations

The following recommendations can help you create effective designs.

Use RTL Code

By using register transfer level (RTL) code and avoiding (when possible) instantiating specific components, you can create designs with the following characteristics.

Note: In certain cases, instantiating optimized modules, such as Logi-BLOX modules, is beneficial with RTL.

- Readable code
- Faster and simpler simulation
- Portable code for migration to different device families
- Reuse of code in future designs

Carefully Select Design Hierarchy

Selecting the correct design hierarchy is advantageous for the following reasons.

- Improves simulation and synthesis results
- Modular designs are easier to debug and modify
- Allows parallel engineering (a team of engineers can work on different parts of the design at the same time)
- Improves the placement and routing of your design by reducing routing congestion and improving timing
- Allows for easier code reuse in the current design, as well as in future designs

Functional Simulation of your Design

Use functional or RTL simulation to verify the syntax and functionality of your design. Use the following recommendations when simulating your design.

- Typically with larger hierarchical HDL designs, you should perform separate simulations on each module before testing your entire design. This makes it easier to debug your code.
- Once each module functions as expected, create a test bench to verify that your entire design functions as planned. You can use the test bench again for the final timing simulation to confirm that your design functions as expected under worst-case delay conditions.

Synthesizing and Optimizing your Design

This section includes recommendations for compiling your designs to improve your results and decrease the run time.

Note: Refer to your synthesis tool documentation for more information on compilation options and suggestions.

Creating an Initialization File

Before you can compile your design, you must create an initialization file to specify compiler defaults, and to point to the applicable implementation libraries. Refer to your synthesis tool documentation for information on creating this file.

Creating a Compile Run Script

The next step is to create a compile run script for iterative design compilations, and to use as a reference for the steps in the synthesis process. Many commonly-used synthesis tools have this capability. If you are a new user, you may want to use the graphical user interface to compile your design instead of using a run script. However, the iterative design compilation process can be tedious with the graphical interface. A run script can speed up the design process.

Compiling Your Design

Use the recommendations in this section to successfully compile your design.

Modifying your Design

You may need to modify your code to successfully compile your design because certain design constructs that are effective for simulation may not be as effective for synthesis. The synthesis syntax and code set may differ slightly from the simulator syntax and code set.

Compiling Large Designs

Older versions of synthesis tools required incremental design compilations to decrease run times. Some or all levels of hierarchy were compiled with separate compile commands and saved as output or database files. The output netlist or compiled database file for each module was read during synthesis of the top level code. This method is not necessary with new synthesis tools, which can handle large designs from the top down. The 5,000 gates per module rule of thumb no longer applies with the new synthesis tools. Refer to your synthesis tool documentation for details.

Saving Compiled Design as XNF or EDIF

After your design is successfully compiled, save it as an XNF or EDIF file for input to the Xilinx software.

Setting Timing Constraints

You can define timing specifications for your design in the User Constraints File (UCF). The UCF gives you tight control of the overall specifications by giving you access to more types of constraints; the ability to define precise timing paths; and the ability to prioritize signal constraints. Furthermore, you can group signals together to simplify timing specifications. Some synthesis tools translate certain synthesis constraints to Xilinx implementation constraints. The translated constraints are placed in a special TIMESPEC component. For more information on timing specifications in the UCF file, refer to the *Quick Start Guide for Xilinx Alliance Series*, the *Libraries Guide*, and the Answers Database on the Xilinx Web site.

Naming and Grouping Signals Together

You can name and group signals with TNMs (Timing Names) or with TIMEGRPs (Time groups). TNMs and TIMEGRPs are placed on these start and end points: ports, registers, latches, or synchronous RAMs. The new specification, TPSYNC, allows you to define an asynchronous node for a timing specification.

TNMs

Timing Names are used to identify a port, register, latch, RAM, or groups of these components for timing specifications. TNMs are specified from a UCF with the following syntax.

```
INST Instance_Name TNM=TNM_Name;
```

Instance_Name is the name given to the port, register, latch, or RAM in your design. The instance names for any port or instantiated component are provided by you in your HDL code. Inferred flip-flops and latch names can usually be determined from the log files. TNM_NAME is the arbitrary name you give the timing group.

You can include several of these statements in the UCF file with a common TNM_NAME to group elements for a timing specification as follows.

```
NET DATA TNM=INPUT_PORTS;  
NET SELECT TNM=INPUT_PORTS;
```

The above example takes two ports, DATA and SELECT, and gives them the common timing name INPUT_PORTS.

TIMEGRPs

Time Groups are another method for specifying a group of components for timing specifications.

Time groups use existing TNMs or TIMEGRPs to create new groups or to define new groups based on the output net that the group sources. There are several methods to create TIMEGRPs in the UCF file, as follows.

```
TIMEGRP TIMEGRP_Name=TNM1:TNM2;  
TIMEGRP TIMEGRP_Name=TNM3:EXCEPT:TNM4;
```

The Xilinx software recognizes the following global timing names.

- FFS — All flip-flops in your design
- PADS — All external ports in your design
- RAMS — All synchronous RAMs in your design
- LATCHES — All latches in your design

The following time group specifies the group name, FAST_FFS, which consists of all flip-flops in your design except for the ones with the TNM or TIMEGRP SLOW_FFS attribute.

```
TIMEGRP FAST_FFS=FFS:EXCEPT:SLOW_FFS;
```

TPSYNC Specification

In the latest version of the Xilinx software, you can define any node as a source or destination for a timing specification with the TPSYNC keyword. In synthesis designs, it is usually difficult to identify the net names for asynchronous paths of inferred logic. These net names can change from compile to compile, so it is not recommended to use this specification with inferred logic. However, with instantiated logic, the declared SIGNAL or WIRE name usually remains intact in the netlist and does not change from compile to compile. Some synthesis tools can preserve the signal/net name defined in the RTL through the optimization process. Check with your synthesis vendor for this capability. The UCF syntax is as follows.

```
NET Net_Name TPSYNC=TPSYNC_Name;
```

In the following NET statement, the TPSYNC is attached to the output net of a 3-state buffer, BUS3STATE. If a TPSYNC is attached to a net, then the source of the net is considered to be the endpoint (in this case, the 3-state buffer itself). The subsequent TIMESPEC statement can use the TPSYNC name just as it uses a TNM name.

```
NET BUS3STATE TPSYNC=bus3 ;  
  
TIMESPEC TSNewSpc3=FROM:PAD(ENABLE_BUS):TO:bus3:20ns ;
```

Specifying Timing Constraints

After your design signals are specified with TNMs, TIMEGRPs, or global timing names, you can place a specification on the design paths. There are a few methods for specifying these timing paths and different specifications have different priorities.

Note: Current versions of the Xilinx implementation tools have improved methods for entering timing constraints. Refer to the Xilinx documentation for your version of the place and route tools for the latest constraints commands and styles.

Period Constraint

The Period constraint specifies a clock period or clock speed on a net or clock port. The Xilinx tools attempt to meet all Pad to Setup requirements, as well as all Clock to Setup delays for registers clocked by the specified clock net. This is equivalent to a create clock type of command in a synthesis tool script.

```
NET Clock_Name PERIOD = Clock_Period ;
```

or

```
NET Clock_Name TNM=TNM_Name;
```

```
TIMESPEC TIMESPEC_Name = PERIOD:TNM_Name:Clock_Period ;
```

The following example specifies that the CLOCK port has a period of 50ns. All input paths to flip-flops clocked with this port are designated to operate at 50ns.

```
NET CLOCK PERIOD = 50;
```

FROM:TO Style Constraint

Specific paths can be specified with a FROM:TO style timing specification. These constraints are specified using global timing names, TNMs, TIMEGRPs, or TPSYNCS to connect the source and destination of the timing path, as well as the desired maximum delay of the path. An equivalent synthesis tool command is a set max delay type of command. A UCF example follows.

```
TIMESPEC TIMESPEC_Name =  
FROM:Source_Name:TO:Desination_Name:Delay_Value ;
```

TIMESPEC_Name is specified with the TS identifier followed by a number, such as TS01.

The following example specifies a new timespec with the identifier TS01 so that all paths that are sourced by a port and end at a register grouped with the name DATA_FLOPS have a delay less than 30ns.

```
TIMESPEC TS01 = FROM:PADS:TO:DATA_FLOPS:30;
```

Offset Constraint

Note: The OFFSET constraint must be used with the clock PERIOD constraint.

The OFFSET constraint can be applied to ports defined in your code. It defines the delay of a signal relative to a clock, and is only valid for registered data paths. The OFFSET constraint specifies the signal delay external to the chip, allowing the implementation tools to automatically adjust relevant internal delays (CLK buffer and distribution delays) to accommodate the external delay specified with this constraint. This constraint is equivalent to the set input delay and set output delay type of commands in your synthesis tool.

```
NET Port_Name OFFSET = {IN | OUT} Time {BEFORE | AFTER}
Clock_Name ;
```

IN | OUT specifies that the offset is calculated with respect to an input IOB or an output IOB.

For a bidirectional IOB, the IN | OUT syntax lets you specify the flow of data (input or output) on the IOB. BEFORE | AFTER indicates whether data is to arrive (input) or leave (output) the device before or after the clock input.

The following example specifies that the data on the output port, DATA_OUT, arrive on the output pin 20ns after the edge of the clock signal, CLOCK, arrives.

```
NET DATA_OUT OFFSET = OUT 20 AFTER CLOCK;
```

Ignoring Timing Paths

When a timespec is issued for a path that is not timing-critical, you can specify to ignore this path for one or all timing specifications. A TIG (Timing IGNore) can be specified on these particular nets. The synthesis tool equivalent is the Set False Path command. The UCF syntax is as follows.

```
NET Signal_Name TIG=TIMESPEC_Name ;
```

To ignore all timing constraints for a signal:

```
NET Signal_Name TIG;
```

To ignore an entire timing constraint:

```
TIG=TIMESPEC_Name ;
```

In the following example, the SLOW_PATH net is set to ignore the timing constraint with the name TS01.

```
NET SLOW_PATH TIG=TS01;
```

Controlling Signal Skew

You can control the maximum allowed skew in your designs. The maximum skew (MAXSKEW) is the difference between the longest and shortest driver-to-load connection delays for a given net. The maximum and minimum delays are determined using worst case maximum delay values for each path. While this specification cannot guarantee that this maximum skew value is achieved in the actual device, it allows the software to minimize the amount of skew on the specified signal. This specification is useful for high-fanout nets when all available global buffers have been used for other critical signals. An example of the UCF syntax for this specification follows.

```
NET Signal_Name MAXSKEW=Skew_Value ;
```

The following example specifies that the CLOCK_ENABLE signal should not have a skew value greater than 4ns.

```
NET CLOCK_ENABLE MAXSKEW=4;
```

Timing Constraint Priority

Timing constraints can be assigned priorities when paths are overlapped by multiple timing constraints. Priorities can be directly specified to a timing constraint as follows.

```
TIMESPEC TIMESPEC_Name = FROM Group1 TO Group2  
Delay_Value PRIORITY Priority_Level;
```

The lower the priority_level, the higher the precedence.

The following example sets a timespec where the source is a time group labeled THESE_FFS and the destination is labeled THOSE_FFS, with a delay value of 25ns and a priority level of 2.

```
TIMESPEC TS04=FROM:THESE_FFS:TO:THOSE_FFS:25 PRIORITY 2;
```

However, timing constraints have an inherent precedence that is based on the type of constraint and the site description provided to the tools. If two constraints are of the same priority and cover the same path, then the last constraint in the constraint file overrides any other constraints that overlap.

Inherent timing constraint priority is shown in the following table.

Note: You cannot assign a priority to override inherent timing constraint priority. You can set priorities for different timing within the same constraint type.

Table 3-1 Precedence of Constraints

| Across Constraint Sources | |
|----------------------------------|---|
| Highest Priority | Physical Constraint File (PCF) |
| | User Constraint File (UCF) |
| Lowest Priority | Input Netlist / Netlist Constraint File (NCF) |
| Within Constraint Sources | |
| Highest Priority | TIG (Timing Ignore) |
| | FROM:USER1:THRU:USER_T:TO:USER2 Specification (USER1 and USER2 are user-defined groups) |
| | FROM:USER1:THRU:USER_T:TO:FFS Specification or FROM:FFS:THRU:USER_T:TO:USER2 Specification (FFS is any pre-defined group) |
| | FROM:FFS:THRU:USER_T:TO:FFS Specification |
| | FROM:USER1:TO:USER2 Specification |
| | FROM:USER1:TO:FFS Specification or FROM:FFS:TO:USER2 Specification |
| | FROM:FFS:TO:FFS specification |
| | Period specification |
| Lowest Priority | “Allpaths” type constraints |

Evaluating Design Size and Performance

Your design should meet the following requirements.

- Design must function at the specified speed
- Design must fit in the targeted device

After your design is compiled, you can determine preliminary device utilization and performance with your synthesis tool's reporting options. After your design is mapped by the Xilinx tools, you can determine the actual device utilization. At this point in the design flow, you should verify that your design is large enough to incorporate any future changes or additions, and that your design will perform as specified.

Using your Synthesis Tool to Estimate Device Utilization and Performance

Use your synthesis tool's area and timing reporting options to estimate device utilization and performance. After compiling, use the report area command to obtain a report of device resource utilization. Some synthesis tools provide area reports automatically. Refer to your synthesis tool documentation for correct command syntax.

Note: See the "Report Files" appendix for sample report files from various synthesis vendors.

This report lists the compiled cells in your design, as well as information on how your design is mapped in the FPGA. These reports are generally accurate for the XC4000 and Spartan family because the synthesis tool creates the logic from your code and maps your design into the FPGA. However, these reports are different for the various synthesis tools. Some reports specify the minimum number of CLBs required, while other reports specify the "unpacked" number of CLBs to make an allowance for routing. For an accurate comparison, you should compare reports from the Xilinx place and route tool after implementation. Also, any instantiated components, such as Logi-BLOX modules, EDIF files, XNF files, or other components that your synthesis tool does not recognize during compilation are not included in the report file. If you include these components in your design, you must include the logic area used by these components when estimating design size. Also, sections of your design may get

trimmed during the mapping process, and may result in a smaller design.

Using the Timing Report Command

Use your synthesis tool's timing report command to obtain a report with estimated data path delays. Refer to your synthesis vendor's documentation for command syntax.

Note: See the "Report Files" appendix for sample report files from various synthesis vendors.

This report is based on the logic level delays from the cell libraries and estimated wire-load models for your design. This report is an estimate of how close you are to your timing goals; however, it is not the actual timing for your design. An accurate report of your design's timing is only available after your design is placed and routed. This timing report does not include information on any instantiated components, such as LogiBLOX modules, EDIF files, XNF files, or other components that are not recognized by your synthesis tool during compilation.

Determining Actual Device Utilization and Pre-routed Performance

To determine if your design fits the specified device, you must map it with the Xilinx Map program. The generated report file *design_name.mrp* contains the implemented device utilization information. You can run the Map program from the Design Manager or from the command line.

Using the Design Manager to Map your Design

Use the following steps to map your design using the Design Manager.

Note: For more information on using the Design Manager, see the *Design Manager/Flow Engine Reference/User Guide*.

1. To start the Design Manager, enter the following command.

```
dsgnmgr
```

2. To create a new project, select the XNF or EDIF file generated by your synthesis tool as your input file from the **File** → **New Project** menu command.
 3. To start design implementation, click the Implement toolbar button or select **Design** → **Implement**.
The Implement dialog box appears.
 4. If necessary, select a part in the dialog box.
 5. Select the Options button in the Implement dialog box.
The Options dialog box appears.
 6. Select the Produce Logic Level Timing Report option in the Optional Targets field. This option creates a timing report prior to place and route, but after map.
 7. Select the type of report you want from the Program Options Templates field by clicking the Edit Template button next to the Implementation drop-down list.
The Implementation Template dialog box appears.
 8. Select the Timing tab. Select the type of report you want to create. The default is Design Performance Summary.
 9. Use the Implementation Template dialog box tabs (Optimize & Map, Place & Route, or Interface) to select any other options applicable to your design. Select **OK** to exit the Implementation Template dialog box.
- Note:** Xilinx recommends using the default Map options for your designs. Also, do not use the guided map option with your synthesized designs.
10. Select Run in the Implement dialog box to begin implementing your design.
 11. When the Flow Engine is displayed, stop the processing of your design after mapping by selecting **Setup** → **Stop After** or by selecting the Set Target toolbar button.
The Stop After dialog box appears.
 12. Select Map and select **OK**.
 13. After the Flow Engine is finished mapping your design, select **Utilities** → **Report Browser** to view the map report.

Double-click the report icon that you want to view. The map report includes a Design Summary section that contains the device utilization information.

14. View the Logic Level Timing Report with the Report Browser. This report shows the performance of your design based on logic levels and best-case routing delays.
15. At this point, you may want to start the Timing Analyzer from the Design Manager to create a more specific report of design paths.
16. Use the Logic Level Timing Report and any reports generated with the Timing Analyzer or the Map program to evaluate how close you are to your performance and utilization goals. Use these reports to decide whether to proceed to the place and route phase of implementation, or to go back and modify your design or implementation options to attain your performance goals. You should have some slack in routing delays to allow the place and route tools to successfully complete your design. Use the verbose option in the Timing Analyzer to see block-by-block delay. The timing report of a mapped design (before place and route) shows block delays, as well as estimated routing delays.

Using the Command Line to Map your Design

1. Translate your design as follows.

```
ngdbuild design_name.xnf
```
2. Map your design as follows.

```
map design_name.ngd
```
3. Use a text editor to view the Device Summary section of the *design_name.mrp* map report. This section contains the device utilization information.
4. Run a timing analysis of the logic level delays from your mapped design as follows.

```
trce [options] design_name.ngd
```

Note: For available options, enter only the trce command at the command line without any arguments.

Use the Trace reports to evaluate how close you are to your performance goals. Use the report to decide whether to proceed

to the place and route phase of implementation, or to go back and modify your design or implementation options to attain your performance goals. You should have some slack in routing delays to allow the place and route tools to successfully complete your design.

The following is the Device Summary section of a Map report.

```
Design Summary
-----
Number of errors:          0
Number of warnings:       3
Number of CLBs:           39 out of 100 39%
  CLB Flip Flops:         32
  4 input LUTs:           66
  3 input LUTs:           5
Number of bonded IOBs:    30 out of 61 49%
  IOB Flops:              0
  IOB Latches:            0
Number of secondary CLKS: 1 out of 4 25%
Number of oscillators:    1
Number of STARTUPs:       1
Number of READCLKs:       1
Number of READBACKs:      1
Number of MD0 pads:       1
Number of MD1 pads:       1
Total equivalent gate count for design: 1538
Additional JTAG gate count for IOBs: 1536
```

The following is a sample Logic Level Timing Report.

```
-----
Xilinx TRACE, Version M1.4.12
Copyright (c) 1995-1997 Xilinx, Inc. All rights reserved.

Design file:                map.ncd
Physical constraint file:   demo_board.pcf
Device,speed:               xc4003e,-2 (x1_0.86 PRELIMINARY)
Report level:               summary report
-----

=====
Timing constraint: NET "FAST_CLOCK" PERIOD = 15.200 nS HIGH 50.000 % ;
  1 item analyzed, 0 timing errors detected.
  Minimum period is 5.585ns.
-----
```

```
=====  
Timing constraint: NET "control_logic/SLOW_CLOCK" PERIOD = 121.600 nS  
HIGH 50.000 % ;  
677 items analyzed, 0 timing errors detected.  
Minimum period is 17.295ns.  
-----  
  
All constraints were met.  
  
Timing summary:  
-----  
  
Timing errors: 0 Score: 0  
  
Constraints cover 811 paths, 0 nets, and 232 connections (73.2% coverage)  
  
Design statistics:  
Minimum period: 17.295ns (Maximum frequency: 57.820MHz)  
  
Analysis completed Tue Jan 27 12:07:59 1998  
-----
```

Evaluating your Design for Coding Style and System Features

At this point, if you are not satisfied with your design performance, you can re-evaluate your code and make any necessary improvements. Modifying your code and selecting different compiler options can dramatically improve device utilization and speed.

Tips for Improving Design Performance

This section includes ways of improving design performance by modifying your code and by incorporating FPGA system features. Most of these techniques are described in more detail in this manual.

Modifying your Code

You can improve design performance with the following design modifications.

- Reduce levels of logic to improve timing

- Redefine hierarchical boundaries to help the compiler optimize design logic
- Pipeline
- Logic replication
- Use of LogiBLOX modules
- Resource sharing

Using FPGA System Features

After correcting any coding style problems, use any of the following FPGA system features in your design to improve resource utilization and to enhance the speed of critical paths.

Note: Each device family has a unique set of system features. Review the current version of the *The Programmable Logic Data Book* for the system features available for the device you are targeting.

- Use global set/reset and global tri-state nets to reduce routing congestion and improve design performance
- Use clock enables
- Place the highest fanout signals on the global buffers
- Modify large multiplexers to use tri-state buffers
- Use one-hot encoding for large or complex state machines
- Use I/O registers when applicable
- Use I/O decoders when applicable
- Use I/O multiplexers when applicable

Using Xilinx-specific Features of your Synthesis Tool

Most synthesis tools have special options for the Xilinx-specific features listed in the previous section. Refer to your synthesis tool documentation for help on using Xilinx-specific features.

Placing and Routing your Design

Note: For more information on placing and routing your design, refer to the *Development System User Guide* and the *Development System Reference Guide* online manuals.

The overall goal when placing and routing your design is fast implementation and high-quality results. However, depending on the situation and your design, you may not always accomplish this goal, as described in the following examples.

- Earlier in the design cycle, run time is more important than the quality of results, and later in the design cycle, the converse is usually true.
- During the day, you may want the tools to quickly process your design while you are waiting for the results. However, you may be less concerned with a quick run time, and more concerned about the quality of results when you run your designs for an extended period of time (during the night or weekend).
- If the targeted device is highly utilized, the routing may become congested, and your design may be difficult to route. In this case, the placer and router may take longer to meet your timing requirements.
- If design constraints are rigorous, it may take longer to correctly place and route your design, and meet the specified timing.

Decreasing Implementation Time

The options you select for the placement and routing of your design directly influence the run time. Generally, these options decrease the run time at the expense of the best placement and routing for a given device. Select your options based on your required design performance.

Note: If you are using the command line, the appropriate command line option is provided in the following procedure.

Use the following steps to decrease implementation time in the Design Manager.

1. Select **Design** → **Implement**

The Implement dialog box appears.

2. Select the Options button in the Implement dialog box.

The Options dialog box appears.

3. Select the Edit Template button next to the Implementation drop-down list in the Program Options Templates field. The Implementation Template dialog box appears.
4. Select the Place & Route tab.
5. Set options in this dialog box as follows.

- Place & Route Effort Level

Generally, you can reduce placement times by selecting a less CPU-intensive algorithm for placement. You can set the placement level from 1 (fastest run time) to 5 (best results) with the default equal to 2. Use the `-l` switch at the command line to perform the same function.

Note: In some cases, poor placement with a lower placement level setting can result in longer route times.

- Router Options

You can limit router iterations to reduce routing times. However, this may prevent your design from meeting timing requirements, or your design may not completely route. From the command line, you can control router passes with the `-i` switch.

- Use Timing Constraints During Place and Route

You can improve run times by not specifying some or all timing constraints. This is useful at the beginning of the design cycle during the initial evaluation of the placed and routed circuit. To disable timing constraints in the Design Manager, deselect the Use Timing Constraints During Place and Route button. To disable timing constraints at the command line, use the `-x` switch with PAR.

6. Select **OK** to exit the Implementation Template dialog box.
7. Select any applicable options in the Options dialog box.
8. Select **OK**.
9. Select Run in the Implement dialog box to begin implementing your design.

Improving Implementation Results

Conversely, you can select options that increase the run time, but produce a better design. These options generally produce a faster design at the cost of a longer run time. These options are useful when you run your designs for an extended period of time (overnight or over the weekend).

Multi-Pass Place and Route Option

Use this option to place and route your design with several different cost tables (seeds) to find the best possible placement for your design. This optimal placement results in shorter routing delays and faster designs. This option works well when the router passes are limited (with the `-i` option). After an optimal cost table is selected, use the re-entrant routing feature to finish the routing of your design. Select this option from the Design menu in the Design Manager, and specify this option at the command line with the `-n` switch.

Turns Engine Option

This option is a Unix-only feature that works with the Multi-Pass Place and Route option to allow parallel processing of placement and routing on several Unix machines. The only limitation to how many cost tables are concurrently tested is the number of workstations you have available. To use this option in the Design Manger, specify a node list when selecting the Multi-Pass Place and Route option. To use this feature at the command line, use the `-m` switch to specify a node list, and the `-n` switch to specify the number of place and route iterations.

Note: For more information on the turns engine option, refer to the *Xilinx Development System Reference Guide*.

Re-entrant Routing Option

Use the re-entrant routing option to further route an already routed design. The router reroutes some connections to improve the timing or to finish routing unrouted nets. You must specify a placed and routed design (.ncd) file for the implementation tools. This option is best used when router iterations are initially limited, or when your design timing goals are close to being achieved.

- From the Design Manager

To initiate a re-entrant route from the Design Manager interface, follow these steps.

1. From the Design Manager, select the placed and routed design revision for the re-entrant option.
2. Select **Tools** → **Flow Engine** to start the Flow Engine from the Design Manager.
3. From the Flow Engine menu, select **Setup** → **Advanced**.
4. In the Advanced dialog box that is displayed, select the Allow Re-entrant Routing option.
5. Select the appropriate options in the Re-entrant Route Options field.
6. Select **OK**.
7. The Place and Route icon in the Flow Engine is replaced with the Re-entrant Route icon. If this step is completed, use the Step Back button until the Re-entrant Route icon no longer indicates completed.
8. Select Run to complete the re-entrant routing.

- From the Command Line

To initiate a re-entrant route from the command line, you can run PAR with the `-k` and `-p` options, as well as any other options you want to use for the routing process. You must either specify a unique name for the post re-entrant routed design (.ncd) file or use the `-w` switch to overwrite the previous design file, as shown in the following examples.

```
par -k -p other_options design_name.ncd new_name.ncd
par -k -p -w other_options design_name.ncd design.ncd
```

Cost-Based Clean-up Option

This option specifies clean-up passes after routing is completed to substitute more appropriate routing options available from the initial routing process. For example, if several local routing resources are used to transverse the chip and a longline is available, the longline is substituted in the clean-up pass. The default value of cost-based clean-up passes is 1. To change the default value, use the Template

Manager in the Design Manager, or the `-c` switch at the command line.

Delay-Based Clean-up Option

This option specifies clean-up passes after routing is completed to substitute more appropriate routing options to reduce delays. The default number of passes for delay-based clean-up is 0. You can change the default in the Design Manager in the Implementation Options window, or at the command line with the `-d` switch.

Guide Option (not recommended)

This option is generally not recommended for synthesis-based designs. Re-synthesizing modules can cause the signal and instance names in the resulting netlist to be significantly different from those in earlier synthesis runs. This can occur even if the source level code (Verilog or VHDL) contains only a small change. Because the guide process is dependent on the names of signals and comps, synthesis designs often result in a low match rate during the guiding process. Generally, this option does not improve implementation results.

Timing Simulation of your Design

Note: Refer to the “Simulating Your Design” chapter for more information on design simulation.

Timing simulation is important in verifying the operation of your circuit after the worst-case placed and routed delays are calculated for your design. In many cases, you can use the same test bench that you used for functional simulation to perform a more accurate simulation with less effort. You can compare the results from the two simulations to verify that your design is performing as initially specified. The Xilinx tools create a VHDL or Verilog simulation netlist of your placed and routed design, and provide libraries that work with many common HDL simulators.

Downloading to the Device and In-system Debugging

After you have verified the functionality and timing of your placed and routed design, you can create a design data file to download for

in-system verification. The design data or bitstream (.bit) file is created from the placed and routed .ncd file. In the Design Manager, use the Configuration step in the Flow Engine to create this file. From the command line, run BitGen on your placed and routed .ncd file to create the .bit file as follows.

```
bitgen [options] design.ncd
```

Use the .bit file with the XChecker cable and the Hardware Debugger to download the data to your device. You can run the Hardware Debugger from the Design Manager, or from the command line as follows.

```
hwdebugr design.bit
```

The Hardware Debugger allows you to download the data to the FPGA using your computer's serial port. The Hardware Debugger can also synchronously or asynchronously probe external or internal nodes in the FPGA. Waveforms can be created from this data and correlated to the simulation data for true in-system verification of your design.

Creating a PROM File for Stand-Alone Operation

After verifying that the FPGA works in the circuit, you can create a PROM file from the .bit file to program a PROM or other data storage device. You can then use this file to program the FPGA in-circuit during normal operation.

Use the Prom File Formatter to create the PROM file, or from the command line use PROMGen. You can run the Prom File Formatter from the Design Manager, or from the command line as follows.

```
promfmtr design.bit
```

Run PROMGen from the command line by typing the following.

```
promgen [options] design.bit
```

Note: For more information on using these programs, refer to the *Xilinx Development System Reference Guide*.

Designing FPGAs with HDL

Xilinx FPGAs provide the benefits of custom CMOS VLSI and allow you to avoid the initial cost, time delay, and risk of conventional masked gate array devices. In addition to the logic in the CLBs and IOBs, the XC4000 family and XC5200 family FPGAs contain system-oriented features such as the following.

- Global low-skew clock or signal distribution network
- Wide edge decoders (XC4000 family only)
- On-chip RAM and ROM (XC4000 family and Spartan)
- IEEE 1149.1 — compatible boundary scan logic support
- Flexible I/O with Adjustable Slew-rate Control and Pull-up/Pull-down Resistors
- 12-mA sink current per output and 24-mA sink per output pair
- Dedicated high-speed carry-propagation circuit

You can use these device characteristics to improve resource utilization and enhance the speed of critical paths in your HDL designs. The examples in this chapter are provided to help you incorporate these system features into your HDL designs. This chapter includes the following sections.

- “Using Global Low-skew Clock Buffers”
- “Using Dedicated Global Set/Reset Resource”
- “Encoding State Machines”
- “Using Dedicated I/O Decoders”
- “Instantiating LogiBLOX Modules”
- “Implementing Memory”

- “Implementing Boundary Scan (JTAG 1149.1)”
- “Implementing Logic with IOBs”
- “Implementing Multiplexers with Tristate Buffers”
- “Using Pipelining”
- “Design Hierarchy”

Using Global Low-skew Clock Buffers

For designs with global signals, use global clock buffers to take advantage of the low-skew, high-drive capabilities of the dedicated global buffer tree of the target device. When you use the Insert Pads or equivalent command, your synthesis tool automatically inserts a BUFG generic clock buffer whenever an input signal drives a clock signal. The Xilinx implementation software automatically selects the clock buffer that is appropriate for your specified design architecture. If you want to use a specific global buffer, you must instantiate it. Many synthesis tools automatically insert I/O pins and clock buffers. Also, some synthesis tools limit I/O and global buffers. Refer to your synthesis tool documentation for detailed information.

You can instantiate an architecture-specific buffer if you understand the architecture and want to specify how the resources should be used. Each XC4000E/L and Spartan device contains four primary and four secondary global buffers that share the same routing resources. XC4000EX/XL/XV devices have sixteen global buffers; each buffer has its own routing resources. XC5200 devices have four dedicated global buffers in each corner of the device.

XC4000 EX/XL/XV devices have two different types of global buffer, Global Low-Skew Buffers (BUFGLS) and Global Early Buffers (BUFGE). Global Low-Skew Buffers are standard global buffers that should be used for most internal clocking or high fanout signals that must drive a large portion of the device. There are eight BUFGLS buffers available, two in each corner of the device. The Global Early Buffers are designed to provide faster clock access, but CLB access is limited to one quadrant of the device. I/O access is also limited. Similarly, there are eight BUFGEs, two in each corner of the device.

Because Global Early and Global Low-Skew Buffers share a single pad, a single IPAD can drive a BUFGE, BUFGLS, or both in parallel. The parallel configuration is especially useful for clocking the fast

capture latches of the device. Since the Global Early and Global Low-Skew Buffers share a common input, they cannot be driven by two different signals.

You can use the following criteria to help select the appropriate global buffer for a given design path.

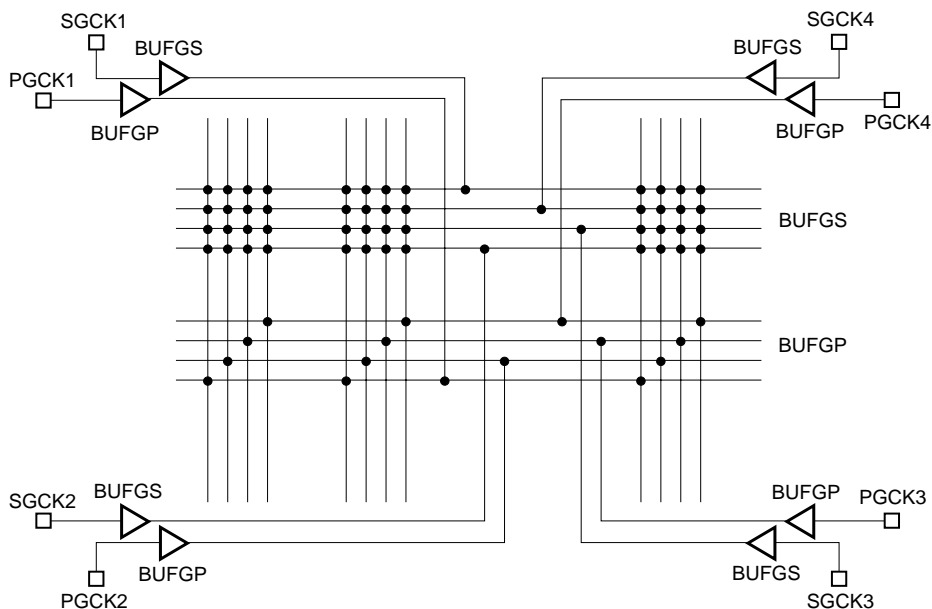
- The simplest option is to use a Global Low-Skew Buffer.
- If you want a faster clock path, use a BUFG. Initially, the software will try to use a Global Low-Skew Buffer. If timing requirements are not met, a BUFGE is automatically used if possible.
- If a single quadrant of the chip is sufficient for the clocked logic, and timing requires a faster clock than the Global Low-Skew Buffer, use a Global Early Buffer.

Note: For more information on using the XC4000 EX/XL/XV device family global buffers, refer to the online version of *The Programmable Logic Data Book* or the Xilinx web site at <http://www.xilinx.com>.

For XC4000E/L and Spartan devices, you can use secondary global buffers (BUFGS) to buffer high-fanout, low-skew signals that are sourced from inside the FPGA. To access the secondary global clock buffer for an internal signal, instantiate the BUFGS cell. You can use primary global buffers (BUFGP) to distribute signals applied to the FPGA from an external source. Internal signals can be globally distributed with a primary global buffer, however, the signals must be driven by an external pin.

Some synthesis tools limit I/O or BUFG resources. For example, BUFG does not synthesize to more than eight instances depending on the selected device architecture. However, some tools do not use all your available resources. Compiling modules separately may also result in resource over-utilization. Check with your synthesis vendor.

XC4000E devices have four primary (BUFGP) and four secondary (BUFGS) global clock buffers that share four global routing lines, as shown in the following figure.



X4987

Figure 4-1 Global Buffer Routing Resources (XC4000E, Spartan)

These global routing resources are only available for the eight global buffers. The eight global nets run horizontally across the middle of the device and can be connected to one of the four vertical longlines that distribute signals to the CLBs in a column. Because of this arrangement only four of the eight global signals are available to the CLBs in a column. These routing resources are “free” resources because they are outside of the normal routing channels. Use these resources whenever possible. You may want to use the secondary buffers first because they have more flexible routing capabilities.

You should use the global buffer routing resources primarily for high-fanout clocks that require low skew, however, you can use them to drive certain CLB pins, as shown in the following figure. In addition, you can use these routing resources to drive high-fanout clock enables, clear lines, and the clock pins (K) of CLBs and IOBs.

In the figure shown, the C pins drive the input to the H function generator, Direct Data-in, Preset, Clear, or Clock Enable pins. The F and G pins are the inputs to the F and G function generators, respectively.

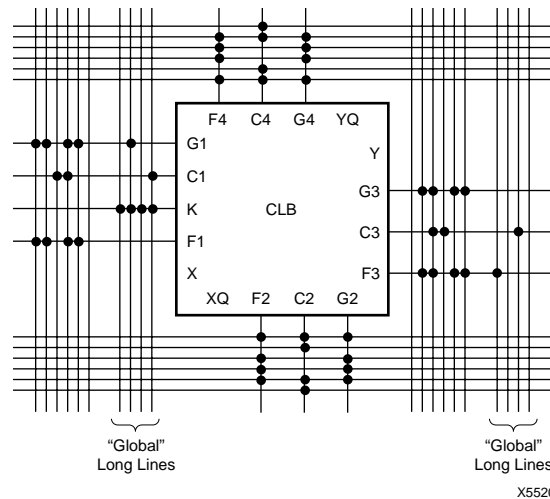


Figure 4-2 Global Longlines Resource CLB Connections

If your design does not contain four high-fanout clocks, use these routing resources for signals with the next highest fanout. To reduce routing congestion, use the global buffers to route high-fanout signals. These high-fanout signals include clock enables and reset signals (*not* global reset signals). Use global buffer routing resources to reduce routing congestion; enable routing of an otherwise unroutable design; and ensure that routing resources are available for critical nets.

Xilinx recommends that you assign up to four secondary global clock buffers to the four signals in your design with the highest fanout (such as clock nets, clock enables, and reset signals). Clock signals that require low skew have priority over low-fanout non-clock signals. You can source the signals with an input buffer or a gate internal to the design. Generate internally sourced clock signals with a register to avoid unwanted glitches. The synthesis tool can insert global clock buffers or you can instantiate them in your HDL code.

Note: Use Global Set/Reset resources when applicable. Refer to the “Using Dedicated Global Set/Reset Resource” section in this chapter for more information.

Inserting Clock Buffers

Many synthesis tools automatically insert a secondary global clock buffer on all input ports that drive a register’s clock pin or a gated clock signal. Refer to your synthesis tool documentation for information on disabling the automatic insertion of clock buffers, and how to specify which ports have clock buffers.

Instantiating Global Clock Buffers

You can instantiate global buffers in your code as described in this section.

Instantiating Buffers Driven from a Port

You can instantiate global buffers and connect them to high-fanout ports in your code rather than inferring them from a synthesis tool script. If you do instantiate global buffers, verify that the Pad parameter is not specified for the buffer.

Instantiating Buffers Driven from Internal Logic

Some synthesis tools require you to instantiate a global buffer in your code to use the dedicated routing resource if a high-fanout signal is sourced from internal flip-flops or logic (such as a clock divider or multiplexed clock), or if a clock is driven from the internal oscillator or non-dedicated I/O pin. The following VHDL and Verilog examples instantiate a BUFGS for an internal multiplexed clock circuit. A Set Dont Touch or equivalent attribute is added to the instantiated component to prevent further optimization by the synthesizer.

- VHDL Example

```

-----
-- CLOCK_MUX.VHD Version 1.1                --
-- This is an example of an instantiation of --
-- global buffer (BUFGS) from an internally  --
-- driven signal, a multiplexed clock.      --
-- March 1998                               --
-----

library IEEE;
use IEEE.std_logic_1164.all;

entity clock_mux is
    port (DATA, SEL:                in  STD_LOGIC;
          SLOW_CLOCK, FAST_CLOCK:   in  STD_LOGIC;
          DOUT:                      out STD_LOGIC);
end clock_mux;

architecture XILINX of clock_mux is

    signal CLOCK:          STD_LOGIC;
    signal CLOCK_GBUF:    STD_LOGIC;

    component BUFGS
        port (I: in  STD_LOGIC;
              O: out STD_LOGIC);
    end component;

begin

    Clock_MUX: process (SEL, FAST_CLOCK, SLOW_CLOCK)
    begin
        if (SEL = '1') then
            CLOCK <= FAST_CLOCK;
        else
            CLOCK <= SLOW_CLOCK;
        end if;
    end process;

    GBUF_FOR_MUX_CLOCK: BUFGS
        port map (I => CLOCK,
                 O => CLOCK_GBUF);

    Data_Path: process (CLOCK_GBUF)
    begin

```

```
        if (CLOCK_GBUF'event and CLOCK_GBUF='1') then
            DOUT <= DATA;
        end if;
    end process;
```

```
end XILINX;
```

- **Verilog Example**

```
////////////////////////////////////
// CLOCK_MUX.V Version 1.1 //
// This is an example of an instantiation of //
// global buffer (BUFGS) from an internally //
// driven signal, a multiplied clock. //
// March 1998 //
////////////////////////////////////

module clock_mux (DATA, SEL, SLOW_CLOCK, FAST_CLOCK,
                 DOUT);

    input DATA, SEL;
    input SLOW_CLOCK, FAST_CLOCK;
    output DOUT;

    reg CLOCK;
    wire CLOCK_GBUF;
    reg DOUT;

    always @ (SEL or FAST_CLOCK or SLOW_CLOCK)
    begin
        if (SEL == 1'b1)
            CLOCK <= FAST_CLOCK;
        else
            CLOCK <= SLOW_CLOCK;
        end

    BUFGS GBUF_FOR_MUX_CLOCK (.O(CLOCK_GBUF),
                              .I(CLOCK));

    always @ (posedge CLOCK_GBUF)
        DOUT = DATA;

endmodule
```

Using Dedicated Global Set/Reset Resource

XC4000 and Spartan devices have a dedicated Global Set/Reset (GSR) net that you can use to initialize all CLBs and IOBs. When the GSR is asserted, every flip-flop in the FPGA is simultaneously preset or cleared. You can access the GSR net from the GSR pin on the STARTUP block or the GSRIN pin of the STARTBUF (VHDL).

Since the GSR net has dedicated routing resources that connect to the Preset or Clear pin of the flip-flops, you do not need to use general purpose routing or global buffer resources to connect to these pins. If your design has a Preset or Clear signal that affects every flip-flop in your design, use the GSR net to increase design performance and reduce routing congestion.

The XC5200 family has a dedicated Global Reset (GR) net that resets all device registers. As in the XC4000 and Spartan devices, the STARTUP or STARTBUF (VHDL) block must be instantiated in your code in order to access this resource. The XC3000A devices also have dedicated Global Reset (GR) that is connected to a dedicated device pin (see device pinout). Since this resource is always active, you do not need to do anything to activate this feature.

For XC4000, Spartan, and XC5200 devices, the Global Set/Reset (GSR or GR) signal is, by default, set to active high (globally resets device when logic equals 1). If you are using an older version of a synthesis tool, for an active low reset, instantiate an inverter in your code to invert the global reset signal. The inverter is absorbed by the STARTUP block and does not use any device resources (function generators). For older versions of synthesis tools, although the inverted signal may be behaviorally described in your code, Xilinx recommends instantiating the inverter to prevent the mapping of the inverter into a CLB function generator, and subsequent delays to the reset signal and unnecessary use of device resources. Also make sure you put a Don't Touch attribute on the instantiated inverter before compiling your design. If you do not add this attribute, the inverter may get mapped into a CLB function generator. Most new synthesis tools automatically insert the STARTUP block and the previous steps are not required.

Note: For more information on simulating the Global Set/Reset, see the “Simulating Your Design” chapter.

Startup State

Note: See the “Simulating Your Design” chapter for more information on STARTUP and STARTBUF.

The GSR pin on the STARTUP block or the GSRIN pin on the STARTBUF block drives the GSR net and connects to each flip-flop’s Preset and Clear pin. When you connect a signal from a pad to the STARTUP block’s GSR pin, the GSR net is activated. Because the GSR net is built into the silicon it does not appear in the pre-routed netlist file. When the GSR signal is asserted High (the default), all flip-flops and latches are set to the state they were in at the end of configuration. When you simulate the routed design, the gate simulator translation program correctly models the GSR function.

Note: For the XC3000 family and the XC5200 family, all flip-flops and latches are reset to zero after configuration.

Preset vs. Clear (XC4000, Spartan)

The XC4000 family flip-flops are configured as either preset (asynchronous set) or clear (asynchronous reset). Automatic assertion of the GSR net presets or clears each flip-flop. You can assert the GSR pin at any time to produce this global effect. You can also preset or clear individual flip-flops with the flip-flop’s dedicated Preset or Clear pin. When a Preset or Clear pin on a flip-flop is connected to an active signal, the state of that signal controls the startup state of the flip-flop. For example, if you connect an active signal to the Preset pin, the flip-flop starts up in the preset state. If you do not connect the Clear or Preset pin, the default startup state is a clear state. To change the default to preset, assign an INIT=S attribute to the flip-flop.

I/O flip-flops and latches do not have individual Preset or Clear pins. The default value of these flip-flops and latches is clear. To change the default value to preset, assign an INIT=S attribute.

Refer to your synthesis tool documentation for information on changing the initial state of registers that do not use the Preset or Clear pins.

Increasing Performance with the GSR/GR Net

Many designs contain a net that initializes most of the flip-flops in the design. If this signal can initialize *all* the flip-flops, you can use the

GSR/GR net. You should always include a net that initializes your design to a known state.

To ensure that your HDL simulation results at the RTL level match the synthesis results, write your code so that every flip-flop and latch is preset or cleared when the GSR signal is asserted. The synthesis tool cannot infer the GSR/GR net from HDL code. To utilize the GSR net, you must instantiate the STARTUP or STARTBUF block (VHDL), as shown in the “No_GSR Implemented with Gates” figure.

Design Example without Dedicated GSR/GR Resource

In the following VHDL and Verilog designs, the RESET signal initializes all the registers in the design; however, it does not use the dedicated global resources. The RESET signal is routed using regular routing resources. These designs include two 4-bit counters. One counter counts up and is reset to all zeros on assertion of RESET and the other counter counts down and is reset to all ones on assertion of RESET. The “No_GSR Implemented with Gates” figure shows the No_GSR design implemented with gates.

- No GSR VHDL Example

```
-- NO_GSR Example
-- The signal RESET initializes all registers
-- May 1997

library IEEE;
use IEEE.std_logic_1164.all;
use IEEE.std_logic_unsigned.all

entity no_gsr is
port (CLOCK: in STD_LOGIC;
      RESET: in STD_LOGIC;
      UPCNT: out STD_LOGIC_VECTOR (3 downto 0);
      DNCNT: out STD_LOGIC_VECTOR (3 downto 0));
end no_gsr;

architecture SIMPLE of no_gsr is

signal UP_CNT: STD_LOGIC_VECTOR (3 downto 0);
signal DN_CNT: STD_LOGIC_VECTOR (3 downto 0);

begin
```

```
UP_COUNTER: process (CLOCK, RESET)
begin
    if (RESET = '1') then
        UP_CNT <= "0000";
    elsif (CLOCK'event and CLOCK = '1') then
        UP_CNT <= UP_CNT + 1;
    end if;
end process;

DN_COUNTER: process (CLOCK, RESET)
begin
    if (RESET = '1') then
        DN_CNT <= "1111";
    elsif (CLOCK'event and CLOCK = '1') then
        DN_CNT <= DN_CNT - 1;
    end if;
end process;

UPCNT <= UP_CNT;
DNCNT <= DN_CNT;

end SIMPLE;
```

- **No GR VHDL Example**

```
-- NO_GR.VHD Example
-- The signal RESET initializes all registers
-- Without the use of the dedicated Global Reset
-- routing
-- December 1997

library IEEE;
use IEEE.std_logic_1164.all;
use IEEE.std_logic_unsigned.all;

entity no_gr is
port (CLOCK: in STD_LOGIC;
      RESET: in STD_LOGIC;
      UPCNT: out STD_LOGIC_VECTOR (3 downto 0);
      DNCNT: out STD_LOGIC_VECTOR (3 downto 0));
end no_gr;

architecture XILINX of no_gr is

signal UP_CNT: STD_LOGIC_VECTOR (3 downto 0);
```

```

signal DN_CNT: STD_LOGIC_VECTOR (3 downto 0);

begin
  UP_COUNTER: process (CLOCK, RESET)
  begin
    if (RESET = '1') then
      UP_CNT <= "0000";
    elsif (CLOCK'event and CLOCK = '1') then
      UP_CNT <= UP_CNT + 1;
    end if;
  end process;

  DN_COUNTER: process (CLOCK, RESET)
  begin
    if (RESET = '1') then
      DN_CNT <= "1111";
    elsif (CLOCK'event and CLOCK = '1') then
      DN_CNT <= DN_CNT - 1;
    end if;
  end process;

  UPCNT <= UP_CNT;
  DNCNT <= DN_CNT;

end XILINX;

```

- **No GSR Verilog Example**

```

/* NO_GSR Example
 * The signal RESET initializes all registers
 * December 1997 */

module no_gsr ( CLOCK, RESET, UPCNT, DNCNT);

input CLOCK, RESET;
output [3:0] UPCNT;
output [3:0] DNCNT;

reg [3:0] UPCNT;
reg [3:0] DNCNT;

always @ (posedge CLOCK or posedge RESET) begin
  if (RESET) begin
    UPCNT = 4'b0000;
    DNCNT = 4'b1111;
  end
end

```

```
        end else begin
            UPCNT = UPCNT + 1'b1;
            DNCNT = DNCNT - 1'b1;
        end
    end
endmodule
```

- **No GR Verilog Example**

```
/* NO_GR.V Example
 * The signal RESET initializes all registers
 * Aug 1997 */

module no_gr ( CLOCK, RESET, UPCNT, DNCNT);

input  CLOCK, RESET;
output [3:0] UPCNT;
output [3:0] DNCNT;

reg [3:0] UPCNT;
reg [3:0] DNCNT;

always @ (posedge CLOCK or posedge RESET) begin
    if (RESET) begin
        UPCNT = 4'b0000;
        DNCNT = 4'b1111;
    end else begin
        UPCNT = UPCNT + 1'b1;
        DNCNT = DNCNT - 1'b1;
    end
end

endmodule
```

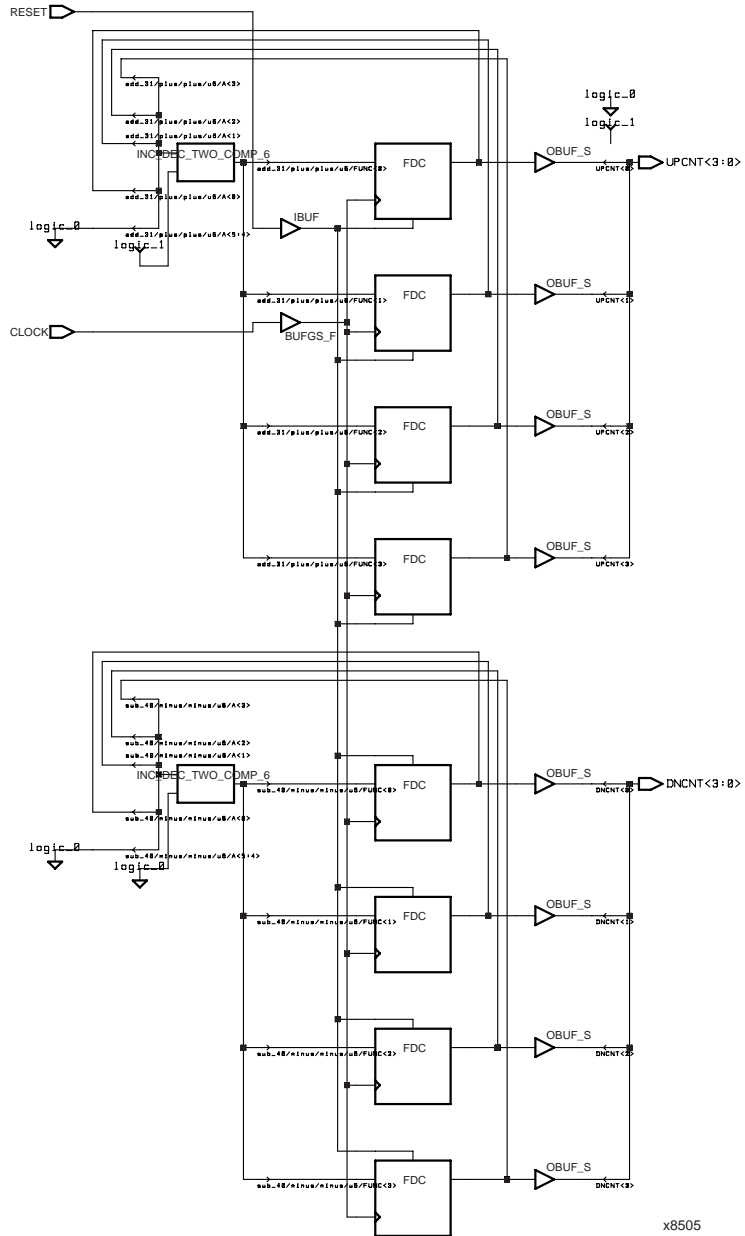


Figure 4-3 No_GSR Implemented with Gates

Design Example with Dedicated GSR/GR Resource

To reduce routing congestion and improve the overall performance of the reset net in the No_GSR and No_GR designs, use the dedicated GSR or GR net instead of the general purpose routing. Instantiate the STARTUP, STARTBUF, ROC, or TOC block in your design and use the GSR or GR pin on the STARTUP block (or the GSRIN pin on the STARTBUF block) to access the global reset net. This is not necessary with many synthesis tools. If you fully define the behavior of the GSR or GR net, the tool infers a STARTUP block. The modified designs (Use_GSR and Use_GR) are included at the end of this section. The Use_GSR design implemented with gates is shown in the “Use_GSR Implemented with Gates” figure.

In XC4000 and Spartan designs, on assertion of the GSR net, flip-flops return to a clear (or Low) state by default. You can override this default by describing an asynchronous preset in your code, or by adding the INIT=“1” or equivalent attribute to the flip-flop (described later in this section).

In XC5200 family designs, the GR resets all flip-flops in the device to a logic zero. If a flip-flop is described as asynchronous preset to a logic 1, the synthesis tool automatically infers a flip-flop with a synchronous preset, and the Xilinx software puts an inverter on the input and output of the device to simulate a preset.

The Use_GSR and Use_GR designs explicitly state that the down-counter resets to all ones, therefore, asserting the reset net causes this counter to reset to a default of all zeros. You can use one of the following two methods to prevent this reset to zeros.

- Remove the comment characters from the last few lines of code in the Use_GSR or Use_GR design. These lines of code correctly describe the behavior of the design (in response to the assertion of reset). However, when you synthesize the design, the Preset pins on the flip-flops that form the down-counter are used and the Clear pins on the flip-flops that form the up-counter are used. Using these pins defeats the purpose of using the GSR or GR net.
- Attach the INIT = “1” or equivalent attribute to the down-counter flip-flops.

The synthesizer may do this if necessary depending on your code’s initialization state when the reset is applied. Refer to your synthesis tool documentation for more information on assigning

attributes. This command allows you to override the default clear (or Low) state when your code does not specify a preset condition. However, because attributes are assigned outside the HDL code, the code no longer accurately represents the behavior of the design.

Xilinx recommends removing the comment characters from the last few lines of the Use_GSR or Use_GR code when you perform an RTL simulation and attaching the INIT=S attribute to the relevant flip-flops when you synthesize the design.

The STARTUP or STARTBUF block must not be optimized during the synthesis process. Add the appropriate attribute to prevent optimization before compiling your design.

- Use GSR VHDL Example (XC4000 family)

```
-- USE_GSR.VHD Example
-- The signal RESET is connected to the
-- GSRIN pin of the STARTBUF block
-- May 1997

library IEEE;
library UNISIM;
use IEEE.std_logic_1164.all;
use IEEE.std_logic_unsigned.all;
use UNISIM.all;

entity use_gsr is
port ( CLOCK: in STD_LOGIC;
      RESET: in STD_LOGIC;
      UPCNT: out STD_LOGIC_VECTOR (3 downto 0);
      DNCNT: out STD_LOGIC_VECTOR (3 downto 0));
end use_gsr;

architecture XILINX of use_gsr is

component STARTBUF
  port (GSRIN: in STD_LOGIC);
        GSROUT: out STD_LOGIC);
end component;

signal RESET_INT: STD_LOGIC;
signal UP_CNT: STD_LOGIC_VECTOR (3 downto 0);
signal DN_CNT: STD_LOGIC_VECTOR (3 downto 0);
```

```
begin

    U1: STARTBUF port map(GSRIN=>RESET,
                          GSROUT=>RESET_INT);

    UP_COUNTER: process(CLOCK, RESET_INT)
    begin
        if (RESET_INT = '1') then
            UP_CNT <= "0000";
        elsif CLOCK'event and CLOCK = '1') then
            UP_CNT <= UP_CNT - 1;
        end if;
    end process;

    DN_COUNTER: (CLOCK, RESET_INT)
    begin
        if (RESET_INT = '1') then
            DN_CNT <= "1111";
        elsif CLOCK'event and CLOCK = '1') then
            DN_CNT <= DN_CNT - 1;
        end if;
    end process;

    UPCNT <= UP_CNT;
    DNCNT <= DN_CNT;

end XILINX;
```

- Use GR VHDL Example

```
-----
-- USE_GR.VHD Version 1.0                               --
-- Xilinx HDL Synthesis Design Guide                    --
-- The signal RESET initializes all registers          --
-- Using the global reset resources since              --
-- STARTBUF block was added                            --
-- December 1997                                       --
-----

library IEEE;
library UNISIM;
use IEEE.std_logic_1164.all;
use IEEE.std_logic_unsigned.all;
use UNISIM.all;
```



```

entity use_gr is
port ( CLOCK: in STD_LOGIC;
      RESET: in STD_LOGIC;
      UPCNT: out STD_LOGIC_VECTOR (3 downto 0);
      DNCNT: out STD_LOGIC_VECTOR (3 downto 0));
end use_gr;

architecture XILINX of use_gr is

component STARTBUF
  port (GSRIN: in STD_LOGIC;
        GSROUT: out STD_LOGIC);
end component;

signal RESET_INT: STD_LOGIC;
signal UP_CNT: STD_LOGIC_VECTOR (3 downto 0);
signal DN_CNT: STD_LOGIC_VECTOR (3 downto 0);

begin

  U1: STARTBUF port map(GSRIN=>RESET,
                       GSROUT=>RESET_INT);

  UP_COUNTER: process(CLOCK, RESET_INT)
  begin
    if (RESET_INT = '1') then
      UP_CNT <= "0000";
    elsif (CLOCK'event and CLOCK = '1') then
      UP_CNT <= UP_CNT + 1;
    end if;
  end process;

  DN_COUNTER: process(CLOCK, RESET_INT)
  begin
    if (RESET_INT = '1') then
      DN_CNT <= "1111";
    elsif (CLOCK'event and CLOCK = '1') then
      DN_CNT <= DN_CNT - 1;
    end if;
  end process;

  UPCNT <= UP_CNT;
  DNCNT <= DN_CNT;

```

```
end XILINX;
```

- Use GSR Verilog Example

```
////////////////////////////////////  
// USE_GSR.V Version 1.0 //  
// The signal RESET initializes all registers //  
// Using the global reset resources (STARTUP) //  
// December 1997 //  
////////////////////////////////////  
  
module use_gsr ( CLOCK, RESET, UPCNT, DNCNT);  
  
input CLOCK, RESET;  
output [3:0] UPCNT;  
output [3:0] DNCNT;  
  
reg [3:0] UPCNT;  
reg [3:0] DNCNT;  
  
STARTUP U1 (.GSR(RESET));  
  
always @ (posedge CLOCK or posedge RESET) begin  
    if (RESET) begin  
        UPCNT = 4'b0000;  
        DNCNT = 4'b1111;  
    end else begin  
        UPCNT = UPCNT + 1'b1;  
        DNCNT = DNCNT - 1'b1;  
    end  
end  
  
endmodule
```

- Use GR Verilog Example

```
////////////////////////////////////  
// USE_GR.V Version 1.0 //  
// The signal RESET initializes all registers //  
// Using the global reset resources since //  
// STARTUP block instantiation was added //  
// December 1997 //  
////////////////////////////////////  
  
module use_gr ( CLOCK, RESET, UPCNT, DNCNT);
```

```
input CLOCK, RESET;
output [3:0] UPCNT;
output [3:0] DNCNT;

reg [3:0] UPCNT;
reg [3:0] DNCNT;

STARTUP U1 (.GR(RESET));

always @ (posedge CLOCK or posedge RESET) begin
    if (RESET) begin
        UPCNT = 4'b0000;
        DNCNT = 4'b1111;
    end else begin
        UPCNT = UPCNT + 1'b1;
        DNCNT = DNCNT - 1'b1;
    end
end

endmodule
```

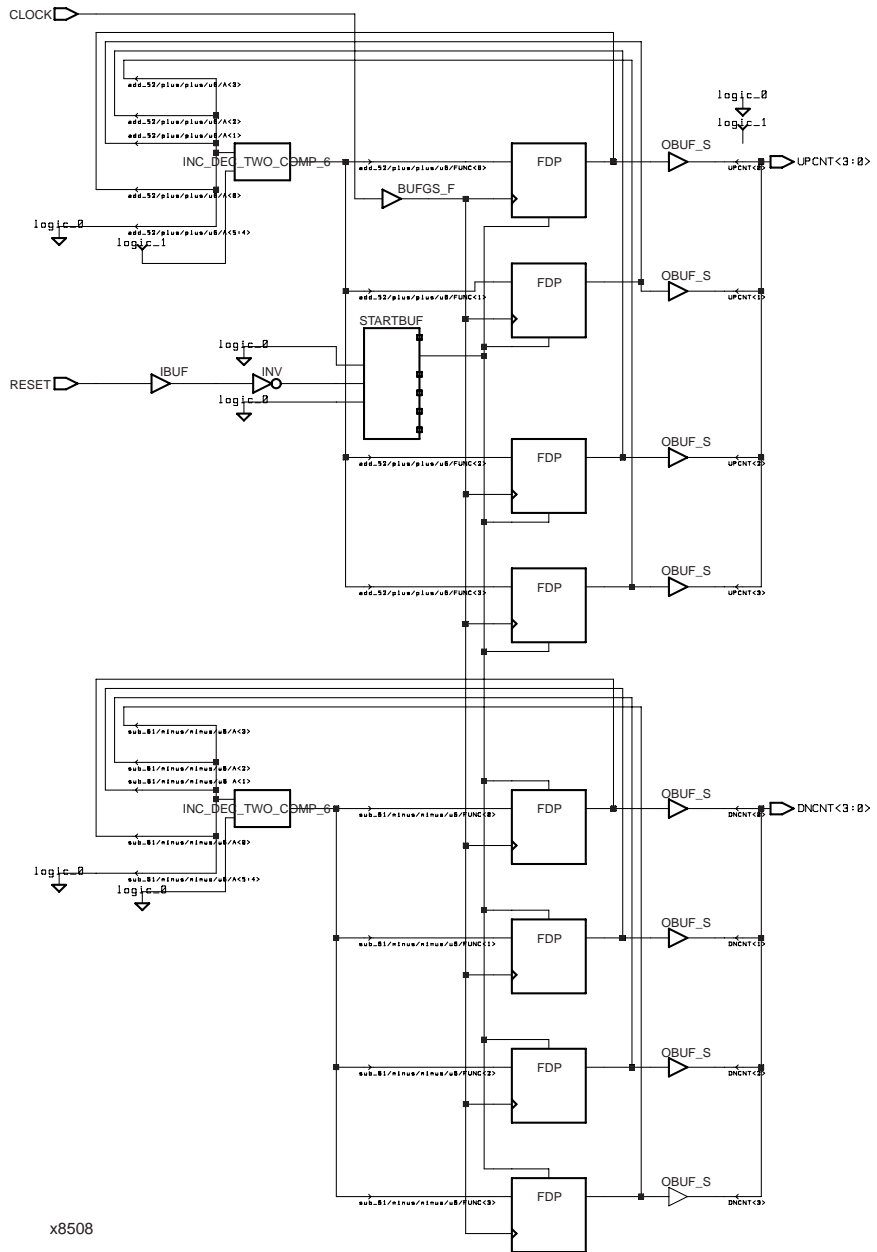


Figure 4-4 Use_GSR Implemented with Gates

Design Example with Active Low GSR/GR Signal

The Active_Low_GSR design is identical to the Use_GSR design except an INV is instantiated and connected between the RESET port and the STARTUP block. Also, a Set Don't Touch (or equivalent) attribute is added to the synthesis tool script for both the INV and STARTUP, or STARTBUF (VHDL) symbols. By instantiating the inverter, the global set/reset signal is now active low (logic level 0 resets all FPGA flip-flops). The inverter is absorbed into the STARTUP block in the device and no CLB resources are used to invert the signal. This is not necessary with many synthesis tools. If all registers and latches are described in the RTL code as reset or set, then a GSR or GR is inferred. Some tools also give you the option to select any signal as the GSR or GR net. This allows you to correct problems if the RTL code does not completely describe the GSR/GR behavior. However, the RTL code will not match the place and route behavior because not all registers are described as set or reset with the GSR/GR signal. Some tools provide a report of the inferred registers that are missing the GSR/GR behavior, and allow you to change the RTL behavior. VHDL and Verilog Active_Low_GSR designs are shown following.

- Active Low GSR VHDL Example

```
-----
-- ACTIVE_LOW_GSR.VHD Version 1.0                --
-- The signal RESET is inverted before being      --
-- connected to the GSRIN pin of the STARTBUF    --
-- The inverter will be absorbed by the STARTBUF  --
-- September 1997                                 --
-----

library IEEE;
library UNISIM;
use IEEE.std_logic_1164.all;
use IEEE.std_logic_unsigned.all;
use UNISIM.all;

entity active_low_gsr is
    port ( CLOCK: in STD_LOGIC;
          RESET: in STD_LOGIC;
          UPCNT: out STD_LOGIC_VECTOR (3 downto 0);
          DNCNT: out STD_LOGIC_VECTOR (3 downto 0));
end active_low_gsr;
```

architecture XILINX of active_low_gsr is

```
component INV
    port (I: in  STD_LOGIC;
          O: out STD_LOGIC);
end component;

component STARTBUF
    port (GSRIN: in  STD_LOGIC;
          GSROUT: out STD_LOGIC);
end component;

signal RESET_NOT:      STD_LOGIC;
signal RESET_NOT_INT: STD_LOGIC;
signal UP_CNT:         STD_LOGIC_VECTOR (3 downto 0);
signal DN_CNT:         STD_LOGIC_VECTOR (3 downto 0);

begin

U1: INV port map(I => RESET, O => RESET_NOT);

U2: STARTBUF port map(GSRIN=>RESET_NOT,
                     GSROUT=>RESET_NOT_INT);

UP_COUNTER: process(CLOCK, RESET_NOT_INT)
begin
    if (RESET_NOT_INT = '1') then
        UP_CNT <= "0000";
    elsif (CLOCK'event and CLOCK = '1') then
        UP_CNT <= UP_CNT + 1;
    end if;
end process;

DN_COUNTER: process(CLOCK, RESET_NOT_INT)
begin
    if (RESET_NOT_INT = '1') then
        DN_CNT <= "1111";
    elsif (CLOCK'event and CLOCK = '1') then
        DN_CNT <= DN_CNT - 1;
    end if;
end process;

UPCNT <= UP_CNT;
DNCNT <= DN_CNT;
```

end XILINX;

- Active Low GSR Verilog Example

```

//////////////////////////////////////////////////////////////////
// ACTIVE_LOW_GSR.V Version 1.0                               //
// The signal RESET is inverted before being                 //
// connected to the GSR pin of the STARTUP bloc             //
// The inverter will be absorbed by STARTUP in M1           //
// September 1997                                           //
//////////////////////////////////////////////////////////////////

module active_low_gsr ( CLOCK, RESET, UPCNT, DNCNT);

    input          CLOCK, RESET;
    output [3:0]  UPCNT;
    output [3:0]  DNCNT;

    wire          RESET_NOT;
    reg [3:0]     UPCNT;
    reg [3:0]     DNCNT;

    INV U1 (.O(RESET_NOT), .I(RESET));

    STARTUP U2 (.GSR(RESET_NOT));

    always @ (posedge CLOCK or posedge RESET_NOT)
    begin
        if (RESET_NOT)
        begin
            UPCNT = 4'b0000;
            DNCNT = 4'b1111;
        end
        else
        begin
            UPCNT = UPCNT + 1'b1;
            DNCNT = DNCNT - 1'b1;
        end
    end
endmodule

```

Encoding State Machines

The traditional methods used to generate state machine logic result in highly-encoded states. State machines with highly-encoded state variables typically have a minimum number of flip-flops and wide combinatorial functions. These characteristics are acceptable for PAL and gate array architectures. However, because FPGAs have many flip-flops and narrow function generators, highly-encoded state variables can result in inefficient implementation in terms of speed and density.

One-hot encoding allows you to create state machine implementations that are more efficient for FPGA architectures. You can create state machines with one flip-flop per state and decreased width of combinatorial logic. One-hot encoding is usually the preferred method for large FPGA-based state machine implementation. For small state machines (fewer than 8 states), binary encoding may be more efficient. To improve design performance, you can divide large (greater than 32 states) state machines into several small state machines and use the appropriate encoding style for each.

Three design examples are provided in this section to illustrate the three coding methods (binary, enumerated type, and one-hot) you can use to create state machines. All three examples contain the same Case statement. To conserve space, the complete Case statement is only included in the binary encoded state machine example; refer to this example when reviewing the enumerated type and one-hot examples.

Some synthesis tools allow you to add an attribute, such as `type_encoding_style`, to your VHDL code to set the encoding style. This is a synthesis vendor attribute (not a Xilinx attribute). Refer to your synthesis tool documentation for information on attribute-driven state machine synthesis.

Note: The bold text in each of the three examples indicates the portion of the code that varies depending on the method used to encode the state machine.

Using Binary Encoding

The state machine bubble diagram in the following figure shows the operation of a seven-state machine that reacts to inputs A through E as well as previous-state conditions. The binary encoded method of

coding this state machine is shown in the VHDL and Verilog examples that follow. These design examples show you how to take a design that has been previously encoded (for example, binary encoded) and synthesize it to the appropriate decoding logic and registers. These designs use three flip-flops to implement seven states.

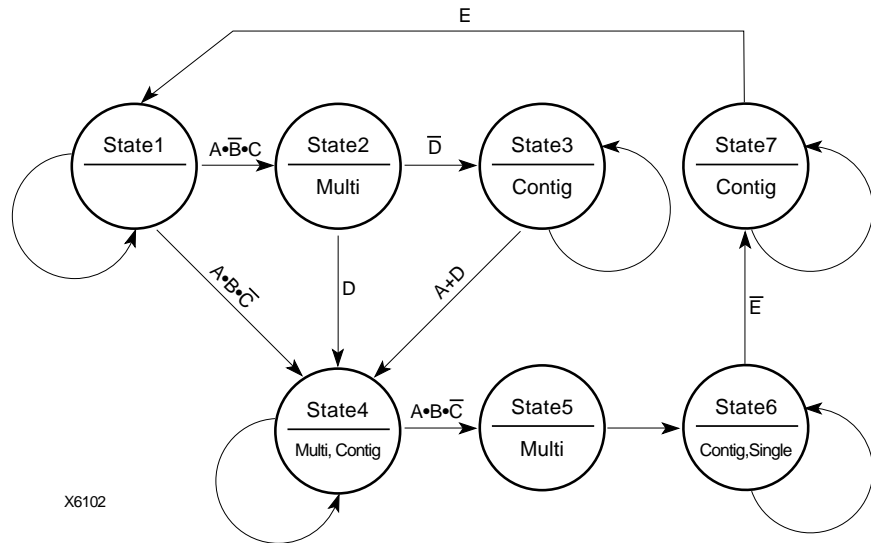


Figure 4-5 State Machine Bubble Diagram

Binary Encoded State Machine VHDL Example

```

-----
-- BINARY.VHD Version 1.0                --
-- Example of a binary encoded state machine --
-- May 1997                               --
-----

Library IEEE;
use IEEE.std_logic_1164.all;
use IEEE.std_logic_unsigned.all;

entity binary is
    port (CLOCK, RESET : in STD_LOGIC;
          A, B, C, D, E: in BOOLEAN;
          SINGLE, MULTI, CONTIG: out STD_LOGIC);
end binary;

```

architecture BEHV of binary is

```
type STATE_TYPE is (S1, S2, S3, S4, S5, S6, S7);
attribute ENUM_ENCODING: STRING;
attribute ENUM_ENCODING of STATE_TYPE:type is "001 010 011 100 101 110
111";
```

```
signal CS, NS: STATE_TYPE;
```

```
begin
```

```
    SYNC_PROC: process (CLOCK, RESET)
    begin
        if (RESET='1') then
            CS <= S1;
        elsif (CLOCK'event and CLOCK = '1') then
            CS <= NS;
        end if;
    end process; --End REG_PROC
```

```
    COMB_PROC: process (CS, A, B, C, D, E)
    begin
        case CS is
            when S1 =>
                MULTI <= '0';
                CONTIG <= '0';
                SINGLE <= '0';
                if (A and not B and C) then
                    NS <= S2;
                elsif (A and B and not C) then
                    NS <= S4;
                else
                    NS <= S1;
                end if;
            when S2 =>
                MULTI <= '1';
                CONTIG <= '0';
                SINGLE <= '0';
                if (not D) then
                    NS <= S3;
                else
                    NS <= S4;
                end if;
            when S3 =>
                MULTI <= '0';
```

```

        CONTIG <= '1';
        SINGLE <= '0';
        if (A or D) then
            NS <= S4;
        else
            NS <= S3;
        end if;
    when S4 =>
        MULTI <= '1';
        CONTIG <= '1';
        SINGLE <= '0';
        if (A and B and not C) then
            NS <= S5;
        else
            NS <= S4;
        end if;
    when S5 =>
        MULTI <= '1';
        CONTIG <= '0';
        SINGLE <= '0';
        NS <= S6;
    when S6 =>
        MULTI <= '0';
        CONTIG <= '1';
        SINGLE <= '1';
        if (not E) then
            NS <= S7;
        else
            NS <= S6;
        end if;
    when S7 =>
        MULTI <= '0';
        CONTIG <= '1';
        SINGLE <= '0';
        if (E) then
            NS <= S1;
        else
            NS <= S7;
        end if;
    end case;
end process; -- End COMB_PROC

end BEHV;

```

Binary Encoded State Machine Verilog Example

```
////////////////////////////////////
// BINARY.V Version 1.0 //
// Example of a binary encoded state machine //
// May 1997 //
////////////////////////////////////

module binary (CLOCK, RESET, A, B, C, D, E,
              SINGLE, MULTI, CONTIG);

input CLOCK, RESET;
input A, B, C, D, E;
output SINGLE, MULTI, CONTIG;

reg SINGLE, MULTI, CONTIG;

// Declare the symbolic names for states
parameter [2:0]
    S1 = 3'b001,
    S2 = 3'b010,
    S3 = 3'b011,
    S4 = 3'b100,
    S5 = 3'b101,
    S6 = 3'b110,
    S7 = 3'b111;

// Declare current state and next state variables
reg [2:0] CS;
reg [2:0] NS;

// state_vector CS

always @ (posedge CLOCK or posedge RESET)
begin
    if (RESET == 1'b1)
        CS = S1;
    else
        CS = NS;
end

always @ (CS or A or B or C or D or D or E)
begin
    case (CS)
        S1 :
            begin
```

```

        MULTI = 1'b0;
        CONTIG = 1'b0;
        SINGLE = 1'b0;
    if (A && ~B && C)
        NS = S2;
    else if (A && B && ~C)
        NS = S4;
    else
        NS = S1;
end
S2 :
begin
    MULTI = 1'b1;
    CONTIG = 1'b0;
    SINGLE = 1'b0;
    if (!D)
        NS = S3;
    else
        NS = S4;
end
S3 :
begin
    MULTI = 1'b0;
    CONTIG = 1'b1;
    SINGLE = 1'b0;
    if (A || D)
        NS = S4;
    else
        NS = S3;
end
S4 :
begin
    MULTI = 1'b1;
    CONTIG = 1'b1;
    SINGLE = 1'b0;
    if (A && B && ~C)
        NS = S5;
    else
        NS = S4;
end
S5 :
begin
    MULTI = 1'b1;
    CONTIG = 1'b0;
    SINGLE = 1'b0;

```

```
        NS = S6;
    end
    S6 :
    begin
        MULTI   = 1'b0;
        CONTIG  = 1'b1;
        SINGLE  = 1'b1;
        if (!E)
            NS = S7;
        else
            NS = S6;
        end
    end
    S7 :
    begin
        MULTI   = 1'b0;
        CONTIG  = 1'b1;
        SINGLE  = 1'b0;
        if (E)
            NS = S1;
        else
            NS = S7;
        end
    end
endcase
end
endmodule
```

Using Enumerated Type Encoding

The recommended encoding style for state machines depends on which synthesis tool you are using. Some synthesis tools encode better than others depending on the device architecture and the size of the decode logic. You can explicitly declare state vectors or you can allow your synthesis tool to determine the vectors. Xilinx recommends that you use enumerated type encoding to specify the states and use the Finite State Machine (FSM) extraction commands to extract and encode the state machine as well as to perform state minimization and optimization algorithms. The enumerated type method of encoding the seven-state machine is shown in the following VHDL and Verilog examples. The encoding style is not defined in the code, but can be specified later with the FSM extraction commands. Alternatively, you can allow your compiler to select the encoding style that results in the lowest gate count when the design is synthesized. Some

synthesis tools automatically find finite state machines and compile without the need for specification.

Note: Refer to the previous VHDL and Verilog Binary Encoded State Machine examples for the complete Case statement portion of the code.

Enumerated Type Encoded State Machine VHDL Example

```

Library IEEE;
use IEEE.std_logic_1164.all;
entity enum is
    port (CLOCK, RESET : in STD_LOGIC;
          A, B, C, D, E: in BOOLEAN;
          SINGLE, MULTI, CONTIG: out STD_LOGIC);
end enum;

architecture BEHV of enum is

    type STATE_TYPE is (S1, S2, S3, S4, S5, S6, S7);

    signal CS, NS: STATE_TYPE;

begin
    SYNC_PROC: process (CLOCK, RESET)
    begin
        if (RESET='1') then
            CS <= S1;
        elsif (CLOCK'event and CLOCK = '1') then
            CS <= NS;
        end if;
    end process; --End SYNC_PROC

    COMB_PROC: process (CS, A, B, C, D, E)
    begin
        case CS is
            when S1 =>
                MULTI <= '0';
                CONTIG <= '0';
                SINGLE <= '0';
            .
            .
            .
        end case;
    end process;
end architecture BEHV;

```

Enumerated Type Encoded State Machine Verilog Example

```
////////////////////////////////////
// ENUM.V Version 1.0 //
// Example of an enumerated encoded state machine //
// May 1997 //
////////////////////////////////////

module enum (CLOCK, RESET, A, B, C, D, E,
            SINGLE, MULTI, CONTIG);

input CLOCK, RESET;
input A, B, C, D, E;
output SINGLE, MULTI, CONTIG;

reg SINGLE, MULTI, CONTIG;

// Declare the symbolic names for states
parameter [2:0]
    S1 = 3'b000,
    S2 = 3'b001,
    S3 = 3'b010,
    S4 = 3'b011,
    S5 = 3'b100,
    S6 = 3'b101,
    S7 = 3'b110;

// Declare current state and next state variables
reg [2:0] CS;
reg [2:0] NS;

// state_vector CS

always @ (posedge CLOCK or posedge RESET)
begin
    if (RESET == 1'b1)
        CS = S1;
    else
        CS = NS;
end

always @ (CS or A or B or C or D or D or E)
begin
    case (CS)
```



```
S1 :  
begin  
MULTI = 1'b0;  
CONTIG = 1'b0;  
SINGLE = 1'b0;  
if (A && ~B && C)  
    NS = S2;  
else if (A && B && ~C)  
    NS = S4;  
else  
    NS = S1;  
end
```

```
.  
. .  
.
```

Using One-Hot Encoding

One-hot encoding allows you to create state machine implementations that are more efficient for FPGA architectures. One-hot encoding is usually the preferred method for large FPGA-based state machine implementation.

The following examples show a one-hot encoded state machine. Use this method to control the state vector specification or when you want to specify the names of the state registers. These examples use one flip-flop for each of the seven states. If you are using FPGA Express, use enumerated type, and avoid using the “when others” construct in the VHDL case statement. This construct can result in a very large state machine.

Note: Refer to the previous VHDL and Verilog Binary Encoded State Machine examples for the complete Case statement portion of the code. See the “Accelerate FPGA Macros with One-Hot Approach” appendix for a detailed description of one-hot encoding and its applications.

One-hot Encoded State Machine VHDL Example

```
Library IEEE;
use IEEE.std_logic_1164.all;
use IEEE.std_logic_unsigned.all;

entity one_hot is
    port (CLOCK, RESET : in STD_LOGIC;
          A, B, C, D, E: in BOOLEAN;
          SINGLE, MULTI, CONTIG: out STD_LOGIC);
end one_hot;

architecture BEHV of one_hot is

    type STATE_TYPE is (S1, S2, S3, S4, S5, S6, S7);
    attribute ENUM_ENCODING: STRING;
    attribute ENUM_ENCODING of STATE_TYPE: type is "0000001 0000010 0000100
    0001000 0010000 0100000 1000000 ";

    signal CS, NS: STATE_TYPE;

begin

    SYNC_PROC: process (CLOCK, RESET)
    begin
        if (RESET='1') then
            CS <= S1;
        elsif (CLOCK'event and CLOCK = '1') then
            CS <= NS;
        end if;
    end process; --End SYNC_PROC

    COMB_PROC: process (CS, A, B, C, D, E)
    begin
        case CS is
            when S1 =>
                MULTI <= '0';
                CONTIG <= '0';
                SINGLE <= '0';
            if (A and not B and C) then
                NS <= S2;
            elsif (A and B and not C) then
                NS <= S4;
            else
                NS <= S1;
            end if;
        end case;
    end process;
end architecture BEHV;
```

```
end if;
.
.
.
```

One-hot Encoded State Machine Verilog Example

```

////////////////////////////////////
// ONE_HOT.V Version 1.0 //
// Example of a one-hot encoded state machine //
// Xilinx HDL Synthesis Design Guide for FPGAs //
// May 1997 //
////////////////////////////////////

module one_hot (CLOCK, RESET, A, B, C, D, E,
               SINGLE, MULTI, CONTIG);

input  CLOCK, RESET;
input  A, B, C, D, E;
output SINGLE, MULTI, CONTIG;

reg SINGLE, MULTI, CONTIG;

// Declare the symbolic names for states
parameter [6:0]
    S1 = 7'b0000001,
    S2 = 7'b0000010,
    S3 = 7'b0000100,
    S4 = 7'b0001000,
    S5 = 7'b0010000,
    S6 = 7'b0100000,
    S7 = 7'b1000000;

// Declare current state and next state variables
reg [2:0] CS;
reg [2:0] NS;

// state_vector CS

always @ (posedge CLOCK or posedge RESET)
begin
    if (RESET == 1'b1)
        CS = S1;
    else
        CS = NS;
end

```

```
always @ (CS or A or B or C or D or D or E)
begin
    case (CS)
        S1 :
            begin
                MULTI = 1'b0;
                CONTIG = 1'b0;
                SINGLE = 1'b0;
                if (A && ~B && C)
                    NS = S2;
                else if (A && B && ~C)
                    NS = S4;
                else
                    NS = S1;
            end
        .
        .
        .
    end
```

Summary of Encoding Styles

In the three previous examples, the state machine's possible states are defined by an enumeration type. Use the following syntax to define an enumeration type.

```
type type_name is (enumeration_literal { , enumeration_literal } );
```

After you have defined an enumeration type, declare the signal representing the states as the enumeration type as follows.

```
type STATE_TYPE is (S1, S2, S3, S4, S5, S6, S7);
```

```
signal CS, NS: STATE_TYPE;
```

The state machine described in the three previous examples has seven states. The possible values of the signals CS (Current_State) and NS (Next_State) are S1, S2, ... , S6, S7.

To select an encoding style for a state machine, specify the state vectors. Alternatively, you can specify the encoding style when the state machine is compiled. Xilinx recommends that you specify an encoding style. If you do not specify a style, your compiler selects a style that minimizes the gate count. For the state machine shown in the three previous examples, the compiler selected the binary encoded style: S1="000", S2="001", S3="010", S4="011", S5="100", S6="101", and S7="110".

You can use the FSM extraction tool to change the encoding style of a state machine. For example, use this tool to convert a binary-encoded state machine to a one-hot encoded state machine.

Note: Refer to your synthesis tool documentation for instructions on how to extract the state machine and change the encoding style.

Comparing Synthesis Results for Encoding Styles

The following table summarizes the synthesis results from the different methods used to encode the state machine in the three previous VHDL and Verilog state machine examples. The results are for an XC4005EPC84-2 device

Note: The Timing Analyzer was used to obtain the timing results in this table.

Table 4-1 State Machine Encoding Styles Comparison (XC4005E-2)

| Comparison | One-Hot | Binary | Enum (One-hot) |
|----------------|--------------------------|-------------|----------------|
| Occupied CLBs | 6 | 9 | 6 |
| CLB Flip-flops | 6 | 3 | 7 |
| PadToSetup | 9.4 ns (3 ^a) | 13.4 ns (4) | 9.6 ns (3) |
| ClockToPad | 15.1 ns (3) | 15.1 ns (3) | 14.9 ns (3) |
| ClockToSetup | 13.0 ns (4) | 13.9 ns (4) | 10.1 ns (3) |

a. The number in parentheses represents the CLB block level delay.

The binary-encoded state machine has the longest ClockToSetup delay. Generally, the FSM extraction tool provides the best results because the compiler reduces any redundant states and optimizes the state machine after the extraction.

Initializing the State Machine

When you use one-hot encoding, add the following lines of code to your design to ensure that the FPGA is initialized to a Set state.

- **VHDL Example**

```
SYNC_PROC: process (CLOCK, RESET)
begin
    if (RESET='1') then
        CS <= s1;
```

- **Verilog Example**

```
always @ (posedge CLOCK or posedge RESET)
begin
    if (RESET == 1'b 1)
        CS = S1;
```

Alternatively, you can assign an INIT=S attribute to the initial state register to specify the initial state. Refer to your synthesis tool documentation for information on assigning this attribute.

In the Binary Encode State Machine example, the RESET signal forces the S1 flip-flop to be preset (initialized to 1) while the other flip-flops are cleared (initialized to 0).

Using Dedicated I/O Decoders

The periphery of XC4000 family devices has four wide decoder circuits at each edge. The inputs to each decoder are any of the IOB signals on that edge plus one local interconnect per CLB row or column. Each decoder generates a High output (using a pull-up resistor) when the AND condition of the selected inputs or their complements is true. The decoder outputs drive CLB inputs so they can be combined with other logic or can be routed directly to the chip outputs.

To implement XC4000 family edge decoders in HDL, you must instantiate edge decoder primitives. The primitive names you can use vary with the synthesis tool you are using. For example, you can instantiate DECODE1_IO, DECODE1_INT, DECODE4, DECODE8, and DECODE16. These primitives are implemented using the dedicated I/O edge decoders. The XC4000 family wide decoder outputs are effectively open-drain and require a pull-up resistor to take the output High when the specified pattern is detected on the decoder inputs. To attach the pull-up resistor to the output signal, you must instantiate a PULLUP component.

The following VHDL example shows how to use the I/O edge decoders by instantiating decode primitives. Each decoder output is a

function of ADR (IOB inputs) and CLB_INT (local interconnects). The AND function of each DECODE output and Chip Select (CS) serves as the source of a flip-flop Clock Enable pin. The four edge decoders in this design are placed on the same device edge. The “Schematic Block Representation of I/O Decoder” figure shows the schematic block diagram representation of this I/O decoder design.

Using Dedicated I/O Decoders VHDL Example

```
--Edge Decoder
--An XC4000 LCA has special decoder circuits at each edge. These decoders
--are open-drained wired-AND gates. When one or more of the inputs (I) are
--Low output(O) is Low. When all of the inputs are High, the output is
--High. A pull-up resistor must be connected to the output node to achieve
--a true logic High.
```

```
Library IEEE;
use IEEE.STD_LOGIC_1164.all;
use IEEE.STD_LOGIC_UNSIGNED.all;

entity io_decoder is
    port (ADR: in std_logic_vector (4 downto 0);
          CS: in std_logic;
          DATA: in std_logic_vector (3 downto 0);
          CLOCK: in std_logic;
          QOUT: out std_logic_vector (3 downto 0));
end io_decoder;

architecture STRUCTURE of io_decoder is

    COMPONENT DECODE1_IO
        PORT ( I: IN std_logic;
              O: OUT std_logic );
    END COMPONENT;

    COMPONENT DECODE1_INT
        PORT ( I: IN std_logic;
              O: OUT std_logic );
    END COMPONENT;

    COMPONENT DECODE4
        PORT ( A3, A2, A1, A0: IN std_logic;
              O: OUT std_logic );
    END COMPONENT;
```

```
COMPONENT PULLUP
  PORT ( O: OUT std_logic );
END COMPONENT;

---- Internal Signal Declarations -----
signal DECODE, CLKEN, CLB_INT: std_logic_vector (3 downto 0);
signal ADR_INV, CLB_INV: std_logic_vector (3 downto 0);
begin

ADR_INV <= not ADR (3 downto 0);
CLB_INV <= not CLB_INT;

----- Instantiation of Edge Decoder: Output "DECODE(0)" -----
  A0: DECODE4 port map (ADR(3), ADR(2), ADR(1), ADR_INV(0), DECODE(0));

  A1: DECODE1_IO port map (ADR(4), DECODE(0));

  A2: DECODE1_INT port map (CLB_INV(0), DECODE(0));

  A3: DECODE1_INT port map (CLB_INT(1), DECODE(0));

  A4: DECODE1_INT port map (CLB_INT(2), DECODE(0));

  A5: DECODE1_INT port map (CLB_INT(3), DECODE(0));

  A6: PULLUP port map (DECODE(0));

----- Instantiation of Edge Decoder: Output "DECODE(1)" -----
  B0: DECODE4 port map (ADR(3), ADR(2), ADR_INV(1), ADR(0), DECODE(1));

  B1: DECODE1_IO port map (ADR(4), DECODE(1));

  B2: DECODE1_INT port map (CLB_INT(0), DECODE(1));

  B3: DECODE1_INT port map (CLB_INV(1), DECODE(1));

  B4: DECODE1_INT port map (CLB_INT(2), DECODE(1));

  B5: DECODE1_INT port map (CLB_INT(3), DECODE(1));

  B6: PULLUP port map (DECODE(1));

----- Instantiation of Edge Decoder: Output "DECODE(2)" -----
  C0: DECODE4 port map (ADR(3), ADR_INV(2), ADR(1), ADR(0), DECODE(2));
```



```

C1: DECODE1_IO port map (ADR(4), DECODE(2));

C2: DECODE1_INT port map (CLB_INT(0), DECODE(2));

C3: DECODE1_INT port map (CLB_INT(1), DECODE(2));

C4: DECODE1_INT port map (CLB_INV(2), DECODE(2));

C5: DECODE1_INT port map (CLB_INT(3), DECODE(2));

C6: PULLUP port map (DECODE(2));

----- Instantiation of Edge Decoder: Output "DECODE(3)" -----
D0: DECODE4 port map (ADR_INV(3), ADR(2), ADR(1), ADR(0), DECODE(3));

D1: DECODE1_IO port map (ADR(4), DECODE(3));

D2: DECODE1_INT port map (CLB_INT(0), DECODE(3));

D3: DECODE1_INT port map (CLB_INT(1), DECODE(3));

D4: DECODE1_INT port map (CLB_INT(2), DECODE(3));

D5: DECODE1_INT port map (CLB_INV(3), DECODE(3));

D6: PULLUP port map (DECODE(3));

-----CLKEN is the AND function of CS & DECODE-----

CLKEN(0) <= CS and DECODE(0);
CLKEN(1) <= CS and DECODE(1);
CLKEN(2) <= CS and DECODE(2);
CLKEN(3) <= CS and DECODE(3);

-----Internal 4-bit counter -----
process (CLOCK)
begin
    if (CLOCK'event and CLOCK='1') then
        CLB_INT <= CLB_INT + 1;
    end if;
end process;

-----"QOUT(0)" Data Register Enabled by "CLKEN(0)"-----
process (CLOCK)
begin

```

```
        if (CLOCK'event and CLOCK='1') then
            if (CLKEN(0) = '1') then
                QOUT(0) <= DATA(0);
            end if;
        end if;
    end process;

-----"QOUT(1)" Data Register Enabled by "CLKEN(1)"-----
    process (CLOCK)
    begin
        if (CLOCK'event and CLOCK='1') then
            if (CLKEN(1) = '1') then
                QOUT(1) <= DATA(1);
            end if;
        end if;
    end process;

-----"QOUT(2)" Data Register Enabled by "CLKEN(2)"-----
    process (CLOCK)
    begin
        if (CLOCK'event and CLOCK='1') then
            if (CLKEN(2) = '1') then
                QOUT(2) <= DATA(2);
            end if;
        end if;
    end process;

-----"QOUT(3)" Data Register Enabled by "CLKEN(3)"-----
    process (CLOCK)
    begin
        if (CLOCK'event and CLOCK='1') then
            if (CLKEN(3) = '1') then
                QOUT(3) <= DATA(3);
            end if;
        end if;
    end process;

end STRUCTURE;
```

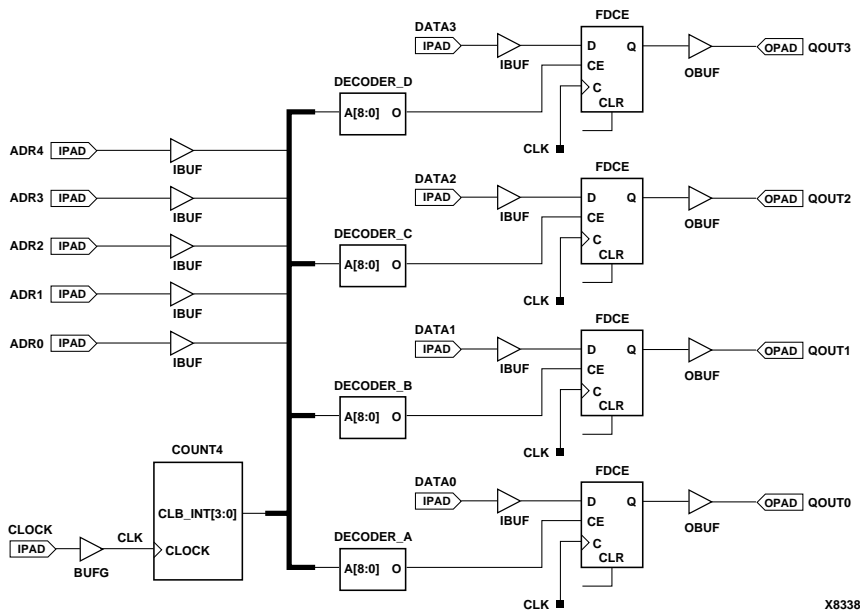


Figure 4-6 Schematic Block Representation of I/O Decoder

Note: In the previous figure, the pull-up resistors are inside the Decoder blocks.

Instantiating LogiBLOX Modules

Note: Refer to the *LogiBLOX Reference/User Guide* for detailed instructions on using LogiBLOX.

Most synthesis tools can infer arithmetic modules from VHDL or Verilog code for these operators: +, -, <, <=, >, >=, =, +1, -1. These adders, subtractors, comparators, incrementers, and decrementers use FPGA dedicated device resources, such as carry logic, to improve the speed and area of designs. For bus widths greater than four, library modules are generally faster unless multiple instances of the same function are compiled together. For more information on the module libraries, refer to your synthesis tool documentation.

If you want to use a module that is not in the module libraries, you can use LogiBLOX to create components that can be instantiated in your code. This is useful for large memory arrays if your synthesis

tool does not infer memory. However, Xilinx recommends properly constraining the synthesis and using the Xilinx-specific module generation capabilities of your tool. A simulation model is also created so that RTL simulation can be performed before your design is compiled.

You can create an instance of an externally defined macro, including a user-defined macro or a Xilinx macro (such as an I/O or flip-flop), by instantiating what some synthesis tool vendors refer to as a “black box” in your HDL code. These black boxes are Verilog empty module descriptions or VHDL component declarations.

Some synthesis tools allow instantiation of higher order Xilinx macros, such as counters and adders from the Unified library. Other synthesis tools provide Xilinx macro libraries that pre-define the Xilinx macros. Without this expansion, macros are not understood by the implementation tools. However, Xilinx does not recommend using these macros. The preferred method is the synthesis tool module expansion, or if you require more control, you can instantiate a LogiBLOX module. If necessary, use these macro libraries only with older schematic-based designs. However, even in these cases, schematic-based netlists are required to expand the macros, which makes the macro library redundant. LogiBLOX modules should also be unnecessary because the synthesis tool should provide equivalent performance. If you find a design in which this is not true, you can use LogiBLOX modules, and contact Xilinx and your synthesis vendor for a solution.

LogiBLOX is a graphical tool that allows you to select from several arithmetic, logic, I/O, sequential, and data storage modules for inclusion in your HDL design. Use LogiBLOX to instantiate the modules listed in the following table.

Table 4-2 LogiBLOX Modules

| Module | Description |
|-------------------|--|
| Arithmetic | |
| Accumulator | Adds data to or subtracts it from the current value stored in the accumulator register |
| Adder/Subtractor | Adds or subtracts two data inputs and a carry input |
| Comparator | Compares the magnitude or equality of two values |
| Counter | Generates a sequence of count values |

Table 4-2 LogiBLOX Modules

| Module | Description |
|------------------------------------|---|
| Logic | |
| Constant | Forces a constant value onto a bus |
| Decoder | Routes input data to 1-of-n lines on the output port |
| Multiplexer | Type 1, Type 2 - Routes input data on 1-of-n lines to the output port |
| Simple Gates | Type 1, Type 2, Type 3 - Implements the AND, INVERT, NAND, NOR, OR, XNOR, and XOR logic functions |
| Tristate | Creates a tri-stated internal data bus |
| I/O | |
| Bi-directional Input/Output | Connects internal and external pin signals |
| Pad | Simulates an input/output pad |
| Sequential | |
| Clock Divider | Generates a period that is a multiple of the clock input period |
| Counter | Generates a sequence of count values |
| Shift Register | Shifts the input data to the left or right |
| Storage | |
| Data Register | Captures the input data on active clock transitions |
| Memory: ROM, RAM, SYNC_RAM, DP_RAM | Stores information and makes it readable |

Using LogiBLOX in HDL Designs

- Before using LogiBLOX, verify the following.
 - Xilinx software is correctly installed
 - Environment variables are set correctly
 - Your display environment variable is set to your machine's display

2. To run LogiBLOX, enter the following command.

lbgui

The LogiBLOX Setup Window appears after the LogiBLOX module generator is loaded. This window allows you to name and customize the module you want to create.

3. Select the Vendor tab in the Setup Window. Select your synthesis tool in the Vendor Name field to specify the correct bus notation for connecting your module.

Select the Project Directory tab. Enter the directory location of your project in the LogiBLOX Project Directory field.

Select the Device Family tab. Select the target device for your design in the Device Family field.

Select the Options tab and select the applicable options for your design as follows.

- Simulation Netlist

This option allows you to create simulation netlists of the selected LogiBLOX module in different formats. You can choose one or more of the outputs listed in the following table.

Table 4-3 Simulation Netlist Options

| Option | Description |
|----------------------------|---|
| Behavioral VHDL netlist | Generates a simulation netlist in behavioral VHDL; output file has a .vhd extension. |
| Gate level EDIF netlist | Generates a simulation netlist in EDIF format; output file has an .edn extension. |
| Structural Verilog netlist | Generates a simulation netlist in structural Verilog; output file has a .v extension. |

- Component Declaration

This option creates instantiation templates in different formats that can be copied into your design. You can select none, one, or both of the following options.

Table 4-4 Component Declaration Options

| Option | Description |
|------------------|---|
| VHDL template | Generates a LogiBLOX VHDL component declaration/instantiation template that is copied into your VHDL design when a LogiBLOX module is instantiated. The output file has a .vhi extension. |
| Verilog template | Generates a LogiBLOX Verilog module definition/instantiation template that is copied into your Verilog design when a LogiBLOX module is instantiated. The output file has a .vei extension. |

- Implementation Netlist

Select NGO File to generate an implementation netlist in Xilinx NGD binary format. You must select this option when instantiating LogiBLOX symbols in an HDL design. The output file has an .ngo extension and can be used as input to NGDBuild.

- LogiBLOX DRC

Select the Stop Process on Warning option to stop module processing if any warning messages are encountered during the design process.

For example, if you have a Verilog design, and you are simulating with Verilog-XL, select Structural Verilog netlist, Verilog template, NGO File, and Stop Process on Warning. For a VHDL design and simulating with VSS, select Behavioral VHDL, VHDL template, NGO File, and Stop Process on Warning.

Select OK.

4. Enter a name in the Module Name field in the Module Selector Window.

Select a base module type from the Module Type field.

Select a bus width from the Bus Width field.

Customize your module by selecting pins and specifying attributes.

After you have completed module specification, select OK.

This initiates the generation of a component instantiation declaration, a behavioral model, and an implementation netlist.

5. Copy the module declaration/instantiation into your design. The template file created by LogiBLOX is *module_name.vhi* (VHDL) or *module_name.vei* (Verilog), and is saved in the project directory as specified in the LogiBLOX setup.
6. Complete the signal connections of the instantiated module to the rest of your design.

Note: For more information on simulation, refer to the “Simulating Your Design” chapter.

7. Create an implementation script. Add the appropriate attribute to the instantiated LogiBLOX module to prevent synthesis of this module. Compile your design.

Also, if you have a Verilog design, use a remove design type of command to make the LogiBLOX netlist unavailable before writing the .xnf or .edif netlist.

Note: If you do not use a remove design type of command, the netlist file may be empty. If this occurs, the Xilinx software will trim this module/component and all connected logic. Refer to your synthesis tool documentation for the correct command and syntax.

8. Compile your design and create a .xnf or .edif file. You can safely ignore the following type of warning messages.

```
Warning: Can't find the design in the library WORK.  
(LBR-1)
```

```
Warning: Unable to resolve reference LogiBLOX_name in  
design_name. (LINK-5)
```

```
Warning: Design design_name has 1 unresolved references.  
For more detailed information, use the "link" command.  
(UID-341)
```


9. Implement your design with the Xilinx tools. Verify that the .ngo file created by LogiBLOX is in the same project directory as the netlist.

You may get the following warnings during the NGDBuild and mapping steps. These messages are issued if the Xilinx software can not locate the corresponding .ngo file created by LogiBLOX.

```
Warning: basnu - logical block LogiBLOX_instance_name of
type LogiBLOX_name is unexpanded. Logical Design DRC
complete with 1 warning(s).
```

If you get this message, you will get the following message during mapping.

```
ERROR:basnu - logical block LogiBLOX_instance_name of
type LogiBLOX_name is unexpanded. Errors detected in
general drc.
```

If you get these messages, first verify that the .ngo file created by LogiBLOX is in the project directory. If the file is there, verify that the module is properly instantiated in the code.

10. To simulate your post-layout design, convert your design to a timing netlist and use the back-annotation flow applicable to your synthesis tool.

Note: For more information on simulation, refer to the “Simulating Your Design” chapter.

Instantiating a LogiBLOX “Black Box” Component

The VHDL example in this section shows how to instantiate a LogiBLOX “black box” component.

VHDL Example

```
entity top is
    port (clk, rst, en, data: in bit; q: out bit);
end top;

architecture structural of top is
    -- Declare the black_box as a boolean attribute
    attribute black_box: boolean;
    -- Declare the black_box_pad_pin as a string attribute
    attribute black_box_pad_pin: string;
```

```
-- In this example, GIZMO is a user macro that I created
-- in a schematic editor, and
-- that I want to directly instantiate in my VHDL design
-- as a black box. Create a component declaration.
component GIZMO
    port(Q: out bit; D, C, CLR: in bit);
end component;
-- Set the black_box attribute on GIZMO to be "true".
attribute black_box of GIZMO: component is true;

-- In this example, MYBUF is a user I/O macro that I created
-- in a schematic editor, and
-- that I want to directly instantiate in my VHDL design
-- as a black box. Create a component declaration.
component MYBUF
    port(O: out bit; I: in bit);
end component;
-- Set the black_box_pad_pin attribute on MYBUF to
-- the pin that interfaces with the external world, "I".
attribute black_box_pad_pin of MYBUF: component is "I";

signal data_core: bit;

begin
    -- Instantiate an MYBUF. Here we connect
    -- data to I and data_core to O.
    data_pad: MYBUF port map (O => data_core, I => data);

    -- Instantiate a GIZMO. Here we connect q to Q,
    -- data_core to D, clk to C, and rst to CLR.
    my_gizmo: GIZMO port map (Q => q, D => data_core,
        C => clk, CLR => rst);
end structural;
```

Implementing Memory

XC4000E/EX/XL and Spartan FPGAs provide distributed on-chip RAM or ROM. CLB function generators can be configured as ROM (ROM16X1, ROM32X1); level-sensitive RAM (RAM16X1, RAM32X1); edge-triggered, single-port (RAM16X1S, RAM32X1S); or dual-port (RAM16x1D) RAM. The edge-triggered capability simplifies system timing and provides better performance for RAM-based designs. This distributed RAM can be used for status registers, index registers, counter storage, constant coefficient multipliers, distributed

shift registers, LIFO stacks, latching, or any data storage operation. The dual-port RAM simplifies FIFO designs.

Note: For more information on XC4000 family RAM, refer to the Xilinx Web site or the current release of *The Programmable Logic Data Book*.

Implementing XC4000 and Spartan ROMs

ROMs can be implemented as follows.

- Use RTL descriptions of ROMs
- Instantiate 16x1 and 32x1 ROM primitives
- Use LogiBLOX to implement any other ROM size

VHDL and Verilog examples of an RTL description of a ROM follow.

RTL Description of a ROM VHDL Example

```
--
-- Behavioral 16x4 ROM Example
--      rom_rtl.vhd
--

library IEEE;
use IEEE.std_logic_1164.all;

entity rom_rtl is
    port (ADDR: in INTEGER range 0 to 15;
          DATA: out STD_LOGIC_VECTOR (3 downto 0));
end rom_rtl;

architecture XILINX of rom_rtl is

    subtype ROM_WORD is STD_LOGIC_VECTOR (3 downto 0);
    type ROM_TABLE is array (0 to 15) of ROM_WORD;
    constant ROM: ROM_TABLE := ROM_TABLE'(
        ROM_WORD("0000"),
        ROM_WORD("0001"),
        ROM_WORD("0010"),
        ROM_WORD("0100"),
        ROM_WORD("1000"),
        ROM_WORD("1100"),
        ROM_WORD("1010"),
        ROM_WORD("1001"),
```

```
        ROM_WORD'("1001"),
        ROM_WORD'("1010"),
        ROM_WORD'("1100"),
        ROM_WORD'("1001"),
        ROM_WORD'("1001"),
        ROM_WORD'("1101"),
        ROM_WORD'("1011"),
        ROM_WORD'("1111"));

    begin
        DATA <= ROM(ADDR); -- Read from the ROM
    end XILINX;
```

RTL Description of a ROM Verilog Example

```
/*
 * ROM_RTL.V
 * Behavioral Example of 16x4 ROM
 */

module rom_rtl(ADDR, DATA) ;
input [3:0] ADDR ;
output [3:0] DATA ;

reg [3:0] DATA ;

// A memory is implemented
// using a case statement

always @(ADDR)
begin
    case (ADDR)
        4'b0000 : DATA = 4'b0000 ;
        4'b0001 : DATA = 4'b0001 ;
        4'b0010 : DATA = 4'b0010 ;
        4'b0011 : DATA = 4'b0100 ;
        4'b0100 : DATA = 4'b1000 ;
        4'b0101 : DATA = 4'b1000 ;
        4'b0110 : DATA = 4'b1100 ;
        4'b0111 : DATA = 4'b1010 ;
        4'b1000 : DATA = 4'b1001 ;
        4'b1001 : DATA = 4'b1001 ;
        4'b1010 : DATA = 4'b1010 ;
        4'b1011 : DATA = 4'b1100 ;
        4'b1100 : DATA = 4'b1001 ;
    endcase
end
```

```
        4'b1101 : DATA = 4'b1001 ;
        4'b1110 : DATA = 4'b1101 ;
        4'b1111 : DATA = 4'b1111 ;
    endcase
end

endmodule
```

When using an RTL description of a ROM, the synthesis tool creates ROMs from random logic gates that are implemented using function generators.

Another method for implementing ROMs is instantiating the 16x1 or 32x1 ROM primitives. To define the ROM value, use the Set Attribute or equivalent command to set the INIT property on the ROM component.

Note: Refer to your synthesis tool documentation for the correct syntax.

This type of command writes the ROM contents to the netlist file so the Xilinx tools can initialize the ROM. The INIT value should be specified in hexadecimal values. See the VHDL and Verilog RAM examples in the following section for examples of this property using a RAM primitive.

Implementing XC4000 Family RAMs

Do not use RTL descriptions of RAMs in your code because they do not compile efficiently and can cause combinatorial loops. The exception to this is if your synthesis tool can infer memory. In this case, you must follow a strict coding style. Refer to your vendor's documentation for more information.

You can implement RAMs as follows.

- Instantiate 16x1 and 32x1 RAM primitives (RAM16X1, RAM32X1, RAM16X1S, RAM32X1S, RAM16X1D)
- Use LogiBLOX to implement any other RAM size
- Some synthesis tools can infer RAMs from your code

When implementing RAM in XC4000 and Spartan designs, Xilinx recommends using the synchronous write, edge-triggered RAM (RAM16X1S, RAM32X1S, or RAM16X1D) instead of the asynchronous-write RAM (RAM16X1 or RAM32X1) to simplify write timing and increase RAM performance.

Examples of an instantiation of edge-triggered RAM primitives are provided in the following VHDL and Verilog designs. As with ROMs, initial RAM values can be specified from the command line. The INIT property value is specified in hexadecimal values. Refer to your synthesis tool documentation for the correct command and syntax.

An Exemplar™ example of a RAM inference (ram.vhd) is also included in this section. Check with your synthesis tool vendor for the availability of this feature.

Instantiating RAM VHDL Example

```
-----  
-- RAM_PRIMITIVE.VHD          --  
-- Example of instantiating 4  --  
-- 16x1 synchronous RAMs      --  
-- HDL Synthesis Design Guide for FPGAs --  
-- May 1997                    --  
-----  
  
library IEEE;  
use IEEE.std_logic_1164.all;  
  
entity ram_primitive is  
  
    port ( DATA_IN, ADDR : in STD_LOGIC_VECTOR(3 downto 0);  
          WE, CLOCK      : in STD_LOGIC;  
          DATA_OUT      : out STD_LOGIC_VECTOR(3 downto 0));  
  
end ram_primitive;  
  
architecture STRUCTURAL_RAM of ram_primitive is  
  
    component RAM16X1S  
        port (D, A3, A2, A1, A0, WE, WCLK : in STD_LOGIC;  
              O : out STD_LOGIC);  
    end component;  
  
end architecture;
```

```

begin

    RAM0 : RAM16X1S port map (O => DATA_OUT(0), D => DATA_IN(0),
                              A3 => ADDR(3), A2 => ADDR(2),
                              A1 => ADDR(1), A0 => ADDR(0),
                              WE => WE, WCLK => CLOCK);

    RAM1 : RAM16X1S port map (O => DATA_OUT(1), D => DATA_IN(1),
                              A3 => ADDR(3), A2 => ADDR(2),
                              A1 => ADDR(1), A0 => ADDR(0),
                              WE => WE, WCLK => CLOCK);

    RAM2 : RAM16X1S port map (O => DATA_OUT(2), D => DATA_IN(2),
                              A3 => ADDR(3), A2 => ADDR(2),
                              A1 => ADDR(1), A0 => ADDR(0),
                              WE => WE, WCLK => CLOCK);

    RAM3 : RAM16X1S port map (O => DATA_OUT(3), D => DATA_IN(3),
                              A3 => ADDR(3), A2 => ADDR(2),
                              A1 => ADDR(1), A0 => ADDR(0),
                              WE => WE, WCLK => CLOCK);

end STRUCTURAL_RAM;

```

Instantiating RAM Verilog Example

```

////////////////////////////////////
// RAM_PRIMITIVE.V //
// Example of instantiating 4 //
// 16x1 Synchronous RAMs //
// HDL Synthesis Design Guide for FPGAs //
// August 1997 //
////////////////////////////////////

module ram_primitive (DATA_IN, ADDR, WE, CLOCK, DATA_OUT);

input [3:0] DATA_IN, ADDR;
input WE, CLOCK;
output [3:0] DATA_OUT;

RAM16X1S RAM0 (.O(DATA_OUT[0]), .D(DATA_IN[0]), .A3(ADDR[3]),
              .A2(ADDR[2]), .A1(ADDR[1]), .A0(ADDR[0]),
              .WE(WE), .WCLK(CLOCK));

```

```
RAM16X1S RAM1 (.O(DATA_OUT[1]), .D(DATA_IN[1]), .A3(ADDR[3]),
               .A2(ADDR[2]), .A1(ADDR[1]), .A0(ADDR[0]),
               .WE(WE), .WCLK(CLOCK));

RAM16X1S RAM2 (.O(DATA_OUT[2]), .D(DATA_IN[2]), .A3(ADDR[3]),
               .A2(ADDR[2]), .A1(ADDR[1]), .A0(ADDR[0]),
               .WE(WE), .WCLK(CLOCK));

RAM16X1S RAM3 (.O(DATA_OUT[3]), .D(DATA_IN[3]), .A3(ADDR[3]),
               .A2(ADDR[2]), .A1(ADDR[1]), .A0(ADDR[0]),
               .WE(WE), .WCLK(CLOCK));

endmodule
```

Inferring RAM VHDL Example

```
library ieee;
use ieee.std_logic_1164.all;
library exemplar;
use exemplar.exemplar_1164.all;
library exemplar;
use exemplar.exemplar.all;

package my_pkg is
    type MEM_WORD is array (6 downto 0) of elbit_vector (1 downto 0);
end my_pkg;
=20
library exemplar;
use exemplar.exemplar.all;
use work.my_pkg.all;

entity mem is
    port (dio      : inout elbit_vector (1 downto 0);=20
          meme, we, inclk, outclk : in bit;
          addr     : integer range 6 downto 0;
          ro       : out bit);
    attribute clock_node : boolean;
    attribute clock_node of inclk : signal is TRUE;
    attribute clock_node of outclk : signal is TRUE;

    =20
end mem;

architecture behav of mem is
    signal mem : MEM_WORD;
```



```

    signal d_int : elbit_vector (1 downto 0);
begin

process (inclk)
begin
    if (inclk'event and inclk =3D '1') then
        if (meme =3D '1' and we =3D '1') then
            mem(addr) <=3D dio;
        end if;
    end if;
end process;
process (outclk)
begin
    if (outclk'event and outclk =3D '1') then
        d_int <=3D mem (addr);
    end if;
end process;
dio <=3D d_int when (meme =3D '1' and we =3D '0') else "ZZ";
end behav;

```

Using LogiBLOX to Implement Memory

If you must instantiate memory, use LogiBLOX to create a memory module larger than 32X1 (16X1 for Dual Port). Implementing memory with LogiBLOX is similar to implementing any module with LogiBLOX except for defining the Memory initialization file. Use the following steps to create a memory module.

Note: Refer to the “Using LogiBLOX in HDL Designs” section for more information on using LogiBLOX.

1. Before using LogiBLOX, verify the following.
 - Xilinx software is correctly installed
 - Environment variables are set correctly
 - Your display environment variable is set to your machine’s display
2. To run LogiBLOX, enter the following command.

lbgui

The LogiBLOX Setup Window appears after the LogiBLOX module generator is loaded. This window allows you to name and customize the module you want to create.

3. Select the Vendor tab in the Setup Window. Select your synthesis tool in the Vendor Name field to specify the correct bus notation for connecting your module.

Select the Project Directory tab. Enter the directory location of your project in the LogiBLOX Project Directory field.

Select the Device Family tab. Select the target device for your design in the Device Family field.

Select the Options tab and select the applicable options for your design.

Select OK.

4. Enter a name in the Module Name field in the Module Selector Window.

Select the Memories module type from the Module Type field to specify that you are creating a memory module.

Select a width (any value from 1 to 64 bits) for the memory from the Data Bus Width field.

In the Details field, select the type of memory you are creating (ROM, RAM, SYNC_RAM, or DP_RAM).

Enter a value in the Memory Depth field for your memory module.

Note: Xilinx recommends (this is not a requirement) that you select a memory depth value that is a multiple of 16 because this is the memory size of one lookup table.

5. If you want the memory module initialized to all zeros on power up, you do not need to create a memory file (Mem File). However, if you want the contents of the memory initialized to a value other than zero, you must create and edit a memory file. Enter a memory file name in the Mem File field and click on the Edit button. Continue with the following steps.

Note: Some memory modules can only be initialized to zero. Refer to the *Xilinx Programmable Logic Data Book* for more information.

- a) A memory template file in a text editor is displayed. This file does not contain valid data, and must be edited before you can use it. The data values specified in the memory file Data Section define the contents of the memory. Data values are

specified sequentially, beginning with the lowest address in the memory, as defined.

- b) Specify the address of a data value. The default radix of the data values is 16. If more than one radix definition is listed in the memory file header section, the last definition is the radix used in the Data Section.

The following definition defines a 16-word memory with the contents 6, 4, 5, 5, 2, 7, 5, 3, 5, 5, 5, 5, 5, 5, 5, 5, starting at address 0. Note that the contents of locations 2, 3, 6, and 8 through 15 are defined via the default definition. Two starting addresses, 4 and 7, are given.

```
depth 16
default 5
data 6,4,
4: 2, 7,
7: 3
```

- c) After you have finished specifying the data for the memory module, save the file and exit the editor.
6. Click the OK button. Selecting OK generates a component instantiation declaration, a behavioral model, and an implementation netlist.
 7. Copy the HDL module declaration/instantiation into your HDL design. The template file created by LogiBLOX is *module_name.vhi* for VHDL and *module_name.vei* for Verilog, and is saved in the project directory as specified in the LogiBLOX setup.
 8. Complete the signal connections of the instantiated LogiBLOX memory module to the rest of your HDL design, and complete initial design coding.
 9. Perform a behavioral simulation on your design. For more information on behavioral simulation, refer to the “Simulating Your Design” chapter.
 10. Create an implementation script. Add a Set Don’t Touch or equivalent attribute to the instantiated LogiBLOX memory module, and compile your design.

Also, if you have a Verilog design, use a remove design type of command before writing the .xnf or .edif netlist.

Note: If you do not use this type of command, the netlist file may be empty. If this occurs, the Xilinx software will trim this module/component and all connected logic. Refer to your synthesis tool documentation for the correct syntax.

11. Compile your design and create a .xnf or .edif file. You can safely ignore the following type of warning messages.

Warning: Can't find the design in the library *WORK*. (LBR-1)

Warning: Unable to resolve reference *LogiBLOX_name* in *design_name*. (LINK-5)

Warning: Design *design_name* has 1 unresolved references. For more detailed information, use the "link" command. (UID-341)

12. Implement your design with the Xilinx tools. Verify that the .ngo file created by LogiBLOX is in the same project directory as the netlist.

You may get the following warnings during the NGDBuild and mapping steps. These messages are issued if the Xilinx software can not locate the corresponding .ngo file created by LogiBLOX.

Warning: basnu - logical block *LogiBLOX_instance_name* of type *LogiBLOX_name* is unexpanded. Logical Design DRC complete with 1 warning(s).

If you get this message, you will get the following message during mapping.

ERROR:basnu - logical block *LogiBLOX_instance_name* of type *LogiBLOX_name* is unexpanded. Errors detected in general drc.

If you get these messages, first verify that the .ngo file created by LogiBLOX is in the project directory. If the file is there, verify that the module is properly instantiated in the code.

13. To simulate your post-layout design, convert your design to a timing netlist and use the back-annotation flow applicable to your synthesis tool.

Note: For more information on simulation, refer to the "Simulating Your Design" chapter.

Implementing Boundary Scan (JTAG 1149.1)

Note: Refer to the *Development System User Guide* for a detailed description of the XC4000/XC5200 boundary scan capabilities.

XC4000, Spartan, and XC5200 FPGAs contain boundary scan facilities that are compatible with IEEE Standard 1149.1. Xilinx devices support external (I/O and interconnect) testing and have limited support for internal self-test.

You can access the built-in boundary scan logic between power-up and the start of configuration. Optionally, the built-in logic is available after configuration if you specify boundary scan in your design. During configuration, a reduced boundary scan capability (sample/preload and bypass instructions) is available.

In a configured FPGA device, the boundary scan logic is enabled or disabled by a specific set of bits in the configuration bitstream. To access the boundary scan logic after configuration in HDL designs, you must instantiate the boundary scan symbol, BSCAN, and the boundary scan I/O pins, TDI, TMS, TCK, and TDO.

The XC5200 BSCAN symbol contains three additional pins: RESET, UPDATE, and SHIFT, which are not available for XC4000 and Spartan. These pins represent the decoding of the corresponding state of the boundary scan internal state machine. If this function is not used, you can leave these pins unconnected in your HDL design.

Instantiating the Boundary Scan Symbol

To incorporate the boundary scan capability in a configured FPGA using synthesis tools, you must manually instantiate boundary scan library primitives at the source code level. These primitives include TDI, TMS, TCK, TDO, and BSCAN. The following VHDL and Verilog examples show how to instantiate the boundary scan symbol, BSCAN, into your HDL code. Note that the boundary scan I/O pins are not declared as ports in the HDL code. The schematic for this design is shown in the “Bnd_scan Schematic” figure.

You *must* assign a Set Don't Touch or equivalent attribute to the net connected to the TDO pad before using the Insert Pads (or equivalent) and compile commands. Otherwise, the TDO pad is removed by the compiler. In addition, you do not need IBUFs or OBUFs for the

TDI, TMS, TCK, and TDO pads. These special pads connect directly to the Xilinx boundary scan module.

Boundary Scan VHDL Example

```
library IEEE;
use IEEE.std_logic_1164.all;
use IEEE.std_logic_unsigned.all;

entity bnd_scan is
    port (TDI_P, TMS_P, TCK_P : in STD_LOGIC;
          LOAD_P, CE_P, CLOCK_P, RESET_P: in
            STD_LOGIC;
          DATA_P: in STD_LOGIC_VECTOR(3 downto 0);
          TDO_P: out STD_LOGIC;
          COUT_P: out STD_LOGIC_VECTOR(3 downto 0));
end bnd_scan;

architecture XILINX of bnd_scan is

    component BSCAN
        port (TDI, TMS, TCK :in STD_LOGIC;
              TDO: out STD_LOGIC);
    end component;

    component TDI
        port (I: in STD_LOGIC);
    end component;

    component TMS
        port (I: in STD_LOGIC);
    end component;

    component TCK
        port (I: in STD_LOGIC);
    end component;

    component TDO
        port (O: out STD_LOGIC);
    end component;

    component count4
        port (LOAD, CE, CLOCK, RST: in STD_LOGIC;
              DATA: in STD_LOGIC_VECTOR (3 downto 0);
              COUT: out STD_LOGIC_VECTOR (3 downto 0));
    end component;
```

```

end component;

-- Defining signals to connect BSCAN to Pins --
signal TCK_NET : STD_LOGIC;
signal TDI_NET : STD_LOGIC;
signal TMS_NET : STD_LOGIC;
signal TDO_NET : STD_LOGIC;

begin

    U1: BSCAN port map (TDO => TDO_NET,
                       TDI => TDI_NET,
                       TMS => TMS_NET,
                       TCK => TCK_NET);

    U2: TDI port map (I =>TDI_NET);

    U3: TCK port map (I =>TCK_NET);

    U4: TMS port map (I =>TMS_NET);

    U5: TDO port map (O =>TDO_NET);

    U6: count4 port map (LOAD => LOAD_P,
                       CE => CE_P,
                       CLOCK => CLOCK_P,
                       RST => RESET_P,
                       DATA => DATA_P,
                       COUT => COUT_P);

end XILINX;

```

Boundary Scan Verilog Example

```

////////////////////////////////////
// BND_SCAN.V //
// Example of instantiating the BSCAN symbol in //
// activating the Boundary Scan circuitry //
// Count4 is an instantiated .v file of a counter //
// September 1997 //
////////////////////////////////////

module bnd_scan (LOAD_P, CLOCK_P, CE_P, RESET_P,
                DATA_P, COUT_P);

    input          LOAD_P, CLOCK_P, CE_P, RESET_P;

```

```

input  [3:0] DATA_P;
output [3:0] COUT_P;

wire      TDI_NET, TMS_NET, TCK_NE, TDO_NET;

BSCAN U1 (.TDO(TDO_NET), .TDI(TDI_NET), .TMS(TMS_NET), .TCK(TCK_NET));

TDI U2 (.I(TDI_NET));

TCK U3 (.I(TCK_NET));

TMS U4 (.I(TMS_NET));

TDO U5 (.O(TDO_NET));

count4 U6 (.LOAD(LOAD_P), .CLOCK(CLOCK_P), .CE(CE_P),
            .RST(RESET_P), .DATA(DATA_P), .COUT(COUT_P));

endmodule

```

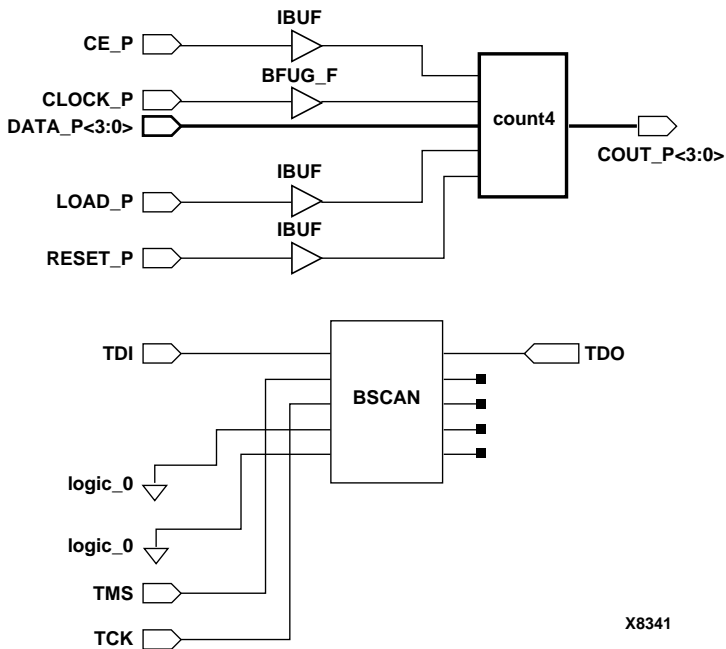


Figure 4-7 Bnd_scan Schematic

Implementing Logic with IOBs

You can move logic that is normally implemented with CLBs to IOBs. By moving logic from CLBs to IOBs, additional logic can be implemented in the available CLBs. Using IOBs also improves design performance by increasing the number of available routing resources.

The XC4000 and Spartan devices have different IOB functions. The following sections provide a general description of the IOB function in XC4000E/EX/XL/XV and Spartan devices. A description of how to manually implement additional I/O features is also provided.

XC4000E/EX/XL/XV and Spartan IOBs

You can configure XC4000E/EX/XL/XV and Spartan IOBs as input, output, or bidirectional signals. You can also specify pull-up or pull-down resistors, independent of the pin usage.

These various buffer and I/O structures can be inferred from commands executed in a script or in your synthesis tool. The Set Port Is Pad (or equivalent) command in conjunction with the Insert Pads (or equivalent) command creates the appropriate buffer structure according to the direction of the specified port in the HDL code. You can add attributes to these commands to further control pull-up, pull-down, and clock buffer insertion, as well as slew-rate control. Some tools operate on I/Os by selecting a chip level (inserts I/O) or module level (no I/O) synthesis. Also, you can add synthesis tool attributes, such as BUFFER_SIG, to ports in your VHDL code to control insertion of I/Os.

Inputs

The buffered input signal that drives the data input of a storage element can be configured as either a flip-flop or a latch. Additionally, the buffered signal can be used in conjunction with the input flip-flop or latch, or without the register.

To avoid external hold-time requirements, IOB input flip-flops and latches have a delay block between the external pin and the D input. You can remove this default delay by instantiating a flip-flop or latch with a NODELAY attribute. The NODELAY attribute decreases the setup-time requirement and introduces a small hold time.

If an IOB or register is instantiated in your HDL code, you may not be able to use the Set Port Is Pad (or equivalent) command on that port. Doing so may automatically infer a buffer on that port and create an invalid double-buffer structure. This varies with the tool you are using. Check with your synthesis vendor to see if partial instantiation interferes with automatic I/O insertion or the use of IOB registers.

Registers that connect to an input or output pad and require a Direct Clear or Preset pin are not implemented by the synthesis tool in the IOB. The VHDL emulation of GSR or GR on these registers prevents them from being pulled into the IOB. The VHDL emulation of GSR/GR through direct clear or preset pins is described in the “Simulating Your Design” chapter. If GSR/GR behavior is not completely described, automatic inferencing of GSR/GR does not occur. In this case, instantiate STARTBUF in VHDL, and fully describe the GSR/GR behavior except for registers that you want in the IOB. In VHDL, these registers do not initialize pre-route, but do indicate X's until the first data is registered. However, they do initialize properly during back-annotation. Verilog models initialize properly and do not interfere with the automatic use of IOB registers instead of CLB registers.

Outputs

The output signal that drives the programmable tristate output buffer can be a registered or a direct output. The register is a positive-edge triggered flip-flop and the clock polarity can be inverted inside the IOB. (Xilinx software automatically optimizes any inverters into the IOB.) The XC4000 and Spartan output buffers can sink 12 mA. Two adjacent outputs can be inter-connected externally to sink up to 24mA.

Note: Most FPGA synthesis tools can optimize flip-flops attached to output pads into the IOB. However, some of these tools cannot optimize flip-flops into an IOB configured as a bidirectional pad. Refer to your synthesis tool documentation for more information.

Slew Rate

Refer to your synthesis tool documentation for information on configuring I/O's, including how to control slew rate.

Pull-ups and Pull-downs

XC4000 and Spartan devices have programmable pull-up and pull-down resistors available in the I/O regardless of whether it is configured as an input, output, or bi-directional I/O. By default, all unused IOBs are configured as an input with a pull-up resistor. The value of the pull-ups and pull-downs vary depending on operating conditions and device process variances but should be approximately 50 K Ohms to 100 K Ohms. If a more precise value is required, use an external resistor. Refer to your synthesis tool documentation for information on how to specify internal pull-up or pull-down I/O resistors.

XC4000EX/XL/XV Output Multiplexer/2-Input Function Generator

A function added to XC4000EX/XL/XV families is a two input multiplexer connected to the IOB output allowing the output clock to select either the output data or the IOB clock enable as the output pad. This allows you to share output pins between two signals, effectively doubling the number of device outputs without requiring a larger device or package. Additionally, this multiplexer can be configured as a two-input function generator allowing you to implement any 2-input logic function in the IOB thus freeing up additional logic resources in the device and allowing for very fast pin-to-pin paths.

To use the output multiplexer (OMUX), you must instantiate it in your code. See the following VHDL and Verilog examples. Instantiation of the other types of two-input output primitives (such as OAND2, OOR2, and OXOR2) are similar to these examples.

Note: Since the OMUX uses the IOB output clock and clock enable routing structures, the output flip-flop (OFD) can not be used within the same IOB. The input flip-flop (IFD) can be used if the clock enable is not used.

- Output Multiplexer VHDL Example

```
-----
-- OMUX_EXAMPLE.VHD                      --
-- Example of OMUX instantiation          --
-- For an XC4000EX/XL/XV device          --
-- HDL Synthesis Design Guide for FPGAs --
-- August 1997                            --
-----
```

```
library IEEE;
use IEEE.std_logic_1164.all;

entity omux_example is

    port (DATA_IN: in STD_LOGIC_VECTOR (1 downto 0);
          SEL: in STD_LOGIC;
          DATA_OUT: out STD_LOGIC);

end omux_example;

architecture XILINX of omux_example is

    component OMUX2
        port (D0, D1, S0 : in STD_LOGIC;
              O : out STD_LOGIC);
    end component;

begin

    DUEL_OUT: OMUX2 port map (O=>DATA_OUT,
                              D0=>DATA_IN(0), D1=>DATA_IN(1), S0=>SEL);

end XILINX;
```

- **Output Multiplexer Verilog Example**

```
////////////////////////////////////
// OMUX_EXAMPLE.V //
// Example of instantiating an OMUX2 //
// in an XC4000EX/XL IOB //
// HDL Synthesis Design Guide for FPGAs //
// August 1997 //
////////////////////////////////////

module omux_example (DATA_IN, SEL, DATA_OUT) ;

input [1:0] DATA_IN ;
input SEL ;
output DATA_OUT ;

OMUX2 DUEL_OUT (.O(DATA_OUT), .D0(DATA_IN[0]),
                .D1(DATA_IN[1]), .S0(SEL));

endmodule
```

XC5200 IOBs

XC5200 IOBs consist of an input buffer and an output buffer that can be configured as an input, output, or bi-directional I/O. The structure of the XC5200 is similar to the XC4000 IOB except the XC5200 does not contain a register/latch. The XC5200 IOB has a programmable pull-up or pull-down resistor, and two slew rate control modes (Fast and Slow) to minimize bus transients. The input buffer can be globally configured to TTL or CMOS levels, and the output buffer can sink or source 8.0 mA.

I/O buffer structures (as with the XC4000 IOBs) can be inferred from your synthesis tool script with the Set Port Is Pad (or equivalent) command in conjunction with the Insert Pads (or equivalent) command. Controlling pull-up and pull-down insertion and slew rate control are performed as previously described for the XC4000 IOB.

The XC5200 IOB also contains a delay element so that an input signal that is directly registered or latched can have a guaranteed zero hold time at the expense of a longer setup time. You can disable this (equivalent to NODELAY in XC4000) by instantiating an IBUF_F buffer for that input port. This only needs to be done for ports that connect directly to the D input of a register in which a hold time can be tolerated.

Bi-directional I/O

You can create bi-directional I/O with one or a combination of the following methods.

- Behaviorally describe the I/O path
- Structurally instantiate appropriate IOB primitives
- Create the I/O using LogiBLOX

Xilinx FPGA IOBs consist of a direct input path into the FPGA through an input buffer (IBUF) and an output path to the FPGA pad through a tri-stated buffer (OBUFT). The input path can be registered or latched; the output path can be registered. If you instantiate or behaviorally describe the I/O, you must describe this bi-directional path in two steps. First, describe an input path from the declared INOUT port to a logic function or register. Second, describe an output path from an internal signal or function in your code to a tri-stated

output with a tri-state control signal that can be mapped to an OBUFT.

You should always describe the I/O path at the top level of your code. If the I/O path is described in a lower level module, your synthesis tool may incorrectly create the I/O structure.

Inferring Bi-directional I/O

This section includes VHDL and Verilog examples that show how to infer a bi-directional I/O. In these examples, the input path is latched by a CLB latch that is gated by the active high READ_WRITE signal.

The output consists of two latched outputs with an AND and OR, and connected to a described tri-state buffer. The active low READ_WRITE signal enables the tri-state gate.

- **Inferring a Bi-directional Pin VHDL Example**

```
-----  
-- BIDIR_INFER.VHD                                     --  
-- Example of inferring a Bi-directional pin          --  
-- August 1997                                       --  
-----  
  
Library IEEE;  
use IEEE.STD_LOGIC_1164.all;  
use IEEE.STD_LOGIC_UNSIGNED.all;  
  
entity bidir_infer is  
  
    port (DATA :          inout STD_LOGIC_VECTOR(1 downto 0);  
          READ_WRITE : in   STD_LOGIC);  
  
end bidir_infer;  
  
architecture XILINX of bidir_infer is  
  
    signal LATCH_OUT : STD_LOGIC_VECTOR(1 downto 0);  
  
begin  
  
    process(READ_WRITE, DATA)  
    begin  
  
        if (READ_WRITE = '1') then
```

```

        LATCH_OUT <= DATA;
    end if;

end process;

process(READ_WRITE, LATCH_OUT)
    begin

        if (READ_WRITE = '0') then
            DATA(0) <= LATCH_OUT(0) and LATCH_OUT(1);
            DATA(1) <= LATCH_OUT(0) or LATCH_OUT(1);
        else
            DATA(0) <= 'Z';
            DATA(1) <= 'Z';
        end if;

    end process;

end XILINX;

```

- **Inferring a Bi-directional Pin Verilog Example**

```

//////////////////////////////////////////////////////////////////
// BIDIR_INFERR.V Version 1.1                                     //
// This is an example of an inference of a bi-directional signal. //
// Note: Logic description of port should always be on top-level //
//       code when using Synopsys Compiler and verilog.         //
// March 1998                                                    //
//////////////////////////////////////////////////////////////////

module bidir_infer (DATA, READ_WRITE);

    input      READ_WRITE ;
    inout [1:0] DATA ;

    reg  [1:0] LATCH_OUT ;

    always @ (READ_WRITE or DATA)
        begin
            if (READ_WRITE == 1'b1)
                LATCH_OUT <= DATA;
            end

    assign DATA[0] = READ_WRITE ? 1'bZ : (LATCH_OUT[0] & LATCH_OUT[1]);
    assign DATA[1] = READ_WRITE ? 1'bZ : (LATCH_OUT[0] | LATCH_OUT[1]);

```

endmodule

Instantiating Bi-directional I/O

Instantiating the bi-directional I/O gives you more control over the implementation of the circuit; however, as a result, your code is more architecture-specific and usually more verbose. The VHDL and Verilog examples in this section are identical to the examples in the “Inferring Bi-directional I/O” section; however, since there is more control over the implementation, an input latch is specified rather than the CLB latch inferred in the previous examples. The following examples are a more efficient implementation of the same circuit.

When instantiating I/O primitives, do not specify the Set Port Is Pad (or equivalent) command on the instantiated ports to prevent the I/O buffers from being inferred by your synthesis tool. This precaution also prevents the creation of an illegal structure.

- Instantiation of a Bi-directional Pin VHDL Example

```
-----  
-- BIDIR_INSTANTIATE.VHD      --  
-- Example of an instantiation --  
-- of a Bi-directional pin    --  
-- August 1997                 --  
-----  
  
Library IEEE;  
use IEEE.STD_LOGIC_1164.all;  
use IEEE.STD_LOGIC_UNSIGNED.all;  
  
entity bidir_instantiate is  
  
    port (DATA :          inout STD_LOGIC_VECTOR(1 downto 0);  
          READ_WRITE : in   STD_LOGIC);  
  
end bidir_instantiate;  
  
architecture XILINX of bidir_instantiate is  
  
    signal LATCH_OUT : STD_LOGIC_VECTOR(1 downto 0);  
    signal DATA_OUT : STD_LOGIC_VECTOR(1 downto 0);  
    signal GATE      : STD_LOGIC;  
  
    component ILD_1
```



```

        port (D, G : in  STD_LOGIC;
              Q   : out STD_LOGIC);
end component;

component OBUFT_S
    port (I, T : in  STD_LOGIC;
          O   :   out STD_LOGIC);
end component;

begin

DATA_OUT(0) <= LATCH_OUT(0) and LATCH_OUT(1);
DATA_OUT(1) <= LATCH_OUT(0) or  LATCH_OUT(1);

GATE <= not READ_WRITE;

INPUT_PATH_0 : ILD_1
    port map (D => DATA(0), G => GATE,
              Q => LATCH_OUT(0));

INPUT_PATH_1 : ILD_1
    port map (D => DATA(1), G => GATE,
              Q => LATCH_OUT(1));

OUPUT_PATH_0 : OBUFT_S
    port map (I => DATA_OUT(0), T => READ_WRITE,
              O => DATA(0));

OUPUT_PATH_1 : OBUFT_S
    port map (I => DATA_OUT(1), T => READ_WRITE,
              O => DATA(1));

end XILINX;

```

- **Instantiation of a Bi-directional Pin Verilog Example**

```

////////////////////////////////////
// BIDIR_INSTANTIATE.V                //
// This is an example of an instantiation //
// of a bi-directional port.           //
// August 1997                          //
////////////////////////////////////

module bidir_instantiate (DATA, READ_WRITE);

    input      READ_WRITE ;

```

```
inout [1:0] DATA ;

reg    [1:0] LATCH_OUT ;
wire  [1:0] DATA_OUT ;
wire          GATE ;

assign GATE = ~READ_WRITE;

assign DATA_OUT[0] = LATCH_OUT[0] & LATCH_OUT[1];
assign DATA_OUT[1] = LATCH_OUT[0] | LATCH_OUT[1];

// I/O primitive instantiation

ILLD_1 INPUT_PATH_0 (.Q(LATCH_OUT[0]), .D(DATA[0]), .G(GATE));

ILLD_1 INPUT_PATH_1 (.Q(LATCH_OUT[1]), .D(DATA[1]), .G(GATE));

OBUFT_S OUPUT_PATH_0 (.O(DATA[0]), .I(DATA_OUT[0]), .T(READ_WRITE));

OBUFT_S OUPUT_PATH_1 (.O(DATA[1]), .I(DATA_OUT[1]), .T(READ_WRITE));

endmodule
```

Using LogiBLOX to Create Bi-directional I/O

You can use LogiBLOX to create I/O structures in an FPGA. LogiBLOX gives you the same control as instantiating I/O primitives, and is usually less verbose. LogiBLOX is especially useful for bused I/O ports.

Note: Refer to the “Using LogiBLOX in HDL Designs” section section, for details on creating, instantiating, and compiling LogiBLOX modules.

Do not use the Set Port Is Pad (or equivalent) command on LogiBLOX-created ports. Also, when designing with Verilog, you must issue a Remove Design or equivalent command before writing out the .xnf files from your synthesis tool.

The following VHDL and Verilog examples show how to instantiate bi-directional I/O created with LogiBLOX. These examples produce the same results as the examples in the “Instantiating Bi-directional I/O” section.

- Using LogiBLOX to Create a Bi-directional Port VHDL Example

```

-----
-- BIDIR_LOGIBLOX.VHD          --
-- Example of using LogiBLOX   --
-- to create a Bi-directional port --
-- August 1997                 --
-----

-- LogiBLOX BIDI Module "bidir_io_from_lb"
-- Created by LogiBLOX version M1.3.7
--   on Mon Sep  8 13:14:02 1997
-- Attributes
--   MODTYPE = BIDI
--   BUS_WIDTH = 2
--   IN_TYPE = LATCH
--   OUT_TYPE = TRI
-----

Library IEEE;
use IEEE.STD_LOGIC_1164.all;
use IEEE.STD_LOGIC_UNSIGNED.all;

entity bidir_logiblox is

    port (DATA :          inout STD_LOGIC_VECTOR(1 downto 0);
          READ_WRITE : in   STD_LOGIC);

end bidir_logiblox;

architecture XILINX of bidir_logiblox is

    signal LATCH_OUT : STD_LOGIC_VECTOR(1 downto 0);
    signal DATA_OUT : STD_LOGIC_VECTOR(1 downto 0);

    -----
    -- Component Declaration
    -----

    component bidir_io_from_lb
        PORT( O:      IN      STD_LOGIC_VECTOR(1 DOWNTO 0);
              OE:    IN      STD_LOGIC;
              IGATE: IN      STD_LOGIC;
              IQ:    OUT     STD_LOGIC_VECTOR(1 DOWNTO 0);
              P:     INOUT   STD_LOGIC_VECTOR(1 DOWNTO 0));
    end component;

```

```

begin

DATA_OUT(0) <= LATCH_OUT(0) and LATCH_OUT(1);
DATA_OUT(1) <= LATCH_OUT(0) or LATCH_OUT(1);

-----
-- Component Instantiation
-----
BIDIR_BUSSED_PORT : bidir_io_from_lb
    port map (O => DATA_OUT, OE => READ_WRITE,
              IGATE => READ_WRITE, IQ => LATCH_OUT, P => DATA);

end XILINX;

```

- Using LogiBLOX to Create a Bi-directional Port Verilog Example

```

////////////////////////////////////
// BIDIR_LOGIBLOX.V //
// This is an example of using LogiBLOX //
// to create a bi-directional port. //
// August 1997 //
////////////////////////////////////

//-----
// LogiBLOX BIDI Module "bidir_io_from_lb"
// Created by LogiBLOX version M1.3.7
// on Mon Sep 8 17:10:15 1997
// Attributes
// MODTYPE = BIDI
// BUS_WIDTH = 2
// IN_TYPE = LATCH
// OUT_TYPE = TRI
//-----

module bidir_logiblox (DATA, READ_WRITE);

    input READ_WRITE ;
    inout [1:0] DATA ;

    reg [1:0] LATCH_OUT ;
    wire [1:0] DATA_OUT ;

    assign DATA_OUT[0] = LATCH_OUT[0] & LATCH_OUT[1];
    assign DATA_OUT[1] = LATCH_OUT[0] | LATCH_OUT[1];

```

```
// LogiBLOX instantiation

bidir_io_from_lb BIDIR_BUSSED_PORT
    (.O(DATA_OUT), .OE(READ_WRITE), .P(DATA),
     .IQ(LATCH_OUT), .IGATE(READ_WRITE));

endmodule

module bidir_io_from_lb (O, OE, P, IQ, IGATE);
    input  [1:0] O;
    input          OE;
    input          IGATE;
    inout  [1:0] P;
    output [1:0] IQ;
endmodule
```

Specifying Pad Locations


Although Xilinx recommends allowing the software to select pin locations to ensure the best possible pin placement in terms of design timing and routing resources, sometimes you must define the pad locations prior to placement and routing. You can assign pad locations either from your synthesis tool's script prior to writing out the netlist file, or from a User Constraints File (UCF). Use one or the other method, but not both. Refer to your synthesis tool documentation for the correct syntax for configuring your I/O with the PLOC property. Also, refer to *The Programmable Logic Data Book* or the Xilinx Web site for the pad locations for your device and package.

Moving Registers into the IOB


Note: XC5200 devices do not have input and output flip-flops.

IOBs contain an input register or latch and an output register. IOB inputs can be register or latch inputs as well as direct inputs to the device array. Registers without a direct reset or set function can be moved into IOBs. Moving registers or latches into IOBs may reduce the number of CLBs used and decreases the routing congestion. In addition, moving input registers and latches into the IOB reduces the external setup time, as shown in the following figure.

Input Register



Output Register



X4974

Figure 4-8 Moving Registers into the IOB

Although moving output registers into the IOB may increase the internal setup time, it may reduce the clock-to-output delay, as shown in this figure. Most FPGA synthesis tools automatically move registers into IOBs if the Preset, Clear, and Clock Enable pins are *not* used.

Use `-pr` Option with Map

Use the `-pr` (pack registers) option when running MAP. The `-pr {i | o | b}` (input | output | both) option specifies to the MAP program to move registers into IOBs under the following circumstances.

1. The input of the register must be connected to an input port, or the Q pin must be connected to an output port. For the

XC4000EX/XL/XV this applies to non-I/O latches, as well as flip-flops.

2. IOBs must have input or output flip-flops. XC5200 devices do not have IOB flip-flops.
3. The flip-flop does not use an asynchronous set or reset signal.
4. In XC4000, Spartan, and XC3000 devices, a flop/latch is not added to an IOB if it has a BLKNM or LOC conflict with the IOB.
5. In XC4000 or Spartan devices, a flop/latch is not added to an IOB if its control signals (clock or clock enable) are not compatible with those already defined in the IOB. This occurs when a flip-flop (latch) is already in the IOB with different clock or clock enable signals, or when the XC4000EX/XL/XV output MUX is used in the same IOB.
6. In XC4000EX/XV devices, if a constant 0 or 1 is driven on the IOPAD, a flip-flop/latch with a CE is not added to the input side of the IOB.

Using Unbonded IOBs (XC4000E/EX/XL/XV and Spartan Only)

In some package/device pairs, not all pads are bonded to a package pin. You can use these unbonded IOBs and the flip-flops inside them in your design by instantiating them in the HDL code. You can implement shift registers with these unbonded IOBs. The VHDL and Verilog examples in this section show how to instantiate unbonded IOB flip-flops in a 4-bit shift register in an XC4000 device.

Note: The synthesis tool compilers cannot infer unbonded primitives. Refer to your synthesis tool documentation for a list of library primitives that can be used for instantiations.

4-bit Shift Register Using Unbonded I/O VHDL Example

```
-----
-- UNBONDED_IO.VHD Version 1.0           --
-- XC4000 LCA has unbonded IOBs which have --
-- storage elements that can be used to build --
-- shift registers.                       --
-- Below is a 4-bit Shift Register using  --
```

```
-- Unbonded IOB Flip Flops --
-- Xilinx HDL Synthesis Design Guide for FPGAs --
-- May 1997 --
-----

Library IEEE;
use IEEE.std_logic_1164.all;
use IEEE.std_logic_unsigned.all;

entity unbonded_io is
    port (A, B: in STD_LOGIC;
          CLK: in STD_LOGIC;
          Q_OUT: out STD_LOGIC);
end unbonded_io;

architecture XILINX of unbonded_io is

    component IFD_U -- Unbonded Input FF with INIT=Reset
        port (Q: out std_logic;
              D, C: in std_logic);
    end component;

    component IFDI_U -- Unbonded Input FF with INIT=Set
        port (Q: out std_logic;
              D, C: in std_logic);
    end component;

    component OFD_U -- Unbonded Output FF with INIT=Reset
        port (Q: out std_logic;
              D, C: in std_logic);
    end component;

    component OFDI_U -- Unbonded Output FF with INIT=Set
        port (Q: out std_logic;
              D, C: in std_logic);
    end component;

    --- Internal Signal Declarations ----
    signal U_Q : STD_LOGIC_VECTOR (3 downto 0);
    signal U_D : STD_LOGIC;

begin
    U_D <= A and B;
    Q_OUT <= U_Q(0);
end;
```



```

U3: OFD_U port map (Q => U_Q(3),
                   D => U_D,
                   C => CLK);

U2: IFDI_U port map (Q => U_Q(2),
                   D => U_Q(3),
                   C => CLK);

U1: OFDI_U port map (Q => U_Q(1),
                   D => U_Q(2),
                   C => CLK);

U0: IFD_U port map (Q => U_Q(0),
                   D => U_Q(1),
                   C => CLK);

end XILINX;

```

4-bit Shift Register Using Unbonded I/O Verilog Example

```

//////////////////////////////////////////////////////////////////
// UNBONDED.V //
// XC4000 family has unbonded IOBs which have //
// storage elements that can be used to build //
// functions like shift registers. //
// Below is a 4-bit Shift Register using Unbonded //
// IOB Flip Flops //
// HDL Synthesis Design Guide for FPGAs //
// May 1997 //
//////////////////////////////////////////////////////////////////

module unbonded_io (A, B, CLK, Q_OUT);

input A, B, CLK;
output Q_OUT;

wire[3:0] U_Q;
wire U_D;

assign U_D = A & B;
assign Q_OUT = U_Q[0];

OFD_U U3 (.Q(U_Q[3]), .D(U_D), .C(CLK));

```

```
IFDI_U U2 (.Q(U_Q[2]), .D(U_Q[3]), .C(CLK));  
OFDI_U U1 (.Q(U_Q[1]), .D(U_Q[2]), .C(CLK));  
IFD_U U0 (.Q(U_Q[0]), .D(U_Q[1]), .C(CLK));  
  
endmodule
```

Implementing Multiplexers with Tristate Buffers

A 4-to-1 multiplexer is efficiently implemented in a single XC4000 or Spartan family CLB. The six input signals (four inputs, two select lines) use the F, G, and H function generators. Multiplexers that are larger than 4-to-1 exceed the capacity of one CLB. For example, a 16-to-1 multiplexer requires five CLBs and has two logic levels. These additional CLBs increase area and delay. Xilinx recommends that you use internal tristate buffers (BUFTs) to implement large multiplexers.

Large multiplexers built with BUFTs have the following advantages.

- Can vary in width with only minimal impact on area and delay
- Can have as many inputs as there are tristate buffers per horizontal longline in the target device
- Have one-hot encoded selector inputs

This last point is illustrated in the following VHDL and Verilog designs of a 5-to-1 multiplexer built with gates. Typically, the gate version of this multiplexer has binary encoded selector inputs and requires three select inputs (SEL<2:0>). The schematic representation of this design is shown in the “5-to-1 MUX Implemented with Gates” figure.

Some synthesis tools include commands that allow you to switch between multiplexers with gates or with tristates. Check with your synthesis vendor for more information.

The VHDL and Verilog designs provided at the end of this section show a 5-to-1 multiplexer built with tristate buffers. The tristate buffer version of this multiplexer has one-hot encoded selector inputs and requires five select inputs (SEL<4:0>). The schematic representation of these designs is shown in the “5-to-1 MUX Implemented with BUFTs” figure.

Mux Implemented with Gates VHDL Example

```
-- MUX_GATE.VHD
-- 5-to-1 Mux Implemented in Gates
-- May 1997

library IEEE;
use IEEE.std_logic_1164.all;
use IEEE.std_logic_arith.all;

entity mux_gate is
port (SEL: in STD_LOGIC_VECTOR (2 downto 0);
      A,B,C,D,E: in STD_LOGIC;
      SIG: out STD_LOGIC);
end mux_gate;

architecture RTL of mux_gate is
begin
  SEL_PROCESS: process (SEL,A,B,C,D,E)
  begin
    case SEL is
      when "000" => SIG <= A;
      when "001" => SIG <= B;
      when "010" => SIG <= C;
      when "011" => SIG <= D;
      when others => SIG <= E;
    end case;
  end process SEL_PROCESS;
end RTL;
```

Mux Implemented with Gates Verilog Example

```
/* MUX_GATE.V
 * May 1997 */

module mux_gate (A,B,C,D,E,SEL,SIG);

input A,B,C,D,E;
input [2:0] SEL;
output SIG;
reg SIG;


  always @ (A or B or C or D or SEL)
  case (SEL)
    3'b000:
```

```

        SIG=A;
    3'b001:
        SIG=B;
    3'b010:
        SIG=C;
    3'b011:
        SIG=D;
    3'b100:
        SIG=E;
default: SIG=A;
    endcase

endmodule

```



X6229

Figure 4-9 5-to-1 MUX Implemented with Gates

Mux Implemented with BUFTs VHDL Example

```

-- MUX_TBUF.VHD
-- 5-to-1 Mux Implemented in 3-State Buffers
-- May 1997

library IEEE;
use IEEE.std_logic_1164.all;
use IEEE.std_logic_arith.all;

entity mux_tbuf is
port (SEL: in STD_LOGIC_VECTOR (4 downto 0);
      A,B,C,D,E: in STD_LOGIC;
      SIG: out STD_LOGIC);
end mux_tbuf;

architecture RTL of mux_tbuf is
begin

    SIG <= A when (SEL(0)='0') else 'Z';

```

```
SIG <= B when (SEL(1)='0') else 'Z';
SIG <= C when (SEL(2)='0') else 'Z';
SIG <= D when (SEL(3)='0') else 'Z';
SIG <= E when (SEL(4)='0') else 'Z';

end RTL;
```

Mux Implemented with BUFTs Verilog Example

```
/* MUX_TBUF.V
 * May 1997 */

module mux_tbuf (A,B,C,D,E,SEL,SIG);

input A,B,C,D,E;
input [4:0] SEL;
output SIG;
reg SIG;

always @ (SEL or A)
begin
    if (SEL[0]==1'b0)
        SIG=A;
    else
        SIG=1'bz;
    end

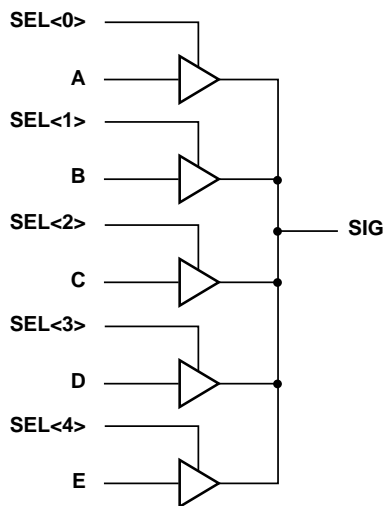
always @ (SEL or B)
begin
    if (SEL[1]==1'b0)
        SIG=B;
    else
        SIG=1'bz;
    end

always @ (SEL or C)
begin
    if (SEL[2]==1'b0)
        SIG=C;
    else
        SIG=1'bz;
    end

always @ (SEL or D)
begin
```

```
    if (SEL[3]==1'b0)
        SIG=D;
    else
        SIG=1'bz;
    end

    always @ (SEL or E)
    begin
        if (SEL[4]==1'b0)
            SIG=E;
        else
            SIG=1'bz;
        end
    end
endmodule
```



X6228

Figure 4-10 5-to-1 MUX Implemented with BUFTs

A comparison of timing and area for a 5-to-1 multiplexer built with gates and tristate buffers in an XC4005EPC84-2 device is provided in the following table. When the multiplexer is implemented with tristate buffers, no CLBs are used and the delay is smaller.

Table 4-5 Timing/Area for 5-to-1 MUX (XC4005EPC84-2)

| Timing/Area | Using BUFTs | Using Gates |
|-------------|--------------------------|---------------------------|
| Timing | 15.31 ns (1 block level) | 17.56 ns (2 block levels) |
| Area | 0 CLBs, 5 BUFTs | 3 CLBs |

Using Pipelining

You can use pipelining to dramatically improve device performance. Pipelining increases performance by restructuring long data paths with several levels of logic and breaking it up over multiple clock cycles. This method allows a faster clock cycle and, as a result, an increased data throughput at the expense of added data latency. Because the Xilinx FPGA devices are register-rich, this is usually an advantageous structure for FPGA designs because the pipeline is created at no cost in terms of device resources. Because data is now on a multi-cycle path, special considerations must be used for the rest of your design to account for the added path latency. You must also be careful when defining timing specifications for these paths.

Some synthesis tools have limited capability for constraining multi-cycle paths, or translate these constraints to Xilinx implementation constraints. Check your synthesis tool documentation for information on multi-cycle paths. If your tool cannot translate the constraint but can synthesize to a multi-cycle path, you can add the constraint to the UCF file.

Before Pipelining

In the following example, the clock speed is limited by the clock-to-out-time of the source flip-flop; the logic delay through four levels of logic; the routing associated with the four function generators; and the setup time of the destination register.




Figure 4-11 Before Pipelining

After Pipelining

This is an example of the same data path in the previous example after pipelining. Because the flip-flop is contained in the same CLB as the function generator, the clock speed is limited by the clock-to-out time of the source flip-flop; the logic delay through one level of logic; one routing delay; and the setup time of the destination register. In this example, the system clock runs faster than in the previous example.




Figure 4-12 After Pipelining

Design Hierarchy

HDL designs can either be synthesized as a flat module or as many small modules. Each methodology has its advantages and disadvantages, but as higher density FPGAs are created, the advantages of hierarchical designs outweigh any disadvantages.

Advantages to building hierarchical designs are as follows.

- Easier and faster verification/simulation
- Allows several engineers to work on one design at the same time

- Speeds up design compilation
- Reduces design time by allowing design module re-use for this and future designs.
- Allows you to produce designs that are easier to understand
- Allows you to efficiently manage the design flow

Disadvantages to building hierarchical designs are as follows.

- Design mapping into the FPGA may not be as optimal across hierarchical boundaries; this can cause lesser device utilization and decreased design performance
- Design file revision control becomes more difficult
- Designs become more verbose

Most of the disadvantages listed above can be overcome with careful design consideration when choosing the design hierarchy.

Using Synthesis Tools with Hierarchical Designs

By effectively partitioning your designs, you can significantly reduce compile time and improve synthesis results. This section provides recommendations for partitioning your designs.

Restrict Shared Resources to Same Hierarchy Level

Resources that can be shared should be on the same level of hierarchy. If these resources are not on the same level of hierarchy, the synthesis tool cannot determine if these resources should be shared.

Compile Multiple Instances Together

You may want to compile multiple occurrences of the same instance together to reduce the gate count. However, to increase design speed, do not compile a module in a critical path with other instances.

Restrict Related Combinatorial Logic to Same Hierarchy Level

Keep related combinatorial logic in the same hierarchical level to allow the synthesis tool to optimize an entire critical path in a single operation. Boolean optimization does not operate across hierarchical boundaries. Therefore, if a critical path is partitioned across boundaries, logic optimization is restricted. In addition, constraining modules is difficult if combinatorial logic is not restricted to the same level of hierarchy.

Separate Speed Critical Paths from Non-critical Paths

To achieve satisfactory synthesis results, locate design modules with different functions at different levels of the hierarchy. Design speed is the first priority of optimization algorithms. To achieve a design that efficiently utilizes device area, remove timing constraints from design modules.

Restrict Combinatorial Logic that Drives a Register to Same Hierarchy Level

To reduce the number of CLBs used, restrict combinatorial logic that drives a register to the same hierarchical block.

Restrict Module Size

Restrict module size to 100 - 200 CLBs. This range varies based on your computer configuration; the time required to complete each optimization run; if the design is worked on by a design team; and the target FPGA routing resources. Although smaller blocks give you more control, you may not always obtain the most efficient design. For the final compilation of your design, you may want to compile fully from the top down. Check with your synthesis vendor for guidelines.

Register All Outputs

Arrange your design hierarchy so that registers drive the module output in each hierarchical block. Registering outputs makes your design easier to constrain because you only need to constrain the clock period and the ClockToSetup of the previous module. If you have multiple combinatorial blocks at different levels of the hierarchy, you must manually calculate the delay for each module. Also, registering the outputs of your design hierarchy can eliminate any possible problems with logic optimization across hierarchical boundaries.

Restrict One Clock to Each Module or to Entire Design

By restricting one clock to each module, you only need to describe the relationship between the clock at the top level of the design hierarchy and each module clock. By restricting one clock to the entire design, you only need to describe the clock at the top level of the design hierarchy.

Note: See your synthesis tool documentation for more information on optimizing logic across hierarchical boundaries and compiling hierarchical designs.

Simulating Your Design

Xilinx supports functional and timing simulation of HDL designs at the following three points in the HDL design flow.

- Register Transfer Level (RTL) simulation; may include the following.
 - Instantiated UNISIM library components
 - LogiBLOX modules
 - LogiCORE models
- Post-synthesis functional simulation with one of the following.
 - Gate-level UNISIM library componentsor
 - Gate-level pre-route SIMPRIM library components
- Post-implementation back-annotated timing simulation with the following.
 - SIMPRIM library components
 - Standard Delay Format (SDF) file

This chapter includes the following sections.

- “Functional Simulation”
- “Timing Simulation”
- “Using VHDL/Verilog Libraries and Models”
- “Simulating Global Signals”
- “Adapting Schematic Global Signal Methodology for VHDL”

- “Setting VHDL Global Set/Reset Emulation in Functional Simulation”
- “Using Oscillators (VHDL)”
- “Compiling Verilog Libraries”
- “Verilog Simulation”
- “Setting Verilog Global Set/Reset”
- “Setting Verilog Global Tristate (XC4000, Spartan, and XC5200 Outputs Only)”
- “Using Verilog Oscillators (OSC, OSC4, OSC5)”

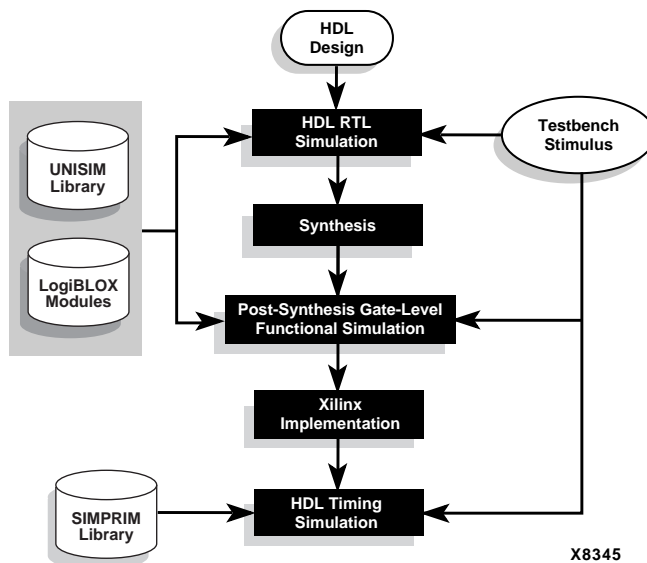


Figure 5-1 Three Simulation Points for HDL Designs

The three primary simulation points can be expanded to allow for two additional post-synthesis simulations, as shown in the following table. These two additional points can be used when the synthesis tool either cannot write VHDL or Verilog, or if the netlist is not in terms of UNISIM components.

Table 5-1 Five Simulation Points in HDL Design Flow

| Simulation | | UNISIM | LogiBLOX Models | SIMPRIM | SDF |
|------------|-------------------------------------|--------|-----------------|---------|-----|
| 1. | RTL | X | X | | |
| 2. | Post-Synthesis | X | X | | |
| 3. | Functional Post-NGDBuild (Optional) | | | X | |
| 4. | Functional Post-MAP (Optional) | | | X | X |
| 5. | Post-Route Timing | | | X | X |

These simulation points are described in detail in the “Functional Simulation” section and the “Timing Simulation” section. The libraries required to support the simulation flows are described in detail in the “Using VHDL/Verilog Libraries and Models” section. The new flows and libraries now support closer functional equivalence of initialization behavior between functional and timing simulations. This is due to the addition of new methodologies and library cells to simulate GSR/GTS behavior.

It is important to address the built-in reset circuitry behavior in your designs starting with the first simulation to ensure that the simulations agree at the three primary points.

If you do not simulate GSR behavior prior to synthesis and place and route, your RTL and possibly post-synthesis simulations will not initialize to the same state as your post-route timing simulation. As a result, your various design descriptions are not functionally equivalent and your simulation results will not match. In addition to the behavioral representation for GSR, you need to add a Xilinx implementation directive. This directive is used to specify to the place and route tools to use the special purpose GSR net that is pre-routed on the chip, and not to use the local asynchronous set/reset pins. Some synthesis tools can identify, from the behavioral description, the GSR net, and will place the STARTUP module on the net to direct the implementation tools to use the global network. However, other synthesis tools interpret behavioral descriptions literally, and will introduce additional logic into your design to implement a function. Without specific instructions to use device global networks, the Xilinx

implementation tools will use general purpose logic and interconnect resources to redundantly build functions already provided by the silicon.

Even if GSR behavior is not described, the actual chip initializes during configuration, and the post-route netlist will have this net that must be driven during simulation. The “Simulating Global Signals” section includes the methodology to describe this behavior, as well as the GTS behavior for output buffers.

Xilinx VHDL simulation supports the VITAL standard. This standard allows you to simulate with any VITAL-compliant simulator, including MTI/Mentor® ModelSim, Synopsys VSS, Active-VHDL, and Cadence.

Built-in Verilog support allows you to simulate with the Cadence Verilog-XL and other compatible simulators. Xilinx HDL simulation supports all current Xilinx FPGA and CPLD devices. Refer to the “Using VHDL/Verilog Libraries and Models” section for the list of supported VHDL and Verilog standards.

Functional Simulation

Functional simulation of HDL designs includes support for FPGA and CPLD architecture features, such as global set/reset, global tristate enable, on-chip oscillators, RAMs, and ROMs.

You can perform functional simulation at the following points in the design flow.

- RTL simulation, including instantiated UNISIM library components, LogiBLOX modules, and LogiCORE models
- Post-synthesis gate-level VHDL or Verilog netlist
- Post-NGDBuild gate-level VHDL or Verilog netlist from implementation after NGDBuild using the SIMPRIM library
- Post-map partial-timing functional simulation with netlist and SDF file from implementation after mapping, and before place and route, using the SIMPRIM library

Register Transfer Level (RTL)

The first simulation point is the RTL-level simulation that allows you to verify or simulate a description at the system or chip level. At this

level, the system or chip is described with high-level RTL language constructs. VHDL and Verilog simulators exercise the design to check its functionality before it is implemented in gates. A test bench is created to model the environment of the system or chip. At this level, the Unified Simulation Library (UNISIM) is used to instantiate cells. You can also instantiate LogiBLOX components if you do not want to rely on the module generation capabilities of your synthesis tool, or if your design requires larger memory structures.

Post-Synthesis (Pre-NGDBuild) Gate-Level Simulation

After the RTL simulation is error-free, the system or chip is synthesized to gates. The test bench is used again to simulate the synthesized result and check its consistency with the original design description.

Gate-level simulation in the Xilinx design flow includes any simulation performed after any of the synthesis, map, or place and route stages. The post-synthesis pre-NGDBuild gate-level simulation is a functional simulation with unit delay timing. The gates are expressed in terms of UNISIM components.

This gate-level functional simulation allows you to directly verify your design after it has been generated by the synthesis tool. If there are differences in the behavior of the original RTL description and the synthesized design, this may indicate a problem with the synthesis tool. Although RTL-level and gate-level simulation may differ even when the synthesis is correct. This is due to the change in the level of abstraction. In general, overall functionality should be the same, but timing differences may occur.

Post-synthesis simulation is synthesis vendor-dependent, and the synthesis tool must write VHDL or Verilog netlists in terms of UNISIM library components. Check with your synthesis vendor for this feature. The library usage guidelines for RTL simulation also apply to post-synthesis pre-NGDBuild gate-level functional simulation. LogiBLOX models remain as behavioral blocks and can be simulated in the same design as structural UNISIM components. You may have to insert library statements into the HDL code.

Post-NGDBuild (Pre-Map) Gate-Level Simulation

The post-NGDBuild (pre-map) gate-level functional simulation is used when it is not possible to simulate the direct output of the synthesis tool. This occurs when the tool cannot write UNISIM-compatible VHDL or Verilog netlists. In this case, NGDBuild translates the EDIF or XNF output of synthesis to SIMPRIM library components. Like post-synthesis, pre-NGDBuild simulation, this simulation allows you to verify that your design has been synthesized correctly, and you can begin to identify any differences due to the lower level of abstraction. Unlike the post-synthesis pre-NGDBuild simulation, there are GSR, GR (global reset), PRLD (preload), and GTS nets that must be initialized, just as for post-map partial timing simulation.

Different simulation libraries are used to support simulation before and after running NGDBuild. Prior to NGDBuild, your design is expressed as a UNISIM netlist containing Unified Library components. After NGDBuild, your design is a netlist containing SIMPRIMs. Although these library changes are fairly transparent, two important considerations are specifying different simulation libraries for pre- and post-implementation simulation, and the different gate-level cells in pre- and post-implementation netlists.

You can pause the implementation tools after reading in your source files; select either a VHDL or Verilog netlist as an output file; and write the file.

Post-Map Partial Timing (CLB and IOB Block Delays)

The third type of post-synthesis functional simulation is post-map partial timing simulation. This gate-level description is also expressed in terms of SIMPRIM components.

The SDF file that is created contains timing numbers for the CLB and IOB mapped blocks. At this point, your design is not routed and does not contain net delays, except for pre-laid out macros, such as cores. Like the post-NGDBuild and full-timing simulation, there are GSR, GR, PRLD, and GTS nets that must be initialized. If you want partial timing delays, you can use the SDF file (this is optional).

You can pause the implementation tools after the Map program stops, and write one of the HDL formats.

Creating a Simulation Test Bench

When you create a test bench for functional simulation, refer to the following table for coding style examples. You should refer to your synthesis vendor's documentation before using these styles because many of the styles cannot be synthesized.

Table 5-2 Common Coding Examples for Verilog/VHDL

| Description | Verilog | VHDL |
|--|---|---|
| Delay or wait 20 nS | <code>'timescale 1 ns/ 100 ps #20</code> | <code>wait for 20 ns;</code> |
| Creation of a free running clock | <code>Initial begin clock = 0; #25 forever #25 clock = ~clock; end</code> | <code>Loop wait for 25 ns; clock <=not (clock); end loop;</code> |
| Print "Text." to screen | <code>\$display("Text");</code> | <code>report "Text."</code> |
| Print value of signal to screen whenever the value changes | <code>\$monitor("%r", \$real time, "%b", clock, "%b" my_signal);</code> | |
| Apply a binary value 1010 to an input bus X. | <code>X = 4'b1010;</code> | <code>X <="1010";</code> |
| Creation of a for loop 0 to 10 | <code>for(x=0; x < 10; x=x+1) begin actions end</code> | <code>for x in 0 to 9 loop actions end loop;</code> |
| Write "X = value" to an output file | <code>\$dumpfile ("file name.dmp"); \$dumpvars (X);</code> | <code>variable TEMP; write (TEMP, "X ="); write (TEMP, X); writeline (filename, TEMP);</code> |
| Wait until X is logic one | <code>wait (X == 1'b1);</code> | <code>wait (X == '1');</code> |
| Wait until X transitions to a logic one | <code>@(posedge X);</code> | <code>wait on X;</code> |

Table 5-2 Common Coding Examples for Verilog/VHDL

| Description | Verilog | VHDL |
|---------------------------------|--|--|
| If-Else construct | <pre>always @ (X) begin if (X = 1) Y = 1'b0 else Y = 1'b1; end</pre> | <pre>process (X) if (X = '1') then Y = '0'; else Y = '1'; end if; end process;</pre> |
| Case construct | <pre>always @(X or A) case (X) 2'b00 : Y = 1'b0); 2'b01 : Y = 1b1; default : Y = A endcase</pre> | <pre>process (X,A) begin case X is when "00" => Y = '0'; when "01" => Y = '1'; when others => Y = A; end case; end process;</pre> |
| Example instantiation of an OFD | <pre>OFD U1 (Q(D_OUT),,D(D_IN), .C(CLOCK));</pre> | <pre>U1: OFD port map (Q => D_OUT), D => D_IN, C => CLOCK);</pre> |

Timing Simulation

Timing simulation is important in verifying the operation of your circuit after the worst-case placed and routed delays are calculated for your design. In many cases, you can use the same test bench that you used for functional simulation to perform a more accurate simulation with less effort. You can compare the results from the two simulations to verify that your design is performing as initially specified. The Xilinx tools create a VHDL or Verilog simulation netlist of your placed and routed design, and provide libraries that work with many common HDL simulators.

Post-Route Full Timing (Block and Net Delays)

After your design is routed using PAR, it can be simulated with the actual block and net timing delays with the same test bench used in the behavioral simulation.

The back-annotation process (NGDAnno) produces a netlist of SIMPRIMs annotated with an SDF file with the appropriate block and net delay data from the place and route process. This netlist has GSR, GR, PRLD, and GTS nets that must be initialized. For more information, refer to the “Simulating Global Signals” section.

Creating a Timing Simulation Netlist

You can create a timing simulation netlist from the Design Manager or from the command line, as described in this section.

From the Design Manager

1. Select **Setup** → **Options** in the Flow Engine.
The Options dialog box appears.
2. Select the Produce Timing Simulation Data button in the Optional Targets field.
3. Select the Edit Template button next to the Implementation drop-down list in the Program Options Templates field.
The Implementation Template dialog box appears.
4. Select the Interface tab.
5. In the Simulation Data Options field, select applicable options as follows.
 - **Format**
Specify the netlist format to use for simulation. The format is usually VHDL or Verilog for synthesis designs.
 - **Correlate Simulation Data to Input Design**
This option enables signal back annotation to the original compiled netlist. By default, this option is on. However, you can turn it off to decrease run time, or if there are problems with the back-annotated simulation data. Port names are not changed by this option; however, internal node names may change.
 - **Bring Out Global Set/Reset Net as a Port**
This option creates an external port in the simulation netlist to allow control of the power-on reset from a port. This

option is not necessary for most simulators, and is off by default.

6. Click **OK** in the Implementation Template dialog box.
7. Click **OK** in the Options dialog box.
8. When you implement your design, the Flow Engine produces timing simulation data files.

Note: If you are using the Verilog-XL simulator, you may want to use the `-ul` switch for the NGD2VER program to automatically add the `uselib` directive to the simulation netlist to point to the location of the simulation libraries. Use the Template Manager to set this switch. Refer to the *Cadence Interface/Tutorial Guide* or <http://www.xilinx.com/techdocs/3167.htm> for more information.

From the Command Line

Note: To display the available options for the programs in this section, enter the program executable name at the prompt without any arguments. For complete descriptions of these options, refer to the *Development System Reference Guide*.

1. Run NGDAnno on your placed and routed `.ncd` file.

For back-annotated output (signals correlated to original netlist), enter the following.

```
ngdanno -p design.pcf design.ncd design.ngm
```

For output that is not back-annotated (faster run time), enter the following.

```
ngdanno design.ncd
```

2. Run the NGD2XXX program for the particular netlist you want to create.

For VHDL, enter the following.

```
ngd2vhd1 [options] design.nga
```

For Verilog, enter the following.

```
ngd2ver [options] design.nga
```

Using VHDL/Verilog Libraries and Models

The five simulation points listed previously require the UNISIM, SIMPRIM, XDW (Xilinx DesignWare), and LogiBLOX libraries. The first point, RTL simulation, is a behavioral description of your design at the register transfer level. RTL simulation is not architecture-specific unless your design contains instantiated UNISIM or LogiBLOX components. To support these instantiations, Xilinx provides a functional UNISIM library and a behavioral LogiBLOX library.

The second point, post-synthesis (pre-NGDDBuild) gate-level simulation uses the UNISIM and XDW libraries. The third, fourth, and fifth points (post-NGDDBuild, post-map, and post-route) use the SIMPRIM library. The following table indicates what library is required for each of the five simulation points.

Table 5-3 Simulation Phase Library Information

| Simulation Point | Compilation Order of Library Required |
|------------------|---|
| RTL | UNISIM (Device Dependent) LogiBLOX (VHDL only, Device Independent) |
| Post-Synthesis | UNISIM (Device Dependent) XDW (VHDL only) |
| Post-NGDDBuild | SIMPRIM (Device Independent) |
| Post-MAP | SIMPRIM (Device Independent) |
| Post-Route | SIMPRIM (Device Independent) |

Adhering to Industry Standards

The standards in the following table are supported by the Xilinx simulation flow.

Table 5-4 Standards Supported by Xilinx Simulation Flow

| Description | Version |
|-----------------------------|--------------------|
| VHDL Language | IEEE-STD-1076-87 |
| Verilog Language | IEEE-STD-1364-95 |
| VITAL Modeling Standard | IEEE-STD-1076.4-95 |
| Standard Delay Format (SDF) | 2.1 |
| Std_logic Data Type | IEEE-STD-1164-93 |

The UNISIM and SIMPRIM libraries adhere to IEEE-STDs. The VHDL library uses the VITAL IEEE-STD-1076.4 standard, and the Verilog library uses the IEEE-STD-1364 standard. By following these standards, Xilinx provides support for customers who use HDL tools from various vendors.

VHDL Initiative Towards ASIC Libraries (VITAL) was created to promote the standardization and interchangeability of VHDL libraries and simulators from various vendors. It also defines a standard for timing back-annotation to VHDL simulators.

Most simulator vendors have agreed to use the IEEE-STD 1076.4 VITAL standard for the acceleration of gate-level simulations. Check with your simulator vendor to confirm that this standard is being followed, and to verify proper settings and VHDL packages for this standard. The simulator may also accelerate IEEE-STD-1164, the standard logic package for types.

VITAL libraries include some overhead for timing checks and back-annotation styles. The UNISIM Library turns these timing checks off for unit delay functional simulation. The SIMPRIM back-annotation library keeps these checks on by default; however, you or your system administrator can turn them off. You must edit and re-compile the SIMPRIM components file after setting the generics.

Locating Library Source Files

The following table provides information on the location of the simulation library source files, as well as the order for a typical compilation.

Table 5-5 Simulation Library Source Files

| Library | Location of Source Files | | Required Libraries | |
|--|---|------------------------------------|--|---|
| | Verilog | VITAL VHDL | Verilog | VITAL VHDL |
| UNISIM 4K Family, Spartan (use UNI4000E for Spartan) | \$XILINX/ verilog/ src/ UNI4000E Or UNI4000X | \$XILINX/ vhdl/src/ unisims | Not required for Verilog-XL; see vendor documentation for other simulators | Required; typical compilation order: unisim_VCOMP.vhd unisim_VPKG.vhd unisim_VITAL.vhd <i>unisim_VCFG4K.vhd</i> (optional) |
| UNISIM 52K Family | \$XILINX/ verilog/ src/ UNI5200 | \$XILINX/ vhdl/src/ unisims | Not required for Verilog-XL; see vendor documentation for other simulators | Required; typical compilation order: unisim_VCOMP52K.vhd unisim_VITAL.vhd unisim_VITAL52K.vhd unisim_VCFG52K.vhd |
| LogiBLOX (Device Independent) | None; uses SIMPRIM library | \$XILINX/ vhdl/src/ logiblox | None; uses SIMPRIM library | Required; typical compilation order: mvlutil.vhd mvlarith.vhd logiblox.vhd |
| SIMPRIM (Device Independent) | \$XILINX/ verilog/ data | \$XILINX/ vhdl/src/ simprims | Not required for Verilog-XL; see vendor documentation for other simulators | Required; typical compilation order: simprim_Vcomponents.vhd simprim_Vpackage.vhd simprim_VITAL.vhd |

Case Sensitivity

Verilog is case sensitive, and is dictated by the targeted synthesis tool. Xilinx provides lower case Unified Library component libraries in its Cadence design kit. These libraries can also be used by other vendors requiring lower case simulation. Cadence Concept schematics require lower case libraries. Upper case versions are available as part of the UNISIM Verilog Library. Instantiations of components for Synopsys

FPGA Compiler synthesis require upper case libraries. Check with your synthesis vendor for the proper case to use with their libraries.

Using the UNISIM Library

The UNISIM Library, used for functional simulation only, contains default unit delays. This library includes all the Xilinx Unified Library components that are inferred by most popular synthesis tools. In addition, the UNISIM Library includes components that are commonly instantiated, such as IOs and memory cells. You should use your synthesis tool's module generators (such as LogiBLOX) to generate higher order functions, such as adders, counters, and large RAMs.

UNISIM Library Structure

The UNISIM library directory structure is different for VHDL and Verilog. There is only one VHDL library for all Xilinx technologies because the implementation differences between architectures are not important for unit delay functional simulation except in a few cases where functional differences occur.

For example, the decode8 in XC4000 devices has a pull-up, and in XC5200 devices, it does not. In these few cases, configuration statements are used to choose between architectures for the components. One library makes it easy to switch between technologies. It is left up to the user and the synthesis tool to use technology-appropriate cells. For Verilog, separate libraries are provided for each technology because Verilog does not have a configuration statement.

Schematic macros are not provided because most schematic vendors provide the lower-level netlist for importing into a synthesis tool. Some synthesis vendors have these macros in their libraries, and can expand them to gates. You can use the HDL output from synthesis to simulate these macros. You can also use a post-NGDBuild or post-Map netlist to simulate netlists with embedded schematic macros. This lower-level netlist for a schematic macro is also required for implementation. The VHDL models for Synopsys DesignWare components are in the Xilinx Synopsys Interface at \$XILINX/synopsys/libraries/sim/src/xdw. Because Verilog versions of the DesignWare components are not currently available, use post-NGDBuild functional simulation instead.

- VHDL UNISIM Library

The VHDL version of the UNISIM library is VITAL-compliant, and can be accelerated, however, it is not back-annotated with an SDF file. The files are in `$XILINX/vhdl/src/unisims`. The source file should be compiled into a library named UNISIM.

- Verilog UNISIM Library

The Verilog version of the UNISIM library may not have to be compiled, depending on the Verilog tool. Because there are a few cells with functional differences between Xilinx devices, a separate library is provided for each supported device. For example, decoders contain pull-ups in some devices and not in others. The libraries are in uppercase only. If needed, lowercase libraries are in the Xilinx Cadence Interface. The libraries are located at `$XILINX/verilog/src/technology`, where *technology* is UNI3000, UNI4000E, UNI4000X, UNI5000, or UNI9000.

When targeting a Spartan device, use the UNI4000E library. Spartan does not support asynchronous RAM (RAM16x1, RAM32x1); wide edge decoders (DECODE4, DECODE8, DECODE16); wired AND (WAND1, WOR2AND); or mode pins (MD0, MD1, MD2).

Note: Verilog reserves the names `buf`, `pullup`, and `pulldown`; the Xilinx versions are changed to `buff`, `pullup1`, `pullup2`, or `pulldown2`, and then mapped to the proper cell during implementation.

Compiling the UNISIM Library

The UNISIM VHDL library (or Verilog library) can be compiled to any physical location. The VHDL source files are found in `$XILINX/vhdl/src/unisims` and are listed here in the order in which they must be compiled.

1. `unisim_VCOMP.vhd` (component declaration file)
2. `unisim_VCOMP52K.vhd` (substitutional component declaration file for XC5200 designs)
3. `unisim_VPKG.vhd` (package file)
4. `unisim_VITAL.vhd` (model file)
5. `unisim_VITAL52K.vhd` (additional model file for XC5200 designs)

6. `unisim_VCFG4K.vhd` (configuration file for XC4K edge decoders)
7. `unisim_VCFG52K.vhd` (configuration file for XC5200 internal decoders)

Note: To use both 4K and 52K, compile them into separate directories as a UNISIM library. Change the mapping of the UNISIM logical name to the appropriate directory for each design.

The uppercase Verilog components are found in individual component files in the following directories.

1. `$XILINX/verilog/src/UNI3000` (Series 3K)
2. `$XILINX/verilog/src/UNI4000E` (Series 4KE, 4KL, Spartan)
3. `$XILINX/verilog/src/UNI4000X` (Series 4KX, 4KXL, 4KXV)
4. `$XILINX/verilog/src/UNI5200` (Series 5200)

Instantiating UNISIM Library Components

You can instantiate UNISIM library components in your design and simulate them during RTL simulation. Your VHDL code must refer to the UNISIM library compiled by you or by your system administrator. The VHDL simulation tool must map the logical library to the physical location of the compiled library. VHDL component declarations are provided in the library and do not need to be repeated in your code. Verilog must also map to the UNISIM Verilog library.

Using the LogiBLOX Library

LogiBLOX is a module generator used for schematic-based design entry of modules such as adders, counters, and large memory blocks. In the HDL flow, you can use LogiBLOX to generate large blocks of memory for instantiation. Refer to the *LogiBLOX Reference/User Guide* for more information.

In addition to the RTL code that results in synthesized logic, you can generate modules such as counters, adders, and large memory arrays with LogiBLOX. You can enter the desired parameters into LogiBLOX and select a VHDL model as output. The VHDL model is at the behavioral level because it allows for quicker simulation times. LogiBLOX is primarily useful for building large memory arrays that cannot be inferred. VHDL models are provided for LogiBLOX

modules from the schematic environment, or for large memory arrays. Most LogiBLOX modules contain registers and require the global set/reset (GSR) initialization. Since the modules do not contain output buffers going off-chip, the global tristate enable (GTS) initialization does not apply.

LogiBLOX models begin as behavioral in VHDL, but are mapped to SIMPRIM structural models in the back-annotated netlist. The behavioral model is also used for any post-synthesis simulation because the module is processed as a “black box” during synthesis. It is important that the initialization behavior is consistent for the behavioral model used for RTL and post-synthesis simulation and for the structural model used after implementation. In addition, the initialization behavior must work with the method used for synthesized logic and cores.

Note: For Verilog, the LogiBLOX model is a structural netlist of SIMPRIM models. Do *not* synthesize this netlist; it is for functional simulation only.

Compiling the LogiBLOX Library

The LogiBLOX library is not a library of modules. It is a set of packages required by the LogiBLOX models that are created “on-the-fly” by the LogiBLOX tool.

You can compile the LogiBLOX VHDL library (or Verilog) to any specified physical location. The VHDL source files are in `$XILINX/vhdl/src/logiblox`, and are listed below in the order in which they must be compiled.

1. `mvlutil.vhd`
2. `mvlarith.vhd`
3. `logiblox.vhd`

The Verilog source files are in `$XILINX/verilog/src/logiblox`.

Instantiating LogiBLOX Modules

LogiBLOX components are simulated with behavioral code. They are not intended to be synthesized, but they can be simulated. The synthesizer processes the components as a “black box”. Implementation uses the NGO file created by LogiBLOX. The source libraries for LogiBLOX packages are in `$XILINX/vhdl/src/logiblox` and

`$XILINX/verilog/src/logiblox`. The actual models are output from the LogiBLOX tool. The package files must be compiled into a library named `logiblox`. The component model from the LogiBLOX GUI should be compiled into your working directory with your design.

Using the LogiCORE Library

In addition to synthesized or generated logic, you can use high-level pre-designed LogiCORE models. These models are high-level VHDL behavioral or RTL models that are mapped to SIMPRIM structural models in the back-annotated netlist. The behavioral model is used for any post-synthesis simulation because synthesis processes the core as a “black box”. As with LogiBLOX models, the initialization behavior must be consistent for the behavioral model used for RTL and post-synthesis simulation and for the structural model used after implementation. In addition, the initialization behavior must work with the method used for synthesized logic and LogiBLOX modules.

The UNISIM VHDL and Verilog libraries can emulate the global set/reset and global tristate network in Xilinx FPGAs. VHDL uses special components for driving the local reset and tristate enable lines, and sending implementation directives to move the nets to the global signal resources. Verilog uses a macro definition.

The local signals emulate the fully routed global signals in a post-routed netlist. Both the VHDL and Verilog post-route netlists use the SIMPRIM Library and have global reset and output tristate enable signals fully routed; they are not emulated.

LogiBLOX and LogiCORE models are at a behavioral level and do not use library components for global signals. However the LogiBLOX model does require the packages that are compiled into the LogiBLOX library.

Simulating Global Signals

Xilinx PLDs have register (flip-flops and latches) set/reset circuitry that pulses at the end of the configuration mode and after power-up. This pulse is automatic and does not need to be programmed. All the flip-flops and latches in a PLD receive this pulse through a dedicated global set/reset (GSR), PRLD, or reset (GR) net. The registers either set or reset, depending on how the registers are defined.

In addition to the set/reset pulse, all output buffers are tristated during configuration mode and after power-up with the dedicated global output tristate enable (GTS) net. The global tri-state and reset net names are provided in the following table.

Table 5-6 Global Reset and Tristate names for Xilinx Devices

| Device Family | Global Reset Name | Global Tristate Name | Default Reset Polarity |
|---------------|-------------------|----------------------|------------------------|
| XC3000 | GR | Not Available | Low |
| XC4000 | GSR | GTS | High |
| XC5200 | GR | GTS | High |
| XC9500 | PRLD | GTS | High |
| SPARTAN | GSR | GTS | High |

These PRLD, GSR, and GR nets require special handling during synthesis, simulation, and implementation to prevent them from being assigned to normally routed nets, which uses valuable routing resources and degrades design performance. The GSR, PRLD, or GR net receives a reset-on-configuration pulse from the initialization controller, as shown in the following figure.

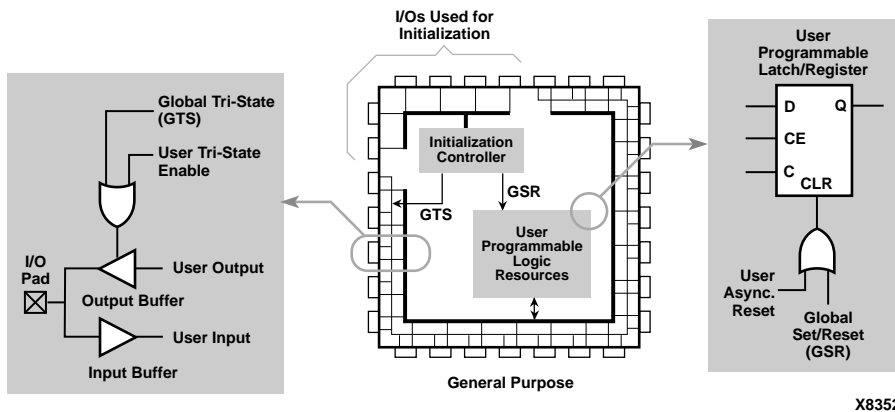


Figure 5-2 Built-in FPGA Initialization Circuitry

This pulse occurs during the configuration or power-up mode of the PLD. However, for ease of simulation, it is usually inserted at time zero of the test bench, before logical simulation is initiated. The pulse

width is device-dependent and can vary widely, depending on process voltage and temperature changes. The pulse is guaranteed to be long enough to overcome all net delays on the reset special-purpose net. The parameter for the pulse width is TPOR, as described in *The Programmable Logic Data Book*.

The tristate-on-configuration circuit shown in the “Built-in FPGA Initialization Circuitry” figure also occurs during the configuration or power-up mode of the PLD. Just as for the reset-on-configuration simulation, it is usually inserted at time zero of the test bench before logical simulation is initiated. The pulse drives all outputs to the tristate condition they are in during the configuration of the PLD. All general-purpose outputs are affected whether they are regular, tristate, or bi-direct outputs during normal operation. This ensures that the outputs do not erroneously drive other devices as the PLD is being configured. The pulse width is device-dependent and can vary widely with process and temperature changes. The pulse is guaranteed to be long enough to overcome all net delays on the GTS net. The generating circuitry is separate from the reset-on-configuration circuit. The pulse width parameter is TPOR, as described in *The Programmable Logic Data Book*. Simulation models use this pulse width parameter for determining HDL simulation for global reset and tristate circuitry (initially developed for schematic design.)

Adapting Schematic Global Signal Methodology for VHDL

There are no global set/reset or output tristate enable pins on the simulation, synthesis, or implementation models of the register components in schematic-based designs. During synthesis, both the global and local reset and tristate-state enable signals are connected to the local pin. Schematic simulators can simulate global signals without a pin. The global signals are represented as internal signals in the schematic simulation model and the test vectors drive the internal global signals directly. If you want complete control of initialization, use registers with asynchronous set/reset to emulate the GSR, even if local set/reset is not required. Synchronous set/reset registers will initialize on their own at time zero. They can be synchronously set after that but cannot emulate GSR behavior after time zero. Some memory components without asynchronous clears will exhibit similar behavior.

In VHDL designs, you must declare as ports any signals that are stimulated or monitored from outside a module. Global GSR and GTS signals are used to initialize the simulation and require access ports if controlled from the test bench. However, the addition of these ports makes the pre- and post-implementation versions of your design different, and your original test bench is no longer applicable to both versions of your design. Since the port lists for the two versions of your design are different, the *socket* in the test bench matches only one of them. To address this issue, five new cells are provided for VHDL simulation: ROC, ROCBUF, TOC, TOCBUF, and STARTBUF.

Verilog can simulate internal signals, and these signals are driven directly from the test bench. However, interpretive Verilog (such as Verilog-XL) and compiled Verilog (such as MTI or NC-Verilog) require a different approach for handling the libraries.

You do not need to incorporate any ports into schematic designs for simulators to mimic the device's global reset (GSR) or global tristate (GTS) networks. Schematic simulators specify these signals on the register model as 'global' to indicate to the simulator that these signals are all connected. These signals are not part of the cell's pin list, do not appear in the netlist, and are not implemented in the resulting design. These global signals are mapped into the equivalent signals in the back-end simulation model. Using this methodology with schematic designs, you can fully simulate the silicon's built-in global networks and implement your design without causing congestion of the general-purpose routing resources and degrading the clock speed.

Setting VHDL Global Set/Reset Emulation in Functional Simulation

When using the VHDL UNISIM library, it is important to control the global signals for reset and output tristate enable. If do not control these signals, your timing simulation results will not match your functional simulation results because the initialization differs.

VHDL simulation does not support test bench driven internal global signals. If the test bench drives the global signal, a port is required. Otherwise, the global net must be driven by a component within the architecture.

Also, the register components do not have pins for the global signals because you do not want to wire to these special pre-laid nets. Instead, you want implementation to use the dedicated network on the chip.

For the HDL synthesis flow, the global reset and output tristate enable signals are emulated with the local reset and tristate enable signals. Special implementation directives are put on the nets to move them to the special pre-routed nets for global signals.

The VHDL UNISIM library uses special components to drive the local reset and tristate enable signals. These components use the local signal connections to emulate the global signal, and also provide the implementation directives to ensure that the pre-routed wires are used.

You can instantiate these special components in the RTL description to ensure that all functional simulations match the timing simulation with respect to global signal initializations.

Global Signal Considerations (VHDL)

The following are important to VHDL simulation, synthesis, and implementation of global signals in FPGAs.

- The global signals have automatically generated pulses that always occur even if the behavior is not described in the front-end description. The back-annotated netlist has these global signals, to match the silicon, even if the source design does not.
- The simulation and synthesis models for registers (flip-flops and latches) and output buffers do not contain pins for the global signals. This is necessary to maintain compatibility with schematic libraries that do not require the pin to model the global signal behavior.
- VHDL does not have a standardized method for handling global signals that is acceptable within a VITAL-compatible library.
- LogiBLOX generates modules that are represented as behavioral models and require a different way to handle the global signal, yet still maintain compatibility with the method used for general user-defined logic and LogiCOREs.
- Intellectual property cores from the LogiCORE product line are represented as behavioral, RTL, or structural models and require

a different way to handle the global signal, yet still maintain compatibility with the method used for general user-defined logic and LogiBLOX.

- The design is represented at different levels of abstraction during the pre- and post-synthesis and implementation phases of the design process. The solutions work for all three levels and give consistent results.
- The place and route tools must be given special directives to identify the global signals in order to use the built-in circuitry instead of the general-purpose logic.

GSR Network Design Cases

When defining a methodology to control a device’s global set/reset (GSR) network, you should consider the following three general cases.

Table 5-7 GSR Design Cases

| Name | Description |
|---------|---|
| Case 1 | Reset-On-Configuration pulse only; no user control of GSR |
| Case 1A | Simulation model ROC initializes sequential elements |
| Case 1B | User initializes sequential elements with ROCBUF model and simulation vectors |
| Case 2 | User control of GSR after Power-on-Reset |
| Case 2A | External Port driving GSR |
| Case 2B | Internal signal driving GSR |
| Case 3 | Don’t Care |

Note: Reset-on-Configuration for PLDs is similar to Power-on-Reset for ASICs except it occurs at power-up and during configuration of the PLD.

Case 1 is defined as follows.

- Automatic pulse generation of the Reset-On-Configuration signal
- No control of GSR through a test bench
- Involves initialization of the sequential elements in a design during power-on, or initialization during configuration of the device

- Need to define the initial states of a design's sequential elements, and have these states reflected in the implemented and simulated design
- Two sub-cases
 - In Case 1A, you do not provide the simulation with an initialization pulse. The simulation model provides its own mechanism for initializing its sequential elements (such as the real device does when power is first applied).
 - In Case 1B, you can control the initializing power-on reset pulse from a test bench without a global reset pin on the FPGA. This case is applicable when system-level issues make your design's initialization synchronous to an off-chip event. In this case, you provide a pulse that initializes your design at the start of simulation time, and possibly provide further pulses as simulation time progresses (perhaps to simulate cycling power to the device). Although you are providing the reset pulse to the simulation model, this pulse is not required for the implemented device. A reset port is not required on the implemented device, however, a reset port is required in the behavioral code through which your reset pulse can be applied with test vectors during simulation.

Using VHDL Reset-On-Configuration (ROC) Cell (Case 1A)

For Case 1A, the ROC (Reset-On-Configuration) instantiated component model is used. This model creates a one-shot pulse for the global set/reset signal. The pulse width is a generic and can be configured to match the device and conditions specified. The ROC cell is in the post-routed netlist and, with the same pulse width, it mimics the pre-route global set/reset net. The following is an example of an ROC cell.

Note: The TPOR parameter from *The Programmable Logic Data Book* is used as the WIDTH parameter.

```

library IEEE;
use IEEE.std_logic_1164.all;
use IEEE.std_logic_unsigned.all;
library UNISIM;
use UNISIM.all;
entity EX_ROC is
    port (CLOCK, ENABLE : in std_logic;
          CUP, CDOWN : out std_logic_vector (3 downto 0));
end EX_ROC;
architecture A of EX_ROC is
    signal GSR : std_logic;
    signal COUNT_UP, COUNT_DOWN : std_logic_vector (3 downto 0);
    component ROC
        port (O : out std_logic);
    end component;
begin
    U1 : ROC port map (O => GSR);

    UP_COUNTER : process (CLOCK, ENABLE, GSR)
    begin
        if (GSR = '1') then
            COUNT_UP <= "0000";
        elsif (CLOCK'event AND CLOCK = '1') then
            if (ENABLE = '1') then
                COUNT_UP <= COUNT_UP + "0001";
            end if;
        end if;
    end process UP_COUNTER;
    DOWN_COUNTER : process (CLOCK, ENABLE, GSR, COUNT_DOWN)
    begin
        if (GSR = '1' OR COUNT_DOWN = "0101") then
            COUNT_DOWN <= "1111";
        elsif (CLOCK'event AND CLOCK = '1') then
            if (ENABLE = '1') then
                COUNT_DOWN <= COUNT_DOWN - "0001";
            end if;
        end if;
    end process DOWN_COUNTER;
    CUP <= COUNT_UP;
    CDOWN <= COUNT_DOWN;
end A;

```

Using ROC Cell Implementation Model (Case 1A)

Complementary to the previous VHDL model is an implementation model that guides the place and route tool to connect the net driven by the ROC cell to the special purpose net.

This cell is created during back-annotation if you do not use the `-gp` or `STARTUP` block options. It can be instantiated in the front end to match functionality with GSR, GR, or PRLD (in both functional and timing simulation.) During back-annotation, the entity and architecture for the ROC cell is placed in your design's output VHDL file. In the front end, the entity and architecture are in the UNISIM Library, requiring only a component instantiation. The ROC cell generates a one-time initial pulse to drive the GR, GSR, or PRLD net starting at time zero for a specified pulse width. You can set the pulse width with a generic in a configuration statement. The default value of the pulse width is 0 ns. This value disables the ROC cell and causes the global set/reset to be held low. (Active low resets are handled within the netlist itself and need to be inverted before using.)

ROC Test Bench (Case 1A)

With the ROC cell you can simulate with the same test bench used in RTL simulation, and you can control the width of the global set/reset signal in your implemented design. ROC cells require a generic `WIDTH` value, usually specified with a configuration statement. Otherwise, a generic map is required as part of the component instantiation. You can set the generic with any generic mapping method. Set the width generic after consulting *The Programmable Logic Data Book* for the particular part and mode implemented. For example, an XC4000E part can vary from 10 ms to 130 ms. Use the `TPOR` parameter in the Configuration Switching Characteristics tables for Master, Slave, and Peripheral modes. The following is the test bench for the ROC example.

```

library IEEE;
use IEEE.std_logic_1164.all;
use IEEE.std_logic_unsigned.all;

library UNISIM;
use UNISIM.all;

entity test_ofex_roc is end test_ofexroc;

architecture inside of test_ofex_roc is

Component ex_roc
Port ( CLOCK, ENABLE: in STD_LOGIC;
      CUP, CDOWN: out STD_LOGIC_VECTOR (3 downto 0));
End component;

.
.
.

```

Begin

```

UUT: ex_roc port map(. . . .);

.
.
.
End inside;

```

The best method for mapping the generic is a configuration in your test bench, as shown in the following example.

```

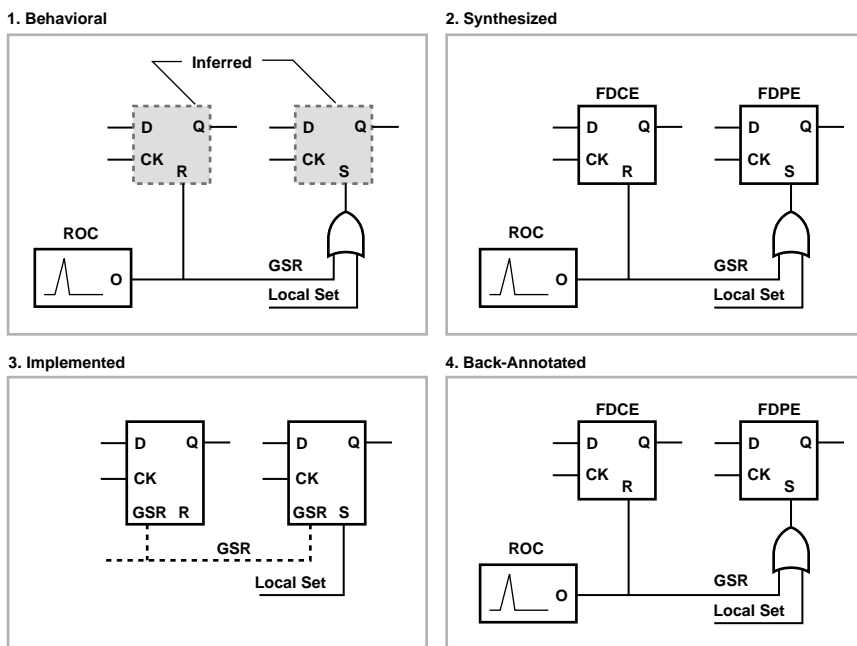
Configuration overall of test_ofexroc is
For inside
    For UUT:ex_roc
        For A
            For U1:ROC use entity UNISIM.ROC (ROC_V)
                Generic map (WIDTH=>52 ns);
            End for;
        End for;
    End for;
End overall;

```

This configuration is for pre-NGDBuild simulation. A similar configuration is used for post-NGDBuild simulation. The ROC, TOC, and OSC4 are mapped to the WORK library, and corresponding architecture names may be different. Review the .vhd file created by NGD2VHDL for the current entity and architecture names for post-NGDBuild simulation.

ROC Model in Four Design Phases (Case 1A)

The following figure shows the progression of the ROC model and its interpretation in the four main design phases.



X8348

Figure 5-3 ROC Simulation and Implementation

- Behavioral Phase

In this phase, the behavioral or RTL description registers are inferred from the coding style, and the ROC cell can be instantiated. If it is not instantiated, the signal is not driven during simulation or is driven within the architecture by code that cannot be synthesized. Some synthesizers infer the local resets that are best

for the global signal and insert the ROC cell automatically. When this occurs, instantiation may not be required unless RTL level simulation is needed. The synthesizer may allow you to select the reset line to drive the ROC cell. Xilinx recommends instantiation of the ROC cell during RTL coding because the global signal is easily identified. This also ensures that GSR behavior at the RTL level matches the behavior of the post-synthesis and implementation netlists.

- Synthesized Phase

In this phase, inferred registers are mapped to a technology and the ROC instantiation is either carried from the RTL or inserted by the synthesis tools. As a result, consistent global set/reset behavior is maintained between the RTL and synthesized structural descriptions during simulation.

- Implemented Phase

During implementation, the ROC is removed from the logical description that is placed and routed as a pre-existing circuit on the chip. The ROC is removed by making the output of the ROC cell appear as an open circuit. Then the implementation tool can trim all the nets driven by the ROC to the local sets or resets of the registers, and the nets are not routed in general purpose routing. All set/resets for the registers are automatically assumed to be driven by the global set/reset net so data is not lost.

- Back-annotated Phase

In this phase, the Xilinx VHDL netlist program assumes all registers are driven by the GSR net; replaces the ROC cell; and rewires it to the GSR nets in the back-annotated netlist. The GSR net is a fully wired net and the ROC cell is inserted to drive it. A similar VHDL configuration can be used to set the generic for the pulse width.

Using VHDL ROCBUF Cell (Case 1B)

For Case 1B, the ROCBUF (Reset-On-Configuration Buffer) instantiated component is used. This component creates a buffer for the global set/reset signal, and provides an input port on the buffer to drive the global set reset line. During the place and route process, this port is removed so it is not implemented on the chip. ROCBUF does not reappear in the post-routed netlist. Instead, you can select an

implementation option to add a global set/reset port to the back-annotated netlist. A buffer is not necessary since the implementation directive is no longer required.

The following example illustrates how to use the ROCBUF in your designs.

```
library IEEE;
use IEEE.std_logic_1164.all;
use IEEE.std_logic_unsigned.all;

library UNISIM;
use UNISIM.all;
entity EX_ROCBUF is
    port (CLOCK, ENABLE, SRP : in std_logic;
          CUP, CDOWN : out std_logic_vector (3 downto 0));
end EX_ROCBUF;
architecture A of EX_ROCBUF is
    signal GSR : std_logic;
    signal COUNT_UP, COUNT_DOWN : std_logic_vector (3 downto 0);
    component ROCBUF
        port (I : in std_logic;
              O : out std_logic);
    end component;
begin
    U1 : ROCBUF port map (I => SRP, O => GSR);
    UP_COUNTER : process (CLOCK, ENABLE, GSR)
    begin
        if (GSR = '1') then
            COUNT_UP <= "0000";
        elsif (CLOCK'event AND CLOCK = '1') then
            if (ENABLE = '1') then
                COUNT_UP <= COUNT_UP + "0001";
            end if;
        end if;
    end process UP_COUNTER;
    DOWN_COUNTER : process (CLOCK, ENABLE, GSR, COUNT_DOWN)
    begin
        if (GSR = '1' OR COUNT_DOWN = "0101") then
            COUNT_DOWN <= "1111";
        elsif (CLOCK'event AND CLOCK = '1') then
            if (ENABLE = '1') then
                COUNT_DOWN <= COUNT_DOWN - "0001";
            end if;
        end if;
    end process DOWN_COUNTER;
```

```

CUP <= COUNT_UP;
CDOWN <= COUNT_DOWN;
end A;

```

ROCBUF Model in Four Design Phases (Case 1B)

The following figure shows the progression of the ROCBUF model and its interpretation in the four main design phases.

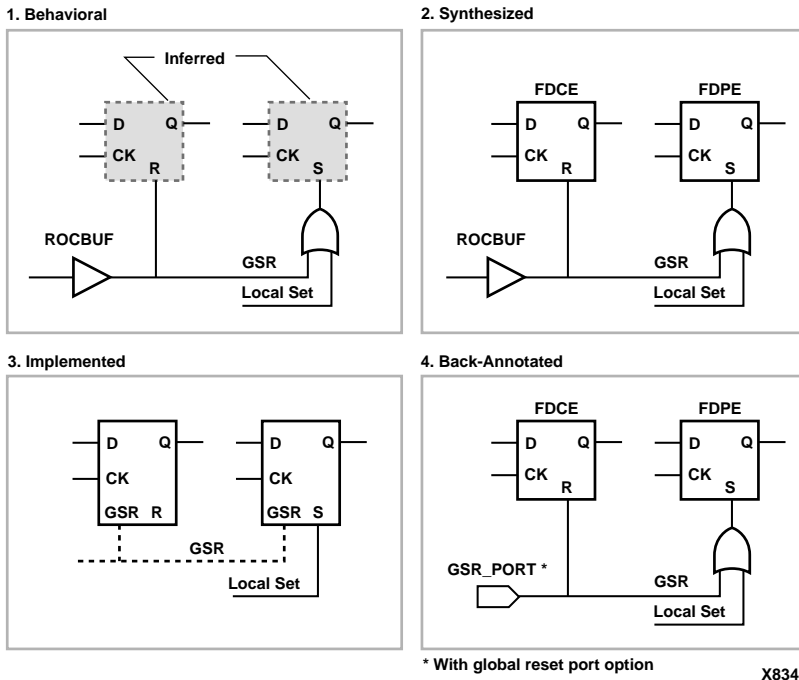


Figure 5-4 ROCBUF Simulation and Implementation

- Behavioral Phase

In this phase, the behavioral or RTL description registers are inferred from the coding style, and the ROCBUF cell can be instantiated. If it is not instantiated, the signal is not driven during simulation, or it is driven within the architecture by code that cannot be synthesized. Use the ROCBUF cell instead of the ROC cell when you want test bench control of GSR simulation. Xilinx recommends instantiating the ROCBUF cell during RTL coding because the global signal is easily identified, and you are

not relying on a synthesis tool feature that may not be available if ported to another tool. This also ensures that GSR behavior at the RTL level matches the behavior of the post-synthesis and implementation netlists.

- **Synthesized Phase**

In this phase, inferred registers are mapped to a technology and the ROCBUF instantiation is either carried from the RTL or inserted by the synthesis tools. As a result, consistent global set/reset behavior is maintained between the RTL and synthesized structural descriptions during simulation.

- **Implemented Phase**

During implementation, the ROCBUF is removed from the logical description that is placed and routed as a pre-existing circuit on the chip. The ROCBUF is removed by making the input and the output of the ROCBUF cell appear as an open circuit. Then the implementation tool can trim the port that drives the ROCBUF input, as well as the nets driven by the ROCBUF output. As a result, nets are not routed in general purpose routing. All set/resets for the registers are automatically assumed to be driven by the global set/reset net so data is not lost. You can use a VHDL netlist tool option to add the port back.

- **Back-annotated Phase**

In this phase, the Xilinx VHDL netlist program starts with all registers initialized by the GSR net, and it replaces the ROC cell it would normally insert with a port if the GSR port option is selected. The GSR net is a fully wired net driven by the added GSR port. A ROCBUF cell is not required because the port is sufficient for simulation, and implementation directives are not required

Using VHDL STARTBUF Block (Case 2A and 2B)

The STARTUP block is traditionally instantiated to identify the GR, PRLD, or GSR signals for implementation if the global reset on tristate is connected to a chip pin. However, this implementation directive component cannot be simulated, and causes warning messages from the simulator. However, you can use the STARTBUF cell instead, which can be simulated. STARTUP blocks are allowed if the warnings can be addressed or safely ignored.

For Cases 2A and 2B, use the STARTBUF cell. This cell provides access to the input and output ports of the STARTUP cell that direct the implementation tool to use the global networks. The input and output port names differ from the names of the corresponding ports of the STARTUP cell. This was done for the following reasons.

- To make the STARTBUF a model that can be simulated with inputs and outputs. The STARTUP cell hangs from the net it is connected to.
- To make one model that works for all Xilinx technologies. The XC4000 and XC5200 families require different STARTUP cells because the XC5200 has a global reset (GR) net and not a GSR.

The mapping to the architecture-specific STARTUP cell from the instantiation of the STARTBUF is done during implementation. The STARTBUF pins have the suffix “IN” (input port) or “OUT” (output port). Two additional output ports, GSROUT and GTSOUT, are available to drive a signal for clearing or setting a design's registers (GSROUT), or for tri-stating your design's I/Os (GTSOUT).

The input ports, GSRIN and GTSIN, can be connected either directly or indirectly via combinational logic to input ports of your design. Your design's input ports appear as input pins in the implemented design. The design input port connected to the input port, GSRIN, is then referred to as the device reset port, and the design input port connected to the input port, GTSIN, is referred to as the device tristate port. The table below shows the correspondence of pins between STARTBUF and STARTUP.

Table 5-8 STARTBUF/STARTUP Pin Descriptions

| STARTBUF Pin Name | Connection Point | XC4000 STARTUP Pin Name | XC5200 STARTUP Pin Name | Spartan |
|--------------------------|--|---|---|---|
| GSRIN | Global Set/ Reset Port of Design | GSR | GR | GSR |
| GTSIN | Global Tristate Port of Design | GTS | GTS | GTS |
| GSROUT | All Registers Asynchronous Set/Reset | Not Available For Simulation Only | Not Available For Simulation Only | Not Available For Simulation Only |

Table 5-8 STARTBUF/STARTUP Pin Descriptions

| STARTBUF Pin Name | Connection Point | XC4000 STARTUP Pin Name | XC5200 STARTUP Pin Name | Spartan |
|-------------------|-------------------------------------|-----------------------------------|-----------------------------------|---------|
| GTSOUT | All Output Buffers Tristate Control | Not Available For Simulation Only | Not Available For Simulation Only | N/A |
| CLKIN | Port or Internal Logic | CLK | CLK | CLK |
| Q2OUT | Port Or Internal Logic | Q2 | Q2 | Q2 |
| Q3OUT | Port Or Internal Logic | Q3 | Q3 | Q3 |
| OUT | Port Or Internal Logic | Q1Q4 | Q1Q4 | Q1Q4 |
| DONEINOUT | Port Or Internal Logic | DONEIN | DONEIN | DONEIN |

Note: Using STARTBUF indicates that you want to access the global set/reset and/or tristate pre-routed networks available in your design's target device. As a result, you must provide the stimulus for emulating the automatic pulse as well as the user-defined set/reset. This allows you complete control of the reset network from the test bench.

The following example shows how to use the STARTBUF cell.

```

library IEEE;
use IEEE.std_logic_1164.all;
use IEEE.std_logic_unsigned.all;
library UNISIM;
use UNISIM.all;
entity EX_STARTBUF is
    port (CLOCK, ENABLE, DRP, DTP : in std_logic;
          CUP, CDOWN : out std_logic_vector (3 downto 0));
end EX_STARTBUF;
architecture A of EX_STARTBUF is
    signal GSR, GSRIN_NET, GROUND, GTS : std_logic;
    signal COUNT_UP, COUNT_DOWN : std_logic_vector (3 downto 0);
    component STARTBUF
        port (GSRIN, GTSIN, CLKIN : in std_logic; GSR_OUT, GTS_OUT,

```

```

        DONEINOUT, Q1Q4OUT, Q2OUT, Q3OUT : out std_logic);
    end component;
begin
    GROUND <= '0';
    GSRIN_NET <= NOT DRP;
    U1 : STARTBUF port map (GSRIN => GSRIN_NET, GTSIN => DTP,
        CLKIN => GROUND, GSROUT => GSR, GTSOUT => GTS);
    UP_COUNTER : process (CLOCK, ENABLE, GSR)
    begin
        if (GSR = '1') then
            COUNT_UP <= "0000";
        elsif (CLOCK'event AND CLOCK = '1') then
            if (ENABLE = '1') then
                COUNT_UP <= COUNT_UP + "0001";
            end if;
        end if;
    end process UP_COUNTER;
    DOWN_COUNTER : process (CLOCK, ENABLE, GSR, COUNT_DOWN)
    begin
        if (GSR = '1' OR COUNT_DOWN = "0101") then
            COUNT_DOWN <= "1111";
        elsif (CLOCK'event AND CLOCK = '1') then
            if (ENABLE = '1') then
                COUNT_DOWN <= COUNT_DOWN - "0001";
            end if;
        end if;
    end process DOWN_COUNTER;
    CUP <= COUNT_UP when (GTS = '0' AND COUNT_UP /= "0000") else "ZZZZ";
    CDOWN <= COUNT_DOWN when (GTS = '0') else "ZZZZ";
end A;
```

GTS Network Design Cases

Just as for the global set/reset net there are three cases for using your device's output tristate enable (GTS) network, as shown in the following table.

Table 5-9 GTS Design Cases

| Name | Description |
|------------------------------|--|
| Case A Case A1 Case A2 | Tristate-On-Configuration only; no user control of GTS Simulation Model TOC Tristates output buffers during configuration or power-up User initializes sequential elements with TOCBUF model and simulation vectors |
| Case B Case B1 Case B2 | User control of GTS after Tristate-On-Configuration External PORT driving GTS Internal signal driving GTS |
| Case C | Don't Care |

Case A is defined as follows.

- Tristating of output buffers during power-on or configuration of the device
- Output buffers are tristated and reflected in the implemented and simulated design
- Two sub-cases
 - In Case A1, you do not provide the simulation with an initialization pulse. The simulation model provides its own mechanism for initializing its sequential elements (such as the real device does when power is first applied).
 - In Case A2, you can control the initializing Tristate-On-Configuration pulse. This case is applicable when system-level issues make your design's configuration synchronous with an off-chip event. In this case, you provide a pulse to tristate the output buffers at the start of simulation time, and possibly provide further pulses as simulation time progresses (perhaps to simulate cycling power to the device). Although you are providing the Tristate-On-Configuration pulse to the simulation model, this pulse is not required for the implemented device. A Tristate-On-Configuration port is not required on the implemented device, however, a TOC port is required in the behavioral code through which your TOC pulse can be applied with test vectors during simulation.

Using VHDL Tristate-On-Configuration (TOC)

The TOC cell is created if you do not use the `-tp` or `STARTUP` block options. The entity and architecture for the TOC cell is placed in the design's output VHDL file. The TOC cell generates a one-time initial pulse to drive the GR, GSR, or PRLD net starting at time '0' for a user-defined pulse width. The pulse width can be set with a generic. The default WIDTH value is 0 ns, which disables the TOC cell and holds the tristate enable low. (Active low tristate enables are handled within the netlist itself; you must invert this signal before using it.)

The TOC cell enables you to simulate with the same test bench as in the RTL simulation, and also allows you to control the width of the tristate enable signal in your implemented design.

The TOC components require a value for the generic WIDTH, usually specified with a configuration statement. Otherwise, a generic map is required as part of the component instantiation.

You may set the generic with any generic mapping method you choose. Set the WIDTH generic after consulting *The Programmable Logic Data Book* for the particular part and mode you have implemented. For example, an XC4000E part can vary from 10 ms to 130 ms. Use the TPOR (Power-On Reset) parameter found in the Configuration Switching Characteristics tables for Master, Slave, and Peripheral modes.

VHDL TOC Cell (Case A1)

For Case A1, use the TOC (Tristate-On-Configuration) instantiated component. This component creates a one-shot pulse for the global Tristate-On-Configuration signal. The pulse width is a generic and can be selected to match the device and conditions you want. The TOC cell is in the post-routed netlist and, with the same pulse width set, it mimics the pre-route Tristate-On-Configuration net.

TOC Cell Instantiation (Case A1)

The following is an example of how to use the TOC cell.

Note: The TPOR parameter from *The Programmable Logic Data Book* is used as the WIDTH parameter in this example.

```
library IEEE;
use IEEE.std_logic_1164.all;
use IEEE.std_logic_unsigned.all;
library UNISIM;
use UNISIM.all;
entity EX_TOC is
    port (CLOCK, ENABLE : in std_logic;
          CUP, CDOWN : out std_logic_vector (3 downto 0));
end EX_TOC;
architecture A of EX_TOC is
    signal GSR, GTS : std_logic;
    signal COUNT_UP, COUNT_DOWN : std_logic_vector (3 downto 0);
    component ROC
        port (O : out std_logic);
    end component;
    component TOC
        port (O : out std_logic);
    end component;
begin
    U1 : ROC port map (O => GSR);
    U2 : TOC port map (O => GTS);
    UP_COUNTER : process (CLOCK, ENABLE, GSR)
    begin
        if (GSR = '1') then
            COUNT_UP <= "0000";
        elsif (CLOCK'event AND CLOCK = '1') then
            if (ENABLE = '1') then
                COUNT_UP <= COUNT_UP + "0001";
            end if;
        end if;
    end process UP_COUNTER;
    DOWN_COUNTER : process (CLOCK, ENABLE, GSR, COUNT_DOWN)
    begin
        if (GSR = '1' OR COUNT_DOWN = "0101") then
            COUNT_DOWN <= "1111";
        elsif (CLOCK'event AND CLOCK = '1') then
            if (ENABLE = '1') then
                COUNT_DOWN <= COUNT_DOWN - "0001";
            end if;
        end if;
    end process DOWN_COUNTER;
    CUP <= COUNT_UP when (GTS = '0' AND COUNT_UP /= "0000") else "ZZZZ";
    CDOWN <= COUNT_DOWN when (GTS = '0') else "ZZZZ";
end A;
```

TOC Test Bench (Case A1)

The following is the test bench for the TOC example.

```

library IEEE;
use IEEE.std_logic_1164.all;
use IEEE.std_logic_unsigned.all;

library UNISIM;
use UNISIM.all;

entity test_ofex_toc is end test_ofextoc;

architecture inside of test_ofex_toc is

Component ex_toc
Port ( CLOCK, ENABLE: in STD_LOGIC;
      CUP, CDOWN: out STD_LOGIC_VECTOR (3 downto 0));
End component;

.
.
.

Begin

UUT: ex_toc port map(. . . .);

.
.
.

End inside;

```

The best method for mapping the generic is a configuration in the test bench, as shown in the following example.

```

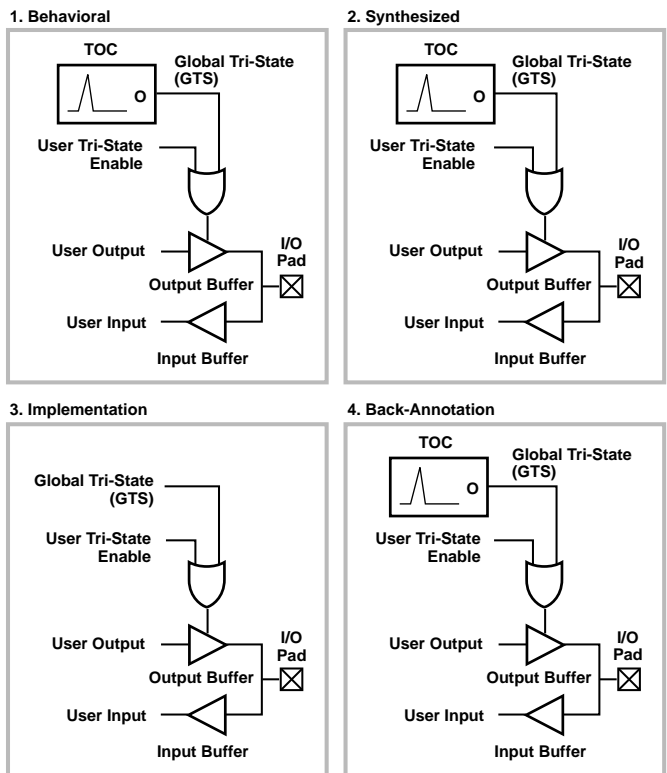
Configuration overall of test_ofextoc is
For inside
    For UUT:ex_toc
        For A
            For U1:TOC use entity UNISIM.TOC (TOC_V)
                Generic map (WIDTH=>52 ns);
            End for;
        End for;
    End for;
End overall;

```

This configuration is for pre-NGDBuild simulation. A similar configuration is used for post-NGDBuild simulation. The ROC, TOC, and OSC4 are mapped to the WORK library, and corresponding architecture names may be different. Review the .vhd file created by NGD2VHDL for the current entity and architecture names for post-NGDBuild simulation.

TOC Model in Four Design Phases (Case A1)

The following figure shows the progression of the TOC model and its interpretation in the four main design phases.



X8350

Figure 5-5 TOC Simulation and Implementation

- Behavioral Phase

In this phase, the behavioral or RTL description of the output buffers are inferred from the coding style. The TOC cell can be instantiated. If it is not instantiated, the GTS signal is not driven during simulation or is driven within the architecture by code that cannot be synthesized. Some synthesizers can infer which of the local output tristate enables is best for the global signal, and will insert the TOC cell automatically so instantiation may not be required unless RTL level simulation is desired. The synthesizer may also allow you to select the output tristate enable line you want driven by the TOC cell. Instantiation of the TOC cell in the RTL description is recommended because you can immediately identify what signal is the global signal, and you are not relying on a synthesis tool feature that may not be available if ported to another tool.

- Synthesized Phase

In this phase, the inferred registers are mapped to a device, and the TOC instantiation is either carried from the RTL or is inserted by the synthesis tools. This results in maintaining consistent global output tristate enable behavior between the RTL and the synthesized structural descriptions during simulation.

- Implemented Phase

During implementation, the TOC is removed from the logical description that is placed and routed because it is a pre-existing circuit on the chip. The TOC is removed by making the input and output of the TOC cell appear as an open circuit. This allows the router to remove all nets driven by the TOC cell as if they were undriven nets. The VHDL netlist program assumes all output tristate enables are driven by the global output tristate enable so data is not lost.

- Back-annotation Phase

In this phase, the VHDL netlist tool re-inserts a TOC component for simulation purposes. The GTS net is a fully wired net and the TOC cell is inserted to drive it. You can use a configuration similar to the VHDL configuration for RTL simulation to set the generic for the pulse width.

Using VHDL TOCBUF (Case B1)

For Case B1, use the TOCBUF (Tristate-On-Configuration Buffer) instantiated component model. This model creates a buffer for the global output tristate enable signal. You now have an input port on the buffer to drive the global set reset line. The implementation model directs the place and route tool to remove the port so it is not implemented on the actual chip. The TOCBUF cell does not reappear in the post-routed netlist. Instead, you can select an option on the implementation tool to add a global output tristate enable port to the back-annotated netlist. A buffer is not necessary because the implementation directive is no longer required.

TOCBUF Model Example (Case B1)

The following is an example of the TOCBUF model.

```
library IEEE;
use IEEE.std_logic_1164.all;
use IEEE.std_logic_unsigned.all;
library UNISIM;
use UNISIM.all;
entity EX_TOCBUF is
    port (CLOCK, ENABLE, SRP, STP : in std_logic;
          CUP, CDOWN : out std_logic_vector (3 downto 0));
end EX_TOCBUF;
architecture A of EX_TOCBUF is
    signal GSR, GTS : std_logic;
    signal COUNT_UP, COUNT_DOWN : std_logic_vector (3 downto 0);
    component ROCBUF
        port (I : in std_logic;
              O : out std_logic);
    end component;
    component TOCBUF
        port (I : in std_logic;
              O : out std_logic);
    end component;
begin
    U1 : ROCBUF port map (I => SRP, O => GSR);
    U2 : TOCBUF port map (I => STP, O => GTS);
    UP_COUNTER : process (CLOCK, ENABLE, GSR)
    begin
        if (GSR = '1') then
            COUNT_UP <= "0000";
        elsif (CLOCK'event AND CLOCK = '1') then
```

```

        if (ENABLE = '1') then
            COUNT_UP <= COUNT_UP + "0001";
        end if;
    end if;
end process UP_COUNTER;
DOWN_COUNTER : process (CLOCK, ENABLE, GSR, COUNT_DOWN)
begin
    if (GSR = '1' OR COUNT_DOWN = "0101") then
        COUNT_DOWN <= "1111";
    elsif (CLOCK'event AND CLOCK = '1') then
        if (ENABLE = '1') then
            COUNT_DOWN <= COUNT_DOWN - "0001";
        end if;
    end if;
end process DOWN_COUNTER;
CUP <= COUNT_UP when (GTS = '0' AND COUNT_UP /= "0000") else "ZZZZ";
CDOWN <= COUNT_DOWN when (GTS = '0') else "ZZZZ";
end A;

```

TOCBUF Model in Four Design Phases (Case B1)

The figure below shows the progression of the TOCBUF model and its interpretation in the four main design phases.

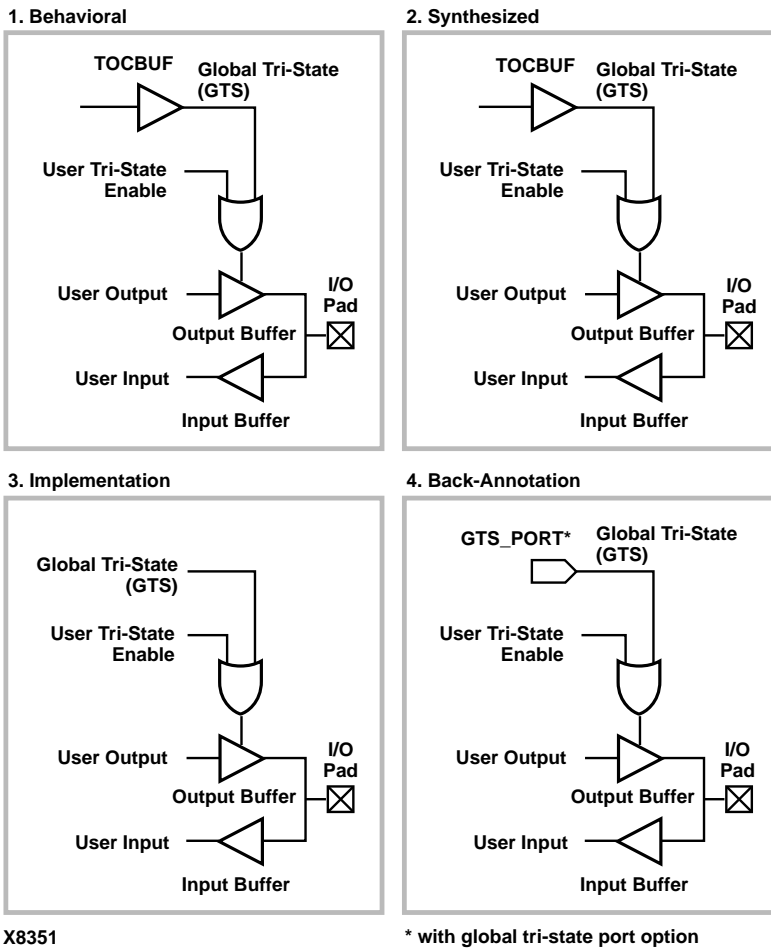


Figure 5-6 TOCBUF Simulation and Implementation

- Behavioral Phase

In this phase, the behavioral or RTL description of the output buffers are inferred from the coding style and may be inserted. You can instantiate the TOCBUF cell. If it is not instantiated, the GTS signal is not driven during simulation or it is driven within the architecture by code that cannot be synthesized. Some synthesizers can infer the local output tristate enables that make the best global signals, and will insert the TOCBUF cell automatically. As a result, instantiation may not be required unless you want RTL level simulation. The synthesizer can allow you to select the output tristate enable line you want driven by the TOCBUF cell. Instantiation of the TOCBUF cell in the RTL description is recommended because you can immediately identify which signal is the global signal and you are not relying on a synthesis tool feature that may not be available if ported to another tool.

- Synthesized Phase.

In this phase, the inferred output buffers are mapped to a device and the TOCBUF instantiation is either carried from the RTL or is inserted by the synthesis tools. This maintains consistent global output tristate enable behavior between the RTL and the synthesized structural descriptions during simulation.

- Implemented Phase

In this phase, the TOCBUF is removed from the logical description that is placed and routed because it is a pre-existing circuit on the chip.

The TOCBUF is removed by making the input and output of the TOCBUF cell appear as an open circuit. This allows the router to remove all nets driven by the TOCBUF cell as if they were undriven nets. The VHDL netlist program assumes all output tristate enables are driven by the global output tristate enable so data is not lost.

- Back-annotated Phase

In this phase, the TOCBUF cell does not reappear in the post-routed netlist. Instead, you can select an option in the implementation tool to add a global output tristate enable port to the back-annotated netlist. A buffer is not necessary because the implementation directive is no longer required. If the option is not selected, the VHDL netlist tool re-inserts a TOCBUF component

for simulation purposes. The GTS net is a fully wired net and the TOCBUF cell is inserted to drive it. You can use a configuration similar to the VHDL configuration used for RTL simulation to set the generic for the pulse width.

Using Oscillators (VHDL)

Oscillator output can vary within a fixed range. This cell is not included in the SIMPRIM library because you cannot drive global signals in VHDL designs. Schematic simulators can define and drive global nets so the cell is not required. Verilog has the ability to drive nets within a lower level module as well. Therefore the oscillator cells are only required in VHDL. After back-annotation, their entity and architectures are contained in your design's VHDL output. For functional simulation, they can be instantiated and simulated with the UNISIM Library.

The period of the base frequency must be set in order for the simulation to proceed, since the default period of 0 ns disables the oscillator. The oscillator's frequency can vary significantly with process and temperature.

Before you set the base period parameter, consult *The Programmable Logic Data Book* for the part you are using. For example, the section in *The Programmable Logic Data Book* for the XC4000 Series On-Chip Oscillator states that the base frequency can vary from 4MHz to 10 MHz, and is nominally 8 MHz. This means that the base period generic "period_8m" in the XC4000E OSC4 VHDL model can range from 250ns to 100ns. An example of this follows.

Oscillator VHDL Example

```
library IEEE;
use IEEE.std_logic_1164.all;
use IEEE.std_logic_unsigned.all;

library UNISIM;
use UNISIM.all;

entity test1 is
port (DATAIN: in STD_LOGIC;
DATAOUT: out STD_LOGIC);
end test1;
```

```
architecture inside of test1 is

signal RST: STD_LOGIC;

component ROC
port(O: out STD_LOGIC);
end component;

component OSC4
port(F8M: out STD_LOGIC);
end component;

signal internalclock: STD_LOGIC;
begin

U0: ROC port map (RST);

U1: OSC4 port map (F8M=>internalclock);

process(internalclock)
begin
if (RST='1') then
DATAOUT <= '0';

elsif(internalclock'event and internalclock='1') then
DATAOUT <= DATAIN;

end if;

end process;

end inside;
```

Oscillator Test Bench

```
library IEEE;
use IEEE.std_logic_1164.all;
use IEEE.std_logic_unsigned.all;

library UNISIM;
use UNISIM.all;

entity test_ofptest1 is end test_ofptest1;

architecture inside of test_ofptest1 is
```

```
component test1
port(DATAIN: in STD_LOGIC;
DATAOUT: out STD_LOGIC);
end component;

signal userdata, userout: STD_LOGIC;

begin

UUT: test1 port map(DATAIN=>userdata,DATAOUT=>userout);

myinput: process
begin
userdata <= '1';
wait for 299 ns;
userdata <= '0';
wait for 501 ns;
end process;

end inside;

configuration overall of test_ofctest1 is
for inside
    for UUT:test1
        for inside
            for U0:ROC use entity UNISIM.ROC(ROC_V)
            generic map (WIDTH=> 52 ns);
            end for;

            for U1:OSC4 use entity UNISIM.OSC4(OSC4_V)
            generic map (PERIOD_8M=> 25 ns);
            end for;
        end for;
    end for;
end for;
end overall;
```

This configuration is for pre-NGDBuild simulation. A similar configuration is used for post-NGDBuild simulation. The ROC, TOC, and OSC4 are mapped to the WORK library, and corresponding architecture names may be different. Review the .vhd file created by NGD2VHDL for the current entity and architecture names for post-NGDBuild simulation.

Compiling Verilog Libraries

For some Verilog simulators, you may need to compile the Verilog libraries before they can be used for design simulations. You do not need to compile the libraries for Verilog-XL because it uses an interpretive compilation of the libraries. Also, other simulators, such as NC-Verilog and Model Technology Modelsim, use a pre-compiled library methodology to speed up the simulation of your designs.

Compiling Libraries for Modelsim

This section includes instructions for compiling the Verilog Libraries for the Modelsim simulator. You can use a similar methodology for other simulators, however, the specific commands will vary.

SIMPRIM Libraries

The SIMPRIM libraries are used for simulations after the following steps in the design flow.

- RTL simulations that include instantiated LogiBLOX modules
- Post-implementation simulations (including post-NGDBuild and post-MAP)
- Timing simulation

To compile the libraries in Modelsim, enter the following commands.

```
vlib $VERILOG_DESTN/simprim
```

```
vmap simprim $VERILOG_DESTN/simprim
```

Note: \$VERILOG_DESTN is the desired location of the compiled libraries.

```
vlog -work simprim $XILINX/verilog/data/*.vmd
```

These commands have equivalent graphical user interface operations. Refer to the Modelsim documentation for details on compiling with the graphical interface.

LogiBLOX Libraries

The LogiBLOX libraries are used for RTL and pre-implementation simulations. Verilog does not have an additional library for

LogiBLOX simulations. The SIMPRIM libraries are used for RTL-instantiated LogiBLOX simulations.

UNISIM Libraries

The UNISIM libraries are used for RTL and post-synthesis simulations.

Because Modelsim uses pre-compiled libraries, Xilinx recommends compiling the UNISIM components that are instantiated in your current design. This is necessary for the GSR signal construct in the libraries. The library must be compiled with the simulation netlist and test bench, or with a special switch specific to your design.

Note: Some UNISIM components include more than one file.

The following command is an example of compiling an FDCE component to your project directory. This only needs to be done once per design or project.

```
vlog +define+GSR_SIGNAL=testbench_module_name.design_instance_name.GSR \  
$XILINX/verilog/src/UNI4000X/FDCE.v $XILINX/verilog/src/UNI4000X/  
U_MUX_2_1.v \  
$XILINX/verilog/src/UNI4000X/U_FD_P_R_S_NO_IO.v
```

In this example, `testbench_module_name` is the module name declared in the test bench file. `Design_instance_name` is the instance name assigned to the instantiated design netlist in the test bench.

You can use a wildcard to compile all library components if your design contains several different types of components, or if you are performing a post-synthesis simulation. This is a much simpler way to compile the libraries, however, it uses more disk space.

```
vlog +define+GSR_SIGNAL=testbench_module_name.design_instance_name.GSR \  
$XILINX/verilog/src/UNI4000X/*.v
```

You can also create a file that includes the components you want to compile, as well as a list of the command line switches you want to use.

```
verilog -f filename
```

The format of this file is the same as the command line. Using a file allows you to break up the commands for readability, and allows you to repeat a command without having to type it in again.

Verilog Simulation

Verilog can use the lower case simulation libraries in the Xilinx Cadence interface. These libraries must be used for netlists from Concept or Composer, or if you are using Synergy for synthesis. Upper case libraries are available in the UNISIM library. These libraries are useful for non-schematic flows that use case-sensitive synthesis tools with upper case libraries.

Setting Verilog Global Set/Reset

At the beginning of an FPGA design simulation, you must toggle the global set/reset signal (GSR in XC4000E/L/X/XL/XV designs), or the GR global reset signal in XC5200, XC3000A/L, or XC3100A/L designs. Toggling the global set/reset emulates the power-on reset of the FPGA. If you do not do this, the flip-flops and latches in your simulation may not function correctly.

The global set/reset net is present in your implemented design even if you do not instantiate the STARTUP block in your design. The function of STARTUP is to give you the option to control the global reset net from an external pin.

If you want to select the global set/reset pulse width so that it reflects the actual amount of time it takes for the chip to go through the reset process when power is supplied to it, refer to *The Programmable Logic Data Book* for the device you are simulating.

The general procedure for specifying global set/reset or global reset during a pre-NGDBuild Verilog UNISIM simulation involves defining the global reset signals with one of the following Verilog macros: GSR_SIGNAL or GR_SIGNAL. This is necessary because these global nets do not exist in the UNISIM libraries, and as a result, the reset of the UNISIM components is controlled by the detection of the GSR_SIGNAL or GR_SIGNAL macros. In addition, you must declare the global set/reset signal either as a Verilog wire or reg. Your choice of wire or reg depends on whether your design contains a STARTUP component.

Note: In the Xilinx software, the Verilog UNISIM library is only used in RTL simulations of your designs. Simulation at other points in the flow use the Verilog SIMPRIM Libraries. This occurs because the FPGA Compiler writes out unexpanded LogiBLOX type modules

that do not have corresponding models. However, there are models for VHDL that VSS can use.

For pre-NGDBuild UNISIM functional simulation, you must set the value of the appropriate Verilog macro (GSR_SIGNAL or GR_SIGNAL) to the name of the GSR or GR net, qualified by the appropriate scope identifiers.

Note: GSR_SIGNAL and GR_SIGNAL are used in the Verilog UNISIM to emulate the global reset signals.

The scope identifiers are a combination of the test module scope and the design instance scope. The scope qualifiers are required because the scope information is needed when the GSR_SIGNAL and GR_SIGNAL macros are interpreted by the Verilog UNISIM simulation models to emulate a global reset signal.

The net name you specify, and whether you specify the net as a Verilog reg or a wire, depends on whether your design includes an instantiated STARTBUF.

Defining GSR in a Test Bench

Use the following steps to define the global set/reset signals in a test bench for your design.

Note: Use the first step if you do not have a STARTBUF in your design, otherwise proceed to the second step.

1. If you do not have a STARTBUF in your design, Xilinx recommends naming the global set/reset net `test.design_instance.GSR` or `test.design_instance.GR` (Verilog is case-sensitive), and the signal should be declared as a Verilog reg data type.

Note: Test refers to the test fixture module name and `design_instance` refers to the designated instance name for the instantiated design netlist within the test bench.

2. If there is a STARTBUF block in your design, and the GSR pin is connected to a net, you should set the value of GSR_SIGNAL to the net connected to the GSR pin on the STARTUP symbol.

The signal you toggle at the beginning of the simulation is the port or signal in your design that is used to control global set/reset. This is usually an external input port in the Verilog netlist,

but it may also be a wire if global reset is controlled by logic internal to your design.

3. When invoking Verilog-XL or Modelsim to run the simulation, Xilinx recommends that you specify the test bench file before the Verilog netlist for your design for the simulation to work properly, as in the following examples.

- Cadence Verilog-XL

For RTL simulation, enter the following.

```
verilog -y $XILINX/verilog/src/UNI4000X
design.stim design.v
```

The path specified with the `-y` switch points the simulator to the UNISIM models and is only necessary if Xilinx primitives are instantiated in your code. When targeting a device family other than the XC4000X family, change the UNI4000X reference in the path to the targeted device family.

For post-implementation simulation, enter the following.

```
verilog design.stim time_sim.v
```

In this example, the same test fixture file is declared first followed by the simulation netlist created by the Xilinx tools. The name of the Xilinx simulation netlist may change depending on how the file was created. It is also assumed that the `-ul` switch was specified during NGD2VER to specify the location of the SIMPRIM libraries.

- Model Technology Modelsim

Because Modelsim is a compiled Verilog simulator, modules must be compiled from the bottom up. This is the opposite of the Verilog-XL recommendation.

For RTL simulation, enter the following.

```
vlog design.v
vlog design.stim
vsim -L uni4000x testbench_module_name
```

This example targets the XC4000X family and assumes the UNI4000X libraries are properly compiled and named uni4000x. For more information on the compilation of the

Modelsim libraries, refer to <http://www.xilinx.com/tech-docs/1923.htm>

For post-implementation simulation, enter the following.

```
vlog time_sim.v
vlog design.stim
vsim -L simprim testbench_module_name
```

This example is based on targeting the SIMPRIM libraries, which have been properly compiled and named simprim. Also, the name of the simulation netlist may change depending on how the file is created.

4. Xilinx recommends giving the name *test* to the main module in the test bench file. This name is consistent with the name of the test bench module that is written later in the design flow by NGD2VER during post-NGDBuild, post-MAP, or post-route simulation. If this naming consistency is maintained, you can use the same test bench file for simulation at all stages of the design flow with minimal modification

For Unified Library functional simulation, you must always define the appropriate Verilog macro (GSR_SIGNAL or GR_SIGNAL) for the global set/reset signal. (This macro is not used in timing simulation when there is a STARTUP block in your design.)

The GSR signal in XC4000E/L/X/XL/XV devices and the GR signal in XC5200 devices is active High, and the GR signal (XC3000A/L and XC3100A/L designs) signal is active Low.

For post-NGDBuild and post-route timing simulation, the test bench template (.tv file) produced by running NGD2VER with the `-tf` option contains most of the code previously described for defining and toggling GSR or GR. However, if you use a signal to control the STARTUP block, you must manually edit the test bench template file (generated by NGD2VER) to specify the signal connected to the GSR or GR pin on the STARTUP block symbol as GSR_SIGNAL (4K) or GR_SIGNAL (5K).

Designs without a STARTUP Block

If you do not have a STARTUP block in your design, you can use the same test bench file with little or no modification at all stages of the design flow, as described in the following examples.

Example 1: XC4000E/L/X/XL/XV and Spartan RTL Functional Simulation (No STARTUP Block)

The following design shows how to drive the GSR signal in a Verilog-XL test bench file at the beginning of a pre-NGDBuild Unified Library functional simulation.

Note: The terms “test bench” and “test fixture” are used synonymously throughout this manual.

You should reference the global set/reset net as GSR in XC4000E/L/X/XL/XV and Spartan designs without a STARTUP block. The Verilog macro defining the global net must be referenced as GSR_SIGNAL because this is how it is modeled in the Verilog UNISIM library.

In the design code, declare GSR as a Verilog wire; however, it is not specified in the port list for the module. Describe GSR to reset or preset every inferred register or latch in your design. GSR does not need to be connected to any instantiated registers or latches, as shown in the following example.

```
module my_counter (CLOCK, COUT, Q, D);

    input  CLOCK, D;

    output Q;
    output [3:0] COUT;

    wire GSR;

    always @(posedge GSR or posedge CLOCK)
        begin
            if (GSR == 1'b1)
                COUT = 4'h0;
            else
                COUT = COUT + 1'b1;
            end

    // Example of an instantiated FDCE
```

```
//  
// If a macro name GSR_SIGNAL is defined in the  
// testbench, CLR does not need to be connected  
// to GSR and flop will still be reset with GSR.  
  
FDCE test_flop (.Q(Q), .D (D), .C(CLOCK),  
               .CE (open), .CLR (open));  
endmodule
```

Because GSR is declared as a floating wire and is not in the port list, the synthesis tool optimizes the GSR signal out of the design. GSR is replaced later by the implementation software for all post-implementation simulation netlists.

In the test fixture file, set a GSR_SIGNAL macro to test.my_counter.GSR (the name of the global set/reset signal, qualified by the name of the design instantiation instance name and the test fixture module name) using the 'define compiler directive, as follows.

```
`define GSR_SIGNAL test.my_counter.GSR
```

GSR_SIGNAL should be toggled high, then low at the beginning of an initial block using the force command.

```
module test;  
  
    `define GSR_SIGNAL test.my_counter.GSR  
  
    count4 my_counter (.CLOCK (CLOCK), .COUT(COUT), .Q(Q), .D(D));  
  
    reg CLOCK, D;  
  
    initial  
    begin  
  
        CLOCK = 0;  
        D = 0;  
  
        force `GSR_SIGNAL = 1; //Here is the Global Reset  
        #100 force `GSR_SIGNAL = 0; //End of Global Reset  
  
        //rest of simulation stimuli  
  
    end  
endmodule
```

In this example, the active high GSR signal in the XC4000 family device is activated by driving it high. 100ns later, it is deactivated by driving it low. (100ns is an arbitrarily chosen value.)

You can use the same test fixture for simulating at other stages in the design flow if this methodology is used.

Example 2: XC5200 RTL Functional Simulation (No STARTUP Block)

For pre-NGDBuild functional simulation, the active High GR net in XC5200 devices should be simulated in the same manner as GSR for XC4000E/L/X/XL/XV and Spartan.

In the design code, GR should be declared as a Verilog wire and not specified in the port list for the module. GR should be described to reset every inferred register or latch in the design. GR does not need to be connected to any instantiated registers or latches.

```

module my_counter (CLOCK, COUT, Q, D);

    input  CLOCK, D;

    output Q;
    output [3:0] COUT;

    wire GR;

    always @(posedge GR or posedge CLOCK)
        begin
            if (GR == 1'b1)
                COUT = 4'h0;
            else
                COUT = COUT + 1'b1;
            end

    // Example of an instantiated FDCE
    //
    // If a macro name GR_SIGNAL is defined in the
    // testbench, CLR does not need to be connected
    // to GR and flop will still be reset with GR.

    FDCE test_flop (.Q(Q), .D (D), .C(CLOCK),
                   .CE (open), .CLR (open));
endmodule

```

In the test fixture file, set a macro called GR_SIGNAL to test.my_counter.GR (the name of the global set/reset signal, qualified by the name of the design instantiation instance name and the test fixture module name) using the 'define compiler directive, as follows.

```
`define GR_SIGNAL test.my_counter.GR
```

GR_SIGNAL should be toggled high, then low at the beginning of an initial block using the force command.

```
module test;

    `define GR_SIGNAL test.my_counter.GR

    count4 my_counter (.CLOCK (CLOCK), .COUT(COUT), .Q(Q), .D(D));

    reg CLOCK, D;

    initial
    begin

        CLOCK = 0;
        D = 0;

        force `GR_SIGNAL = 1; //Here is the Global Reset

        #100 force `GR_SIGNAL = 0; //End of Global Reset

        //rest of simulation stimuli

    end

endmodule
```

In this example, the active high GR signal in the XC5200 family device is activated by driving it high. 100ns later, it is deactivated by driving it low. (100ns is an arbitrarily chosen value.)

You can use the same test fixture for simulating at other stages of the design if this methodology is used.

Example 3: XC3000A/L and XC3100A/L RTL and Post-synthesis Functional Simulation (No STARTUP Block)

Asserting global reset in XC3000A/L and XC3100A/L designs is almost identical to the procedure for asserting global reset in XC5200

designs, except that GR in XC3000A/L and XC3100A/L devices is active Low. (Also note that the STARTUP block is not supported on XC3000A/L and XC3100A/L devices):

```
module test;
reg GR;
`define GR_SIGNAL test.GR;
initial
begin
    `GR_SIGNAL = 0;    // reset the device
    #100 `GR_SIGNAL = 1;
```

Note: The Global Reset (GR) signal in the XC3000A/L architecture is modeled differently in functional simulation netlists and SIMPRIM library-based netlists generated by NGD2VER. In the Verilog Unified Library, GR is modeled as a wire within a global module, while in a SIMPRIM-based netlist, it is always modeled as an external port. As a result, you cannot use the same test bench file for both Unified library simulation and SIMPRIM-based simulation.

Designs with a STARTUP block (XC4000E/L/X/XL/XV, Spartan, and XC5200 Devices Only)

Asserting global set/reset when the STARTUP block is specified in your design is similar to asserting global set/reset without a STARTUP block in your design. However, there are the following two differences.

- The ``define` statement must now specify the name of the net attached to the GSR pin (XC4000E/L/X/XL/XV, Spartan devices) or GR pin (XC5200 devices) on the STARTUP block.


```
`define GSR_SIGNAL net_connected_to_GSR_pin
```
- The signal you toggle is now the external input port that controls the “net_connected_to_GSR_pin” (or “net_connected_to_GR_pin”) on the STARTUP block. If the global reset signal is inverted, as shown in the following figure, it appears in your Verilog netlist as an input port, and can be driven as follows.

```
initial
begin
    GSR_user_control_signal = 0;
    #100 GSR_user_control_signal = 1;
```

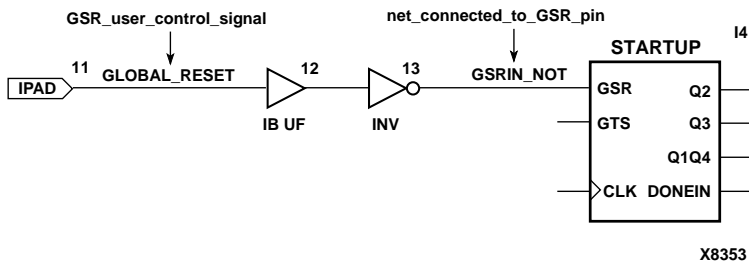


Figure 5-7 Verilog User-Controlled Inverted GSR

Example 1a: XC4000E/L/X/XL/XV and Spartan: RTL and Post-synthesis Simulation (With STARTUP)

The following is an example of driving the global set/reset signal in a test fixture file at the beginning of an RTL or post-synthesis functional simulation when there is a STARTUP block in an XC4000E/L/X/XL/XV or Spartan design.

In the following figure, the mygsr signal is the GSR_user_control_signal. In this case, mygsr is an external user signal that controls GSR. Mygsr sources an IBUF, which in turn sources the gsrin signal. Gsrin represents the net_connected_to_GSR_pin pin that directly sources the GSR pin of the STARTUP block.




Figure 5-8 Verilog User-Controlled GSR

This design allows you to control global set/reset in the device by driving the external mygsr input port. In the test fixture file, mygsr is a Verilog reg in the test module:

```
module test;
    reg mygsr;
```

In addition, for XC4000E/L/X/XL/XV and Spartan designs, a Verilog macro called GSR_SIGNAL must be declared to make the connection between the user logic and the global GSR net embedded in the Unified Library models. This is done by using a ``define` directive to set GSR_SIGNAL to the following.

```
test_module_name.design_instance_name.gsr_pin_signal
```

Gsr_pin_signal corresponds to the name of the signal connected to the GSR pin on the STARTUP block (in this case, gsrin). The scope qualifier in this case also includes the name of the design instance (uut) in anticipation that the net appears as an internal net of the design in the post-NGDBuild, post-Map, and post-route simulations further in the design flow.

The global set/reset control signal should be toggled High, then Low in an initial block.

```
module test;
    reg mygsr;
    `define GSR_SIGNAL test.uut.gsrin;

    initial
    begin
        mygsr = 1;           // reset the device
        #100 mygsr = 0;
```

Example 1b: Post-NGDBuild Functional, Post-Map Timing, and Post-Route Timing Simulation (With STARTUP)

For post-NGDBuild functional simulation, post-Map timing simulation, and post-route timing simulation, the procedure is identical to Unified Library functional simulation, except that you must omit the ``define` statement for GSR_SIGNAL. This is done because the net connections exist in the post-NGDBuild design, and retaining the macro definition causes a possible conflict with these connections. In the following example, the macro definition is commented out to avoid a possible conflict.

```

module test;
  reg mygsr;
  // `define GSR_SIGNAL test.uut.gsrin;
initial
begin
  mygsr = 1;          // reset the device
  #100 mygsr = 0;

```

Example 2a: XC5200: RTL or Post-synthesis Functional Simulation Designs with STARTUP Block

For an XC5200 design with a STARTUP block, the net controlling GR should be simulated in the same manner as for the XC4000E/L/X/XL/XV or Spartan.

Substitute GR_SIGNAL for GSR_SIGNAL, mygr for mygsr, and gr_in for gsr_in in Example 1 to obtain the test fixture fragment for simulating GR in a Verilog RTL or post-synthesis simulation.




Figure 5-9 Verilog User-Controlled Inverted GR

```

module test;
  reg mygr;
  `define GR_SIGNAL test.uut.gr_in;
initial
begin
  mygr = 1;          // reset the device
  #100 mygr = 0;

```

Example 2b: Post-NGDBuild Functional, Post-Map Timing, and Post-Route Timing Simulation (With STARTUP Block)

For post-NGDBuild functional simulation, post-Map timing simulation, and post-route timing simulation, the procedure is identical to Unified Library functional simulation, except that you must omit the ``define` statement for `GR_SIGNAL`. This is done because the net connections exist in the post-NGDBuild design, and retaining the macro definition may cause a conflict with these connections. In the following example the Verilog macro definition is commented out to avoid a possible conflict.

```
module test;
    reg mygr;
    // `define GR_SIGNAL test.uut.gr_in;
initial
begin
    mygr = 1;          // reset the device
    #100 mygr = 0;
```

Example 3: XC3000A/L and XC3100A/L designs

STARTUP is not supported or required in XC3000A/L and XC3100A/L designs. Follow the procedure for XC3000A/L and XC3100A/L designs without STARTUP blocks.

Setting Verilog Global Tristate (XC4000, Spartan, and XC5200 Outputs Only)

XC4000E/L/X/XL/XV, Spartan, and XC5200 devices also have a global control signal (GTS) that tristates all output pins. This allows you to isolate the actual device part during board level testing. You can also tristate the FPGA device outputs during board level simulation to assist in debugging simulation. In most cases, GTS is deactivated so that the outputs are active.

Although the STARTUP component also gives you the option of controlling the global tristate net from an external pin, usually it is used for controlling global reset. In this case, the GTS pin can be left unconnected in the design entry phase, and it will float to its inactive state level. The global tristate net, GTS, is implemented in designs even if a STARTUP block is not instantiated. You can deactivate GTS

by driving it Low in your test fixture file, or by connecting the GTS pin to GND in your input design.

Specifying GTS

The general procedure for specifying GTS is similar to that used for specifying the global set/reset signals, GSR and GR. You define the global tristate signal with the Verilog macro, GTS_SIGNAL. You must declare the global tristate signal either as a Verilog wire or reg. If you do not want to specify GTS for simulation, you do not need to change anything in your design or test bench.

The net name you select, and whether you specify the net as a Verilog reg or a wire, depends on whether you have a STARTUP block instantiated in your design, and if you have a signal connected to the STARTUP GTS pin.

If a STARTUP block is not in your design, Xilinx recommends that you name the global tristate wire `testbench_module.design_instance.GTS`, and declare the signal as a Verilog wire data type in your design netlist.

If there is a STARTUP block in your design and the GTS pin is connected to a net, the value of GTS_SIGNAL should be set to the name of the net connected to the GTS pin on the STARTUP symbol. The signal you toggle at the beginning of simulation is the port or signal in your design that is used to control global tristate. This is usually an external input port in the Verilog netlist, but can be a wire if global tristate is controlled by internal logic in your design.

Xilinx recommends that you name the main module in your test bench file `test` to be consistent with the name of the test fixture module that is written later in the design flow by NGD2VER. If this naming consistency is maintained, you can use the same test fixture file for simulation at all stages in the design flow with minimal modification.

The GTS signal in XC4000E/L/X/XL/XV, Spartan, and XC5200 devices is active High. This macro is not used in timing simulation when there is a STARTUP block in your design and the GTS pin is connected.

For post-NGDBuild and post-route timing simulation, the test bench template (.tv file) produced by NGD2VER with the `-tf` option contains most of the code described previously for defining and

driving GTS. However, in cases where you have a signal controlling the STARTUP block, you must manually edit the test fixture template file generated by NGD2VER to specify the control signal for GTS.

Designs without a STARTUP Block

When you do not have a STARTUP block in your design, you can use the same test fixture file with little or no modification if you use the guidelines in the following example.

XC4000E/L/X/XL/XV, Spartan, and XC5200 RTL Functional Simulation (No STARTUP Block)

The following is an example of how you can drive the GTS signal in a test fixture file at the beginning of a pre-NGDBuild RTL or post-synthesis functional simulation. The global tristate net is named GTS in XC4000E/L/X/XL/XV, Spartan, or XC5200 designs when there is not a STARTUP block. The Verilog macro defining the global net must be named GTS_SIGNAL because this is the name of the predefined macro used to model the global tristate signal in the Xilinx Verilog UNISIM simulation models. Use the following guidelines.

- In your design, declare GTS as a Verilog wire, as follows.

```
module my_design;
  wire GTS;
```

- Set a macro named GTS_SIGNAL to *testbench_module.design_instance.GTS* (the name of the global tristate signal, qualified by the name of the instantiated design instance name and test fixture module), using the ``define` compiler directive, as follows.

```
`define GTS_SIGNAL test.GTS;
```

GTS should be driven Low in an initial block, as follows.

```
module test;
  `define GTS_SIGNAL test.uut.GTS;
  initial
  begin
    `GTS_SIGNAL = 0;
```

In this example, the active High GTS signal is deactivated by driving it Low to activate the outputs of the design.

Designs with a STARTUP block (XC4000E/L/X/XL/XV, Spartan, and XC5200 Devices Only)

Asserting global tristate when the STARTUP block is specified in your design is similar to asserting global tristate without a STARTUP block. There are two primary differences, as follows.

- If the GTS pin on the STARTUP block is connected, the ``define` statement must now set `GTS_SIGNAL` to the name of the net attached to the GTS pin on the STARTUP block, as follows.

```
`define GTS_SIGNAL net_connected_to_GTS_pin
```

- If the GTS pin on the STARTUP block is connected, the signal you drive is now either the external input port or internal signal that controls the `net_connected_to_GTS_pin` on the STARTUP block. If it is an external input, it is in your Verilog netlist as an input port. To tristate your outputs, drive this signal High, and to activate your outputs, drive it Low, as shown here.

```
initial
begin
  GTS_user_control_signal = 1;
  #100 GTS_user_control_signal = 0;
```

Example 1a: XC4000E/L/X/XL/XV, Spartan, and XC5200: RTL or Post-Synthesis Functional Simulation (With STARTUP, GTS Pin Connected)

In the following figure, the design contains a STARTUP block, and the GTS pin on STARTUP is connected to an external input named `mygts`.

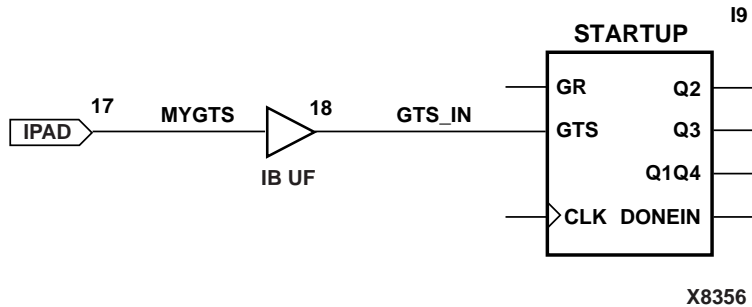


Figure 5-10 Verilog User-Controlled Inverted GTS

The external input, mygts, is declared as a Verilog register. A ``define` directive setting `GTS_SIGNAL` to the name of the net connected to the GTS pin is required to connect the user logic to the global GTS model in the UNISIM simulation models for output buffers (OBUF, OBUFT, and so on). The following is an example of a test fixture.

```

module test;
  reg mygts;
  `define GTS_SIGNAL test.uut.gts_in;
  .
  .
  .
  initial
  begin
    mygts = 1; // if you wish to tristate the
              // device;
    #100 mygts = 0; // deactivate GTS
  end

```

Example 1b: Post-NGDBuild Simulation of GTS (With STARTUP, GTS Pin connected)

For post-route timing simulation, the procedure is similar, except you must omit the `define` statement for `GTS_SIGNAL` because it conflicts with the GTS net driver.

```

module test;
  reg mygts;
  // `define GTS_SIGNAL test.uut.gtsin
  initial

```

```
begin
    mygts = 1; // if you wish to tristate the
              // device;
    #100 mygts = 0; // deactivate GTS
```

Example 2a: XC4000E/L/X/XL/XV, Spartan, and XC5200: Unified Library Simulation (With STARTUP, GTS Pin not connected)

For Unified Library functional simulation, define a wire named GTS, and set the GTS_SIGNAL macro to test.GTS. Toggle GTS_SIGNAL as shown in the following example.

```
module test;
    wire GTS;
    `define GTS_SIGNAL test.GTS
initial
    begin
        force `GTS_SIGNAL = 1; // if you wish to
                              // tristate the
                              // device;
        #100 force `GTS_SIGNAL = 0; // deactivate GTS
```

Example 2b: Post-NGDBuild Simulation of GTS (With STARTUP, GTS Pin not connected)

For post-NGDBuild functional simulation, the actual net exists and must be further qualified by the design instance scope, uut, as shown here.

```
module test;
    // wire GTS;
    // `define GTS_SIGNAL test.GTS
    `define GTS_SIGNAL test.uut.GTS
initial
    begin
        force `GTS_SIGNAL = 1; // if you wish to
                              // tristate the
                              // device;
        #100 force `GTS_SIGNAL = 0; // deactivate GTS
```

Note: For post-route timing simulation, you can use the same test bench.

Using Verilog Oscillators (OSC, OSC4, OSC5)

The OSC (X3000A/L), OSC4 (XC4000E/L/X/XL/XV, Spartan) and OSC5 (XC5200) oscillator components do not have Verilog simulation models associated with them. For OSC, the clock signal frequency is derived from an external crystal-controlled oscillator. The OSC4 and OSC5 are internal oscillators, and are useful in applications in which timing is not critical.

To simulate these oscillators, you must reference the net attached to the output of the oscillator component. For example, for an oscillator output net named `osclk` attached to an oscillator symbol (OSC, OSC4, or OSC5) with a timescale unit of `1ns`, use the Always block to emulate an oscillator with a 10 Mhz clock frequency. Toggle this net at the desired frequency in your Verilog test fixture using the Force command, as shown in the following example.

```
force osclk = 1`b0;
always #100 force osclk = ~osclk;
```


Accelerate FPGA Macros with One-Hot Approach

Steven K. Knapp

Xilinx Inc.
2100 Logic Dr.
San Jose, CA 95124

Reprinted with permission from *Electronic Design*, September 13, 1990. © Penton Publications.

State machines - one of the most commonly implemented functions with programmable logic - are employed in various digital applications, particularly controllers. However, the limited number of flip-flops and the wide combinatorial logic of a PAL device favors state machines that are based on a highly encoded state sequence. For example, each state within a 16-state machine would be encoded using four flip-flops as the binary values between 0000 and 1111.

A more flexible scheme - called one-hot encoding (OHE) - employs one flip-flop per state for building state machines. Although it can be used with PAL-type programmable-logic devices (PLDs), OHE is better suited for use with the fan-in limited and flip-flop-rich architectures of the higher-gate-count field-programmable gate arrays (FPGAs), such as offered by Xilinx, Actel, and others. This is because OHE requires a larger number of flip-flops. It offers a simple and easy-to-use method of generating performance-optimized state-machine designs because there are few levels of logic between flip-flops.

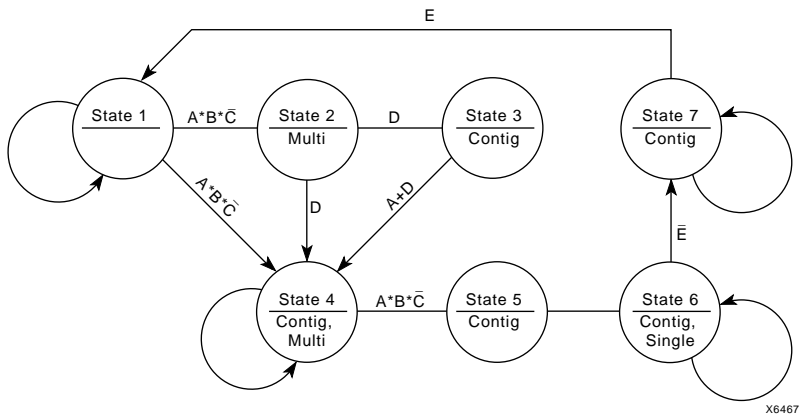


Figure A-1 A TYPICAL STATE MACHINE BUBBLE

(In reference to the figure above) A Typical State Machine Bubble diagram shows the operation of a seven-state state machine that reacts to inputs A through E as well as previous-state conditions.

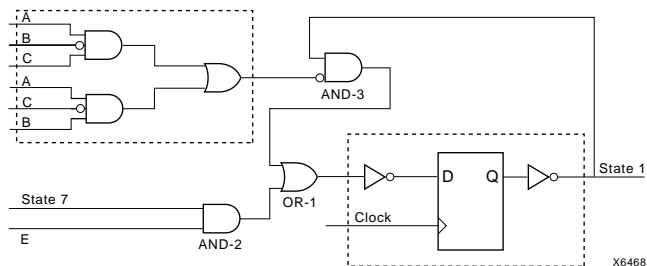


Figure A-2 INVERTERS

(In reference to the figure above) Inverters are required at the D input and the Q output of the state flip-flop to ensure that it powers on in the proper state. Combinatorial logic decodes the operations based on the input conditions and the state feedback signals. The flip-flop will remain in State 1 as long as the conditional paths out of the states are not valid.

A state machine implemented with a highly encoded state sequence will generally have many, wide-input logic functions to interpret the inputs and decode the states. Furthermore, incorporating a highly encoded state machine in an FPGA requires several levels of logic

between clock edges because multiple logic blocks will be needed for decoding the states. A better way to implement state machines in FPGAs is to match the state-machine architecture to the device architecture.

LIMITING FAN-IN

A good state-machine approach for FPGAs limits the amount of fan-in into one logic block. While the one-hot method is best for most FPGA applications, binary encoding is still more efficient in certain cases, such as for small state machines. It's up to the designer to evaluate all approaches before settling on one for a particular application.

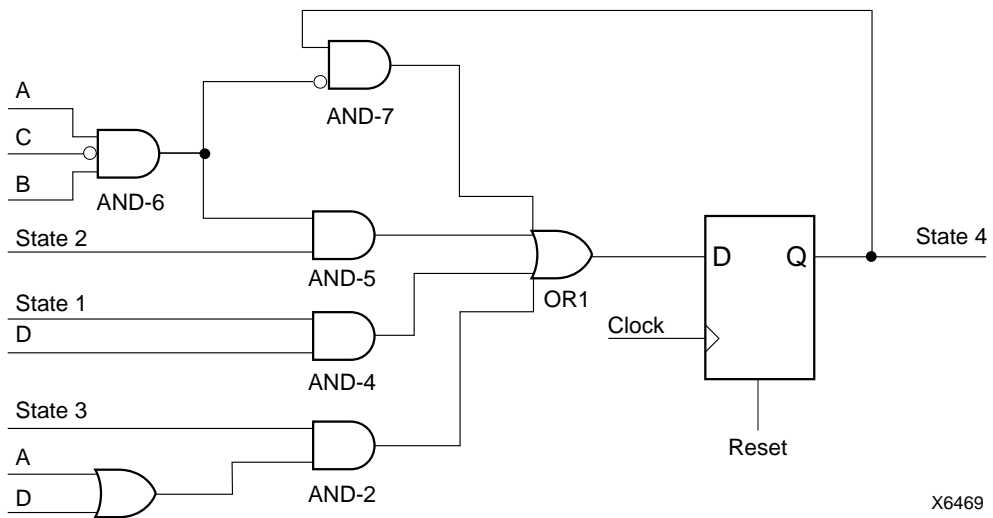


Figure A-3 THE SEVEN STATES

(In reference to the figure above) Of the seven states, the state-transition logic required for State 4 is the most complex, requiring inputs from three other state outputs as well as four of the five condition signals (A - D).

FPGAs are high-density programmable chips that contain a large array of user-configurable logic blocks surrounded by user-programmable interconnects. Generally, the logic blocks in an FPGA have a limited number of inputs. The logic block in the Xilinx XC-3000 series, for instance, can implement any function of five or less inputs. In contrast, a PAL macrocell is fed by each input to the chip and all of

the flip-flops. This difference in logic structure between PALs and FPGAs is important for functions with many inputs: where a PAL could implement a many-input logic function in one level of logic, an FPGA might require multiple logic layers due to the limited number of inputs.

The OHE scheme is named so because only one state flip-flop is asserted, or “hot”, at a time. Using the one-hot encoding method for FPGAs was originally conceived by High-Gate Design - a Saratoga, Calif.-based consulting firm specializing in FPGA designs.

The OHE state machine's basic structure is simple - first assign an individual flip-flop to each state, and then permit only one state to be active at any time. A state machine with 16 states would require 16 flip-flops using the OHE approach; a highly encoded state machine would need just four flip-flops. At first glance, OHE may seem counter-intuitive. For designers accustomed to using PLDs, more flip-flops typically indicates either using a larger PLD or even multiple devices.

In an FPGA, however, OHE yields a state machine that generally requires fewer resources and has higher performance than a binary-encoded implementation. OHE has definite advantages for FPGA designs because it exploits the strengths of the FPGA architecture. It usually requires two or less levels of logic between clock edges than binary encoding. That translates into faster operation. Logic circuits are also simplified because OHE removes much of the state-decoding logic - a one-hot-encoded state machine is already fully decoded.

OHE requires only one input to decode a state, making the next-state logic simple and well-suited to the limited fan-in architecture of FPGAs. In addition, the resulting collection of flip-flops is similar to a shift-register-like structure, which can be placed and routed efficiently inside an FPGA device. The speed of an OHE state machine remains fairly constant even as the number of states grows. In contrast, a highly encoded state machine's performance drops as the states grow because of the wider and deeper decoding logic that's required.

To build the next-state logic for OHE state machine is simple, lending itself to a “cookbook” approach. At first glance, designers familiar with PAL-type devices may be concerned by the number of potential illegal states due to the sparse state encoding. This issue, to be discussed later, can be solved easily.

A typical, simple state machine might contain seven distinct states that can be described with the commonly used circle-and-arc bubble diagrams, see the “A TYPICAL STATE MACHINE BUBBLE” figure. The label above the line in each “bubble” is the state’s name. The labels below the line are the outputs asserted while the state is active. In the example, there are seven states labeled State 1-7. The “arcs” that feed back into the same state are the default paths. These will be true only if no other conditional paths are true.

Each conditional path is labeled with the appropriate logical condition that must exist before moving to the next state. All of the logic inputs are labeled as variables A through E. The outputs from the state machine are called Single, Multi, and Contig. For this example, State 1, which must be asserted at power-on, has a double-inverted flip-flop structure (shaded region of the “INVERTERS” figure)

The state machine in the example was built twice, once using OHE and again with the highly encoded approach employed in most PAL designs. A Xilinx XC3020-100 2000-gate FPGA was the target for both implementations. Though the OHE circuit required slightly more logic than the highly-encoded state machine, the one-hot state machine operated 17% faster (see the table). Intuitively, the one-hot method might seem to employ many more logic blocks than the highly encoded approach. But the highly encoded state machine needs more combinatorial logic to decode the encoded state values.

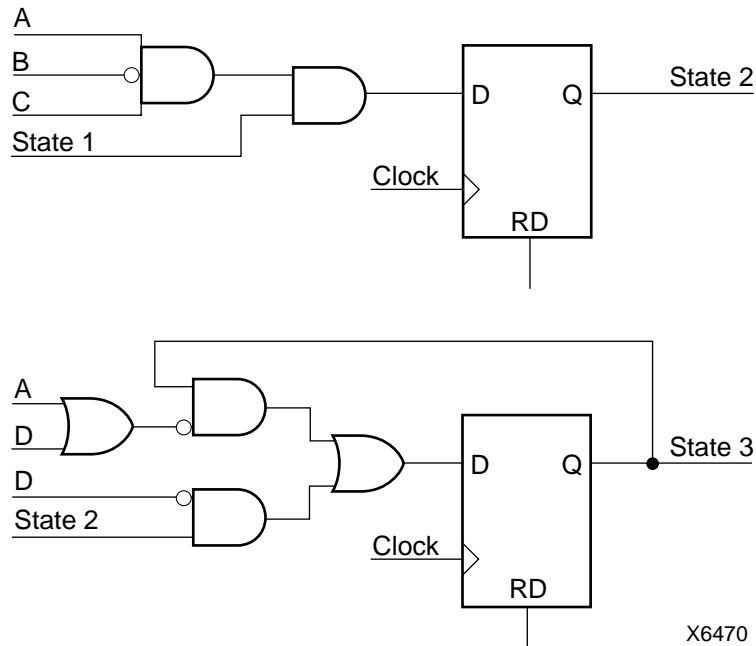


Figure A-4 ONLY A FEW GATES

(In reference to the figure above) Only a few gates are required by States 2 and 3 to form simple state-transition logic decoding. Just two gates are needed by State 2 (top), while four simple gates are used by State 3 (bottom).

The OHE approach produces a state machine with a shift-register structure that almost always outperforms a highly encoded state machine in FPGAs. The one-state design had only two layers of logic between flip-flops, while the highly encoded design had three. For other applications, the results can be far more dramatic. In many cases, the one-hot method yields a state machine with one layer of logic between clock edges. With one layer of logic, a one-hot state machine can operate at 50 to 60 MHz.




Figure A-5 LOOKING NEARLY THE SAME

(In reference to the figure above) Looking nearly the same as a simple shift register, the logic for States 5, 6, and 7 is very simple. This is because the OHE scheme eliminates almost all decoding logic that precedes each flip-flop.

The initial or power-on condition in a state machine must be examined carefully. At power-on, a state machine should always enter an initial, known state. For the Xilinx FPGA family, all flip-flops are reset at power-on automatically. To assert an initial state at power-on, the output from the initial-state flip-flop is inverted. To maintain logical consistency, the input to flip-flop also is inverted.

All other states use a standard, D-type flip-flop with an asynchronous reset input. The purpose of the asynchronous reset input will be discussed later when illegal states are covered.

Once the start-up conditions are set up, the next-state transition logic can be configured. To do that, first examine an individual state. Then count the number of conditional paths leading into the state and add an extra path if the default condition is to remain in the same state. Second, build an OR-gate with the number of inputs equal to the number of conditional paths that were determined in the first step.

Third, for each input of the OR-gate, build an AND-gate of the previous state and its conditional logic. Finally, if the default should remain in the same state, build an AND-gate of the present state and the inverse of all possible conditional paths leaving the present state.

To determine the number of conditional paths feeding State 1, examine the state diagram - State 1 has one path from State 7 whenever the variable E is true. Another path is the default condition, which stays in State 1. As a result, there are two conditional paths feeding State 1. Next, build a 2-input OR-gate - one input for the conditional path from State 7, the other for the default path to stay in State 1 (shown as OR-1 in the "INVERTERS" figure).

The next step is to build the conditional logic feeding the OR-gate. Each input into the OR-gate is the logical AND of the previous state and its conditional logic feeding into State 1. State 7, for example, feeds State 1 whenever E is true and is implemented using the gate called AND-2, in the “INVERTERS” figure. The second input into the OR-gate is the default transition that’s to remain in State 1. In other words, if the current state is State 1, and no conditional paths leaving State 1 are valid, then the state machine should remain in State 1. Note in the state diagram that two conditional paths are leaving State 1, in the “A TYPICAL STATE MACHINE BUBBLE” figure.

The first path is valid whenever $(A*B*C)$ is true, which leads into State 2. The second path is valid whenever $(A*B*C)$ is true, leading into State 4. To build the default logic, State 1 is ANDed with the inverse of all the conditional paths leaving State 1. The logic to perform this function is implemented in the gate labeled AND-3 and the logic elements that feed into the inverting input of AND-3, in the “INVERTERS” figure.




Figure A-6 S-R FLIP-FLOPS

(In reference to the figure above) S-R Flip-Flops offer another approach to decoding the Contig output. They can also save logic blocks, especially when an output is asserted for a long sequence of contiguous states.

State 4 is the most complex state in the state-machine example. However, creating the logic for its next-state control follows the same basic method as described earlier. To begin with, State 4 isn't the initial state, so it uses a normal D-type flip-flop without the inverters. It does, however, have an asynchronous reset input, three paths into the state, and a default condition that stays in State 4. Therefore, four-input OR-gate feeds the flip-flop (OR-1 in the "THE SEVEN STATES" figure).

The first conditional path comes from State 3. Following the methods established earlier, an AND of State 3 and the conditional logic, which is A Ored with D, must be implemented (AND-2 and OR-3 in the "THE SEVEN STATES" figure). The next conditional path is from State 2, which requires an AND of State 2 and variable D (AND-4 in the "THE SEVEN STATES" figure). Lastly, the final conditional path leading into State 4 is from State 1. Again, the State-1 output must be ANDed with its conditional path logic - the logical product, $A*B*C$ (AND-5 and AND-6 in the "THE SEVEN STATES" figure).

Now, all that must be done is to build the logic that remains in State 4 when none of the conditional paths away from State 4 are true. The path leading away from State 4 is valid whenever the product, $A*B*C$, is true. Consequently, State 4 must be ANDed with the inverse of the product, $A*B*C$. In other words, "keep loading the flip-flop with a high until a valid transfer to the next state occurs." The default path logic uses AND-7 and shares the output of AND-6.

| ONE-STATE VS. BINARY ENCODING METHODS | | |
|--|-------------------------------|-------------------------------|
| Method | Number of Logic Blocks | Worst-case performance |
| One-hot | 7.5 | 40 Mhz |
| Binary encoding | 7.0 | 34 Mhz |

Configuring the logic to handle the remaining states is very simple. State 2, for example, has only one conditional path, which comes from State 1 whenever the product $A*B*C$ is true. However, the state machine will immediately branch in one of two ways from State 2, depending on the value of D. There's no default logic to remain in State 2, the "ONLY A FEW GATES" figure. State 3, like States 1 and 4 has a default state, and combines the A, D, State 2, and State 3 feedback to control the flop-flop's D input in the "ONLY A FEW GATES" figure.

State 5 feeds State 6 unconditionally. Note that the state machine waits until variable E is low in State 6 before proceeding to State 7. Again, while in State 7, the state machine waits for variable E to return to true before moving to State 1 in the “LOOKING NEARLY THE SAME” figure.

OUTPUT DEFINITIONS

After defining all of the state transition logic, the next step is to define the output logic. The three output signals - Single, Multi, and Contig - each fall into one of three primary output types:

1. Outputs asserted during one state, which is the simplest case. The output signal Single, asserted only during State 6, is an example.
2. Outputs asserted during multiple contiguous states. This appears simple at first glance, but a few techniques exist that reduce logic complexity. One example is Contig. It's asserted from State 3 to State 7, even though there's a branch at State 2.
3. Outputs asserted during multiple, non-contiguous states. The best solution is usually brute-force decoding of the active states. One such example is Multi, which is asserted during State 2 and State 4.

OHE makes defining outputs easy. In many cases, the state flip-flop is the output. For example, the Single output also is the flip-flop output for State 6; no additional logic is required. The Contig output is asserted throughout States 3 through 7. Though the paths between these states may vary, the state machine will always traverse from State 2 to a point where Contig is active in either State 3 or State 4.

There are many ways to implement the output logic for the Contig output. The easiest method is to decode States 3, 4, 5, 6, and 7 with a 5-input OR gate. Any time the state machine is in one of these states, Contig will be active. Simple decoding works best for this state machine example. Decoding five states won't exceed the input capability of the FPGA logic block.

ADDITIONAL LOGIC

However, when an output must be asserted over a longer sequence of states (six or more), additional layers of decoding logic would be required. Those additional logic layers reduce the state machine's performance.

Employing S-R flip-flops gives designers another option when decoding outputs over multiple, contiguous states. Though the basic FPGA architecture may not have physical S-R flip-flops, most macro-cell libraries contain one built from logic and D-type flip-flops. Using S-R flip-flops is especially valuable when an output is active for six or more contiguous states.

The S-R flip-flop is set when entering the contiguous states, and reset when leaving. It usually requires extra logic to look at the state just prior to the beginning and ending state. This approach is handy when an output covers multiple, non-contiguous states, assuming there are enough logic savings to justify its use.

In the example, States 3 through 7 can be considered contiguous. Contig is set after leaving State 2 for either States 3 or 4, and is reset after leaving State 7 for State 1. There are no conditional jumps to states where Contig isn't asserted as it traverses from State 3 or 4 to State 7. Otherwise, these states would not be contiguous for the Contig output.

The Contig output logic, built from an S-R flip-flop, will be set with State 2 and reset when leaving State 7 in the "S-R FLIP-FLOPS" figure. As an added benefit, the Contig output is synchronized to the master clock. Obvious logic reduction techniques shouldn't be overlooked either. For example, the Contig output is active in all states except for States 1 and 2. Decoding the states where Contig isn't true, and then asserting the inverse, is another way to specify Contig.

The Multi output is asserted during multiple, non-contiguous states - exclusively during States 2 and 4. Though States 2 and 4 are contiguous in some cases, the state machine may traverse from State 2 to State 4 via State 3, where the Multi output is unasserted. Simple decoding of the active states is generally best for non-contiguous states. If the output is active during multiple, non-contiguous states over long sequences, the S-R flip-flop approach described earlier may be useful.

One common issue in state-machine construction deals with preventing illegal states from corrupting system operation. Illegal states exist in areas where the state machine's functionality is undefined or invalid. For state machines implemented in PAL devices, the state-machine compiler software usually generates logic to prevent or to recover from illegal conditions.

In the OHE approach, an illegal condition will occur whenever two or more states are active simultaneously. By definition, the one-hot method makes it possible for the state machine to be in only one state at a time. The logic must either prevent multiple, simultaneous states or avoid the situation entirely.

Synchronizing all of the state-machine inputs to the master clock signal is one way to prevent illegal states. “Strange” transitions won't occur when an asynchronous input changes too closely to a clock edge. Though extra synchronization would be costly in PAL devices, the flip-flop-rich architecture of an FPGA is ideal.

Even off-chip inputs can be synchronized in the available input flip-flops. And internal signals can be synchronized using the logic block's flip-flops (in the case of the Xilinx LCAs). The extra synchronization logic is free, especially in the Xilinx FPGA family where every block has an optional flip-flop in the logic path.

RESETTING STATE BITS

Resetting the state machine to a legal state, either periodically or when an illegal state is detected, give designers yet another choice. The Reset Direct (RD) inputs to the flip-flops are useful in this case. Because only one state bit should be set at any time, the output of a state can reset other state bits. For example, State 4 can reset State 3.

If the state machine did fall into an illegal condition, eventually State 4 would be asserted, clearing State 3. However, State 4 can't be used to reset State 5, otherwise the state machine won't operate correctly. To be specific, it will never transfer to State 5; it will always be held reset by State 4. Likewise, State 3 can reset State 2, State 5 can reset State 4, etc. - as long as one state doesn't reset a state that it feeds.

This technique guarantees a periodic, valid condition for the state machine with little additional overhead. Notice, however, that State 1 is never reset. If State 1 were “reset”, it would force the output of State 1 high, causing two states to be active simultaneously (which, by definition, is illegal).

Report Files

This appendix includes report files from various synthesis vendors. To reduce the size of this appendix, some of the files are truncated (indicated by a series of dots) where information is repeated.

Synplicity®

Synplicity report files include the following.

- Synthesis information on generated state machines and inserted items, such as clock buffers
- Predicted timing performance, including maximum frequency of worst case paths
- Area usage information for IOBs, carry logic, registers, FMAPs, HMAPs, and CLBs

Note: The report file in this section is for a design compiled with the VHDL compiler. A report file for a design compiled with the Verilog compiler is essentially the same.

```
Content-Type: text/plain; charset="us-ascii"
Content-Disposition: attachment; filename="atm_chip1_ed.log"
```

```
$ Start of Compile
#Mon Jan 12 07:59:16 1998
```

```
Synplify VHDL Compiler, version 3.0b, built Dec 17 1997
Copyright (C) 1994-1997, Synplicity Inc. All Rights Reserved
```

```
VHDL syntax check successful!
```

```
Compiler output is up to date. No re-compile necessary
```

```
Synthesizing work.acc_chip.schematic
```

@N:"e:\customer\atm\lcl_int1.vhd":73:6:73:7|Trying to extract state machine for register wr_state_o

Extracted state machine for register wr_state_o

State machine has 4 reachable states with original encodings of:

11
10
01
00

@N:"e:\customer\atm\lcl_int1.vhd":73:6:73:7|Trying to extract state machine for register rd_state_o

Extracted state machine for register rd_state_o

State machine has 5 reachable states with original encodings of:

00001
00010
00100
01000
10000

Post processing for work.acc_chip.schematic

@END

Process took 0.371 seconds realtime, 0.371 seconds cputime

Synplify Xilinx Technology Mapper, version 3.0b, built Dec 21 1997

Copyright (C) 1994-1997, Synplicity Inc. All Rights Reserved

Setting fanout limit to 100

@N:"e:\customer\atm\token1.vhd":103:8:103:9|Found counter in view:work.TOKEN(vhdl_rtl) inst token_cntr[5:0]

@N:"e:\customer\atm\rx_agen1.vhd":61:6:61:7|Found counter in view:work.RX_ADD_GEN(vhdl_rtl) inst rx_20_add[18:0]

@N:"e:\customer\atm\rx_agen1.vhd":61:6:61:7|Found counter in view:work.RX_ADD_GEN(vhdl_rtl) inst rx_10_add[18:0]

@N:"e:\customer\atm\slw_clk1.vhd":21:10:21:11|Found counter in view:work.SLOW_CLOCKS(vhdl_rtl) inst slow_counter[16:0]

@N:"e:\customer\atm\mngmnt1.vhd":998:4:998:5|Found counter in view:work.MNGMNT(vhdl_rtl) inst tx_frame_cntr[11:0]

@N:"e:\customer\atm\mngmnt1.vhd":931:4:931:5|Found counter in view:work.MNGMNT(vhdl_rtl) inst rx_frame_cntr[11:0]

Clock Buffers:

Inserting Clock buffer for port CLK_20M_SMP_I,TNM=CLK_20M_SMP_I

Inserting Clock buffer for port CPU_WR_N_I,TNM=CPU_WR_N_I

Inserting Clock buffer for port CLK_40M_P_I,TNM=CLK_40M_P_I

Inserting Clock buffer for port CLK_40M_S2_I,TNM=CLK_40M_S2_I

Net buffering Report:

No nets needed buffering.

Timing reports

Delay - This is the delay from a start point
such as a register or primary input.

Slack - If this value is negative then it indicates
the size of the timing violation.

FO - This is the estimated fanout or loading
used in calculating net delays

Requested default timing:

Frequency=15.0 MHz, Period=66.7 ns

Estimated result:

Frequency=16.7 MHz, Period=60.0 ns
Slack on longest path: 6.7 ns

Timing Information for Longest Paths:

Instance: Z_40ACCCHIPZ_32TXD2PZ_41.base0_reg20[12], cell DFFRE
D I -Z_40ACCCHIPZ_32TXD2PZ_41.sync_tx20_0_base0_reg20_4[12]Delay=60.4,
Slack=6.7

Instance: Z_40ACCCHIPZ_32TXD2PZ_41.base0_reg20[18], cell DFFRE
D I -Z_40ACCCHIPZ_32TXD2PZ_41.sync_tx20_0_base0_reg20_4[18]Delay=59.8,
Slack=7.3

Instance: Z_40ACCCHIPZ_32TXD2PZ_41.base0_reg20[15], cell DFFRE
D I -Z_40ACCCHIPZ_32TXD2PZ_41.sync_tx20_0_base0_reg20_4[15]Delay=59.0,
Slack=8.1

Instance: Z_40ACCCHIPZ_32TXD2PZ_41.base3_reg20[15], cell DFFRE
D I - Z_40ACCCHIPZ_32TXD2PZ_41.N_1956Delay=58.6, Slack=8.5

Instance: Z_40ACCCHIPZ_32TXD2PZ_41.base0_reg20[11], cell DFFRE
D I -Z_40ACCCHIPZ_32TXD2PZ_41.sync_tx20_0_base0_reg20_4[11]Delay=58.4,
Slack=8.7

Instance: Z_40ACCCHIPZ_32TXD2PZ_41.base3_reg20[13], cell DFFRE
D I -Z_40ACCCHIPZ_32TXD2PZ_41.sync_tx20_3_base3_reg20_4[13]Delay=58.3,
Slack=8.8

Instance: Z_40ACCCHIPZ_32TXD2PZ_41.base1_reg20[13], cell DFFRE
D I -Z_40ACCCHIPZ_32TXD2PZ_41.sync_tx20_1_base1_reg20_4[13]Delay=58.3,
Slack=8.8

.
. .
. .
. .

Resource Usage Report

Mapping to part: 40125xvbg560-1

I/O primitives:

| | |
|--------------|---------|
| OBUF | 98 uses |
| IBUF | 43 uses |
| OUTFF INIT=R | 14 uses |
| INFF INIT=R | 10 uses |

| | |
|------|--------|
| BUFG | 4 uses |
|------|--------|

Carry primitives used for arithmetic functions:

| | |
|------------|----------|
| INC-FG-1 | 11 uses |
| INC-FG-CI | 121 uses |
| FORCE-0 | 14 uses |
| ADD-FG-CI | 183 uses |
| EXAMINE-CI | 55 uses |
| FORCE-1 | 30 uses |
| ADD-G-F1 | 4 uses |
| SUB-FG-CI | 180 uses |
| DEC-FG-0 | 3 uses |
| DEC-FG-CI | 12 uses |

Register bits not including I/Os: 1258

Logic Mapping Summary:

FMAPs: 3059 of 9248 (34%)

HMAPs: 647 of 4624 (14%)

Total packed CLBs: 1530 of 4624 (34%)

(Packed CLBs is determined by the larger of three quantities:

Registers / 2, HMAPs, or FMAPs / 2.)

Mapper successful!

Process took 169.984 seconds realtime, 169.984 seconds cputime

Exemplar Logic

The area section of the Exemplar Logic report includes area utilization information for FMAPs, HMAPs, CLBs, IO buffers, and IOB registers. The timing section of the report lists the predicted timing of the critical path, and shows the number of levels of logic and buffers.

Summary Report

Content-Type: text/plain; charset="us-ascii"
 Content-Disposition: attachment; filename="ram.sum"

Cell: mem View: behav Library: work

Total accumulated area :

| | |
|-------------------------------------|----|
| Number of FG Function Generators : | 4 |
| Number of Packed CLBs : | 1 |
| Number of IBUF : | 9 |
| Number of OBUF : | 1 |
| Number of IOB Output Flip Flops : | 2 |
| Number of ports : | 10 |
| Number of nets : | 24 |
| Number of instances : | 17 |
| Number of references to this view : | 0 |

| Cell | Library | References | Total Area |
|----------|---------|------------|--------------------------|
| F2_LUT | xi4ex | 2 x 1 | 2 FG Function Generators |
| OFDTX | xi4ex | 2 x 1 | 2 IOB Output Flip Flops |
| OBUF | xi4ex | 1 x 1 | 1 OBUF |
| GND | xi4ex | 1 x 1 | 1 GND |
| IBUF | xi4ex | 9 x 1 | 9 IBUF |
| RAM16x1S | xi4ex | 2 x 1 | 2 FG Function Generators |

Using wire table: 4052ex-3_avg

Slack Table at End Points

| End points | Slack | Arrival | | Required | |
|----------------------------|-------|---------|-------|----------|------|
| | | rise | fall | rise | fall |
| ix328_l5_1_l0_l0_0_l0_l0/D | : n/a | 20.02 | 20.92 | n/a | n/a |
| ix328_l5_0_l0_l0_0_l0_l0/D | : n/a | 20.02 | 20.92 | n/a | n/a |
| dio(0)/ | : n/a | 15.90 | 16.80 | n/a | n/a |

Synthesis and Simulation Design Guide

| | | | | | | |
|-------------------------------|---|-----|-------|-------|-----|-----|
| dio(1)/ | : | n/a | 15.90 | 16.80 | n/a | n/a |
| ix328_l5_1_l0_l0_0_l0_l0/WE | : | n/a | 11.80 | 11.80 | n/a | n/a |
| ix328_l5_0_l0_l0_0_l0_l0/WE | : | n/a | 11.80 | 11.80 | n/a | n/a |
| ix328_reg_q(1)_O1/D | : | n/a | 9.12 | 9.12 | n/a | n/a |
| ix328_reg_q(0)_O1/D | : | n/a | 9.12 | 9.12 | n/a | n/a |
| ro/ | : | n/a | 8.55 | 8.55 | n/a | n/a |
| ix328_l5_0_l0_l0_0_l0_l0/WCLK | : | n/a | 7.06 | 7.06 | n/a | n/a |

Critical Path Report

Critical path #1, (unconstrained path)

| NAME | GATE | ARRIVAL | LOAD |
|----------------------------|----------|----------|--------------------|
| we/ | | 0.00 up | 2.32 |
| ix351/O | IBUF | 4.12 up | 2.94 |
| ix328_nx4/O | F2_LUT | 8.86 up | 2.94 |
| ix328_reg_q(0)_O1/O | OFDTX | 16.80 dn | 2.94 |
| ix328_reg_q(0)_I1/O | IBUF | 18.60 dn | 2.32 |
| ix328_l5_0_l0_l0_0_l0_l0/D | RAM16x1S | 20.92 dn | 0.00 |
| data arrival time | | 20.92 | |
| data required time | | | not specified |
| data required time | | | not specified |
| data arrival time | | | 20.92 |
| | | | unconstrained path |

Critical path #2, (unconstrained path)

| NAME | GATE | ARRIVAL | LOAD |
|----------------------------|----------|----------|------|
| meme/ | | 0.00 up | 2.32 |
| ix352/O | IBUF | 4.12 up | 2.94 |
| ix328_nx4/O | F2_LUT | 8.86 up | 2.94 |
| ix328_reg_q(0)_O1/O | OFDTX | 16.80 dn | 2.94 |
| ix328_reg_q(0)_I1/O | IBUF | 18.60 dn | 2.32 |
| ix328_l5_0_l0_l0_0_l0_l0/D | RAM16x1S | 20.92 dn | 0.00 |
| data arrival time | | 20.92 | |

```

data required time                               not specified
-----
data required time                               not specified
data arrival time                               20.92
-----
                                           unconstrained path
-----

```

```

.
.
.
.

```

Log Report

```

Content-Type: text/plain; charset="us-ascii"
Content-Disposition: attachment; filename="ram.log"

```

```

C:\Program Files\Exemplar Logic\Galileo 4.2\bin\win32\gc.exe \
F:/rel4.2/example/ram.vhd F:/rel4.2/example/ram.edf -input_format=VHDL \
-target=xi4ex -output_format=EDIF -area -effort=quick \
-edif_timing_file=F:/rel4.2/example/ram.tim -encoding=OneHot -
wire_tree=Worst \
-nocontrol -vhdl_93 -process=3 -wire_table=4052ex-3_avg -chip
-----

```

```

Galileo - V4.2 (build 2.01, compiled Dec 19 1997 at 16:48:48)
Copyright 1990-1996 Exemplar Logic, Inc. All rights reserved.

```

```

Checking Security ...
Info: setting encoding to OneHot
Info: setting process to 3
Info: setting wire_tree to Worst
Info: setting wire_table to 4052ex-3_avg
--
-- Welcome to Galileo
-- Run By massoumi@WACO
-- Run Started On Tue Jan 13 11:05:09 Pacific Daylight Time 1998
--
-- read -format VHDL {F:/rel4.2/example/ram.vhd}
-- Reading file C:\PROGRA~1\EXEMPL~1\GALILE~1.2\data\standard.vhd for
unit standard
-- Loading package standard into library std

```

```
-- Reading vhdl file F:/rel4.2/example/ram.vhd into library work
-- Reading file C:\PROGRA~1\EXEMPL~1\GALILE~1.2\data\std_1164.vhd for
unit std_logic_1164
-- Loading package std_logic_1164 into library ieee
-- Reading file C:\PROGRA~1\EXEMPL~1\GALILE~1.2\data\ex_1164.vhd for unit
exemplar_1164
-- Loading package exemplar_1164 into library exemplar
-- Reading file C:\PROGRA~1\EXEMPL~1\GALILE~1.2\data\exemplar.vhd for
unit exemplar
-- Loading package exemplar into library exemplar
-- Loading package my_pkg into library work
-- Loading entity mem into library work
-- Loading architecture behav of mem into library work
"F:/rel4.2/example/ram.vhd",line 20: Warning, output ro is never assigned
a value.
-- Compiling root entity mem(behav)
-- Reading target technology xi4ex
Reading library file `C:\PROGRA~1\EXEMPL~1\GALILE~1.2\lib\xi4ex.syn`...
Library version = 0.9
Delays assume: Process=3
-- Pre Optimizing Design .work.mem.behav
INFO: Using Ram Cell ram_io_inclock_outclock_2_3_8.
-- Read Module Generators
-- Reading module generator description from file
C:\PROGRA~1\EXEMPL~1\GALILE~1.2\data\modgen\xi4e.vhd
-- Reading vhdl file C:\PROGRA~1\EXEMPL~1\GALILE~1.2\data\modgen\xi4e.vhd
into library OPERATORS
-- Modgen File xi4e.vhd Version 4.20
-- Resolving Modgen With modgen_select "small"
-- Start module generator resolving for design .work.mem.behav
-- Resolving function ram_io with module generator
ram_io_2_3_8_true_true_false from file xi4e.vhd
-- optimize -target xi4ex -effort quick -chip -area
-- Start optimization for design .work.mem.behav
Using wire table: 4052ex-3_avg

      Pass      Area      Delay      DFFs  PIs    POs  --CPU--
          (FGs)      (ns)
1          2          17          2     7     3    00:00
Info, setting outputs in top level view 'behav' to fast.
Using wire table: 4052ex-3_avg
-- Start timing optimization for design .work.mem.behav

Latest arrival time at primary output: 20.9 ns
Latest arrival time at register input: 20.9 ns
```

Forcing timing constraints at all end points: 18.8 ns

Initial Timing Optimization Statistics:

```

Most Critical Slack      :      -2.1
Sum of Negative Slacks  :      -4.2
Longest Path            :      20.9 ns
Area                    :          2.0
  
```

Final Timing Optimization Statistics:

```

Most Critical Slack      :      -2.1
Sum of Negative Slacks  :      -4.2
Longest Path            :      20.9 ns
Area                    :          2.0
  
```

Total time taken : 0 cpu secs

Using wire table: 4052ex-3_avg

-- Start timespec generation for design .work.mem.behav

Cell: mem View: behav Library: work

```

Number of ports :          10
Number of nets  :          24
Number of instances :       17
Number of references to this view :    0
  
```

Total accumulated area :

```

Number of FG Function Generators :    4
Number of Packed CLBs :              1
Number of IBUF :                      9
Number of OBUF :                      1
Number of IOB Output Flip Flops :    2
  
```

-- Writing file F:/rel4.2/example/ram.edf

-- CPU time taken for this run was 29.55 sec

-- Run ended On Tue Jan 13 11:05:37 Pacific Daylight Time 1998
-- Galileo run successfully completed. Goodbye !

Synopsys FPGA Express™

FPGA Express reports include information on the following.

- Synthesis options used
- Primitives used
- Required and predicted frequency of the clocks
- Critical paths

```
=====  
Chip fd32ce-Optimized  
=====
```

Summary Information:

```
-----  
Type: Optimized implementation  
Source: fd32ce, up to date  
Status: 0 errors, 0 warnings, 1 messages  
Export: not exported since last optimization
```

Target Information:

```
-----  
Vendor: Xilinx  
Family: XC4000  
Device: 4085XLBG560  
Speed: xl-09
```

Chip Parameters:

```
-----  
Optimize for: Speed  
Optimization effort: High  
Frequency: 50 MHz  
Is module: No  
Keep io pads: No  
Number of flip-flops: 32  
Number of latches: 0
```

Chip Design Hierarchy:

```
-----  
fd32ce: defined in /home/elayda/drcust/stapp2/test1/fd32ce.v
```


Primitive reference count:

```

-----
BUFG          1
IBUF          2
INFF         32
OBUF         32
STARTUP       1
  
```

Clocks:

```

-----

```

| Period (ns) | Rise (ns) | Fall (ns) | Required Freq (MHz) | Estimated Freq (MHz) | Signal |
|----------------|--------------|--------------|---------------------------|----------------------------|-------------|
| 20 | 0 | 10 | 50.00 | n/a | default |
| n/a | n/a | n/a | n/a | 100.00 | clock_BUFGE |

Timing Groups:

```
-----
```

| Name | Description |
|------------------|---------------------------------------|
| (I) | Input ports |
| (O) | Output ports |
| (RC,clock_BUFGE) | Clocked by rising edge of clock_BUFGE |

Timing Path Groups:

```
-----
```

| >From | To | Required Delay (ns) | Estimated Delay (ns) |
|------------------|------------------|---------------------------|----------------------------|
| (I) | (RC,clock_BUFGE) | 20.00 | 8.42 |
| (RC,clock_BUFGE) | (O) | 20.00 | 9.08 |

Input Port Timing:

```
-----
```

| Port Name | Required Delay (ns) | Estimated Slack (ns) | To-Group |
|--------------|---------------------------|----------------------------|------------------|
| clock | 11.58 | 11.58 | (RC,clock_BUFGE) |
| ce | n/a | n/a | (RC,clock_BUFGE) |
| reset | n/a | n/a | (RC,clock_BUFGE) |

| | | | |
|-------------|-------|-------|-------------------|
| data_in<31> | 12.10 | 12.10 | (RC,clock_BUFGeD) |
| data_in<30> | 12.10 | 12.10 | (RC,clock_BUFGeD) |
| data_in<29> | 12.10 | 12.10 | (RC,clock_BUFGeD) |
| data_in<28> | 12.10 | 12.10 | (RC,clock_BUFGeD) |
| data_in<27> | 12.10 | 12.10 | (RC,clock_BUFGeD) |
| data_in<26> | 12.10 | 12.10 | (RC,clock_BUFGeD) |
| data_in<25> | 12.10 | 12.10 | (RC,clock_BUFGeD) |
| data_in<24> | 12.10 | 12.10 | (RC,clock_BUFGeD) |
| data_in<23> | 12.10 | 12.10 | (RC,clock_BUFGeD) |
| data_in<22> | 12.10 | 12.10 | (RC,clock_BUFGeD) |
| data_in<21> | 12.10 | 12.10 | (RC,clock_BUFGeD) |
| data_in<20> | 12.10 | 12.10 | (RC,clock_BUFGeD) |
| data_in<19> | 12.10 | 12.10 | (RC,clock_BUFGeD) |
| data_in<18> | 12.10 | 12.10 | (RC,clock_BUFGeD) |
| data_in<17> | 12.10 | 12.10 | (RC,clock_BUFGeD) |
| data_in<16> | 12.10 | 12.10 | (RC,clock_BUFGeD) |
| data_in<15> | 12.10 | 12.10 | (RC,clock_BUFGeD) |
| data_in<14> | 12.10 | 12.10 | (RC,clock_BUFGeD) |
| data_in<13> | 12.10 | 12.10 | (RC,clock_BUFGeD) |
| data_in<12> | 12.10 | 12.10 | (RC,clock_BUFGeD) |
| data_in<11> | 12.10 | 12.10 | (RC,clock_BUFGeD) |
| data_in<10> | 12.10 | 12.10 | (RC,clock_BUFGeD) |
| data_in<9> | 12.10 | 12.10 | (RC,clock_BUFGeD) |
| data_in<8> | 12.10 | 12.10 | (RC,clock_BUFGeD) |
| data_in<7> | 12.10 | 12.10 | (RC,clock_BUFGeD) |
| data_in<6> | 12.10 | 12.10 | (RC,clock_BUFGeD) |
| data_in<5> | 12.10 | 12.10 | (RC,clock_BUFGeD) |
| data_in<4> | 12.10 | 12.10 | (RC,clock_BUFGeD) |
| data_in<3> | 12.10 | 12.10 | (RC,clock_BUFGeD) |
| data_in<2> | 12.10 | 12.10 | (RC,clock_BUFGeD) |
| data_in<1> | 12.10 | 12.10 | (RC,clock_BUFGeD) |
| data_in<0> | 12.10 | 12.10 | (RC,clock_BUFGeD) |

Output Port Timing:

| Port Name | Required Delay (ns) | Estimated Slack (ns) | From-Group |
|--------------|---------------------------|----------------------------|-------------------|
| | | | |
| data_out<31> | 20.00 | 10.92 | (RC,clock_BUFGeD) |
| data_out<30> | 20.00 | 10.92 | (RC,clock_BUFGeD) |
| data_out<29> | 20.00 | 10.92 | (RC,clock_BUFGeD) |
| data_out<28> | 20.00 | 10.92 | (RC,clock_BUFGeD) |
| data_out<27> | 20.00 | 10.92 | (RC,clock_BUFGeD) |
| data_out<26> | 20.00 | 10.92 | (RC,clock_BUFGeD) |

| | | | |
|--------------|-------|-------|-------------------|
| data_out<25> | 20.00 | 10.92 | (RC,clock_BUFGeD) |
| data_out<24> | 20.00 | 10.92 | (RC,clock_BUFGeD) |
| data_out<23> | 20.00 | 10.92 | (RC,clock_BUFGeD) |
| data_out<22> | 20.00 | 10.92 | (RC,clock_BUFGeD) |
| data_out<21> | 20.00 | 10.92 | (RC,clock_BUFGeD) |
| data_out<20> | 20.00 | 10.92 | (RC,clock_BUFGeD) |
| data_out<19> | 20.00 | 10.92 | (RC,clock_BUFGeD) |
| data_out<18> | 20.00 | 10.92 | (RC,clock_BUFGeD) |
| data_out<17> | 20.00 | 10.92 | (RC,clock_BUFGeD) |
| data_out<16> | 20.00 | 10.92 | (RC,clock_BUFGeD) |
| data_out<15> | 20.00 | 10.92 | (RC,clock_BUFGeD) |
| data_out<14> | 20.00 | 10.92 | (RC,clock_BUFGeD) |
| data_out<13> | 20.00 | 10.92 | (RC,clock_BUFGeD) |
| data_out<12> | 20.00 | 10.92 | (RC,clock_BUFGeD) |
| data_out<11> | 20.00 | 10.92 | (RC,clock_BUFGeD) |
| data_out<10> | 20.00 | 10.92 | (RC,clock_BUFGeD) |
| data_out<9> | 20.00 | 10.92 | (RC,clock_BUFGeD) |
| data_out<8> | 20.00 | 10.92 | (RC,clock_BUFGeD) |
| data_out<7> | 20.00 | 10.92 | (RC,clock_BUFGeD) |
| data_out<6> | 20.00 | 10.92 | (RC,clock_BUFGeD) |
| data_out<5> | 20.00 | 10.92 | (RC,clock_BUFGeD) |
| data_out<4> | 20.00 | 10.92 | (RC,clock_BUFGeD) |
| data_out<3> | 20.00 | 10.92 | (RC,clock_BUFGeD) |
| data_out<2> | 20.00 | 10.92 | (RC,clock_BUFGeD) |
| data_out<1> | 20.00 | 10.92 | (RC,clock_BUFGeD) |
| data_out<0> | 20.00 | 10.92 | (RC,clock_BUFGeD) |

Critical Path Timing:

| Cell Type | Arrival Time (ns) | Required Time (ns) | Fanout Count | Pin-Name |
|--------------|-------------------------|--------------------------|-----------------|-----------------------|
| | | | | |
| port | 9.08 | 20.00 | 0 | /fd32ce-Optimized/dat |
| obuf | 9.08 | 20.00 | 0 | /fd32ce-Optimized/C17 |
| obuf | 4.38 | 15.30 | 1 | /fd32ce-Optimized/C17 |
| inff | 1.30 | 12.22 | 1 | /fd32ce-Optimized/dat |
| inff | 0.00 | 10.92 | 32 | /fd32ce-Optimized/dat |

Index

A

- active_low_gsr design, 4-23
- after xx ns statement, 2-2
- arithmetic functions
 - gate reduction, 2-37
 - ordering and grouping, 2-3
 - resource sharing, 2-32
- ASIC
 - comparing to FPGA, 1-3, 2-1
- asynchronous reset pin, 2-38
- asynchronous set pin, 2-38

B

- barrel shifter design, 2-18
- bi-directional I/O, 4-71
 - inferring, 4-72
 - instantiating, 4-74
 - using LogiBLOX, 4-76
- binary encoded state machine, 4-26
- boundary scan, 4-63
 - instantiating in HDL, 4-63
- BSCAN, 4-63
- BUFGP, 4-3
- BUFGS, 4-3
- BUFT *see* tristate buffer

C

- capitalization style in code, 2-4
- case sensitivity, 5-13
- case statement, 2-3
 - comparing to if statement, 2-55

- design example, 2-58
- syntax, 2-47
- when to use, 2-47

CLB

- XC4000, 2-39
- clear pin, 2-38, 4-9, 4-16
- clock buffers
 - inserting, 4-6
 - instantiating, 4-2, 4-6
- clock enable pin, 2-38, 2-42
- combinatorial feedback loop, 2-29
- comments in code, 2-13
- compile run script, 3-4
- compiling large designs, 3-5
- compiling your design, 3-4, 3-5
- conditional expression, 2-27
- constants, 2-8
- constraint precedence, 3-11
- cost-based clean-up option, 3-22
- creating readable code, 2-10

D

- D register, 2-29
 - design, 2-7
- decoders, 4-40
- delay-based clean-up option, 3-23
- design compiling, 3-4
- design entry, 3-3
- design flow
 - description, 3-1
 - diagram, 3-2

- using the Design Manager, 3-2
- design hierarchy, 3-3, 4-90, 4-91
- design performance, 3-17
- device downloading, 3-23
- directory tree structure, 1-11
- disk space requirements, 1-6
- downloading files, 1-7, 1-8
- downloading to the device, 3-23

E

- EDIF file, 3-5
- else statement, 2-29
- entering your design, 3-3
- enumerated type encoded state machine, 4-32
- Exemplar Logic report files, B-4
- extracting downloaded files, 1-10

F

- Field Programmable Gate Array *see* FPGA
- file transfer protocol, 1-8
- Finite State Machine, 4-32, 4-39
 - changing encoding style, 4-39
 - extraction commands, 4-32
- flip-flop, 2-30
- formatting styles, 2-4
- FPGA
 - comparing to ASIC, 1-3, 2-1
 - creating with HDLs, 4-1
 - global clock buffer, 4-2
 - system features, 1-4, 4-1
- FPGA Express report files, B-10
- from
 - to style timing constraint, 3-8
- FSM *see* Finite State Machine
- functional simulation, 1-2, 3-4, 5-1, 5-4
 - comparing to synthesis, 2-1

G

- gate reduction
 - definition, 2-37

- gated clocks, 2-42
- global clock buffer, 4-2
- global longlines, 4-5
- global set/reset, 4-9, 5-23, 5-51
 - increasing performance, 4-10
 - STARTUP block, 4-9
 - test bench, 5-52
- global signals, 5-18, 5-22
- GSR *see* global set/reset
- GSRIN, 4-9
- GTS, 5-35, 5-64
- guide option, 3-23

H

- hardware description language *see* HDL
- HDL
 - also see* Verilog
 - also see* VHDL
 - coding for FPGAs, 4-1
 - converting to gates, 1-2
 - definition, 1-1
 - designing FPGAs, 1-2, 1-3
 - FPGA system features, 4-1
 - boundary scan, 4-63
 - global clock buffer, 4-2
 - global set/reset, 4-9
 - I/O decoders, 4-40
 - implementing logic with IOBs, 4-67
 - on-chip RAM, 4-52
 - implementing registers, 2-27
 - schematic entry design hints, 2-17
 - hdl_resource_allocation command, 2-34
 - hdlin_check_no_latch command, 2-30
 - hierarchy in designs, 1-4
 - high-density design flow, 3-1
 - hold-time requirement, 2-28, 4-67

I

- I/O decoder, 4-40

- if statement, 2-30
 - comparing to case statement, 2-55
 - design example, 2-56
 - registers, 2-30
 - syntax, 2-45
 - when to use, 2-46
- if-case statement
 - design example, 2-51
- if-else statement, 2-3, 2-45
- ignore timing paths, 3-9
- indenting HDL code, 2-10
- INIT=S attribute, 4-10, 4-16, 4-40
- initialization statement, 2-4
- Insert Pads command, 4-2
- installation
 - design examples, 1-5
 - directory tree structure, 1-11
 - disk space requirements, 1-6
 - downloading files, 1-7, 1-8
 - extracting downloaded files, 1-10
 - file transfer protocol, 1-8
 - internet site, 1-7, 1-8
 - memory requirements, 1-5
 - tactical software, 1-5
- internet site, 1-7
- IOB
 - implementing logic, 4-67
 - moving registers, 4-79, 4-80
 - unbonded, 4-81
- J**
- JTAG 1149.1, 4-63
- L**
- labeling in code, 2-7
- latch
 - combinatorial feedback loop, 2-28
 - comparing speed and area, 2-31
 - converting to register, 2-29
 - D flip-flop, 2-30
 - D latch implemented with gates, 2-28
 - hdlin_check_no_latch command, 2-30
 - implementing in HDL, 2-27
 - inference, 2-46
 - latch count, 2-30
 - RAM primitives, 2-30
- libraries, 5-11
- LogiBLOX
 - bi-directional I/O, 4-76
 - implementing memory, 4-59
 - instantiating modules, 4-45
 - libraries, 5-11, 5-16, 5-49
- LogiCORE
 - library, 5-18
- M**
- mapping your design
 - using design manager, 3-13
 - using the command line, 3-15
- maxskew, 3-10
- memory
 - implementing in HDL, 4-52
 - requirements, 1-5
- Modelsim simulator, 5-49
- multi-pass place and route option, 3-21
- multiplexer
 - comparing gates and tristate buffer, 4-89
 - implementing with gates, 4-86
 - implementing with tristate buffer, 4-84
 - resource sharing, 2-32
- N**
- named association, 2-9
- naming conventions, 2-5, 2-6
- nested if statement, 2-48
- no_gsr design, 4-11
- NODELAY attribute, 4-67
- O**
- offset constraint, 3-9
- OMUX, 4-69

one-hot encoded state machine, 4-35
oscillators, 5-46
output multiplexer, 4-69

P

pad location, 4-79
period constraint, 3-8
pipelining, 4-89
placing and routing your design, 3-18
port declarations, 2-14
positional association, 2-9
post-route full timing simulation, 5-8
post-synthesis simulation, 5-5
-pr option, 4-80
preset pin, 2-38, 4-9, 4-10, 4-16
priority-encoded logic, 2-48, 2-55
PROM file, 3-24
pull-downs, 4-69
pull-ups, 4-69

R

RAMs
 inferring, 4-58
 instantiating, 4-56
re-entrant routing option, 3-21
register
 clear pin, 2-38
 converting latch to register, 2-29
 D register, 2-29
 if statement, 2-30
 implementing in HDL, 2-27
 inference, 2-39
 moving into IOB, 4-79
 preset pin, 2-38
report files, B-1
report_timing command, 3-13
reset on configuration, 5-24
reset on configuration buffer, 5-29
resource sharing
 CLB count, 2-37
 definition, 2-31

delay, 2-37
design examples, 2-32
disabling, 2-34

ROC, 5-24
ROCBUF, 5-29
ROMs, 4-53
RTL simulation, 4-11, 4-17, 5-1, 5-4
 definition, 1-2

S

schematic entry design hints, 2-17
set don't touch attribute, 4-63
signal skew, 3-10
signals, 2-15
SIMPRIM libraries, 5-11, 5-49
simulating your design, 5-1
simulation
 creating a test bench, 5-7
 functional, 5-4
 global signals, 5-18
 industry standards, 5-11
 library source files, 5-12
 post-map, 5-6
 post-NGDBuild, 5-6
 post-synthesis, 5-5
 timing, 5-8
simulation diagram, 5-2
slew rate, 4-68
software requirements, 1-5
Spartan
 IOB, 4-67
STARTBUF, 4-9, 5-32
STARTUP block, 4-9, 4-17
startup state, 4-10
state machine, 4-26
 binary encoded, 4-26
 bubble diagram, 4-27
 encoding style summary, 4-38
 enumerated type encoded, 4-32
 design example, 4-33, 4-35
 enumeration type, 4-38

- initializing, 4-39
- limiting fan-in, A-3
- one-hot encoded, 4-35, A-1
- one-state vs. binary encoded, A-9
- resetting state bits, A-12
- seven states, A-3
- std_logic data type, 2-13
- Synopsys
 - creating compile run script, 3-4
- Synplicity report files, B-1
- synthesis
 - comparing to simulation, 2-1

T

- TCK pin, 4-63
- TDI pin, 4-63
- TDO pin, 4-63
- technical support, 1-14
- test bench, 5-7
- TIG, 3-9
- TIMEGRPs, 3-6
- timing
 - constraint precedence, 3-11
 - constraint priority, 3-10
 - constraints, 1-5, 3-5, 3-8
 - requirements, 1-5
 - simulation, 3-23, 5-1, 5-8
 - simulation netlist
 - Command Line, 5-10
 - Design Manager, 5-9
- TMS pin, 4-63
- TNMs, 3-6
- TOC, 5-37
- TOCBUF, 5-42
- TPSYNC keyword, 3-7
- tristate buffer
 - comparing to gates, 4-89
 - implementing multiplexer, 4-84
- tristate enable, 5-35
- tristate on configuration, 5-37
- tristate on configuration buffer, 5-42

- turns engine option, 3-21

U

- UCF, 4-79
- unbonded IOBs, 4-81
- UNISIM libraries, 5-11, 5-50
- UNISIM library, 5-14
- use_gsr design, 4-17
- user constraints file, 4-79

V

- variables, 2-15
- Verilog
 - capitalization style, 2-5
 - constants for opcode, 2-9
 - definition, 1-3
 - global set/reset, 5-51
 - GTS, 5-63
 - libraries, 5-49
 - oscillators, 5-69
 - parameters for constants, 2-9
 - register inference, 2-41
 - simulation, 5-51
- VHDL
 - after xx ns statement, 2-2
 - also see* HDL
 - arithmetic functions, 2-3
 - capitalization style, 2-5
 - case statement, 2-3
 - coding styles, 2-4
 - constants, 2-8
 - constants for opcode, 2-8
 - definition, 1-3
 - if-else statement, 2-3
 - initialization statement, 2-4
 - naming identifiers, 2-6
 - naming packages, 2-6
 - naming types, 2-6
 - register inference, 2-39
 - simulation, 2-1
 - std_logic data type, 2-13

- synthesis, 2-1
- variables for constants, 2-8
- wait for xx ns statement, 2-2
- Xilinx naming conventions, 2-5

VHSIC Hardware Description Language
see VHDL

W

- wait for xx ns statement, 2-2

X

XC4000

- CLB, 2-39

- IOB, 4-67

XC5200

- IOB, 4-71

XDW libraries, 5-11

Xilinx internet site, 1-8

XNF file, 3-5

xor_sig design, 2-16

xor_var design, 2-17