

UltraFAST™ 

UltraFast Design Methodology Guide for the Vivado Design Suite

UG949 (v2015.3) November 23, 2015

Revision History

The following table shows the revision history for this document.

Date	Version	Revision
11/23/2015	2015.3	Added Using IP Core Containers , updated Logic Simulation , added information on core container files to in Recommended Source Files to Manage , Minimum Sets of Source Files to Manage , and Vivado Design Suite Source Types , added information on core containers in Managing IP Sources in Chapter 2, Using the Vivado Design Suite . Updated Table 3-1 , added SLR Utilization Considerations , and added SLR Crossing for Wide Buses in Chapter 3, Board and Device Planning . Updated UltraScale Device Clocking and updated Pipelining Considerations in Chapter 4, Design Creation . Added information on <code>report_design_analysis</code> in Clock Skew and Uncertainty in Chapter 5, Implementation . Updated Figure 2-1 .
06/01/2015	2015.1	Reorganized and updated Chapter 2, Using the Vivado Design Suite , including adding new Source Management and Revision Control Recommendations section. Updated I/O Planning Design Flows in Chapter 3, Board and Device Planning . Reorganized and updated Chapter 4, Design Creation . Updated Timing Closure in Chapter 5, Implementation , including moving common design bottlenecks information to the <i>Vivado Design Suite User Guide: Design Analysis and Closure Techniques</i> (UG906). Reorganized and updated Chapter 6, Configuration and Debug , including adding links to various additional resources.
10/14/2014	2014.3	Modified IP Flows related sections. Minor fixes/clarifications based on specific feedback/suggestions.
04/02/2014	2014.1	Condensed "power" section. Major revamp of "Vivado Design Suite Flows", and "Configuration and Debug" chapters. Fixed specific typos, heading name/levels and minor changes.
12/18/2013	2013.4	Removed checklist appendix. These links have been replaced with a checklist version that is available in Documentation Navigator.
11/25/2013	2013.3	Fixed errors in table of contents.
10/27/2013	2013.3	Fixed incorrect hyperlinks.
10/23/2013	2013.3	Initial Xilinx release.

Table of Contents

Chapter 1: Introduction

About This Guide	5
Guide Contents	5
Guide Applicability and References	6
Need for Design Methodology	6
Design Methodology Checklist	7
Design Process	8
Rapid Validation	11
Accessing Documentation and Training	12

Chapter 2: Using the Vivado Design Suite

Overview of Using the Vivado Design Suite	15
Vivado Design Suite Use Models	19
Configuring IP	23
Creating IP Subsystems with IP Integrator	28
Packaging Custom IP and IP Subsystems	32
Creating Custom Peripherals	33
Logic Simulation	34
Synthesis, Implementation, and Design Analysis	43
Source Management and Revision Control Recommendations	43
Upgrading Designs and IP to the Latest Vivado Design Suite Release	60

Chapter 3: Board and Device Planning

Overview of Board and Device Planning	62
PCB Layout Recommendations	62
Clock Resource Planning and Assignment	65
I/O Planning Design Flows	67
FPGA Power Aspects and System Dependencies	87
Worst Case Power Analysis Using Xilinx Power Estimator (XPE)	98
Configuration	102

Chapter 4: Design Creation

Overview of Design Creation	104
-----------------------------------	-----

Defining a Good Design Hierarchy	105
RTL Coding Guidelines	108
Clocking Guidelines	157
Working With Intellectual Property (IP)	204
Working with Constraints	210
 Chapter 5: Implementation	
Overview of Implementation	255
Synthesis	255
Synthesis Attributes	259
Bottom Up Flow	261
Moving Past Synthesis	263
Implementing the Design	266
Timing Closure	280
Power	327
 Chapter 6: Configuration and Debug	
Overview of Configuration and Debug	338
Configuration	338
Debugging	344
 Appendix A: Baselining and Timing Constraints Validation Procedure	
Introduction	358
Procedure	358
 Appendix B: Additional Resources and Legal Notices	
Xilinx Resources	360
Solution Centers	360
References	360
Training Resources	363
Please Read: Important Legal Notices	363

Introduction

About This Guide

Xilinx® programmable devices have capacities of multi-million Logic Cells (LC), and integrate an ever-increasing share of today's complex electronic systems, including:

- Embedded subsystems
- Analog and digital processing
- High-speed connectivity
- Network processing

In order to create such complex systems within short design cycles, designers synthesize many large blocks of logic from RTL, and reuse Intellectual Property (IP) modules from Xilinx or third parties.

Given the complexity of this process, it is critical to adopt a set of best practices collectively called the UltraFast™ Design Methodology, a set of best practices that maximize productivity for both system integration and design implementation.

Guide Contents

This guide discusses a design methodology process to follow in order to achieve an efficient and quicker design implementation, and to derive the maximum value from Xilinx devices and tools.

In most cases, this guide tells you the reasoning behind its recommendations. By understanding that reasoning, you can appreciate the potential consequences of deviating from the recommended methodology, and take appropriate precautions.

Guide Applicability and References

Although this guide is primarily for use with the Xilinx Vivado® Design Suite, most of the conceptual information in this guide can be leveraged for use with the Xilinx ISE® Design Suite as well. This guide provides high-level information, design guidelines, and design decision trade-offs.

This guide includes references to other documents such as the *Vivado Design Suite User Guides*, *Vivado Design Suite Tutorials*, and *Quick-Take Video Tutorials*. This guide is not a replacement for those documents. You should still refer to those documents for detailed, current information, including descriptions of tool use and design methodology. For a more complete listing of reference documents, see [Appendix B, Additional Resources and Legal Notices](#).

At various places, the guide gives the Vivado tools command for a specific task. Run the command with `-help` for detailed information (including example usage).

Need for Design Methodology

Advanced algorithms used in today's increasingly complex electronic products are stretching the boundaries of density, performance, and power. This creates many challenges for the design teams to hit the target release window within their allocated budget. The UltraFast Design Methodology allows project managers to:

- Accelerate time to market, thus increasing product revenue and market share.
- Formulate an accurate estimate of the project schedule and cost, reducing risk.

This guide is a collection of best practices covering aspects related to board planning, design creation, IP integration, design implementation and closure techniques, programming, and hardware debug. These best practices and recommendations have been gathered from a large pool of expert users over the past several years. The recommendations in this guide will help you succeed as they have for many of Xilinx customers.

Vivado Design Suite is also automating part of the UltraFast Design Methodology by providing:

- DRC rules that provide guidance on HDL code and XDC constraints so engineers can improve the quality of their design earlier in the flows and avoid problems downstream when iterations would be costlier.
- Proven templates for specific HDL code and XDC constraints that enable optimal-by-construction code.

These DRC checks and HDL/XDC templates are part of the overall UltraFast Design Methodology, in the sense that having a clear DRC and code that adheres to the templates provided with the Vivado Design Suite will make it easier for you to complete your design in a timely manner.

Design Methodology Checklist

To take full advantage of the UltraFast Design Methodology, use this guide in partnership with the Design Methodology Checklist. The checklist includes common questions and recommended actions to consider during the design process starting with planning and continuing through all subsequent stages of design. The checklist questions highlight typical areas in which design decisions are likely to have downstream ramifications and draw attention to issues that are often unknown or ignored.



VIDEO: For a demonstration of the checklist, see the [Vivado Design Suite QuickTake Video: Introducing the UltraFast Design Methodology Checklist](#).

Most checklist questions also provide links to content in this guide or other Xilinx documentation. These references offer guidance on addressing the design concerns raised by the questions.

Documentation Navigator ships with the Vivado Design Suite (see [Using the Documentation Navigator](#)). To access the checklist feature, use Documentation Navigator version 2013.4 or later. From within Documentation Navigator, use these steps to begin using the Design Methodology Checklist:

1. Click the **Design Hub View** tab.
2. At the top of the menu on the left side, click **Create Design Checklist**.
3. Fill out the information in the New Design Checklist Dialog and click **OK**.
4. The new checklist opens. Tabs across the top of the checklist provide navigation as shown in the following figure. The Title Page tab provides some basic information on using the checklist. Click the other tabs to see the checklist questions and guidance.



Figure 1-1: Design Methodology Checklist Tabs in Documentation Navigator

A spreadsheet version of the [UltraFast Design Methodology Checklist](#) is also available.

Design Process

The steps in the design process are shown in the following figure. These steps are usually overlapping in time. Sometimes, the process might also return to a previous step, resulting in iterations.

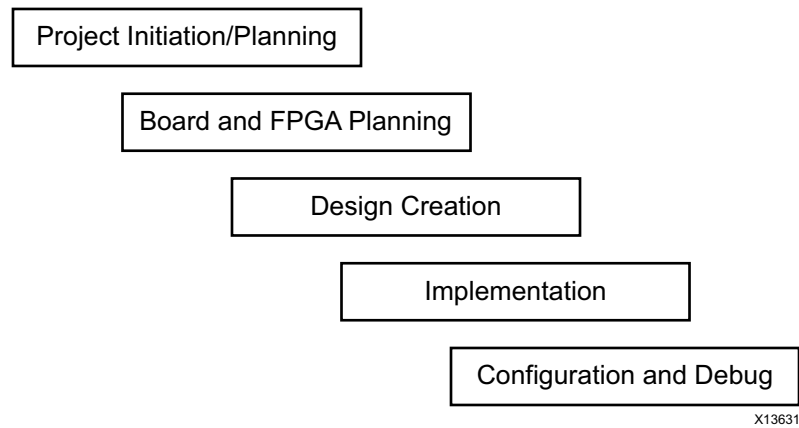


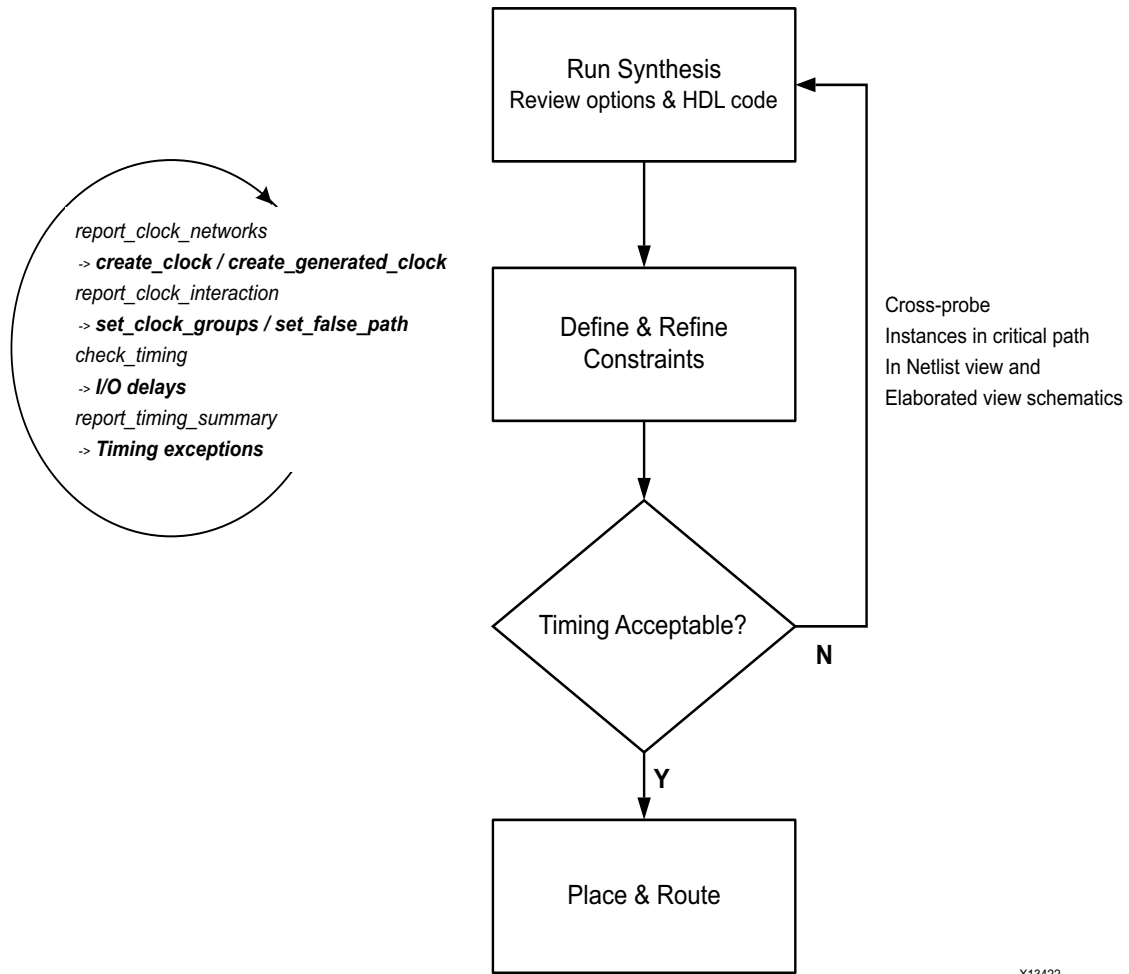
Figure 1-2: Steps in Design Process

The checklist and this methodology guide are organized as per the design phases. As you enter each design phase, Xilinx recommends that you review the corresponding tab in the checklist and the chapter in this methodology guide.

This guide demonstrates the importance of monitoring design budgets (such as area, power, timing, etc.) and correcting designs appropriately from early stages. There is a lot of importance given to creating correct timing constraints for the system, before entering the implementation phase. Because the Vivado tools use timing-driven algorithms throughout, the design must be properly constrained from the beginning of the design flow.

Specifying correct timing requires (among other things) analyzing the relationship between each master clock plus their related generated clocks in the design. Unlike ISE (UCF), in the Vivado tools (XDC) each clock interaction is timed, unless explicitly declared as asynchronous or false-path. Timing analysis should be performed after synthesis and timing should be met with the right constraints at each implementation stage before proceeding to the next.

Overall timing and implementation convergence is accelerated by following this recommendation along with using the interactive analysis environment of the Vivado Design Suite. Further acceleration can be achieved by combining the above with the HDL design guidelines in this guide. The following figure gives some details of this high-level methodology.

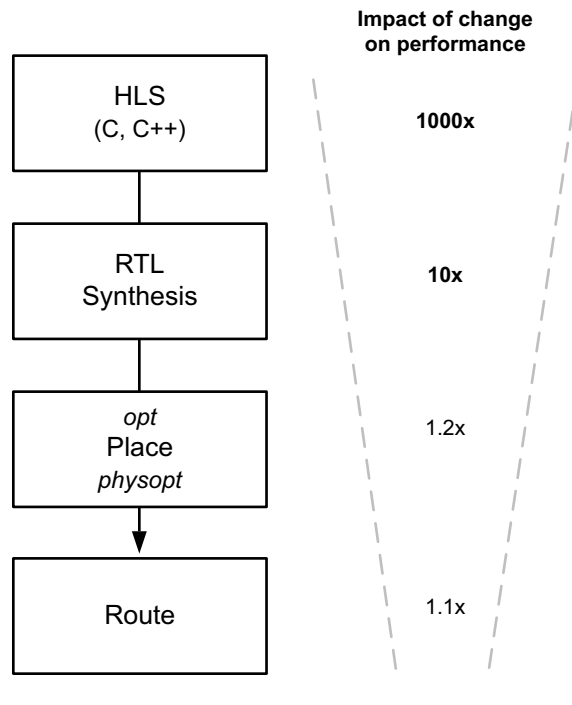


X13422

Figure 1-3: Design Methodology for Rapid Convergence

The synthesis portion of the design flow can be considered complete when the design goals are met with a positive margin (or a relatively small negative margin). For example, if post-synthesis timing is not met (or not close to being met), placement and routing results are not likely to meet timing. You may still go ahead with the rest of the flow; occasionally implementation tools may be able to close timing if they can allocate the best resources to the failing paths. Even if the timing is not met you will have a more accurate understanding of the negative slack magnitude. Having a better understanding of the post-implementation negative slack helps determine how much you need to improve the post-synthesis worst negative slack (WNS) when you come back to synthesis with improvements to HDL and constraints.

As shown in the following figure, early stages in the design flow (C, C++, and HDL synthesis) have a much higher impact on design performance, density, and power than the later stages.



X13423

Figure 1-4: Impact of Design Changes Throughout the Flow

Accordingly, if the design does not meet its timing goals, Xilinx recommends that you revisit the steps of synthesis and its inputs (including HDL and constraints) rather than iterating for a solution only in the implementation stages.

Because it is important to be correct from the beginning and to pay attention to design goals from the early stages, this guide also provides guidelines for RTL, clock, pin, and PCB planning. Properly defining and validating the design at each design stage helps alleviate timing closure, routing closure, and power usage issues during subsequent stages of implementation.

Rapid Validation

This guide introduces the concept of rapid validation of specific aspects of system architecture and micro-architecture choices. This concept can be applied in two different contexts.

In the context of system design, the I/O bandwidth is validated in-system, before even implementing the core of the design. For more information, see [Interface Bandwidth Validation in Chapter 3](#). This step may highlight the need to revise system architecture and interface choices, before finalizing on I/Os.

In the context of design implementation, baselining is used to write the simplest set of constraints, which can identify internal device timing challenges. For more information, see [Baselining the Design in Chapter 5](#). This process may identify the need to revise RTL micro-architecture choices, before moving to the implementation phase.

As part of establishing a good design methodology it is important to establish exactly how you plan to interact with the Vivado Design Suite. It has a flexible use model to accommodate various development flows and different types of designs. [Chapter 2, Using the Vivado Design Suite](#), discusses various use models supported by the Vivado tools. This will help you decide on your use model. Subsequent chapters will help you understand more details on aspects of methodology and techniques related to:

- Timing constraints definition and validation
- I/O and clock planning within the device
- Selecting and configuring IP
- Creating IP subsystems
- Packaging custom IP
- Logic Simulation
- Design rule checking (DRC)
- Power analysis and optimization
- Timing closure flows
- Hardware validation (debug core insertion and configuration)



RECOMMENDED: *Follow the design methodology recommendations discussed in this guide to obtain the most out of Xilinx devices while consuming the least amount of your time and effort.*

Accessing Documentation and Training

Access to the right information at the right time is critical for timely design closure and overall design success. Reference guides, user guides, tutorials, and videos get you up to speed as quickly as possible with the Vivado Design Suite. This section lists some of the sources for documentation and training.

Using the Documentation Navigator

The Vivado Design Suite ships with the Xilinx Documentation Navigator shown in the following figure, which provides an environment to access and manage the entire set of Xilinx software and hardware documentation, training, and support materials. Documentation Navigator allows you to view current and past Xilinx documentation. The documentation display can be filtered based on release, document type, or design task. When coupled with a search capability, you can quickly find the right information. Methodology Guides appear as one of the filters under Document Types, which allows you to reach any of Methodology Guides almost instantaneously.

Documentation Navigator scans the Xilinx website to detect and provide documentation updates. The Update Catalog feature alerts you to available updates, and gives details about the documents that are involved. Xilinx recommends that you always update the catalog when alerted to keep it current. You can establish and manage local documentation catalogs with specified documents.

The Documentation Navigator has a tab called the Design Hub View. Design hubs are collections of documentation related by design activity, such as Applying Design Constraints, Synthesis and Implementation, and Programming and Debug. Documents and videos are organized in each hub in order to simplify the learning curve for that area. Each hub contains a Getting Started Section, a Support Resources section with an FAQ for that flow, as well as Additional Learning Materials. For new users, the Getting Started section provides a good place to start. For those already familiar with the flow, Key Concepts and the FAQ may be of particular interest to gain expertise.

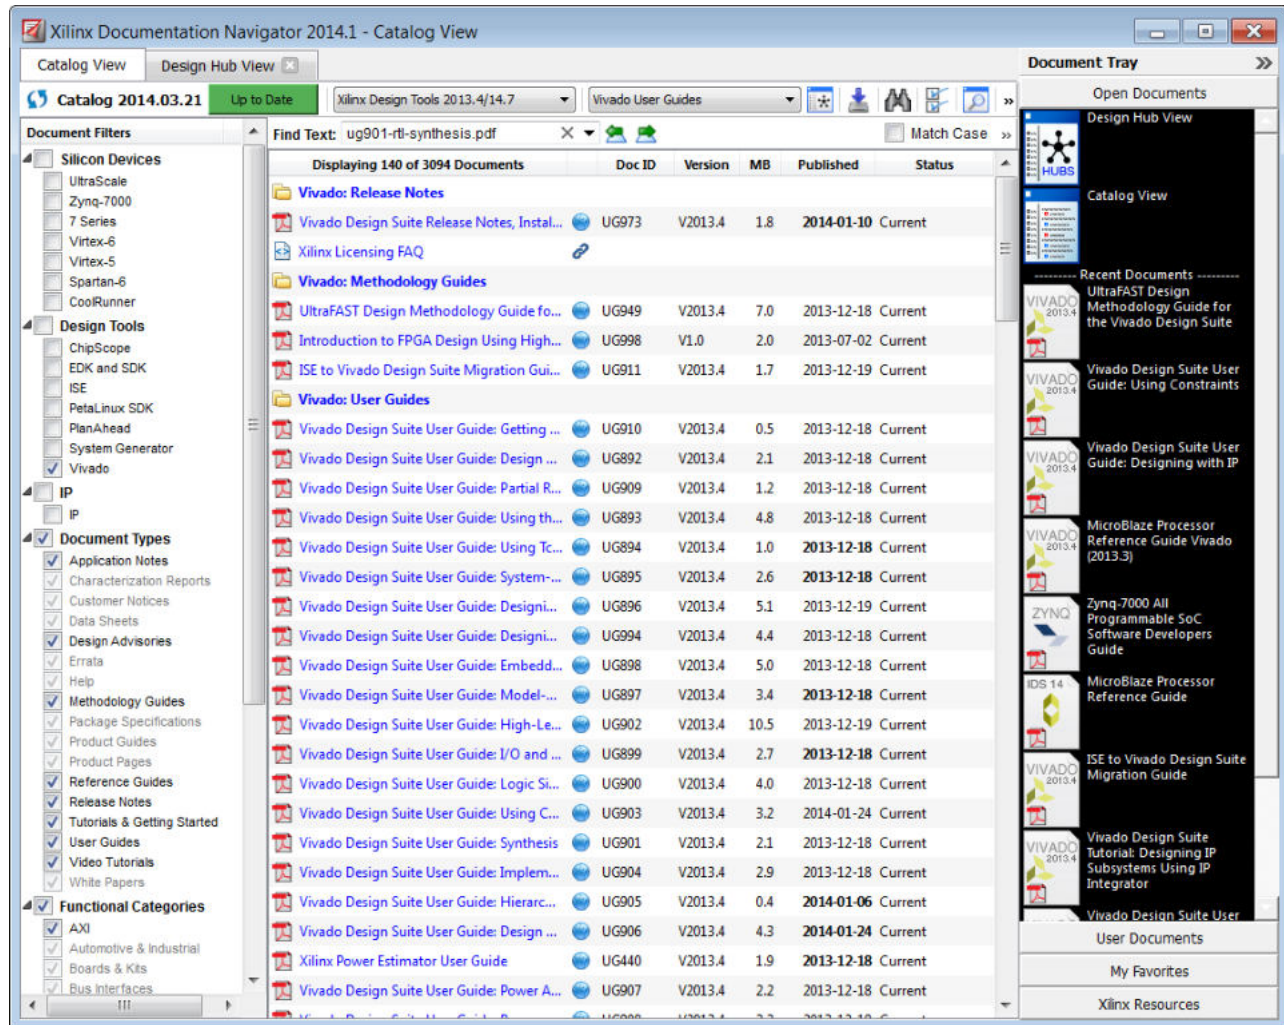


Figure 1-5: Xilinx Documentation Navigator Catalog Viewer

Accessing the QuickTake Video Tutorials

Xilinx QuickTake video tutorials provide guidance on using the features of the Vivado Design Suite. These tutorials are short and succinct training tools. They can be viewed from the [Vivado Design Suite Video Tutorials](#) page on the Xilinx website or the Xilinx [YouTube](#) channel and can be downloaded locally.



TIP: Download the clips locally if connection speed interferes with viewing quality. The QuickTake video tutorials are also available through Documentation Navigator, as shown in the following figure.

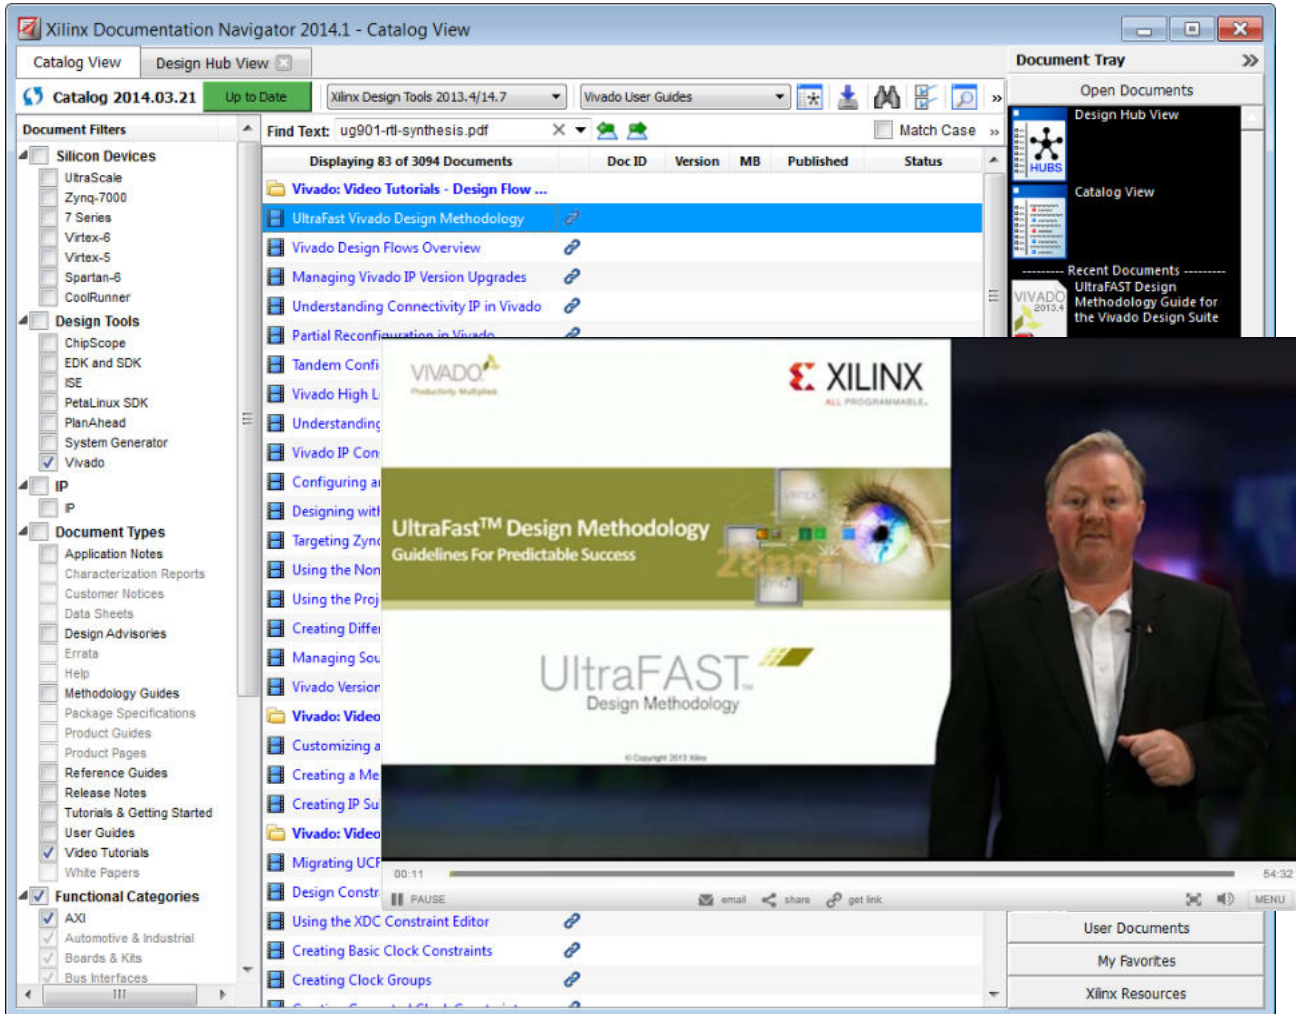


Figure 1-6: Accessing QuickTake Video Tutorials Using Documentation Navigator

In addition to QuickTake Videos, Xilinx also provides written tutorials with examples. These tutorials provide example designs and step-by-step instructions to perform specific design tasks. The [Vivado Design Suite Tutorials](#) are also available on the Xilinx website or through Documentation Navigator. In addition, you can also register for training classes offered by Xilinx or its partners.

Using the Vivado Design Suite

Overview of Using the Vivado Design Suite

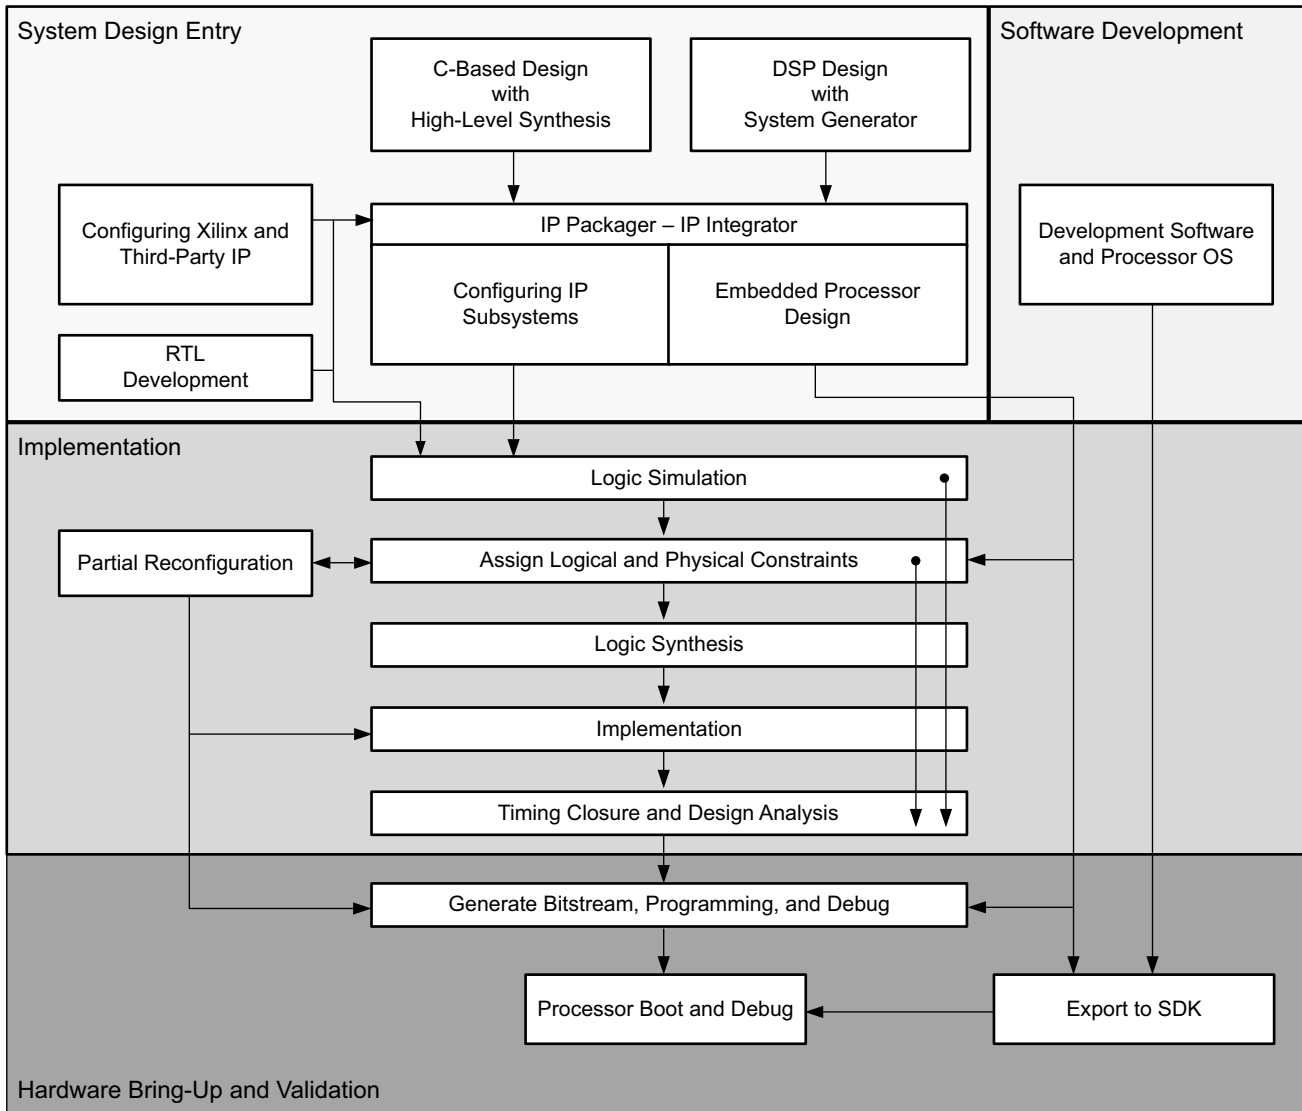
The chapter provides details about various use models, features, and options when using the Vivado® Design Suite. This includes preparation and management of the design sources and IP. For details about configuring and running the implementation tools and performing design analysis, see [Chapter 4, Design Creation](#) and [Chapter 5, Implementation](#).

You can use the Vivado Design Suite for different types of designs, such as hierarchical RTL designs, netlist-based designs, or I/O planning designs. The tool flow and related features might vary depending on the type of design. This document details the HDL-based FPGA design flow where RTL sources, IP cores, or third-party synthesized netlists are compiled through implementation, and you can then use the results to generate a bitstream to program and debug the FPGA.

The following documents and video tutorials provide additional information about Vivado Design Suite flows:

- [Vivado Design Suite QuickTake Video: Vivado Design Flows Overview](#)
- *Vivado Design Suite User Guide: Design Flows Overview* (UG892) [Ref 5]
- *Vivado Design Suite Tutorial: Design Flows Overview* (UG888) [Ref 29]
- [Xilinx Video Training: UltraFast™ Vivado Design Methodology](#)

The following figure shows the Vivado Design Suite design flow.



X15150-110315

Figure 2-1: Vivado Design Suite Design Flow

RTL to Bitstream Design Flow

You can specify RTL source files to create a project and use these sources for RTL code development, analysis, synthesis and implementation. Vivado synthesis and implementation support multiple source file types, including Verilog, VHDL, SystemVerilog, and XDC. This document focuses on proper coding and design techniques for defining hierarchical RTL sources and Xilinx® design constraints (XDC), as well as providing information on using specific features of the Vivado Design Suite, and techniques for performance improvement of the programmed design. For information on creating and working with an RTL project, see this [link](#) in the *Vivado Design Suite User Guide: System-Level Design Entry* (UG895) [Ref 8].

The Vivado Design Suite also supports the use of third-party synthesized netlists, including EDIF or structural Verilog. However, IP cores from the Vivado IP Catalog must be synthesized using Vivado synthesis, and are not supported for synthesis with a third-party synthesis tool. There are a few exceptions to this requirement, such as the memory IP for 7 series devices. Refer to the data sheet for a specific IP for more information.

By default, the Vivado Design Suite uses an out-of-context (OOC) design flow to synthesize IP cores from the Xilinx IP Catalog and block designs from the Vivado IP integrator. You can also choose to synthesize specific modules of a hierarchical RTL design as OOC modules. This OOC flow lets you synthesize, implement, and analyze design modules of a hierarchical design, IP cores, or block designs, out of the context of, or independent from the top-level design. The OOC flow is an efficient technique for supporting hierarchical team design, synthesizing and implementing IP and IP subsystems, and managing modules of large complex designs. For more information on the out-of-context design flow, see this [link](#) in the *Vivado Design Suite User Guide: Design Flows Overview* (UG892) [Ref 5].

Vivado Design Suite features behavioral, functional, and timing simulation of the design at different stages of the design flow. You can also use third-party simulators that can be integrated into and launched from the Vivado IDE. Refer to the *Vivado Design Suite User Guide: Logic Simulation* (UG900) [Ref 11] for more information.

When the synthesized netlist is available, Vivado implementation provides all the features necessary to optimize, place and route the netlist onto the available device resources of the target part. Vivado implementation works to satisfy the logical, physical, and timing constraints of the design. After implementing your design, you can analyze the results with a variety of timing and power analysis features, and run the design in hardware by programming the Xilinx device and debugging the design in the Vivado hardware manager.

The Vivado Design Suite supports several other design flows, as described in the following sections. Each of these flows is derived from the RTL-to-bitstream flow, so the implementation and analysis techniques described here are also applicable to other design flows and design types.

Embedded Processor Design Flow

A slightly different tool flow is needed when creating an embedded processor design. Because the processor requires software to boot and run effectively, the software design flow must work in unison with the hardware design flow. Data handoff between the hardware and software flows, and validation across these two domains is critical for success.

Creating an embedded processor hardware design involves the Vivado IP Integrator feature of the Vivado Design Suite. In the IP Integrator environment, you can instantiate, configure, and assemble the processor core and its interfaces. The IP Integrator enforces rules-based connectivity and provides design assistance. After it is compiled through implementation, the hardware design is exported to the Xilinx Software Development Kit (SDK) for use in the software development and validation. Simulation and debug features allow you to simulate and validate the design across the two domains.

The embedded processor design flow is described in the following resources:

- *Vivado Design Suite User Guide: Embedded Processor Hardware Design* (UG898) [Ref 10]
- *Vivado Design Suite Tutorial: Embedded Processor Hardware Design* (UG940) [Ref 33]
- *UltraFast Embedded Design Methodology Guide* (UG1046) [Ref 36]
- [Vivado Design Suite QuickTake Video: Designing with Vivado IP Integrator](#)
- [Vivado Design Suite QuickTake Video: Targeting Zynq Using Vivado IP Integrator](#)

High-Level C-based Synthesis Flow

The C-based High-Level Synthesis (HLS) tools within the Vivado Design Suite enable you to describe various DSP functions in the design using C, C++, System C, and OpenCL™ API languages. You create and validate the C code with the Vivado HLS tools. Use of higher level languages allows you to abstract algorithmic descriptions, data type, specification, etc. You can create “what-if” scenarios using various parameters to optimize design performance and device area.

HLS lets you simulate the generated RTL directly from its design environment using C-based test benches. HLS automatically uses several optimized libraries and supports floating-point arithmetic through math.h. C-to-RTL synthesis transforms the C-based design into an RTL module that can be packaged and implemented as part of a larger RTL design, or instantiated into an IP Integrator block design.

The HLS tool flow and features are described in the following resources:

- *Vivado Design Suite User Guide: High-Level Synthesis* (UG902) [Ref 17]
- *Vivado Design Suite Tutorial: High-Level Synthesis* (UG871) [Ref 28]
- Vivado High-Level Synthesis video tutorials available from the [Vivado Design Suite Video Tutorials](#) page on the Xilinx website

Partial Reconfiguration Design Flow

Partial Reconfiguration allows portions of a running Xilinx device to be reconfigured in real-time, changing the features and functions of the running design. The reconfigurable modules must be properly planned to ensure they function as needed for maximum performance. The Partial Reconfiguration flow requires a rather strict design process to ensure that the reconfigurable modules are designed properly to enable glitchless operation during partial bitstream updates. This includes reducing the number of interface signals into the reconfigurable module, floorplanning device resources, and pin placement; as well as adhering to special partial reconfiguration DRCs. The device programming method must also be properly planned to ensure the configuration I/O pins are assigned appropriately.

The partial reconfiguration tool flow and features are described in the following resources:

- *Vivado Design Suite User Guide: Partial Reconfiguration* (UG909) [\[Ref 25\]](#)
- *Vivado Design Suite Tutorial: Partial Reconfiguration* (UG947) [\[Ref 34\]](#)
- [Vivado Design Suite QuickTake Video: Partial Reconfiguration in Vivado Design Suite](#)

Vivado Design Suite Use Models

Just as the Vivado Design Suite supports many different design flows, the tools support several different use models depending on how you want to manage your design and interact with the Vivado tools. This section will help guide you through some of the decisions that you must make about the use model you want to use for interacting with the Vivado tools.

Some of these decisions include:

- Are you a script or command-based user; or do you prefer working through a graphical user interface (GUI)?
- Do you want to configure IP cores for use with a single design project; or establish a remote repository to reuse configured IP cores across multiple projects?
- Do you want the Vivado Design Suite to manage the design sources, status, and results by using a project structure; or would you prefer to quickly create and manage a design yourself?
- Are you managing your source files inside a revision control system?
- Are you using third-party tools for synthesis or simulation?



RECOMMENDED: Before beginning your first FPGA design with the Vivado tools, see the Vivado Design Suite User Guide: Getting Started (UG910) [Ref 12] and the Vivado Design Suite User Guide: Design Flows Overview (UG892) [Ref 5].

Understanding Project and Non-Project Software Use Models

The Vivado Design Suite has two primary use models: Project Mode and Non-Project Mode. Both use models can be run through the Vivado IDE, although the Vivado IDE offers many benefits for the Project Mode. Both use models can be run through Tcl commands or scripts, although Tcl commands are the simplest way to run the Non-Project Mode.

Note: You can use a Tcl script based flow, but also use the Vivado IDE when needed to create physical and timing constraints, or view results during design analysis.

Using Project Mode

In Project Mode, Vivado tools automatically manage your design flow and design data. The key advantage of using Project Mode is design process automation with push-button implementation.



TIP: The key advantage of Project Mode is that the Vivado Design Suite manages the entire design process, including dependency management, report generation, data storage, etc.

The project tracks and reports on the design status, source file dependencies, and implementation results. When working in Project Mode, the Vivado Design Suite creates a directory structure on disk in order to manage design source files, either locally or remotely, and manage changes and updates to the source files. The project infrastructure is also used to manage the automated synthesis and implementation runs, track run status, and store synthesis and implementation results and reports. For example:

- If you modify an HDL source after synthesis, the Vivado Design Suite identifies the current results as out-of-date, and prompts you for resynthesis.
- If you modify design constraints, the Vivado tools prompt you to either re-synthesize, re-implement, or both.
- After routing is completed, the Vivado tool automatically generates timing and power reports.
- The entire design flow can be run with a single click within the Vivado IDE.



IMPORTANT: Certain operating systems (for example, Microsoft Windows) restrict the number of characters (such as 256) that can be used for the file path and file name. If your operating system has such a limitation, Xilinx recommends that you create projects closer to the drive root to keep file paths and names as short as possible.

Using Non-Project Mode

In Non-Project Mode, you manage design sources and the design process yourself using Tcl scripts. The key advantage is that you have full control over each step of the flow. You can generate design checkpoints and reports at will. Each implementation step can be tailored to meet specific design challenges, and you can analyze results after each design step. Design sources are generally accessed from their current locations, such as a revision control system, rather than copying them into a local project directory structure. As the design flow progresses, the representation of the design is retained in memory that is allocated to the Vivado Design Suite. In other words, all of the design is stored in-memory and updated in real-time throughout the design flow.



TIP: *In Non-Project Mode, a project infrastructure is created in memory to process the design, but the project is not written to disk.*

In Non-Project Mode, each design step is controlled using Tcl commands. For example:

- If you modify an HDL file after synthesis, you must remember to rerun synthesis to update the in-memory netlist.
- If you want a timing report after routing, you must explicitly generate the timing report when routing completes.
- Design parameters and implementation options are set using Tcl commands.
- You can save design checkpoints and create reports at any stage of the design process using Tcl.

You must write reports or design checkpoints to save the in-memory design as it progresses. The design checkpoint (DCP) refers to a file that is an exact representation of the in-memory design. You can save a design checkpoint after each step in the design flow, such as post synthesis, post optimization, post placement. The DCP file can be read back into the Vivado Design Suite to restore the design to the state captured in the checkpoint file.

You can also open a DCP in the Vivado IDE to perform interactive constraints assignment and design analysis. Because you are viewing the active design in memory, any changes are automatically passed forward in the flow. You can also save updates to new constraint files or design checkpoints for future runs.

While most Non-Project Mode features are also available in Project Mode, some Project Mode features are not available in Non-Project Mode. These features include source file and run results management, saving design and tool configuration, design status, and IP integration. On the other hand, you can use Non-Project mode to skip certain processes, thereby reducing the memory footprint of the design, and saving disk space related to projects.

Working with the Vivado Integrated Design Environment (IDE)

The Vivado Integrated Design Environment (IDE) can be used in both Project Mode and Non-Project Mode. The features displayed in the IDE vary depending on how and when you invoke the IDE. For more information on the Vivado IDE, see the *Vivado Design Suite User Guide: Using the Vivado IDE* (UG893) [Ref 6].

The Vivado IDE enables analysis and constraints assignment throughout the design process by opening designs in memory. Opening a design loads the design netlist at that particular stage of the design flow, assigns the constraints to the design, and applies the design to the target device. This process allows you to visualize and interact with the design at each design stage, as shown in the following figure.

When using Project Mode, the Vivado IDE provides an interface called Flow Navigator, to assemble, implement, and validate your design and IP. In addition, the Vivado IDE supports a push-button design flow that manages all design sources, configuration, and results.

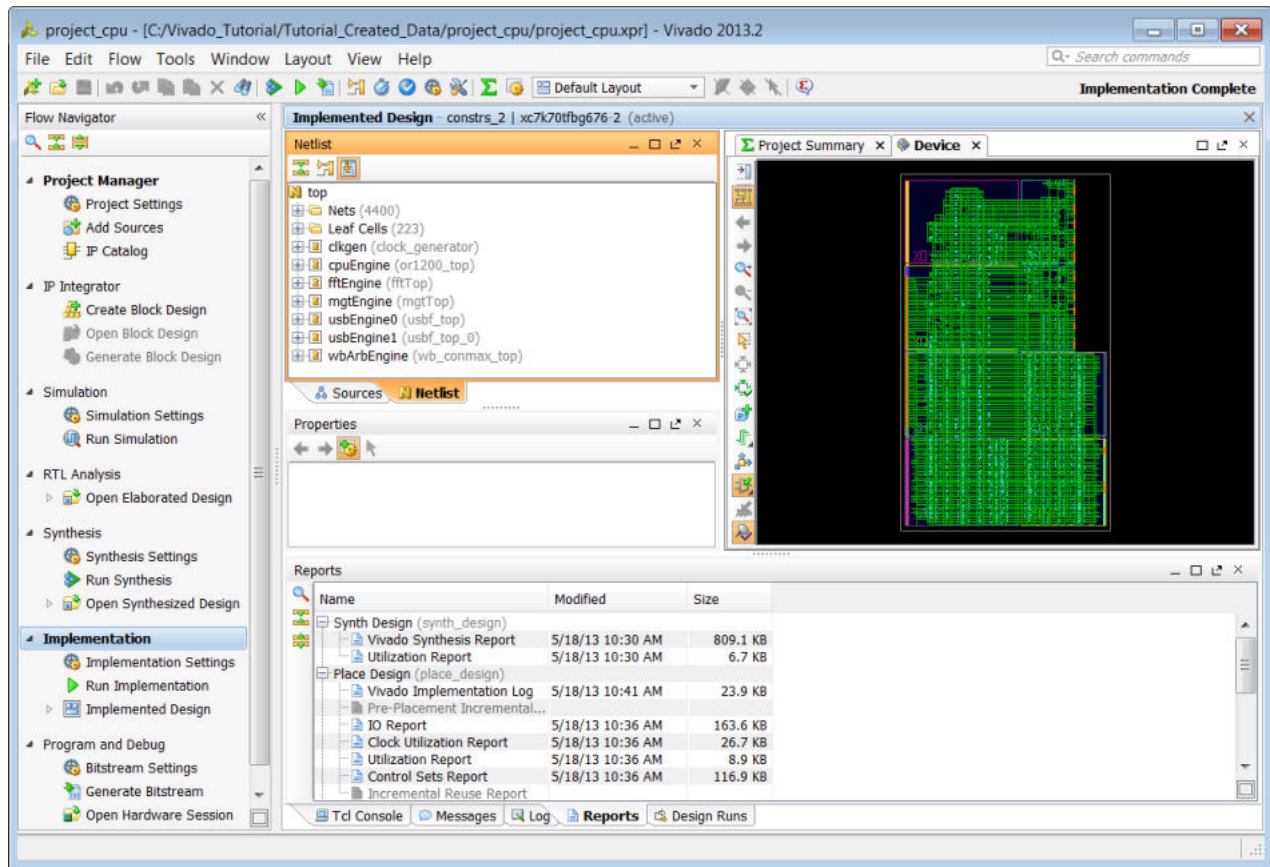


Figure 2-2: Opening the Implemented Design in the Vivado IDE

For more information on the Vivado IDE, see the *Vivado Design Suite User Guide: Using the Vivado IDE* (UG893) [Ref 6].

Working with Tcl

All supported design flows and use models can be run using Tcl commands. You can use either individual Tcl commands or saved scripts of Tcl commands. You can use Tcl scripts to run the entire design flow, including design analysis and reporting, or to run parts of the design flow, such as design creation and synthesis.

If you prefer working directly with Tcl commands, you can interact with your design using a Vivado Design Suite Tcl shell, using the Tcl Console from within the Vivado IDE. For more information about using Tcl and Tcl scripting, see the *Vivado Design Suite User Guide: Using Tcl Scripting* (UG894) [Ref 7]. For a step-by-step tutorial that shows how to use Tcl in the Vivado tools, see the *Vivado Design Suite Tutorial: Design Flows Overview* (UG888) [Ref 29].

When working with Tcl, you can still take advantage of the interactive GUI-based analysis and constraint definition capabilities of the Vivado IDE. From a Tcl script you can save design checkpoints at any time after synthesis, and then open the checkpoints later in the Vivado IDE.

Configuring IP

The Vivado IDE provides an IP-centric design flow that lets you add IP modules to your design from various design sources. The Vivado IDE also provides an extensible IP Catalog that includes Xilinx-delivered Plug-and-Play IP. The IP Catalog can be extended by adding any of the following:

- Modules from System Generator for DSP designs (MATLAB[®] from Simulink[®] algorithms)
- Vivado High-Level Synthesis (HLS) designs (C/C++ algorithms)
- Third-party IP
- Designs packaged as IP using the Vivado IP Packager

The IP Catalog provides a consistent and easy access to Xilinx IP, including building blocks, wizards, connectivity, DSP, embedded, AXI infrastructure and video IP all from a single common repository, regardless of the end application being developed. The IP Catalog also lists Alliance Partner IP with links for licensing information.

The majority of the IP available through the IP Catalog can be used in either an RTL project or in an IP Integrator block design. There is a small number of IP that are designed for use only in IP Integrator as well as a small number designed for use in an RTL project. A message alerts you if you attempt to use an IP in an unsupported flow.

The Vivado IP Catalog displays either **Included** or **Purchase** under the License column in the IP Catalog. The following definitions apply to IP offered by Xilinx:

- **Included:** The Xilinx End User License Agreement includes Xilinx LogiCORE™ IP cores that are licensed within the Xilinx Vivado Design Suite software tools at no additional charge.
- **Purchase:** The Core License Agreement applies to fee-based Xilinx LogiCORE IP, and the Core Evaluation License Agreement applies to the evaluation of fee-based Xilinx LogiCORE IP.

The IP with a license type of Purchase have different possible functionality. To learn more about licensing for a particular IP, select the IP in the catalog, right-click, and select **License Status**. The following table shows the possible status of the Purchase IP.

Table 2-1: Purchase IP Status

Status	Rights Policy	Description
Bought	Full	Licenses customer to fully elaborate the core with full functionality
Hardware_evaluation	Evaluation	Licenses customer to fully elaborate the core with restricted functionality (for example, hardware timeout circuitry)
Design_Linking	Simulation	Licenses customer to elaborate the core for simulation purposes only

This license status information is available for IP cores used in a project using Report IP Status by selecting **Tools > Report > Report IP Status**. For additional information on how to obtain IP licenses, see the [Xilinx IP Licensing](#) page on the Xilinx website.

Xilinx and its partners provide additional IP cores that are not shipped as part of the default Vivado IP Catalog. For more information on the available IP, see the [Intellectual Property](#) page on the Xilinx website.

The following sections discuss how to configure, use, and manage IP cores.

Using the Vivado IP Catalog

IP is best configured using the IP Catalog features of the Vivado IDE, which make it easy to browse, configure, generate output products, and validate the IP. The IP Catalog and configuration wizards of the Vivado IDE make the job easy. You can also access the IP documentation, such as product guides, change logs, and answer records (if applicable) directly from the IP Catalog.

There are Tcl equivalent commands that enable scripting of IP customization, but not all Tcl parameters for IP configuration are documented. If scripting is desired, you can use the Vivado Design Suite journal file to create a script after you have used the IDE to configure the IP and generate output products.

After using the configuration wizards to customize the IP, a Xilinx Core Instance (.xci) file is created. This file contains all the customization options for the IP. From this file the tool can generate all output products for the IP. These output products consist of HDL for synthesis and simulation, constraints, possibly a test bench, C modules, example designs, etc. The tool creates these files based upon the customization options used.

Because IP is generated for specific logic devices, naming the IP with descriptive names may help identify them later.

For more information, see:

- [Updating IP](#)
- [Managing IP](#) (explains the available options for storing and using the IP configurations)
- *Vivado Design Suite User Guide: Designing with IP* (UG896) [\[Ref 9\]](#)

Generating IP Output Products

IP output products are created to enable synthesis, simulation, and implementation tools to use the specific configuration of the IP. While generating output products, a directory structure is set up to store the various output products associated with the IP. The folders and files are fairly self-explanatory and should be left intact.

Using IP Core Containers

To facilitate interactions with revision control systems, you can store IP configuration files and output products in a single, binary IP core container file rather than a directory structure. The Vivado Design Suite interacts directly with the IP core container files. For more information on using IP core containers, see the *Vivado Design Suite User Guide: Designing with IP* (UG896) [\[Ref 9\]](#). For more information on revision control interaction with core containers, see [Managing IP Sources](#).

Netlisting Options

You have an option on the level of data you wish to create for IP when generating output products. The options you choose affect how the design is implemented.

- Using the default setting of **Out-of-Context** (OOC) will perform logic synthesis on just the specific customized version of the IP during generation of output products to create a synthesized design checkpoint file (.dcp). This netlist is then used during implementation along with the constraints that the IP delivers. Synthesis of the IP just by itself is referred to as Out-of-Context (OOC) synthesis.

This method is also required when using third-party synthesis tools for the rest of the design and a module stub file is created for them to infer a black box for the IP. Functional simulation models are also produced in both Verilog and VHDL. If you do not have a mixed language simulation tool then the functional netlist is available. When

using the Vivado Design Suite to launch a third-party simulator either HDL or a functional netlist will be used as appropriate.

- Selecting **Global** will generate the RTL and XDC IP source files for the IP. These are then used along with the rest of the design sources during top-level synthesis and implementation of the design.



TIP: You can also use OOC synthesis results for implementation (for module analysis or to preserve timing). Implementing an OOC module requires additional constraints (for example, HD.CLK_SRC) to ensure accurate timing results. For more information refer to Vivado Design Suite User Guide: Hierarchical Design (UG905) [Ref 20].

Xilinx recommends that for each IP you customize you should generate all available output products, including a synthesized design checkpoint. Doing so provides you with a complete representation of the IP that can be archived or placed in revision control. If future Vivado Design Suite versions do not include that IP, or if the IP has changed in undesirable ways (such as interface changes), you have all the output products required to simulate, and to use for synthesis and implementation with future Vivado Design Suite releases.

IP Constraints

Most of the IP cores include XDC constraints that are used during synthesis and implementation. These constraints are used automatically either in Project Mode or Non-Project Mode if the IP is used by means of the XCI created during customization. Manually modifying IP constraints to work at the top level can be error prone and tedious.

Many IP cores reference their input clocks in their constraints. These clocks can come either from the user at the top level, or even other IP cores in a design. By default, the Vivado tools process any IP clock creation and any user-defined top-level clock creation early. This process makes these clocks available to the IP cores that require them.

Validating IP

Many of the Xilinx IP delivered in the Vivado IP Catalog have an example design. You can determine if an IP comes with an example design by selecting the IP from the IP Sources area of the Manage IP or RTL project and see if **Open IP Example Design** is selectable, as shown in the following figure. This can also be done via Tcl.

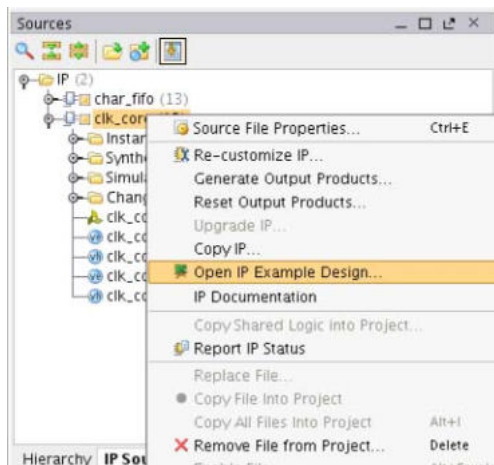


Figure 2-3: Opening an Example Design

When opening the example design a new project is created which uses the user-customized IP. You can refer to the example design to understand a valid usage and connectivity for the specific customization of the IP.

Some IP deliver test benches with the example design, which you can use to validate the customized IP functionality. You can run behavioral, post synthesis, or post-implementation simulations. You can run either functional or timing simulations. In order to perform timing/functional simulations you need to synthesize/implement the example design.

For specific information on simulating an IP, refer to the product guide for the IP. For more detail on simulation, refer to the *Vivado Design Suite User Guide: Logic Simulation* (UG900) [Ref 11]. For more details on working with example designs and IP output products, refer to the *Vivado Design Suite User Guide: Designing with IP* (UG896) [Ref 9].

If you have memory IP in your design, see the following resources:

- For details on simulation, see the *LogiCORE IP UltraScale Architecture-Based FPGAs Memory Interface Solutions Product Guide* (PG150) [Ref 48].
- For an example of simulating memory IP with a MicroBlaze™ processor design, see the *Reference System: Kintex-7 MicroBlaze System Simulation Using IP Integrator* (XAPP1180) [Ref 46].

Using Memory IP

Additional I/O pin planning steps are required when using Xilinx memory IP. After the IP is customized, you then assign the top-level I/O ports to physical package pins in either the elaborated or synthesized design in the Vivado IDE.

All of the ports associated with each memory IP are grouped together into an I/O Port Interface for easier identification and assignment. A Memory Bank/Byte Planner is provided

to assist you with assigning Memory I/O pin groups to Byte lanes on the physical device pins.

For more information, see this [link](#) in the *Vivado Design Suite User Guide: I/O and Clock Planning* (UG899) [Ref 4].

Creating IP Subsystems with IP Integrator

This section discusses how to create and manage Intellectual Property (IP) Subsystems.

Vivado IP Subsystems

The IP Integrator feature of the Vivado Design Suite enables the creation of Block Designs (.bd). These block designs are essentially IP Subsystems containing any number of user-configured IP and connections between them. For packaging and using your own custom IP along with AXI interface, see [Packaging Custom IP and IP Subsystems](#) and [Creating Custom Peripherals](#). The IP Integrator is the interface for connecting IP cores to create domain specific subsystems and designs, including embedded processor-based designs using Zynq®-7000 All Programmable (AP) SoCs and MicroBlaze processors. It can instantiate High-Level Synthesis modules from Vivado HLS, DSP modules from System Generator, and user custom IP made available using the Package IP feature.

For more information, see the following resources:

- [Vivado Design Suite User Guide: Designing IP Subsystems Using IP Integrator](#) (UG994) [Ref 26]
- [Vivado Design Suite QuickTake Video: Designing with Vivado IP Integrator](#)

Building IP Subsystems

IP Subsystems are best configured using the IP Integrator feature of the Vivado IDE. The interactive block design capabilities of the IP Integrator make the job of configuring and assembling groups of IP easy.

If scripting is desired, you can use the Vivado Design Suite journal file after you have used the Vivado IDE to create subsystems. The IP Integrator can create a Tcl script to re-create the current block design in memory. You can also use a combination of GUI, Tcl shell commands, and scripts to create an IP subsystem.

An IP subsystem can be configured so that when it is imported into a design project it can exist as a set of HDL sources and constraints (default) or as a design checkpoint file (.dcp), which contains a synthesized netlist and constraints for the entire subsystem.

Designer Assistance

To expedite the creation of a subsystem or a design, use the Block Automation and Connection Automation features of the IP integrator. The Block Automation feature should be used to configure a basic processor-based design and some complex IP subsystems whereas the Connection Automation feature should be used to make repetitive connections to different pins/ports of the design. IP integrator is also board aware and currently supports all the Xilinx evaluation boards, which means that if you use an evaluation board as the target hardware, IP integrator knows all the components present on that particular board, thereby allowing you to use the Connection Automation feature to connect the I/O ports of the design to an existing component on that target board. Designer assistance also helps with clocks and resets connectivity. Several tabs, such as the Signals tab and the Board tab, aid you with making connections in the block design. Using Designer Assistance not only helps expedite the interconnectivity but also eliminates unintended design errors.

Block Automation

For some complex IP- and processor-based designs, IP integrator offers a feature called Block Automation. This feature allows you to put together a basic processor/IP-based subsystem with commonly used components relatively quickly. After the basic building blocks for an embedded design are put together, you can build upon this basic system by adding other IP from the catalog.

Connection Automation

After the block automation is done and a basic system is in place, you need to connect the design to external I/O pins. The IP integrator Connection Automation feature not only helps make connections to the I/O pins but also helps make connections to different sources on the design itself. Combined with board awareness, the Connection Automation feature can help you connect the ports of the block design to external components present on the target board and create physical constraints for these ports.

Rule-Based Connection Checking

IP Integrator runs basic design rule checks in real time as the design is being assembled. However, there is a potential for something to go wrong during design creation. For example, the frequency on a clock pin may be set incorrectly. The tool can catch these types of errors by running design validation. Validating the design runs design rule checks on the block design and reports warnings and/or errors that are applicable to the design. You can then cross-probe the warnings and/or errors from the messages view into the block diagram. Xilinx recommends validating a board design to catch design errors that would otherwise be unnoticed until later in the design flow.

Running design validation also runs Parameter Propagation on the block design. Parameter Propagation enables an IP to auto-update its parameterization based on its connections in the design. You can package IP with specific propagation rules, and IP Integrator then runs these rules as the block diagram is generated.

You can run design validation by selecting **Tools > Validate Design** or through the Tcl command `validate_bd_design`. The IP integrator canvas also has an icon for Validate Design.

Using the Platform Board Flow

The Vivado Design Suite is board aware and can automatically derive I/O constraints and IP configuration data from the included board files. Through the board files, the Vivado Design Suite knows the various components present on the target boards and can customize and configure an IP to be connected to a particular board component. Several 7 series, Zynq-7000 AP SoC, and UltraScale™ device boards are currently supported. You can download support files for partner-developed boards from the partner websites.

The IP Integrator shows all the component interfaces on the target board in a separate tab called the Board tab. You can use this tab to select the desired components and the Designer Assistance offered by IP Integrator to easily connect your design to the board component of your choice. Interfaces can use the drag-and-drop method individually or as a group to utilize the default IP configurations. Double-clicking the interface uses the interactive method to add the interface with optional IP and configurability when available to the IP Integrator diagram. All the I/O constraints are automatically generated as a part of using this feature.

You can also generate board files for custom boards and add the repository that contains the board file to a project. For more information on generating a custom board file, see this [link](#) in the *Vivado Design Suite User Guide: System-Level Design Entry* (UG895) [Ref 8].

Creating Hierarchical IP Subsystems

IP Integrator can be used to create hierarchical IP subsystems. This feature is useful for designs with a large number of blocks, which could otherwise become harder to manage on the GUI canvas. The tool supports multiple levels of hierarchy, allowing you to group the blocks based on design functionality, thereby keeping the design modular and neat on the IP Integrator canvas.

You can also change the visual aspects of different objects in the design. For example, clocks and resets could be colored differently. You can also choose to show different types of connections while hiding others.

Generating Block Designs

After the block design or IP subsystem has been created, you can generate the design, which includes generation of all source codes, necessary constraints for the IP cores, and the structural netlist of the block design. You can generate the block design by right-clicking on the block design (in the Sources window) and selecting **Generate Output Products** from the pop-up menu. In the Vivado Design Suite Flow Navigator you can also select **IP Integrator > Generate Block Design**. The equivalent Tcl command is:

```
generate_target all [get_files <path to the block design>]
```

After this step is complete, the design is ready to be integrated in a higher level HDL design or to be taken through synthesis and implementation.

Using Out-of-Context Synthesis for Block Designs

Hierarchical design flows enable the partitioning of the design into smaller, more manageable modules that can be processed independently. In the Vivado Design Suite, these flows are based on the ability to synthesize a partitioned module out-of-context (OOC) from the rest of the design.

There are two modes of OOC supported for block designs in the Vivado Design Suite. Select the **Out of context per Block design** when using a third-party synthesis flow for the top-level design. This is a common flow with IP Integrator that creates a design checkpoint (DCP) file. If used as part of the larger design, this block design does not need to be re-synthesized every time you modify other parts of the design (outside of IP Integrator).

Select the **Out of context per IP** option if you want to create an OOC run for each of the IP instantiated on the block design. This option creates a design checkpoint file for each IP in the block design. In this flow, output products for only the IP whose configuration have changed, either because of re-customization or because of an impact from parameter propagation, is generated. This flow improves run-time and can address situations where run-times are a matter of concern, particularly during the early stages of design exploration. IP Caching can be used with this option to prevent IP from regenerating output products if they have not changed.

Creating Remote Block Designs (Recommended)

The reuse feature of IP Integrator facilitates team-based design by allowing you to create block designs outside of a Vivado Design Suite project for reuse by multiple teams. Once you create the design and put it under revision control, multiple teams can reuse the same block design for creating multiple projects.

To create a block design at a remote location, specify the desired location in the drop-down Directory list of the Create Block Design dialog box.

Validating IP Subsystems

After the design is complete, a comprehensive design rule check can be performed on the block design by selecting **Validate Design** from the toolbar in the IP Integrator canvas. This action carries out parameter propagation and checks for any parameter mismatches between different endpoints of the design.

Integrating the Block Design into a Top-Level Design

An IP Integrator block design can be integrated into a higher-level design or it can be the highest level in the design hierarchy. To integrate the IP Integrator design into a higher-level design, instantiate the design in the higher-level HDL file.

You can perform a higher-level instantiation of the block design by selecting the block design in the Vivado IDE Sources window and selecting **Create HDL Wrapper**. This generates a top-level HDL file for the IP Integrator sub-system.

Packaging Custom IP and IP Subsystems

The Vivado IP Packager gives you the ability to create custom IP for delivery in the Vivado IP Catalog. The industry standard IP-XACT format is used for packaging the IP. The location of the packaged IP can be added to the Repository Manager section of the Vivado Design Suite Project Settings. After a repository of one or more IP has been added, the IP core(s) will be shown in the IP Catalog. You can now select and customize the IP visible in the Vivado IP Catalog. Here is an overview of the flow to use the Vivado IP Packager:

1. Use the Create and Package IP wizard to package the HDL and associated data files of the custom IP.
2. Provide the packaged custom IP to a user for the IP.
3. The end-user adds the IP location to the Repository section of the Vivado Design Suite Project Settings.
4. The IP is now visible in the IP Catalog and the end-user can select and customize the IP similar to Xilinx-delivered IP.

The Create and Package IP wizard takes you step-by-step through the IP packaging steps and lets you package IP from a project, a block design, or a specified directory or to create and package a new template AXI4 peripheral.



IMPORTANT: *The directory structure of the custom IP needs to be set up so that all the HDL files forming the IP definition are below the directory level that is being packaged. It is possible for the tool to refer to HDL files at higher directory levels through absolute paths, but this can make the packaged IP non-portable across networks.*

Using the IP Packager allows the end-user to have a consistent experience whether using Xilinx IP, third-party IP, or a custom IP. For more detailed information on creating and packaging IP refer to the *Vivado Design Suite User Guide: Creating and Packaging Custom IP* (UG1118) [Ref 27].



IMPORTANT: *Ensure that the desired list of supported device families is defined properly while creating the custom IP definition. This is especially important if you want your IP to be used with multiple device families.*



TIP: *Before packaging your IP HDL, ensure its correctness by simulating and synthesizing to validate the design.*

During creation stage of packaging a custom IP, another instance of the Vivado IDE might open with `edit_ip` project. This project is a temporary cache and the tools can clear it immediately after packaging the IP.

Creating Custom Peripherals

The Vivado Design Suite requires that all memory-mapped interfaces use an AXI interface. The Vivado Design Suite offers an option in the Create and Package IP Wizard to facilitate the creation of a custom IP that adheres to the AXI interface standard. The Create and Package IP Wizard can generate three types of AXI interfaces:

- AXI4: For memory-mapped interfaces that allow bursts of up to 256 data transfer cycles with a single address phase
- AXI4-Lite: A lightweight, single transaction memory-mapped interface
- AXI4-Stream: An AXI interface that removes the requirement for an address phase and allows unlimited data burst sizes

If you already have the core functionality of your IP, you can use the Create and Package IP Wizard to generate AXI interface logic for your custom IP. The Create and Package IP Wizard can create a template AXI4 peripheral that includes HDL, drivers, test application, Bus Functional Model (BFM), and an example template. After the peripheral has been created, the user design files can be added to complete the custom IP. Refer to the *Vivado Design Suite User Guide: Creating and Packaging Custom IP* (UG1118) [Ref 27] and *Vivado Design Suite User Guide: Designing IP Subsystems Using IP Integrator* (UG994) [Ref 26] for additional details.

The Xilinx LogiCORE™ IP AXI Bus Functional Model (BFM) cores, developed for Xilinx by Cadence Design Systems, support the simulation of customer-designed AXI-based IP. AXI BFM cores support all versions of AXI (AXI3, AXI4, AXI4-Lite, and AXI4-Stream). The BFMs are encrypted Verilog modules. BFM operation is controlled by using a sequence of Verilog tasks contained in a Verilog-syntax text file. This core can be instantiated on a block design

and connected to the custom IP to check AXI functionality. For more information on the BFM Cores, see the *AXI BFM Cores LogiCORE IP Product Guide* (PG129) [Ref 15].

Logic Simulation

The following figure shows all the places where Vivado simulation should be used for functional and timing simulation.

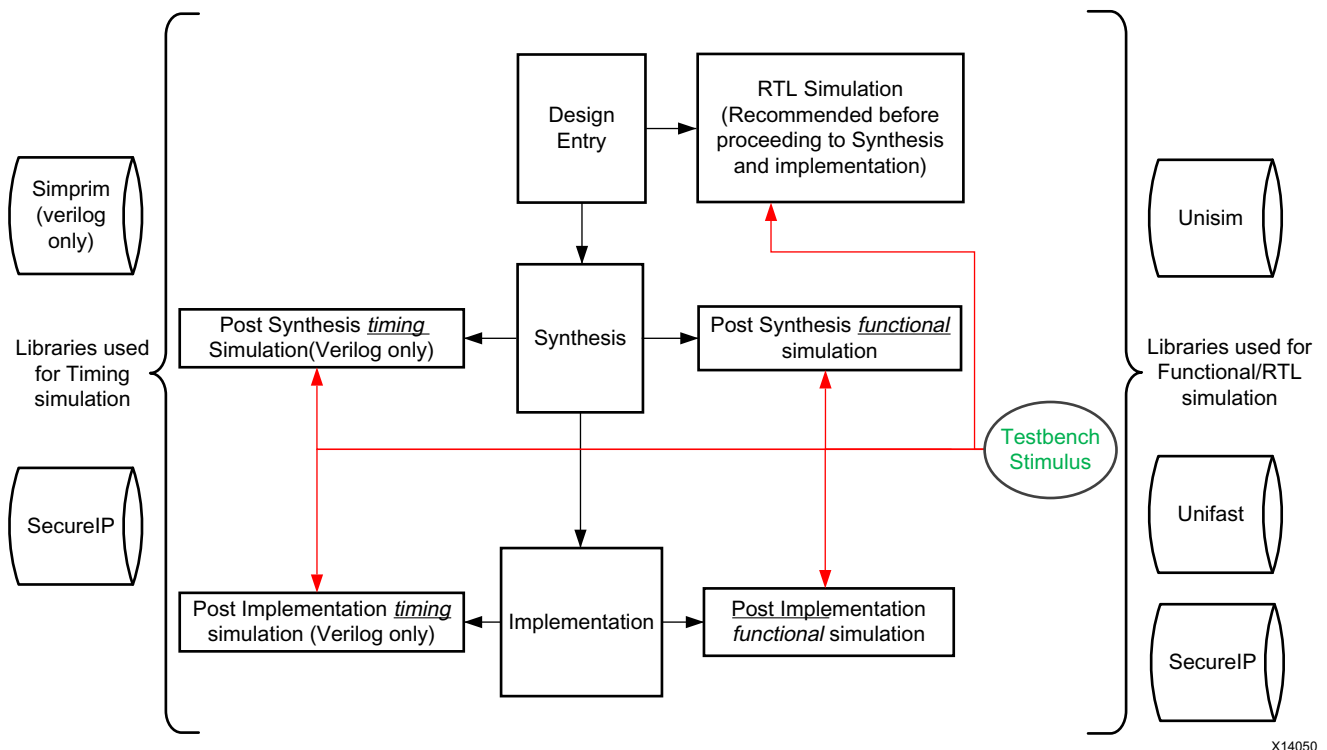


Figure 2-4: Simulation at Various Points in the Design Flow

Functional Simulation Early in the Design Flow

Use functional or register transfer level (RTL) simulation to verify syntax and functionality. This first pass simulation is typically performed to verify the RTL or behavioral code and to confirm that the design is functioning as intended.

With larger hierarchical designs, you can simulate individual IP, block designs, or hierarchical modules before testing your complete design. This simulation process makes it easier to debug your code in smaller portions before examining the larger design.



RECOMMENDED: At this stage, no timing information is provided. Xilinx recommends performing simulation in unit-delay mode to avoid the possibility of a race condition.

Use synthesizable HDL constructs for the initial design creation. Do not instantiate specific components unless necessary. This allows for:

- More readable code
- Faster and simpler simulation
- Code portability (the ability to migrate to different device families)
- Code reuse (the ability to use the same code in future designs)



TIP: *You might need to instantiate components if the components cannot be inferred.*

Instantiation of components might make your design code architecture specific. These instantiated components might include:

- Instantiated UNISIM library components
- Instantiated UniMacro components
- SecureIP

When each module simulates as expected, create a top-level design test bench to verify that your entire design functions as planned. Use the same test bench again for the final timing simulation to confirm that your design functions as expected under worst-case delay conditions.

Using Structural Netlists for Simulation

After you synthesize/implement your design, you can perform netlist simulation in functional or timing mode. The netlist simulation can also help you with the following:

- Identify post-synthesis and post-implementation functionality changes caused by:
 - Synthesis attributes or constraints that create mismatches (such as `full_case` and `parallel_case`)
 - UNISIM attributes applied in the Xilinx Design Constraints (XDC) file
 - Differences in language interpretation between synthesis and simulation
 - Dual-port RAM collisions
 - Missing or improperly applied timing constraints
 - Operation of asynchronous paths
 - Functional issues due to optimization techniques
- Sensitize timing paths declared as false or multi-cycle during STA
- Generate netlist switching activity to estimate power
- Identify X state pessimism

For netlist simulation, use one or more of the libraries shown in the following table.

Table 2-2: Use of Simulation Library

Library Name	Description	VHDL Library Name	Verilog Library Name
UNISIM	Functional simulation of Xilinx primitives	UNISIM	UNISIMS_VER
UNIMACRO	Functional simulation of Xilinx macros	UNIMACRO	UNIMACRO_VER
UNIFAST	Fast simulation library	UNIFAST	UNIFAST_VER

The UNIFAST library is an optional library that you can use during functional simulation to speed up simulation runtime. UNIFAST libraries are supported for 7 series devices only. UltraScale and later device architectures do not support UNIFAST libraries, because all the optimizations are incorporated in the UNISIM libraries by default.



IMPORTANT: You cannot use the UNIFAST model for timing simulations.



RECOMMENDED: Xilinx recommends using the UNIFAST library for initial verification of the design. For a complete verification use the UNISIM library.

For more information on Xilinx simulation libraries, see this [link](#) in the *Vivado Design Suite User Guide: Logic Simulation* (UG900) [Ref 11].

The following table shows the location of the simulation libraries.

Table 2-3: Location of Simulation Library

Library	HDL Type	Location
UNISIM	Verilog	<Vivado_Install_Area>/data/verilog/xsim/unisims
	VHDL	<Vivado_Install_Area>/data/vhdl/xsim/unisims
UNIFAST	Verilog	<Vivado_Install_Area>/data/verilog/src/unifast
	VHDL	<Vivado_Install_Area>/data/vhdl/src/unifast
UNIMACRO	Verilog	<Vivado_Install_Area>/data/verilog/src/unimacro
	VHDL	<Vivado_Install_Area>/data/vhdl/src/unimacro
SECUREIP	Verilog	<Vivado_Install_Area>/data/secureip/<simulator>/<simulator>_secureip_cell.list.f

Note: The Vivado tools allow you to run functional and/or timing simulation at the synthesis and implementation stage. For more information on netlist generation, refer to the `write_verilog` and `write_sdf` commands in the *Vivado Design Suite Tcl Command Reference Guide* (UG835) [Ref 13].

Primitives/elements of the UNISIM library do not have any timing information except the clocked elements. To prevent race conditions during functional simulation, clocked elements have a clock-to-out delay of 100 ps. Waveform views might show spikes and glitches for combinatorial signals, due to lack of any delay in the UNISIM elements.

Timing Simulation

Many users do not run timing simulation due to high run time. If you decide to skip timing simulation, you should make sure of the following:

- Ensure that your STA constraints are absolutely correct. Pay special attention to exceptions.
- Ensure that your netlist is exactly equivalent to what you intended through your RTL. Pay special attention to any inference-related information provided by the synthesis tool.

However, consider using timing simulation because simulation with full timing is the closest method of mimicking hardware behavior. If your design does not work on hardware, it is much easier to debug the failure in simulation, as long as you have a timing simulation that can reproduce the failure.

Running Timing Simulation in the Vivado Design Suite

Xilinx supports timing simulation in Verilog only. You can generate the timing simulation netlist using the `write_verilog` Tcl command. The Verilog system task `$sdf_annotate` within the simulation netlist specifies the name of the standard delay format (SDF) file to be read. The Vivado simulator automatically reads the SDF file when the simulator compiles the Verilog simulation netlist.



TIP: *The Vivado simulator supports mixed-language simulation, which means that if you are a VHDL user, you can generate a Verilog simulation netlist and instantiate it from the VHDL test bench.*

Simulation Time Resolution

Xilinx recommends that you run simulations using a resolution of 1 ps. Some Xilinx primitive components, such as DCM, require a 1 ps resolution to work properly in either functional or timing simulation.



TIP: *Because most of the simulation time is spent in delta cycles, there is no significant simulator performance gain by using coarser resolution with the Xilinx simulation models.*



RECOMMENDED: *Xilinx recommends that you do not use a finer resolution, such as fs. Some simulators might round the numbers, while other simulators might truncate the numbers.*

Simulating Global Set/Reset (GSR)

Xilinx devices have dedicated routing and circuitry that connect to every register in the device. When you assert the dedicated global set/reset (GSR) net, that net is released

immediately after the device is configured. All the flip-flops and latches receive this reset, and reach a known value, depending on the register definitions.

In netlist simulations, the GSR signal is automatically asserted for the first 100 ns to simulate the reset that occurs after configuration. Optionally, you can supply a GSR pulse in pre-synthesis functional simulations, but this is not necessary if the design has a local reset that resets all registers.

For more information on GSR, see this [link](#) in the *Vivado Design Suite User Guide: Logic Simulation* (UG900) [Ref 11].

Disabling X Propagation

When a timing violation occurs during a timing simulation, the default behavior of a latch, register, RAM, or other synchronous elements is to output an X on the element during simulation. This X occurs because the actual output value is not known. On the hardware, the output of the register could show any of the following behavior:

- Retain its previous value
- Update to the new value
- Go metastable, i.e., a definite value is not settled upon until a while after the clocking of the synchronous element

Because this value cannot be determined, and accurate simulation results cannot be guaranteed, the element outputs an X to represent an unknown value. The X output remains until the next clock cycle in which the next clocked value updates the output if another violation does not occur.

The presence of an X output can significantly affect simulation. For example, an X generated by one register can be propagated to others on subsequent clock cycles. This can cause large portions of the design under test to become unknown.

To correct large scale X propagation in the design:

- On a synchronous path, analyze the path and fix any timing problems associated with this or other paths to ensure a properly operating circuit.
- On an asynchronous path, if you cannot otherwise avoid timing violations, disable the X propagation on synchronous elements during timing violations through use of ASYNC_REG attribute. For more information on using the ASYNC_REG constraint see this [link](#) in the *Vivado Design Suite User Guide: Logic Simulation* (UG900) [Ref 11].

When X propagation is disabled, the simulator retains the previous value at the output of the register. In the actual silicon, the register might have a different behavior. Disabling X propagation might yield simulation results that do not match the silicon behavior.



CAUTION! Exercise care when using this option. Use it only if you understand the timing violation and the design has mechanisms to recover the whole circuit from this violation.

Compiling Simulation Libraries

Xilinx recommends that you run library compilation before you begin simulation, especially if your design instantiates VHDL primitives or Xilinx IP. The majority of Xilinx IP is in VHDL so you will see failures related to library binding if you do not compile simulation libraries. You can run the `compile_simlib` Tcl command to compile the Xilinx simulation libraries for the target simulator. You can also issue this command using the Vivado IDE by selecting **Tools > Compile Simulation Libraries**.

Compiling Libraries as a System Administrator

System administrators compiling libraries using the `compile_simlib` Tcl command should compile the libraries in a default location that is accessible to all users.

The following example shows how to compile the libraries for QuestaSim for all devices, libraries, and languages:

1. Launch the Vivado tools.
2. Set up QuestaSim. See the *Vivado Design Suite User Guide: Logic Simulation* (UG900) [Ref 11].
3. Go to the Tcl console, and type the following command:

```
compile_simlib -simulator questa
```

Compiling Libraries as a User

When you run `compile_simlib` as a user, you should compile the libraries for each of your projects. If your project targets a single device, compile the libraries for that specific device only.

The following command shows how to compile libraries for Cadence IES for a design targeting a Kintex® UltraScale device:

```
compile_simlib -simulator ies -directory. -family kintexu
```

Simulation Flow

The Vivado Design Suite supports both integrated simulation, which allows you to run the simulator from within the Vivado IDE, and batch simulation, which allows you to generate a script from the Vivado tools to run simulation on an external verification environment.

Integrated Simulation

The Vivado IDE provides full integration with all supported simulators. In this flow, the simulator is called from within the Vivado IDE, and you can compile and simulate the design easily with a push of a button. For information on the steps involved in setting up the simulation flow, see this [link](#) in the *Vivado Design Suite User Guide: Logic Simulation* (UG900) [Ref 11].

Note: For information on supported simulators, see the *Vivado Design Suite User Guide: Release Notes, Installation, and Licensing* (UG973) [Ref 3].

Xilinx recommends the following for the integrated simulation flow:

1. Make sure the Compiled library location option is set correctly for your simulator in the following files:
 - Mentor simulators: `modelsim.ini`
 - IES simulator: `cds.lib`
 - VCS simulator: `synopsys_sim.setup`
 - Aldec tools: `library.cfg`

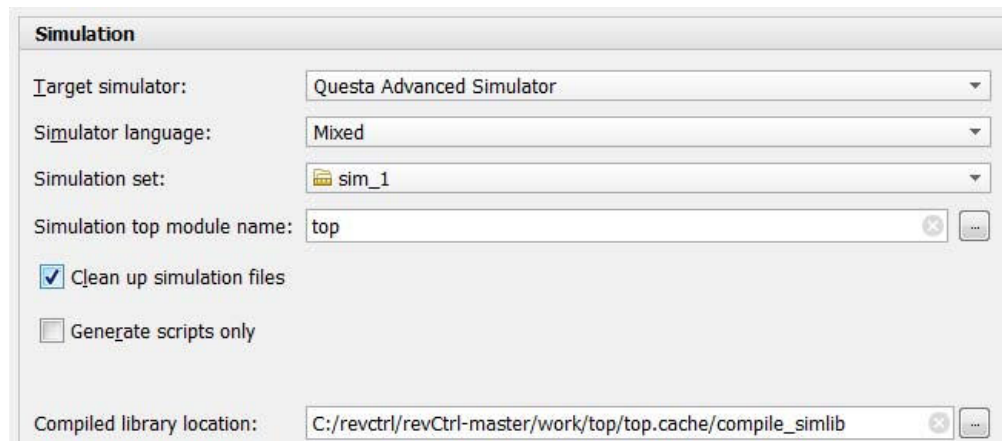


Figure 2-5: Setting Simulation Options

Note: `C:/revctrl/revCtrl-master/work/top/top.cache/compile_simlib` is a project level library in the preceding figure. You can use the same option to point to the libraries compiled by the system administrator.

2. Always select **Mixed** for the Simulator language option if you have a mixed language simulator license. Based on the simulator language selection, you might get an RTL model or a post-synthesis netlist for IP simulation. If you choose **VHDL** or **Verilog** instead of **Mixed**, there might be significant slowdown in simulation. To learn more about this option, see this [link](#) in the *Vivado Design Suite User Guide: Logic Simulation* (UG900) [Ref 11].

3. Do not select **Clean up simulation files** unless you detect data corruption. This option deletes the entire simulation directory and doing so will cause the incremental compile option to stop working because simulators do not have access to the original database to check if the source was previously compiled.
4. Launch integrated simulation from the Vivado IDE or with the `launch_simulation` Tcl command.



IMPORTANT: *The `launch_simulation` command launches integrated simulation for project-based designs. This command does not support Non-Project Mode. For more information, see [Understanding Project and Non-Project Software Use Models](#).*



RECOMMENDED: *If your verification environment has a self-checking test bench, run simulation in batch mode. There is a significant runtime cost when you view simulator waveforms using the integrated simulation.*

Batch Simulation

For batch simulation, the Vivado Design Suite provides the `export_simulation` Tcl command to generate simulation scripts for supported simulators, including the Vivado simulator. You can use the scripts generated by `export_simulation` directly or use the scripts as a reference for building custom simulation scripts. The `export_simulation` command creates separate scripts for each stage of the simulation process (compile, elaborate, and simulate) so that you can easily incorporate the generated scripts in your own verification flow. For more information about generating scripts for batch simulation, see this [link](#) in the *Vivado Design Suite User Guide: Logic Simulation* (UG900) [Ref 11].

Xilinx recommends the following for the batch simulation flow:

- Make sure the Compiled Library Location option is set correctly for your simulator as discussed in [Integrated Simulation](#).
- Use the `export_ip_user_files` Tcl command to export the IP and block design (BD) files needed for use by third-party simulators. For more information, see this [link](#) in the *Vivado Design Suite User Guide: Designing with IP* (UG896) [Ref 9].

Note: The Vivado tools automatically generate the IP and BD files when generating the output products for an IP.
- Create a standalone simulation directory by copying all of the design files needed for simulation using the `export_simulation -export_source_files` Tcl command.

Generating Simulation Scripts for IP

You can generate simulation scripts for individual IP, block designs (BD), or hierarchical modules of your design. This allows you to validate the behavior of the IP or module as a standalone design, which simplifies debugging of the design.

Xilinx IP can contain IEEE P1735 encrypted RTL for synthesis, simulation, and implementation. However, the RTL file encryption makes it impossible for you to visually inspect the library association of a VHDL simulation source file. You can get the library association by querying the properties of a file in the Vivado Design Suite, but Xilinx recommends using the `export_simulation` command to create simulation scripts for encrypted IP.

To export the simulation scripts for a specific design object, use the following command:

```
export_simulation -of_objects [get_files <ip_name.xci>]
```

Generating Simulation Scripts for Top-Level Designs

You can generate simulation scripts for the top-level design using the `export_simulation` Tcl command. Generating simulation scripts for the whole design provides a significant advantage over generating simulation scripts for each IP, BD, or design module. When using `export_simulation` to generate scripts for the top-level design, the Vivado Design Suite uses a filtering algorithm to eliminate duplicate files in the simulation file list. The smaller file list improves compile time significantly.

To create simulation scripts for the top-level of the design, run the following command from the open project or in-memory design:

```
export_simulation
```

Generating Irun Scripts for Cadence Flow

You can also generate a single step script for a Cadence Irun-based flow to simplify the compile, elaborate, and simulate process. To generate an Irun executable script, use the following command:

```
export_simulation -single_step
```

Simulation Flow Summary

1. Run behavioral simulation before proceeding with synthesis and implementation. Issues identified early will save time and money.
2. Always target supported third-party simulator versions. For more information, see the *Vivado Design Suite User Guide: Release Notes, Installation, and Licensing* (UG973) [Ref 3].
3. Use the Xilinx Tcl command `export_simulation` to generate scripts for the selected simulator.
4. Generate simulation scripts for individual IP, BDs, and hierarchical modules as well as for the top-level design.

5. Always set the Target Language to **Mixed** unless you do not have a mixed mode license for your simulator.
6. If you are targeting a 7 series device, use UNIFAST libraries to improve simulation performance.
Note: The UNIFAST libraries are not supported for UltraScale device primitives.
7. Infer logic wherever possible. Instantiating primitives adds significant simulation runtime cost.
8. Make sure incremental compile is turned on when using third-party simulators.
9. Turn off the waveform viewer wherever possible.
10. In the Vivado simulator, turn off debug during `xelab` for a performance boost.
11. In the Vivado simulator, turn on multi-threading to speed up compile time.

Synthesis, Implementation, and Design Analysis

These topics are covered in [Chapter 5, Implementation](#). For additional information, see the following resources:

- *Vivado Design Suite User Guide: Synthesis* (UG901) [\[Ref 16\]](#) or the Synthesis Design Hub
- *Vivado Design Suite User Guide: Implementation* (UG904) [\[Ref 19\]](#) or the Implementation Design Hub
- *Vivado Design Suite User Guide: Design Analysis and Closure Techniques* (UG906) [\[Ref 21\]](#) or the Timing Closure & Design Analysis Design Hub

Source Management and Revision Control Recommendations

Revision Control Methodologies

Overview

The methodologies for source management and revision control can vary depending on user and company preference, as well as the software used to manage Revision Control. This section will describe some of the fundamental methodology choices that design teams need to make in order to manage their active design projects. Specific recommendations on using the Vivado Design Suite with revision control systems are provided later in this section. Throughout this section, the term “manage” refers to the process of checking source versions in and out using a revision control system.

Many design teams use source management systems to manage various design configurations and revisions. There are many readily available systems available, such as Git, RCS, CVS, Subversion, ClearCase, Perforce, Bitkeeper, and many others. No single system is predominant. The Vivado Design Suite can interact with all such systems and is intended to work equally well with each while stopping short of direct integration with any individual tool. For an example of how to use revision control with the Vivado tools and data, see the *Vivado Design Suite Tutorial: Revision Control Systems* (UG1198) [Ref 30].

Most users follow a methodology in which sources are checked out into a local repository as a staging area for local changes. This can be called a *sandbox* or a local working version of the master repository. The purpose of this area is to group and test local changes by one user before pushing them back into the master repository where other users see the latest changes. The sources in this local repository can be modified to complete the design. These modified sources can be checked back into the source control system as new revisions at any time. Design results can also optionally be checked in for revision storage. Many users use a directory structure to store and manage sources and results.

Revision Control Systems use different mechanisms to update only those sources that have been modified since the last check out.

There is a defined subset of Vivado Design Suite files that should be managed using revision control. There are a number of intermediate files generated by the Vivado Design Suite that it would be inefficient and unnecessary to manage. The list shown in [Vivado Design Suite Source Types](#) highlights the key files.

Single User versus Multi-User Access

Design sources and IP are typically stored in a library or repository that can be referenced by any number of designers working on a project. However, using design sources or IP from a common repository can make it subject to modification by anyone using those files in their design. Any modification to a shared source or IP is going to be incorporated by all other designers referencing that same source. This can have unintended consequences if design specific modifications impact multiple designs.

Multi-user access works best if the design sources and IP are managed by a single designer and shared exactly as is to all other designs. These sources could/should be made read-only, which would prevent any one designer from inadvertently updating a common source.

In addition, Vivado Design Suite IP and IP Subsystems target a specific device. Thought should be given when considering sharing them between designs to ensure device compatibility.

Using Read-Only Managed Sources Directly

The Vivado Design Suite allows the use of read-only source files and IP to assemble designs. If sources are not expected to change during the process of the design, they can easily be referenced from their managed read-only locations. If sources need to be modified, they can be checked out into the local working area.

The area where Vivado Design Suite is invoked should have write privileges in order to write log and report files, certain caches, and design results.

Using a Local Working Area

A common methodology for design source management is to make a local copy of the managed sources into an active project working area. Designers can make modifications at will in that area to complete their designs. Sources can be checked back into revision control at any milestone event. Note that all your changes are made in the local copy of the file lying deep inside the project directory structure. So, be sure of the file that you are checking back into revision control system.

This method could be applied only for design sources that are expected to change. Static sources could be referenced from their originally managed read-only locations.

Managing Minimum Sets of Sources

Some design teams want to manage the minimum set of design files needed to reproduce the design. For details, see [Minimum Sets of Source Files to Manage](#).

The Vivado Design Suite enables this capability with the following caveats:

- The output products for IP and IP Integrator Subsystems need to be regenerated, which requires using the same version of software that the IP was originally created with, to get the exact same replica. If all of the IP output products are managed, the IP can be referenced by future software versions.
- It takes significant added time to generate the output products during design compilation.



RECOMMENDED: For best results, Xilinx recommends managing all of the sources listed in the following section, [Recommended Source Files to Manage](#).

Recommended Source Files to Manage

You can recreate a Vivado Design Suite design and associated IP and BD files using a set of source files. To mitigate the risk of recreation using future versions of the Vivado Design Suite and to reduce compile times, Xilinx recommends managing the following source files. Checking in all these files enables the design to be recreated using the current sources and tool configuration settings.

- Scripts
 - Run scripts used to compile the design including any pre and post scripts used when launching runs
 - If modifications were made to the Vivado `init.tcl` file, it must also be checked in
 - Project recreation scripts created with the `write_project_tcl` command
 - Project `.xpr` files can be used to recreate a design project
- RTL and simulation test benches
 - All RTL-based source files including `.include` files
 - Simulation test benches and stimulus files
- Constraints
 - All XDC constraint files and Tcl commands used as constraints
- IP
 - All IP directories and output product files with names intact
 - If using IP core containers, the IP `.xcix` file with all third-party simulation and synthesis files in the `ip_user_files` and `ip_static_files` subdirectories
 - Any `.coe`, `.csv`, `.bmm`, and `.elf` files that are used
- IP Integrator
 - The entire block design location with all of the IP subdirectories, files, and names intact
- HLS
 - All C sources files
 - Packaged HLS IP for use in Vivado synthesis
 - Scripts
- System Generator for DSP
 - The entire System Generator created module directory with all of the subdirectories, files, and names intact

- Vivado Hardware Manager
 - All .bit files needed to program the device
 - The .ltx file containing tool defaults and custom user settings
- Hardware Software Interface
 - The Handoff Design File (.hdf extension) contains all the information required by Xilinx SDK to create the corresponding software project for your design. The HDF file is created when you export your design either through the `write_hwdef` or `write_sysdef` Tcl command or the **File > Export > Export Hardware** command.
- Documentation
 - Any design related docs, specs, reports, etc.

Minimum Sets of Source Files to Manage

A Vivado Design Suite design along with its associated IP and BDs can be recreated using a minimum set of source files. There are limitations to this approach, as described in [Managing Minimum Sets of Sources](#).

Following are the minimum set of files needed to recreate the design:

- Scripts
 - Run scripts used to compile the design including any pre and post scripts used when launching runs
 - Project recreation scripts created with the `write_project_tcl` command
 - Project .xpr files can be used to recreate a design project
- RTL and Test benches
 - All RTL based source files including .include files
 - Simulation test benches and stimulus files
- Constraints
 - All XDC constraint files and Tcl commands used as constraints
- IP
 - The .xci files associated for each IP
 - If using IP core containers, the IP core container .xcix file
- IP Integrator
 - The BD recreation Tcl command created using the `write_bd_tcl` command
 - The .bd file used for each BD
 - The UI directory for graphics locations for blocks and comments

Vivado Design Suite Source Types

The Vivado Design Suite environment references source files that contain design descriptions. Options in the Vivado Design Suite control how you create, use, and manage the source types. These sources include:

- HDL and netlist files: Verilog (.v), SystemVerilog (.sv), VHDL (.vhd), and EDIF (.edf)
- C based source files (.c)
- Tcl files, run scripts, and init.tcl (.tcl)
- Logical and Physical Constraints (.xdc)
- IP core files (.xci)
- IP core container (.xcix)
- IP integrator block design files (.bd)
- Design Checkpoint files (.dcp)
- System Generator subsystems (.sgp)
- Side files for use by related tools (e.g. "do" files for simulator)
- Block Memory Map files (.bmm, .mmi)
- Executable and Linkable Format files (.elf)
- Coefficient files (.coe)
- Archived projects (.zip)

How Vivado Design Suite Recognizes Source Modifications

The Vivado Design Suite uses the time/date stamps on the majority of the source files to indicate whether the source file has been updated. Taking any action on the various sources that may trigger a new time/date stamp may indicate that the design is out of date and needs to be updated.



IMPORTANT: *Ensure that common revision control actions do not alter the file time/date stamps.*

Some Revision Control Systems might alter the timestamps even if there were no changes to the contents of a file, or they might alter the timestamps based on the time of check-in/check-out. To retain the relative time modification ordering of source and generated products, sort your files in reverse order of the timestamps and then perform check-in/check-out. The following is an indicative command (in Unix) to add files to git with timestamp ordering maintained:

```
find . -type f -print0 | xargs -0 ls-rt | xargs git add
```

Defining a Design Source Directory Structure

To effectively distinguish and manage the various types of design sources, Xilinx recommends creating a directory structure for each active design. The various directories store specific types of sources, as shown in the following figure.

Name	Type	Date modified
IP	File folder	1/22/2015 2:42 PM
HLS	File folder	1/16/2015 3:21 PM
Testbenches	File folder	1/16/2015 3:15 PM
Work_Dir	File folder	1/16/2015 12:56 PM
Doc	File folder	1/16/2015 12:44 PM
Constraints	File folder	1/16/2015 12:30 PM
RTL	File folder	1/16/2015 12:28 PM
IPI-BDs	File folder	1/16/2015 11:54 AM
DSP	File folder	1/16/2015 11:54 AM
Scripts	File folder	1/16/2015 11:53 AM

Figure 2-6: Design Source Directory Structure Example

The idea is to categorize and organize the various types of Vivado design sources. Define a naming convention and subdirectory structure that best works for your design and design team. The area itself could be managed by a revision control system or it could be a design specific checked out copy.

Vivado IP and IP Subsystems are device specific, so thought should be given when considering storage directory names. This can help identify them for later use in other designs.

Using an Active Working Area

Xilinx recommends creating an active working directory in order to provide a writable area to create projects, compile the design, write results, and experiment to close the design. This area could also contain the writable versions of the design sources and IP, if desired.

It is good design practice to put some thought into how to name and organize design specific sources, scripts, and results for easy identification. Devise subdirectory and source file naming schemes in order to help understand their contents later. Remove stale or unsuccessful design attempts and stale sources.



CAUTION! *The Windows operating system has a 260-character limit for path lengths, which can affect the Vivado tools. To avoid this issue, use the shortest possible names and directory locations when creating projects, defining IP or managed IP projects, and creating block designs.*

Managing Vivado Design Suite Projects

Using a Vivado Design Suite project can complicate the interaction with a source control system. Certain steps should be followed when using Vivado Design Suite projects with revision control. The project can maintain its own copy of the source files or reference remote sources and provides its own source management. However, there are methods to use Projects with revision control. The following methods are recommended.

Using Remote or Local Sources

When projects are created and managed through the Vivado tools, sources can either be referenced in their original locations or copied local to the project. Both local and remote sources can be interactively manipulated using the Vivado IDE. The text editor can be used to edit the sources, and the results can be analyzed and modified in open designs.

Design sources can be read-only protected, and stored anywhere that is network accessible from the design directory. Read-only sources limit the advantages of the interactive design features of the Vivado IDE by not allowing modifications.

For easiest interfacing with revision control systems when you are using Project Mode, you should create projects using remote design sources and not copy sources into the project directory. Managing sources remotely lets you easily maintain the design sources in the project while using the interactive Vivado IDE capabilities. You can check in new versions to the revision control system as modifications are made to the source files.

Alternately, you can elect to copy sources locally into the project directory which makes the entire design project more portable and self-contained, but it also buries the design source files down inside the project subdirectories.



TIP: *You can save most source modifications with a new name.*

Recreating Projects with only the .xpr file

You can recreate and manage Vivado Design Suite projects with a single file (<project_name>.xpr). The project file, along with the various project source files, are the only files that you need to manage under revision control. The entire project can be recreated by opening this project file, along with its associated source files, provided that the source files are at their original locations and were added as remote sources to the project. If sources are internal to the project, the standalone XPR file will not be able to reproduce the project correctly.

Recreating Projects with the write_project_tcl command

When using projects with source control systems, Xilinx recommends rebuilding the project from scratch using a Tcl-scripted approach. The Vivado `write_project_tcl` command

provides the ability to create a Tcl script that you can use to recreate the current design using the same source files and settings.

```
write_project_tcl <script_file_name>.tcl
```

The resulting script file should be managed with revision control and can be used to recreate the project as follows:

```
source <script_file_name>.tcl
```

Using Vivado Design Suite Script-Based Flows

The easiest way to interact with source control systems is to use the Non-Project scripted flow. The designer can check out the desired sources into a local directory structure or they can be read directly from their managed locations. New source files may also need to be created. Once the files are ready, the `read_<source_type>` Tcl commands pass the files to the Vivado synthesis and implementation commands. The source files remain in their original locations. The checked-out sources can be modified interactively, or with Tcl commands during the design session, using appropriate code editors. A common example of such a modification is a timing constraint change.

Note: Although source files can be read-only protected, this disables them from being modified.

Source files are then checked back into the source control system at the designer's discretion. Design results such as design checkpoints, reports, and bit files can also be checked in for revision management.

Using IP and IP Subsystems

IP and IP Integrator Subsystems are best configured interactively within the Vivado IDE. The IP customization wizards and the interactive IP Integrator canvas make the job very easy. Once the IP or IP Subsystem is configured and the output products are generated, the sources can be read by Vivado Tcl commands. For IP, the `.xci` should be used for the source. For IP Integrator, the `.bd` should be used as the source.

Note: You can also use the design checkpoint files (`.dcp`) for IP or IP Integrator as sources in the script based flow. Xilinx recommends using the `.xci` and `.bd` files, because they are guaranteed to be the latest version.

Managing Scripts and Reports

Designers should also manage any Tcl scripts and modified `.ini` startup files required to recreate the design. Managing the various output report files can also help to identify the design state for future reference.

Managing IP

Xilinx recommends that you create project-specific storage locations for the IP used in the design project. Because IP can be re-configured from any design project that uses the IP, creating design-specific IP prevents unwanted updates by other designers.

Vivado Design Suite IP and IP Subsystems are device specific, so thought should be given when considering sharing them between designs and when setting up management areas.

Vivado Design Suite IP can be configured, stored, and managed using the following methods:

- Create a Manage IP Project (Recommended)
- Project Local IP
- Project Remote IP

Creating Manage IP Projects (Recommended)

With centralized IP management, the customized IP and their output products are stored in a centralized repository outside of the current design or project directory structure. Each IP that is customized will be stored in its own self contained directory. The IP can be referenced from this repository in either Project Mode or Non-Project Mode, either by a Tcl script or by adding it to a project as a remote source.

The Manage IP project creates IP storage locations to configure, validate, and store multiple IP cores. They take advantage of the Vivado IDE to enable IP configuration using the IP Catalog, source management, analysis, and the runs infrastructure to validate and store the results for each IP core. This capability allows IP designers to select IP cores from the IP Catalog, configure and customize them, and then generate output products in the desired design configuration. The environment also allows you to validate the IP by performing synthesis and implementation.



RECOMMENDED: *Xilinx recommends that you manage the entire IP Location directory structure intact, including the `managed_ip_project` directory. This ensures that the Vivado Design Suite maintains the status of the IP runs and output products.*

Project Local IP

The IP can be configured and stored within the Vivado Design Suite project. The IP output products can reside within the project directory structure. Storing the IP output products local to the design project creates a standalone entity for the entire design, which can easily be packaged and shared. This is also useful if the design uses a unique IP configuration that you may not want made available in a shared IP repository.

Using the IP Catalog within a project to customize and add IP is straightforward. The project is self-contained and easily managed in one location. When an IP is not used in multiple projects or by multiple people, this is a simple approach to take. The IP is simply another part of a project that is managed along with all other design data, such as RTL sources and run results.

The entire project directory structure should be managed intact.

Project Remote IP

The Vivado Design Suite also enables individual IP to be configured standalone and remotely for use in both Project Mode and Non-Project Mode. While configuring any Vivado Design Suite IP, you can specify the IP Location, which can be set to a location outside of the project directory as shown in the following figure.

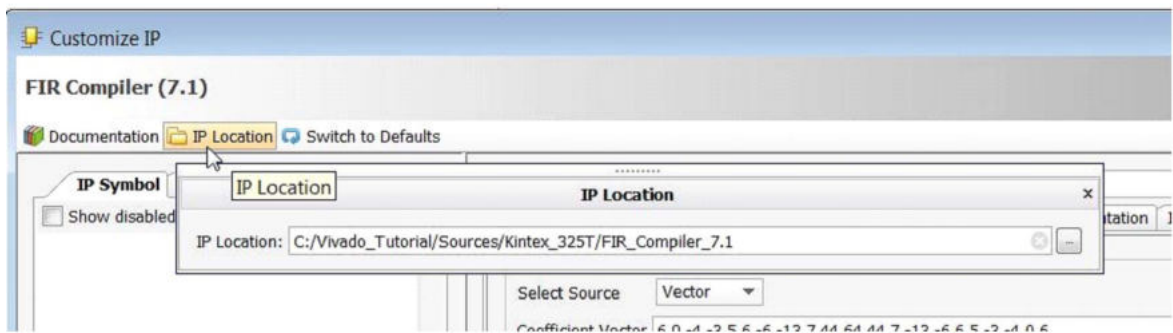


Figure 2-7: Defining a Remote Location for an Individual IP

Setting the IP Location to a directory outside the Vivado Design Suite project writes the IP configuration file (.xci), and the various IP output products including the RTL sources and constraints into in a standalone directory with a name matching the IP customization name.

Remote IP can be used in any number of design projects. However, since the features of the Vivado IDE enable modification and version updates to remote IP, you must take care when modifying or updating shared remote IP, as it may affect other users and designs. Setting up design specific remote IP locations prevent unwanted modification from other designers.

When using Remote IP Locations, follow these guidelines:

- Design an IP storage directory structure that differentiates device types, IP types, and other elements.
- Store example designs outside the IP directory to ensure that they are preserved when IP is updated.

Manage the entire remote IP directory structure intact. You can specify a location, but each IP is always placed in its own directory, named after the IP, just as with in project and manage IP projects.

Managing IP Sources

With Vivado Design Suite IP, each IP core is stored in a separate subdirectory containing the main .xci IP configuration file, along with RTL, XDC, and other related output product files required to synthesize and implement the IP. Xilinx recommends managing all of these files with their directory structure intact in order to mitigate any risk of reproducing the results in future versions of software.

Using IP Core Containers (Recommended)

To facilitate interactions with revision control systems, you can store IP configuration files and output products in a single, binary IP core container file rather than a directory structure. The Vivado Design Suite interacts directly with the IP core container files. When using IP core containers, you only need to manage the IP .xcix file. The .xcix file contains all of the files required for simulation, synthesis, and implementation.

For more information on using IP core containers, see *Vivado Design Suite User Guide: Designing with IP* (UG896) [Ref 9]. For more information on simulating with IP core container files, see the *Vivado Design Suite User Guide: Logic Simulation* (UG900) [Ref 11].

Note: All of the output product source files, including the .xci file or .xcix file, can be read only.



IMPORTANT: Create IP example designs in an unmanaged location to enable compilation and experimentation.

Some IP have additional associated files, such as COE, CSV, ELF, and BMM files. These files need to be managed with revision control systems. In a project flow, these files are registered with the project and might appear in the Design Sources area. Although the files are not imported into the project, they remain at the location they were referenced.

For example, when you specify a .coe file for the FIR Compiler, the file is added to a Coefficient Files folder under the Design Sources. The .coe file location is stored as a relative path in the IP XCI file. When placing the IP and the COE file in revision control, the relative paths need to be maintained. If the relative path needs to be changed, the path is a property of the IP. You can update this path property using the following Tcl command:

```
set_property CONFIG.Coefficient_File {/location/of/coe/file} [get_ips my_FIR_compiler]
```



TIP: You can also update this path property by re-customizing the IP in the GUI and specifying the COE file location.

Note: Archive Project correctly handles these file by copying the files into the project directory structure and updating the path references for the IP. When using the `write_project_tcl` command, the files are either referenced from their current location or imported depending on the options used.

Managing IP Core Containers and External Third-Party Files (Recommended)

When using IP core containers, the following IP output files are duplicated outside of the core container for easier access by third-party tools:

- Instantiation template
- Synthesis stub files
- Simulation source files
- Scripts
- Test benches

You can find these duplicated files in the `ip_user_files` and `ip_static_files` subdirectories for the IP.



RECOMMENDED: *When using third-party tools, Xilinx recommends managing the IP .xcix file as well as these intact external subdirectories.*

Managing All IP Output Products (Recommended)

It is recommend you manage the entire IP directory containing all IP output products with revision control. This gives the flexibility to decide when and if to upgrade the IP at a future point. It also provides better run time as the IP does not have to be regenerated every time it is used.

The Vivado IDE only supports one version of an IP in the IP Catalog. If you upgrade the Vivado IDE and the IP is no longer current, it is still usable. The IP is locked and you are not able to re-customize or generate output products. However, if all the output products are present, the IP can be used.

By default, all IP are synthesized out-of-context from the rest of the design. The synthesized design checkpoint produced is an output product of the IP. To ensure the IP can be used in future versions of the Vivado IDE these files must be present if the default out-of-context flow was used when creating the IP. As with other output products they cannot be created if the IP is not the current version.

Managing the IP .xci File Only

To restore the IP customizations used in a design, you must at a minimum preserve the .xci file for the specific configuration of the IP. From the .xci, the IP can be regenerated using the same Vivado IDE release with which it was created. If you plan to stay with this software version, or plan to always upgrade the IP to the latest version, the IP .xci is sufficient. However, to use an IP core (including its current customization and the RTL, XDC, etc.) with future versions of the Vivado IDE, you should place the entire IP output product directory under revision control and maintain the directory structure.

IP directory structures should be maintained in order to ensure locations to create the output products. Do not store multiple .xci files in a single directory as output products will collide when generated.

Managing the IP .xcix Core Container File Only

When you want to manage a minimum set of files and still include the IP output products needed for simulation and synthesis, use IP core containers. With core containers, you only need to manage the IP .xcix file. You can use the following command to export all of the IP source files needed for simulation, synthesis, and implementation into any unmanaged area for use in any design project:

```
export_ip_user_files
```

Managing IP Integrator Subsystems

IP Integrator Subsystems contain multiple IP, along with their customized parameters, and interconnect. These IP subsystems are created with the IP Integrator and are referred to as Block Designs (BD) within the Vivado IDE.

Creating Remote Block Designs (Recommended)

When you create a Vivado IP Integrator block design (.bd extension), you can specify a remote location. This is the easiest way to manage and store revisions of the block designs. The Vivado tools store all of the files and IP used in the block designs in a directory remote from the project directory, which is similar to [Project Remote IP](#).

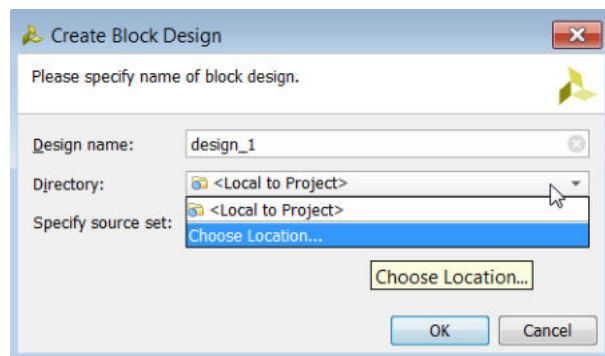


Figure 2-8: Defining a Remote Location for a Block Design

For more information, see this [link](#) on creating block designs in the *Vivado Design Suite User Guide: Designing IP Subsystems Using IP Integrator* (UG994) [Ref 26].

Creating Project-Based Block Designs

By default, Vivado IP Integrator will create and store the block design data inside the local project directory structure. This can make interaction with a revision control system tricky

as the Vivado IDE enables interactive modifications and management of the sources. Storing the IP output products within the design project creates a standalone entity for the entire design, which can easily be packaged and shared. This is also useful if the design uses unique IP configuration that you may not want made available in a shared IP repository.

Managing IP Integrator Sources

With Vivado IP Integrator, each IP core used in the BD is stored in a separate subdirectory under the BD, containing the main .xci IP configuration file, along with RTL, XDC, and other related output product files required to implement each IP. Xilinx recommends managing all of the BD files with the directory structure and file names intact in order to mitigate any risk of reproducing the results on future versions of software.

All of the output product source files including the .bd file can be read only, however it limits the block design capabilities of the Vivado IDE.

You should take care when using IP example designs to ensure they are created in an unmanaged location to enable compilation and experimentation. By default they are created inside the IP subdirectory.

Recreating the Block Design using only the .bd File

At a minimum, the .bd source file associated with the BD should be managed with revision control. Adding the .bd as a project source will trigger the regeneration of the block design and related IP output products. This approach can be time consuming, because all of the IP and block design output products will need to be regenerated and compiled. It also limits the recreation of the block design to the version of software it was created with. Xilinx recommends checking in the entire bd directory. Ensure the BD has been generated successfully with all of the associated IP output products.

Recreating the Block Design Using the write_bd_tcl Command

The Vivado `write_bd_tcl` command provides the ability to create a Tcl script that you can use to recreate the current IP integrator block design (.bd) using the same IP, connectivity, and settings. Ensure the BD has been generated successfully with all of the associated IP output products.

```
write_bd_tcl <script_file_name>.tcl
```

You can use this command with both Project and Non-Project flows to recreate the block design providing the sources are still located in the same places.

```
source <script_file_name>.tcl
```

You should also manage the resulting script file with revision control.

Managing Simulation Sources

Simulation sources are often the majority of the design source files along with test benches and simulation scripts. All of these files should be managed by revision control. They can be exported into a single directory from the Vivado Design Suite.

For more information on exporting files for logic simulation, see [Generating Simulation Scripts for IP](#) and [Generating Simulation Scripts for Top-Level Designs](#).

Managing System Generator Sources

The entire project directory created by the Xilinx System Generator should be managed intact by revision control.

Managing HLS Sources

All of the C based design sources, run scripts, and potentially the output packaged RTL IP should be managed by revision control.

Managing Design Results

Depending on your design team needs, additional files may be desired to be managed. These may include the design bitstream file, implementation results, design checkpoints, log files and reports, custom Tcl commands and run scripts. Manage these files appropriately.

Using Design Checkpoints

The Vivado Design Suite uses Design Checkpoint files (.dcp) to store the current state of the design being processed through the flow. These checkpoints include the netlist, constraints, and design results at the stage when the checkpoint was written.

Checkpoints should be written after each stage of the design process. They are automatically created when using a project to process synthesis and implementation runs.

Checkpoints can be opened in the Vivado IDE for design analysis and constraints modification. Constraint changes can be written to new constraints files for use during the next run through the flow. Checkpoints are images of a design at a given state in the flow. Checkpoints do not contain the entire project or the source files. They can typically be passed on to the remaining design steps, such as Generate Bitstream.

Your design team may elect to also store milestone design checkpoints.

Archiving Designs

The `archive_design` command can package up an entire project into a compressed zip file. The command has several options for storing sources and run results. Essentially the entire project is copied locally in memory, and then zipped into a file on disk while leaving the original project intact. This command also copies any remote sources into the archive. This feature is useful for sending your design description to another person or to store as a self contained entity. You might also need to send your version of `init.tcl` if you are using this file to set specific parameters or variables that affect the design.

For more information, see the following resources:

- *Vivado Design Suite User Guide: System-Level Design Entry* (UG895) [Ref 8]
- [Vivado Design Suite QuickTake Video: Creating Different Types of Projects](#)
- [Vivado Design Suite QuickTake Video: Managing Sources with Projects](#)

Managing Hardware Manager Projects and Sources

Project `.bit` file and `.ltx` files are the primary output files required for using the Vivado Design Suite Debug and Programming features in Vivado hardware manager. Xilinx recommends you manage these files under revision control if you want to use this project in Vivado Lab Edition.

When Using Vivado Design Suite Projects

The `project_name.hw` directory in your Vivado Design Suite project stores information about custom dashboards, trigger, capture conditions, waveform configuration files etc created as part of using the Debug and Programming in the Vivado hardware manager. Xilinx recommends you manage the `project_name.hw` directory in your Vivado Design Suite project under revision control. This also helps if you want to hand off this project to be used in Vivado Lab Edition.

Managing Vivado Lab Edition Sources

Xilinx recommends you manage the project directory that was created for the project in Vivado Lab Edition under revision control. The `hw_*` directories in the Lab Edition project directory stores information about custom dashboards, trigger, capture conditions, waveform configuration files, etc., in the Vivado hardware manager as part of using the Debug and Programming. The `.lpr` file in the Lab Edition project directory is the project file that you need to manage under revision control. The entire project can be recreated by opening this project file, provided that the `hw_*` directory are at their original locations.

Upgrading Designs and IP to the Latest Vivado Design Suite Release

It is highly likely that a new release of the Vivado Design Suite will become available during your design process. You can update to this new release, or stay with your current release. Although Xilinx recommends that you use the latest release, it is not mandatory. You should be aware though, that Xilinx does not support versions prior to the last two major releases. New releases may contain:

- Software bug fixes
- New IP versions
- Updated device files (including speed files for various devices)
- New device offerings from Xilinx
- New software features
- Performance improvements

The Vivado Design Suite project may be upgraded when migrating to a newer release. The project and device file upgrades typically happen automatically without user interaction.

You can elect to update the latest IP version, or lock the IP at the version with which it was configured. To use a locked version of an IP, generate the output products for the IP, and then use those during subsequent software updates. This does prevent you from reconfiguring the IP on the latest release unless you use the software version with which it was originally configured. Although Xilinx recommends that you use the latest IP versions, doing so is not mandatory.

IP Subsystems (created using IP integrator) also require you to manage the IP contained in them. You can either stay with the locked version of the IP output products, or update them to the latest version. All IP included in a particular IP Subsystem must be updated simultaneously.

For more information, see the following resources:

- [Working With Intellectual Property \(IP\)](#)
- *Vivado Design Suite User Guide: Designing with IP* (UG896) [Ref 9]
- *Vivado Design Suite User Guide: Designing IP Subsystems Using IP Integrator* (UG994) [Ref 26]
- [Vivado Design Suite QuickTake Video: Managing Vivado IP Version Upgrades](#)

For information on migrating your design to the new software release, see *Vivado Design Suite User Guide: Release Notes, Installation, and Licensing* (UG973) [Ref 3].

Updating Projects

When opening a project, you are prompted to upgrade if it was last saved with an older version of the Vivado Design Suite. It does allow you to open the project in a read-only manner. Project updates are related to new project formats only. It does not automatically update IP or IP Subsystems. It is recommend to update to the latest project format with each release.

Updating IP

If the IP is updated in a future Vivado Design Suite release, Xilinx recommends that you upgrade. When using the Vivado IDE you can usually upgrade to the latest IP. The Change Log describes the upgrade. If you do not wish to upgrade, you must generate and archive all the output products for an IP. For more information, see [Managing IP](#).

These saved targets may be used, however, they cannot be recustomized using newer version of the Vivado Design Suite. Neither can additional output products be created in a new Vivado IDE release.

For upgrading IP subsystems refer to [Updating IP Subsystems](#).

Updating IP Subsystems

If any of the IP used in the subsystem is updated in a future Vivado Design Suite release, Xilinx recommends that you upgrade the entire block design. If you do not wish to upgrade, you must generate and archive all the output products for an IP. Editing the block design is disabled. For more information, see [Managing IP](#).

These saved targets may be used, however, they cannot be recustomized using newer version of the Vivado Design Suite. Neither can additional output products be created in a new Vivado IDE release.

For more information, see the *Vivado Design Suite User Guide: Designing IP Subsystems Using IP Integrator* (UG994) [[Ref 26](#)].

Board and Device Planning

Overview of Board and Device Planning

Properly planning the FPGA orientation on the board and assigning signals to specific pins can lead to dramatic improvements in overall system performance, power consumption, and design cycle times. Visualizing how the FPGA device interacts physically and logically with the printed circuit board (PCB) enables you to streamline the data flow through the device.

Failing to properly plan the I/O configuration can lead to decreased system performance and longer design closure times. Xilinx highly recommends that you consider I/O planning in conjunction with board planning.

For more information, see the following resources:

- *Vivado Design Suite User Guide: I/O and Clock Planning* (UG899) [Ref 4]
- [Vivado Design Suite QuickTake Video: I/O Planning Overview](#)

PCB Layout Recommendations

The layout of the FPGA device on the board relative to other components with which it interacts can significantly impact the I/O planning.

Aligning with Physical Components on the PCB

The orientation of the FPGA device on the PCB should first be established. Consider the location of fixed PCB components, as well as internal FPGA resources. For example, aligning the GT interfaces on the FPGA package to be as close to the components with which they interface on the PCB will lead to shorter PCB trace lengths and fewer PCB vias.

A sketch of the PCB including the critical interfaces can often help determine the best orientation for the FPGA device on the PCB, as well as placement of the PCB components. Once done, the rest of the FPGA I/O interface can be planned.

High-speed interfaces such as memory can benefit from having very short and direct connections with the PCB components with which they interface. These PCB traces often have to be matched length and not use PCB vias, if possible. In these cases, the package pins closest to the edge of the device are preferred in order to keep the connections short and to avoid routing out of the large matrix of BGA pins.

Power Distribution System

FPGA designers are faced with a unique task when designing a Power Distribution System (PDS). Most other large, dense integrated circuits (such as large microprocessors) come with very specific bypass capacitor requirements. Since these devices are designed only to implement specific tasks in their hard silicon, their power supply demands are fixed, and fluctuate only within a certain range.

FPGA devices do not share this property. FPGA devices can implement an almost infinite number of applications at undetermined frequencies, and in multiple clock domains. For this reason, it is critical that you refer to the *PCB Design Guide* [Ref 35] for your device to fully understand the device PDS.

Key factors to consider during PDS design include:

- Selecting the right voltage regulators to meet the noise and current requirements based on Power Estimation. For more information, see [Power in Chapter 5](#).
- Setting up the XADC power supply (Vrefp and Vrefn pins).
- Running PDN simulation. The recommended amount of decoupling capacitors in the *PCB Design Guide* [Ref 35] for your device assumes worst-case situations, because FPGA devices can implement any arbitrary functionality. Running PDN simulations can help in reducing the amount of decoupling capacitors required to guarantee power supplies that are within the recommended operating range.

For more information on PDN simulation, see the Xilinx White Paper: *Simulating FPGA Power Integrity Using S-Parameter Models* (WP411) [Ref 49].

Specific Considerations for PCB Design

The PCB should be designed considering the fastest signal interfacing with the FPGA device. These high-speed signals are extremely sensitive to trace geometry, vias, loss, and crosstalk. These aspects become even more prominent for multi-layer PCBs. For high-speed interfaces perform a signal integrity simulation. A board re-design with improved PCB material or altered trace geometries may be necessary to obtain the desired performance.

Xilinx recommends going through the following list of items when designing your PCB:

- Review the PCB design checklist for Gigabit Transceivers (GTs).
 - For more information, see the Transceiver User Guide for your device.
 - Run Spice or IBIS-AMI simulations using channel parameters

- Review memory IP and PCIe® design guidelines

For more information, see the respective product guides.

- Follow the proper PCB decoupling capacitor.

For more information, see the *PCB Design Guide* [Ref 35] for your device.

- Run noise analysis.

The Vivado® Design Suite I/O planner can run SSN analysis for a given pinout.

- Run signal integrity analysis.

The Vivado tools can write IBIS files for the design.

- Check to see if there are any issues with overshoot or undershoot due to poor termination.

- Run the built-in Vivado DRC on I/O Pin Planning.

- Run power estimation for the design.

- Make sure you understand total power consumption.
- The Vivado Design Suite has power estimation tools (XPE) that will help analyze power for a given design.

- Determine whether the board has an adequate Power Distribution System (PDS).

- Review schematic recommendations.

For more information, see the *PCB Design Guide* [Ref 35] for your device.

Clock Resource Planning and Assignment

Xilinx recommends that you select clocking resources as one of the first steps of your design, well before pinout selection. Your clocking selections can dictate a particular pinout, and can also direct logic placement for that logic. Proper clocking selections can yield superior results. Consider:

- Constraint creation, particularly in large devices with high utilization in conjunction with clock planning.
- Manual placement of clocking resources if needed for design closure. [Clocking Guidelines in Chapter 4](#), explains more details on clocking resources, if you need to do manual placement.

Selecting Clocking Resources

Making proper decisions for clocking requires some understanding of the target architecture resources, which might be different for different generations of devices.

7 Series Devices

Xilinx® 7 series devices contain thirty-two global clock buffers known as BUFG. Half of these global clock buffers are above the horizontal center of the FPGA device, and the other half are below the horizontal center.

I/Os, PLLs and MMCMs in the top half of the chip can connect only to the sixteen BUFGs above the horizontal center. I/Os, PLLs and MMCMs in the lower half of the chip can connect only to the sixteen BUFGs below the horizontal center.

Thus, when selecting the I/O pins for the clocking resources, balance the clock inputs between the upper and lower half of the device. This is also important when choosing multiple clocks to be muxed by a BUFGCTRL. All associated BUFGCTRL clock inputs should be placed into the same half of the device to allow connectivity to the same BUFGCTRLs. When targeting SSI technology devices a similar rule exists. However, rather than applying to the upper and lower half of a device, it applies to the upper or lower half of the SLR. When choosing between PLLs and MMCMs, use PLLs wherever possible as it provides tighter control of jitter. MMCMs may be used when: (1) the PLLs have been exhausted; or (2) you need advanced features available in MMCM but not in PLL.

BUFG components can meet most clocking requirements for designs with less demanding requirements, such as:

- Number of clocks
- Design performance

BUFG components are easily inferred by synthesis, and have few restrictions, allowing for most general clocking.

However, if clocking demands exceed the capabilities or number of the BUFG component, or if you require better clocking characteristics, Xilinx recommends that you:

1. Analyze the clocking needs against the available clocking resources.
2. Select and control the best resource for the task.

For information on other clocking components, see [Clocking Guidelines in Chapter 4](#).

Single or Multi Region Clock Pin Selection

Based on the interface size, you can decide whether to use a Single Region Clock Capable (SRCC) pin or a Multi Region Clock Capable (MRCC) pin. If your interface spans multiple banks, you must use an MRCC pin, which increases the delay through clock network.

Single ended clocks should be connected to P-side of the differential pair of clocks.

UltraScale Devices

For UltraScale™ devices, clock inputs to the FPGA are generally grouped into the following categories:

- Global Clocks (GC) - These are most of the general purpose clock inputs which can drive PLLs (two per bank) or MMCM (1 per bank) as well as have dedicated access to the Global clocking network. In general it is suggested to locate this clock close to the I/O or logic placement in the array to minimize insertion delay and power; however there is some flexibility in placement due to the clocking structure.
- Byte Lane Clocks (DBC QBC) - These are dedicated clock pins for the I/O bit slices that are generally used with memory interfaces and other high-speed clocking. These are generally assigned by the same IP that is creating the high-speed interface to ensure proper placement for the desired connectivity.
- MGT Reference Clock (MGTREFCLK) - This is the dedicated differential reference clock input to the MGT (GTH or GTY) interfaces. This clock input must be supplied on dedicated pins through an IBUGDS_GTE3 component.
- Configuration clocks (EMCCLK, TCK, I2C_SCLK, CCLK) - These clocks are dedicated clocks associated with configuration, JTAG or SYSMON interfaces. For details, see the *UltraScale Architecture Configuration User Guide* (UG570) [Ref 37].

For more information on UltraScale device clocking, see [Clocking Guidelines in Chapter 4](#).

I/O Planning Design Flows

The Vivado IDE allows you to interactively explore, visualize, assign, and validate the I/O ports and clock logic in your design. The environment ensures correct-by-construction I/O assignment. It also provides visualization of the external package pins in correlation with the internal die pads.

You can visualize the data flow through the device and properly plan I/Os from both an external and internal perspective. Once the I/Os have been assigned and configured through the Vivado IDE, constraints are then automatically created for the implementation tools.

For more information on Vivado Design Suite I/O and clock planning capabilities, see the following resources:

- *Vivado Design Suite User Guide: I/O and Clock Planning* (UG899) [\[Ref 4\]](#)
- *Vivado Design Suite Tutorial: I/O and Clock Planning* (UG935) [\[Ref 31\]](#)
- [Vivado Design Suite QuickTake Video: I/O Planning Overview](#)

Determine When the Final I/O Configuration is Required

The PCB board fabrication schedule often dictates when the final FPGA I/O configuration is required. Whenever possible, perform I/O planning after the initial RTL design has been created and synthesized. The reason for this sequence is that the synthesized netlist is clock aware, and the logic has now been defined at a structural level. This sequencing enables many more clock related DRCs to ensure that the I/O banks and clock logic have been assigned properly.

The design can also be run through implementation to ensure that: (1) all I/O and clock rules are adhered to; and (2) the design successfully generates a bitstream. This is the recommended validation process for a final I/O configuration.

However, not all design cycles allow that much time. Often the I/O configuration has to be defined before you have synthesizable RTL. Although the Vivado tools enable pre-RTL based I/O planning, the level of DRC checks performed are fairly basic. For more information, see the PCB design guide for the selected device and the related I/O hardware documentation. Alternatively, a dummy top-level design with I/O Standards and pin assignments can help perform DRCs related to banking rules.

Pre-RTL I/O Planning

If your design cycle forces you to define the I/O configuration before you have a synthesized netlist, take great care to ensure adherence to all relevant rules. The Vivado tools include a Pin Planning Project environment that allows you to import I/O definitions

using a CSV or XDC format file. You can also create a dummy RTL with just the port directions defined. Availability of port direction makes simultaneous-switching-noise (SSN) analysis more accurate as input and output signals have different contributions to SSN.

I/O ports can also be created and configured interactively. Basic I/O bank DRC rules are provided.

See the *PCB Design Guide* [Ref 35] for your device to ensure proper I/O configuration. For more information, see this [link](#) in the *Vivado Design Suite User Guide: I/O and Clock Planning* (UG899) [Ref 4].

Netlist-Based I/O Planning

The recommended time in the design cycle to assign I/Os and clock logic constraints is after the design has been synthesized. The clock logic paths are established in the netlist for constraint assignment purposes. The I/O and clock logic DRCs are also much more comprehensive.

See the *PCB Design Guide* [Ref 35] for your device to ensure proper I/O configuration. For more information, see this [link](#) in the *Vivado Design Suite User Guide: I/O and Clock Planning* (UG899) [Ref 4].

Defining Alternate Devices

It is often difficult to predict the final device size for any given design during initial planning. Logic can be added or removed during the course of the design cycle, which can result in the need to change the device size.

The Vivado tools enable you to define alternate devices to ensure that the I/O pin configuration defined is compatible across all selected devices, as long as the package is the same.



IMPORTANT: *The device must be in the same package.*

To migrate your design with reduced risk, carefully plan the following at the beginning of the design process: device selection, pinout selection, and design criteria. Take the following into account when migrating to a larger or smaller device in the same package: pinout, clocking, and resource management. For more information, see this [link](#) in the *Vivado Design Suite User Guide: I/O and Clock Planning* (UG899) [Ref 4].

Pin Assignment

Good pinout selection leads to good design logic placement. Poor placement may also create longer routes, causing increased power consumption and reduced performance. These consequences of good pinout selection are particularly true for large FPGA devices. Because some large FPGA devices can span multiple dies, a spread out pinout can cause

related signals to span longer distances. For more information, see this [link](#) in the *Vivado Design Suite User Guide: I/O and Clock Planning* (UG899) [Ref 4].

Using Xilinx Tools in Pinout Selection

Xilinx tools assist in interactive design planning and pin selection. These tools are only as effective as the information you provide them. Tools such as the Vivado design analysis tool can assist pinout efforts. These tools can graphically display the I/O placement, show relationships among clocks and I/O components, and provide Design Rule Check (DRC) capability to analyze pin selection.

If a design version is available, a quick top-level floorplan can be created to analyze the data flow through the device. For more information, see the *Vivado Design Suite User Guide: Design Analysis and Closure Techniques* (UG906) [Ref 21].

Required Information

For the tools to work effectively, you must provide as much information about the I/O characteristics and topologies as possible. You must specify the electrical characteristics, including I/O standard, drive, and slew.

You must also take into account all other relevant information, including clock topology and timing constraints.

Clocking choices in particular can have a significant influence in pinout selection and vice-versa, as discussed in [Selecting Clocking Resources](#).

For more information on specifying the electrical characteristics for an I/O, see this [link](#) in the *Vivado Design Suite User Guide: I/O and Clock Planning* (UG899) [Ref 4].

Pinout Selection

Xilinx recommends careful pinout selection for some specific signals as discussed below.

Interface Data, Address, and Control Pins

Group the same interface data, address, and control pins into the same bank. If you cannot group these components into the same bank, group them into adjacent banks. For Stacked Silicon Interconnect (SSI) technology devices, group all pins of a particular interface into the same SLR.

Interface Control Signals

Place the following interface control signals in the middle of the data buses they control (clocking, enables, resets, and strobes).

Very High Fanout, Design-Wide Control Signals

Place very high fanout, design-wide control signals towards the center of the device.

For SSI technology devices, place the signals in the SLR located in the middle of the SLR components they drive.

Configuration Pins

To design an efficient system, you must choose the FPGA configuration mode that best matches the system requirements. Factors to consider include:

- Using dedicated vs. dual purpose configuration pins.

Each configuration mode dedicates certain FPGA pins and can temporarily use other multi-function pins during configuration only. These multi-function pins are then released for general use when configuration is completed.

- Using configuration mode to place voltage restrictions on some FPGA I/O banks.
- Choosing suitable terminations for different configuration pins.
- Using the recommended values of pullup or pulldown resistors for configuration pins.



RECOMMENDED: *Even though configuration clocks are slow speed, perform signal integrity analysis on the board to ensure clean signals.*

There are several different configuration options. Although the options are flexible, there is often an optimal solution for each system. Consider the following when choosing the best configuration option:

- Setup
- Speed
- Cost
- Complexity

See [Configuration](#). For more information on FPGA configuration options, see *Vivado Design Suite User Guide: Programming and Debugging* (UG908) [Ref 24].

Memory Interfaces

Additional I/O pin planning steps are required when using Xilinx Memory IP. After the IP is customized, you then assign the top-level IP ports to physical package pins in either the elaborated or synthesized design in the Vivado IDE.

All of the ports associated with each Memory IP are group together into an I/O Port Interface for easier identification and assignment. A Memory Bank/Byte Planner is provided to assist you with assigning Memory I/O pin groups to byte lanes on the physical device pins.

For more information, see this [link](#) in the *Vivado Design Suite User Guide: I/O and Clock Planning* (UG899) [Ref 4].

The *Xilinx Zynq-7000 SoC and 7 Series Devices Memory Interface User Guide* (UG586) [Ref 47] and the *LogiCORE IP UltraScale Architecture-Based FPGAs Memory Interface Solutions Product Guide* (PG150) [Ref 48] contain design and pinout guidelines. Be sure that you follow the trace length match recommendations in that Guide, verify that the correct termination is used, and validate the pinout in by running the DRCs after memory IP I/O assignment.

Gigabit Transceivers (GTs)

Gigabit Transceivers (GTs) have specific pinout requirements. You may be able to share reference clocks across multiple GTs. For 7 series devices, the reference clock for a Quad can also be sourced from an adjacent Quad, while for UltraScale devices, the reference clock for a Quad can be sourced from up to two Quads below or from up to two Quads above. For devices that support stacked silicon interconnect (SSI) technology, the reference clock sharing is limited within its own super logic region (SLR). Xilinx recommends that you use the GT wizard to generate the core. For pinout recommendations, see the product guide.

For clock resource balancing, the Vivado placer attempts to constrain loads clocked by GT output clocks (TXOUTCLK or RXOUTCLK) next to the GTs sourcing the clocks. For stacked silicon interconnect (SSI) technology devices, if the GTs are located in the clock regions adjacent to another SLR, the routing resources required for signals entering or exiting SLLs have to compete with the routing resources required by the GT output clock loads. Therefore, GTs located in clock regions next to SLR crossings might reduce the available routing connections to and from the SLL crossings available in those clock regions.

High Speed I/O

HP (high-performance) and HR (high-range) banks have difference in the speed with which they can transmit and receive signals. Depending upon the I/O speed you need, choose between HP or HR banks.

Internal VREF and DCI Cascade Constraints

Based on the settings for DCI Cascade and Internal VREF, you can free up pins to be used for regular I/Os. These settings also ensure that related DRC checks are run to validate the legality of the constraints. For more information, see the *SelectIO Resources User Guide* (UG471) [Ref 38] for your device.

CCIO and CMT Usage

Balance CCIO and CMT usage between the upper and lower halves of the device in order to balance the access to upper and lower BUFG components. For SSI technology devices, balance upper and lower CCIO components or CMT components in an SLR against the other SLR components.

Interface Bandwidth Validation

Create small connectivity designs to validate each interface on the FPGA. These small designs exercise only the specific hardware interface. The internal of the design need not be created yet. A separate design should be created per hardware interface, and should be used to exercise the hardware at full bandwidth at the desired speed. FPGA internal loop-back or simple checkers can be used to verify that the data transmittal is successful at the desired speed. As the FPGA interface on the board is being designed, these designs can be used to validate that the interfaces and the board will be able to work at the desired speed.

These small test designs can be rapidly implemented through the Vivado tools. This flow will also allow for a robust validation of the selected I/O in terms of placement legality and interface timing requirements. Thus, any potential DRC or potential timing issue can be validated as pin locations are being finalized.

For some of the interfaces IP cores, the Vivado tools can provide the test designs; for example, IBERT for SerDes or example design for PCIe.

These same designs can also be used subsequently to systematically validate each hardware component, before working on the bitstream for the whole design.

Designing with SSI Devices

SSI Pinout Considerations

When planning pinouts for components that are located in a particular SLR, place the pins into the same SLR. For example, when using the device DNA information as a part of an external interface, place the pins for that interface in the master SLR in which the DNA_PORT exists. Additional considerations include the following:

- Group all pins of a particular interface into the same SLR.
- For signals driving components in multiple SLRs, place those signals in the middle SLR.
- Balance CCIO or CMT components across SLRs.
- Reduce SLR crossings.

Super Logic Region (SLR)

A Super Logic Region (SLR) is a single FPGA die slice contained in an SSI technology device.

Active Circuitry

Each SLR contains the active circuitry common to most Xilinx FPGA devices. This circuitry includes large numbers of:

- 6-input LUTs
- Registers
- I/O components
- Gigabit Transceivers (GT)
- Block memory
- DSP blocks
- Other blocks

SLR Components

Multiple SLR components are assembled to make up an SSI technology device.

The orientation of the SLR components is stacked vertically onto the interposer.




Figure 3-1: Single SSI SLR

Multiple SLR components are stacked vertically to create the SSI technology devices.

- The bottom SLR is SLR0.
- Subsequent SLR components are incremented as they ascend vertically.

For example, there are four SLR components in the XC7V2000T device. The bottom SLR is SLR0. The SLR directly above SLR0 is SLR1. The SLR directly above SLR1 is SLR2. The top SLR is SLR 3.

The Xilinx tools clearly identify SLR components in the graphical user interface (GUI) and in reports.

SLR Nomenclature

Understanding SLR nomenclature for your target device is important in:

- Pin selection
- Floorplanning
- Analyzing timing and other reports
- Identifying where logic exists and where that logic is sourced or destined

You can use the Vivado Tcl command `get_slrs` to get specific information about SLRs for a particular device. For example, use the following commands:

- `llength [get_slrs]` to obtain the number of SLRs in the device
- `get_slrs -of_objects [get_cells my_cell]` to get the SLR in which `my_cell` is placed




Figure 3-2: Vivado Tools Representation of a XC7V2000T Device

Master Super Logic Region

Every SSI technology device has a single master SLR. See the following table.

Table 3-1: Master SLR Index Reference

Device	Device Viewer
XC7V2000T	SLR1
XC7VX1140T	SLR1
XC7VH580T	SLR0
XC7VH870T	SLR2
XCKU085	SLR0
XCKU115	SLR0
XCVU125	SLR0
XCVU160	SLR1
XCVU190	SLR1
XCVU440	SLR1

The master SLR contains the primary configuration logic that initiates configuration of the device and all other SLR components.

The master SLR is the only SLR that contains dedicated circuitry such as:

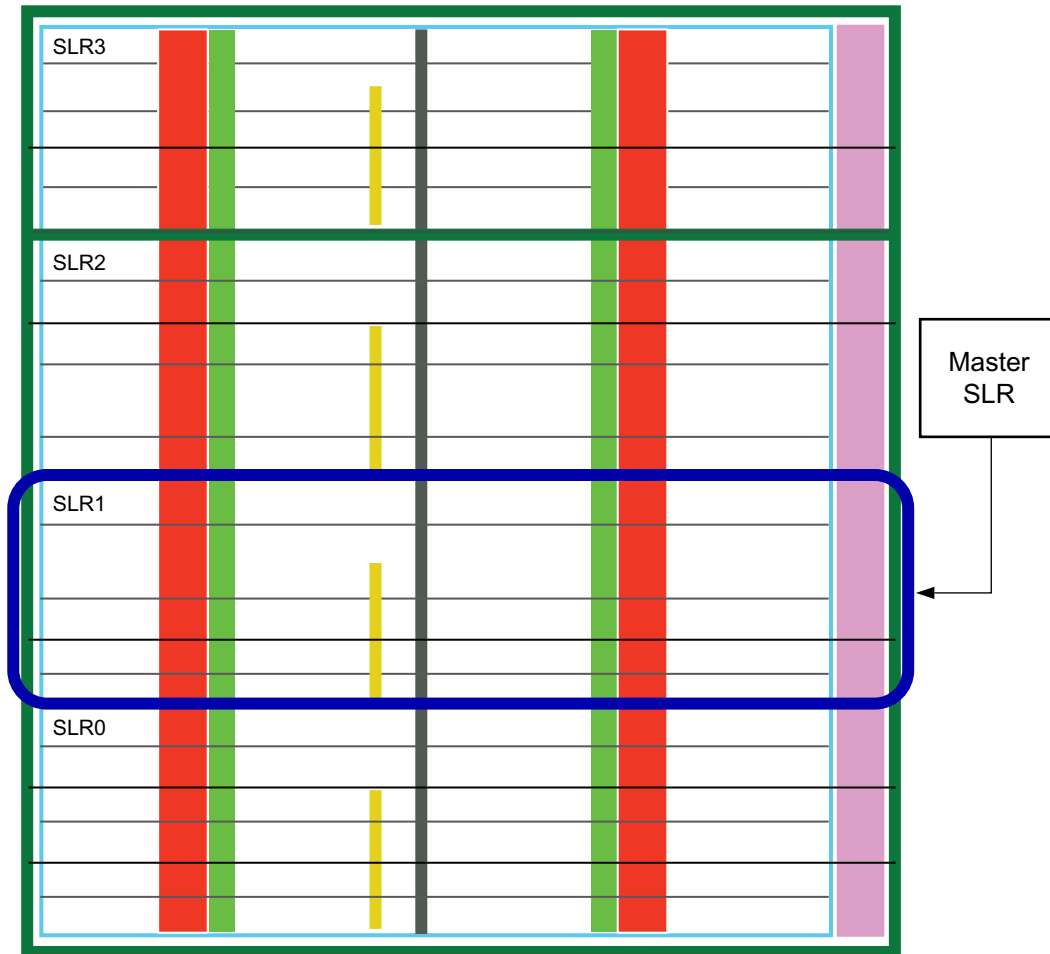
- DNA_PORT
- EFUSE_USER

When using these components, the place and route tools can assign associated pins and logic to the appropriate SLR. In general, no additional intervention is required.



TIP: To query which SLR is the master SLR in the Vivado Design Suite, you can enter the following Tcl command: `get_slrs -filter IS_MASTER`

The following figure shows the master SLR in an XC7V2000T device.



X14689

Figure 3-3: Master SLR in an XC7V2000T Device

Virtex-7 Device Family SLR Components

Two different SLR components are used to create the Virtex[®]-7 device family:

- [XC7V2000T Devices](#)
- [XC7VX1140T and Virtex-7 HT Device Family](#)

XC7V2000T Devices

The XC7V2000T devices share the same type of SLR containing:

- Approximately 500, 000 logic cells
- A mix of the following components:
 - I/O
 - Block RAM
 - DSP blocks
 - GTX Transceivers
 - Other blocks

XC7VX1140T and Virtex-7 HT Device Family

The XC7VX1140T devices and the Virtex-7 HT device family utilize SLR components containing:

- Approximately 290,000 logic cells
- GTH Transceivers
- A larger number of block RAM and DSP components than the XC7V2000T SLR components

Table 3-2: Key Resources Available in Each Virtex-7 SLR Type

	Virtex-7 T SLR	Virtex-7 XT/HT SLR
Logic Cells	488,640	284,800
Slices	76,350	44,500
Block RAM	323	470
DSP Slices	540	840
Clock Regions/MMCM	6	6
I/O	300	300
Transceivers	12	24
Interconnects between SLRs	13,440	10,560

Note: The actual number of bonded out I/O and Transceivers might be different depending on the selected device and package.

UltraScale Device Family SLR Components

Three different SLR components are used to create the UltraScale device family:

- XCKU085 and XCKU115
- XCVU125, XCVU160, XCVU190
- XCVU440

XCKU085 and XCKU115

The XCKU085 and XCKU115 share the same type of SLR containing:

- Approximately 580,000 logic cells
- GTH Transceivers
- A large number of DSP

XCVU125, XCVU160, XCVU190

The XCVU125, XCVU160 and XCVU190 share the same SLR containing:

- Approximately 625,000 Logic Cells
- Several GTH and GTY transceivers
- Several PCIe, Interlaken and Ethernet dedicated blocks

XCVU440

The XCVU440 SLR contains:

- Approximately 1.5 million logic cells
- Several SelectIO™ interface resources
- GTH transceivers

Table 3-3: Key Resources Available in Each UltraScale Device SLR Type

	Kintex® SLRs	XCVU190	XCVU440
Logic Cells	580,440	626,640	1,477,560
CLBs	41,460	44,760	105,540
Block RAM	1,080	1260	840
DSP Slices	2,760	600	960
Clock Regions/MMCM	30	30	45
I/O	624	520	520
GTH Transceivers	32	20	30
GTY Transceivers	0	20	0
Interconnects between SLRs	17,280	17,280	25,920

Note: The actual number of bonded out I/O and transceivers might be different depending on the selected device and package.

Silicon Interposer

The silicon interposer is a passive layer in the SSI technology device.

This layer routes the following between SLR components:

- Configuration
- Global clocking
- General interconnect

The silicon interposer provides:

- Power and ground
- Configuration
- Inter-die connectivity
- Other required connectivity

The active circuitry exists on the SLR. The silicon interposer is bonded to the packaging substrate using Through-Silicon Via (TSV) components. These components connect the circuitry of the FPGA device to the package balls.

The silicon interposer is the conduit between SLR components and the packaging substrate. It connects the following to the device package:

- Power and ground connections
- I/O components
- Gigabit Transceivers (GT)




Figure 3-4: Silicon Interposer

Super Long Line (SLL) Routes

- Super Long Line (SLL) routes provide the general connectivity for signals that cross from one SLR to another.
- SLL routes are located on the interposer.
- SLL routes are connected to the SLR components by microbumps connected directly to the interconnect in the SLR.
- SLL routes connect to the center of Vertical 12 routes in the SLR.

SLL Connectivity

In 7 series devices, each SLL component spans the vertical length of 50 interconnect tiles (equivalent to 50 Slice components). This is exactly the height of one clock region in Xilinx 7 series FPGA devices.

For UltraScale devices, an SLL spans the vertical length of 60 Laguna tiles (equivalent to 60 CLB components). This is exactly the height of one clock region in UltraScale devices.

In 7 series SLR adjacent clock regions, there is one interconnect point connecting to the neighboring SLR at every interconnect tile in the clock region. UltraScale devices differ from 7 series in that the SLLs are connected to a dedicated interface called Laguna. The Laguna

interface has dedicated registers that can be used or bypassed allowing for either high-speed pipelined interfaces or slower speed non-pipelined interfaces to cross the SLRs.

Table 3-4: SLL Components for Each SLR Crossing

Device	SLL Components
7V2000T	13,440
7VX1140T	10,560
KU085	17,280
KU115	17,280
VU125	17,280
VU160	17,280
VU190	17,280
VU440	25,920

The 7VX1140T device has fewer SLL components compared to the 7V2000T, because it has more DSP and block memory columns. These columns displace more interconnect tiles for the same given area. UltraScale devices have a fixed number of SLLs per clock region adjacent to a neighboring SLR. For each clock region, there are 2880 SLLs.





Figure 3-5: Staggered SLLs Crossing in a 7 Series SSI Device

The ratios and gap size between SLR components is for illustration purposes only. The actual gap is comparatively much smaller.

In 7 series devices, the SLL components connect to the SLR at the center point of a Vertical 12 Long Line, which spans 12 interconnect tiles in the SLR. UltraScale devices use the dedicated Laguna interface.

This connectivity provides three optimal places to enter or exit an SLL from SLR to adjacent SLR, and gives additional flexibility to placement with little penalty to performance or power.



X12430

Figure 3-6: Representation of SLL Connectivity in the SLR

Propagation Limitations



TIP: For high-speed propagation across SLRs, consider registering signals that cross SLR boundaries.

SLL signals are the only data connections between SLR components.

The following do not propagate across SLR components:

- Carry chains
- DSP cascades
- Block RAM address cascades
- Other dedicated connections such as DCI cascades and block RAM cascades

The tools normally take this limit on propagation into account. To ensure that designs route properly and meet your design goals, you must also take this limit into account when you

build a very long DSP cascade and manually place such logic near SLR boundaries as well as when you specify a pinout for the design.

SLR Utilization Considerations

The Vivado implementation tools use a special algorithm to partition logic into multiple SLRs. For challenging designs, you can improve timing closure for designs that target SSI technology devices using the following guidelines.

To improve timing closure and compile times, you can use Pblocks to assign logic to each SLR and validate that individual SLRs do not have excessive utilization across all fabric resource types. For example, a design with BRAM utilization of 70% might cause issues with timing closure if the BRAM resources are not balanced across SLRs and one SLR is using over 85% BRAMs.

The following example utilization report for a vu160 shows that the overall BRAM utilization is 56% while the BRAM utilization in SLR0 is 89% (897 out of 1008 available). Timing closure might be more difficult to achieve for the design in SLR0 than with a balanced BRAM utilization across SLRs.

3. BLOCKRAM

```

-----
+-----+-----+-----+-----+-----+
| Site Type | Used | Fixed | Available | Util% |
+-----+-----+-----+-----+-----+
| Block RAM Tile | 1843 | 0 | 3276 | 56.26 |
|   RAMB36/FIFO* | 1820 | 1 | 3276 | 55.56 |
|   FIFO36E2 only | 78 | | | |
|   RAMB36E2 only | 1742 | | | |
|   RAMB18 | 46 | 0 | 6552 | 0.70 |
|   RAMB18E2 only | 46 | | | |
+-----+-----+-----+-----+-----+

```

14. SLR CLB Logic and Dedicated Block Utilization

```

-----
+-----+-----+-----+-----+-----+-----+-----+-----+
| SLR Index | CLBs | (%)CLBs | Total LUTs | Memory LUTs | (%)Total LUTs | Registers | BRAMs | DSPs |
+-----+-----+-----+-----+-----+-----+-----+-----+
| SLR2 | 40109 | 89.61 | 167520 | 156 | 46.78 | 327600 | 512 | 0 |
| SLR1 | 42649 | 95.28 | 205484 | 2297 | 57.38 | 355918 | 434 | 0 |
| SLR0 | 35379 | 97.62 | 163188 | 24 | 56.29 | 313392 | 897 | 0 |
+-----+-----+-----+-----+-----+-----+-----+-----+
| Total | 118137 | | 536192 | 2477 | | 996910 | 1843 | 0 |
+-----+-----+-----+-----+-----+-----+-----+-----+

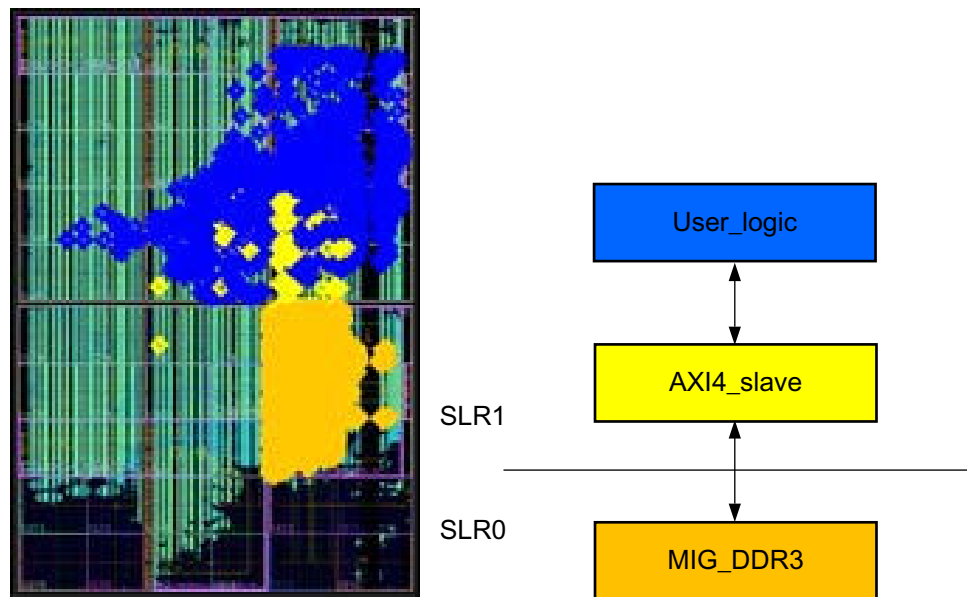
```

Xilinx recommends assigning BRAM and DSP groups to SLR Pblocks to minimize SLR crossings of shared signals. For example, an address bus that fans out to a group of BRAMs that are spread out over multiple SLRs can make timing closure more difficult to achieve, because the SLR crossing incurs additional delay for the timing critical signals.

Device resource location or user I/O selection anchors IP to SLRs (for example, GT, ILKN, PCIe, and CMAC hard IP or memory interface controllers). Pay special attention to hard IP location and pinout selection to avoid data flow crossing SLR boundaries multiple times.

Keep tightly interconnected modules and IPs within the same SLR. If that is not possible, you can add pipeline registers to allow the placer more flexibility to find a good solution despite the SLR crossing between logic groups. Keep critical logic within the same SLR. By ensuring that main modules are properly pipelined at their interfaces, the placer is more likely to find SLR partitions with flip-flop to flip-flop SLR crossings.

In the following figure, a memory interface that is constrained to SLR0 needs to drive user logic in SLR1. An AXI4-Lite slave interface connects to the memory IP backend, and the well-defined boundary between the memory IP and the AXI4-Lite slave interface provides a good transition from SLR0 to SLR1.




X15238-110515

Figure 3-7: Memory Interface in SLR0 Driving User Logic in SLR1

SLR Crossing for Wide Buses

When data flow requirements require that wide buses cross SLRs, use pipelining strategies to improve timing closure and alleviate routing congestion of long resources. For wide buses operating above 250 MHz, Xilinx recommends using at least 3 pipeline stages to cross an SLR: one at the top, one at the bottom, and one in the middle of the SLR. Additional pipeline stages might be required for very high performance buses or when traversing horizontal as well as vertical distances.


The following figure illustrates a worst case crossing for a vu190-2 device. This example starts at an Interlaken hard IP in the bottom left of SLR0 to a packet monitor block assigned to the top right of SLR2. Without pipeline registers for the data bus to and from the packet monitor, the design misses the 300 MHz timing requirement by a wide margin.



X15240-110415

Figure 3-8: Data Path Crossing SLR without Pipeline Flip-Flop

However, adding 7 pipeline stages to aid in the traversal from SLR0 to SLR2 allows the design to meet timing. It also reduces the use of vertical and horizontal long routing resources, as shown in the following figure.



X15239-110415

Figure 3-9: Data Path Crossing SLR with Pipeline Flip-Flop Added

FPGA Power Aspects and System Dependencies

When planning the PCB, you must take power into consideration:

- The FPGA device and the user design create system power supply and heat dissipation requirements.
- Electrical and physical factors can affect the power supply and cooling of the FPGA device, which in turn significantly impacts device performance.

For these reasons, you must understand the power and cooling requirements of the FPGA device. These must be designed on the board.

Power Supply Paths on FPGA Devices

Multiple power supplies are required to power an FPGA device. Some of this power must be provided in a specific sequence. Consider the use of power monitoring or sequencing circuitry to provide the correct power-on sequence to the FPGA device and GTs, as well as

any additional active components on the board. More complex environments may benefit from the use of a microcontroller or system and power management bus such as SMBUS or PMBUS to control the power and reset process. Specific details regarding on/off sequencing can be found in the device data sheet.

The separate sources provide the required power for the different FPGA resources. This allows different resources to work at different voltage levels for increased performance or signal strength, while preserving a high immunity to noise and parasitic effects.

Components of Power Dissipated from the FPGA Device

Three components make up the total required power for each supply source.

- **Device static (leakage) power**

The power required for the device to operate and be available for programming. A large portion of this power is due to leakage in the transistors used to hold the device configuration.

- **Design static power**

The additional continuous power drawn when the device is configured and there is no activity. This includes static current from I/O terminations, clock managers, and other circuits that need power when used, regardless of design activity.

- **Design dynamic power**

The additional power resulting from the design activity. This power varies over time with the design activity. It also depends on the voltage level and the logic and routing resources used.

Power Consumption Paths

The total power supplied to the device flows in then out of the FPGA through multiple paths, including thermal power and off-chip power.

Thermal Power

Thermal power is the power consumed internally within the FPGA. This represents the generation of heat, which contributes to raising the device junction temperature. This heat is then transferred to the environment. The board design must provide heat dissipation paths to ensure that the junction temperature remains within the device operating range.

Off-Chip Power

Off-chip power is the current that flows from the supply source, through the FPGA power pins, then out of the I/Os, and is then dissipated in external board components. The currents

supplied by the FPGA device are generally consumed in off-chip components such as I/O terminations, LEDs, or the I/O buffers of other chips. These do not contribute to raising the device junction temperature of the FPGA device itself. However, power and ground lines must be designed to carry this power.

Power Modes

An FPGA device goes through several power phases from power up to power down with varying power requirements:

- [Power-On](#)
- [Startup Power](#)
- [Standby Power](#)
- [Active Power](#)

Power-On

Power-on power is the transient spike current that occurs when power is first applied to the FPGA device. This current varies for each voltage supply, and depends on the FPGA device construction; the ability of the power supply source to ramp up to the nominal voltage; and the device operating conditions, such as temperature and sequencing between the different supplies.

Spike currents are not a concern in modern FPGA device architectures when the proper power-on sequencing guidelines are followed.

Startup Power

Startup power is the power required during the initial bring-up and configuration of the device. This power generally occurs over a very short period of time and thus is not a concern for thermal dissipation. However, current requirements must still be met. In most cases, the active current of an operating design will be higher and thus no changes are necessary. However, for lower-power designs where active current can be low, a higher current requirement during this time may be necessary. Xilinx Power Estimator (XPE) can be used to understand this requirement. When Process is set to **Maximum**, the current requirement for each voltage rail will be specified to either the operating current or the startup current, whichever is higher. XPE will display the current value in blue if the startup current is the higher value.

Standby Power

Standby power (also called *design static power*) is the power supplied when the device is configured with your design and no activity is applied externally or generated internally.

Standby power represents the minimum continuous power that the supplies must provide while the design operates.

Active Power

Active power (also called *design dynamic power*) is the power required while the device is running your application. Active power includes standby power (all static power), plus power generated from the design activity (design dynamic power). Active power is instantaneous and varies at each clock cycle depending on the input data pattern and the design internal activity.

Environmental Factors Impacting Power

Power depends on several factors beyond the immediate design itself. These are the factors which influence the voltage and the junction temperature of the device, thereby impacting the power dissipation. Such environmental factors impacting power include:

- [Supply Strategies](#)
- [Cooling Strategies](#)

Supply Strategies

Supply strategies include:

- [Regulator Technology](#)
- [Decoupling Network Performance](#)
- [FPGA Device Selection](#)

Regulator Technology

Different regulator technologies exist to balance input-to-output voltage difference, response time, maximum currents, and output voltage accuracy constraints.

Decoupling Network Performance

In addition to supplying the FPGA device during brief high power demand periods, an efficiently designed decoupling circuit reduces current surge requests from the regulator and improves overall regulator consumption.

FPGA Device Selection

Different FPGA device families require different power supply counts and voltage levels. A careful balance among resources, performance, and power is present on all Xilinx FPGA devices. Choose the device that best matches your specific requirements.

Paying excessive attention to one characteristic (for example, performance) can negatively impact another characteristic (for example, power). Selecting a device that supports a lower core voltage or lower voltage I/O interfaces reduces power.

Cooling Strategies

Cooling strategies include:

- [System Environment](#)
- [Heat Sink](#)
- [Package Selection](#)
- [Component Placement](#)

System Environment

The shape and dimension of the system enclosure (together with the ambient air temperature) are the primary factors that impact the transfer of generated heat to the environment.

Heat Sink

The dimension, shape, thermal adhesive, and mounting of the heat sink and the eventual associated forced airflow system determine the amount of heat that can be extracted from the FPGA device.

Package Selection

In addition to cost and signal integrity, the package dimension, material, and connection to the board influence how the generated heat can be transferred to the environment from both the top and bottom level. The larger the contact surface area between the heat sink and board, the lower the thermal resistance.

Component Placement

Component placement relative to the system enclosure and other board material, assembly, and components affects how the heat is transferred to the environment. For example, an obstacle may reduce or redirect the airflow near the FPGA device. Other heat generating components in close proximity to the FPGA device may heat up the air flowing above the device and reduce the heat sink efficiency or transfer heat into the FPGA device by means of the board material.

Power Models Accuracy

The accuracy of the characterization data embedded in the tools evolves over time to reflect the device availability or manufacturing process maturity. This accuracy designation is displayed in the Characterization field. Device family characterization data evolve in the following sequence:

- [Advance Device Designation](#)
- [Preliminary Device Designation](#)
- [Production Device Designation](#)



RECOMMENDED: Use the latest version of Xilinx Power Estimator (XPE) to reflect the latest available data.

Advance Device Designation

Devices with the Advance device designation have data models primarily based on simulation results or measurements from early production device lots. Advance data is typically available within a year of product launch. Advance data is considered relatively stable and conservative, although some under-reporting or over-reporting may occur. Advance data accuracy is considered lower than Preliminary and Production data. Xilinx recommends that you discuss the most recent data with your FAE.

Preliminary Device Designation

The Preliminary device designation is based on complete early production silicon. Almost all the blocks in the device fabric are characterized. The probability of accurate power reporting is improved compared to Advance data.

Production Device Designation

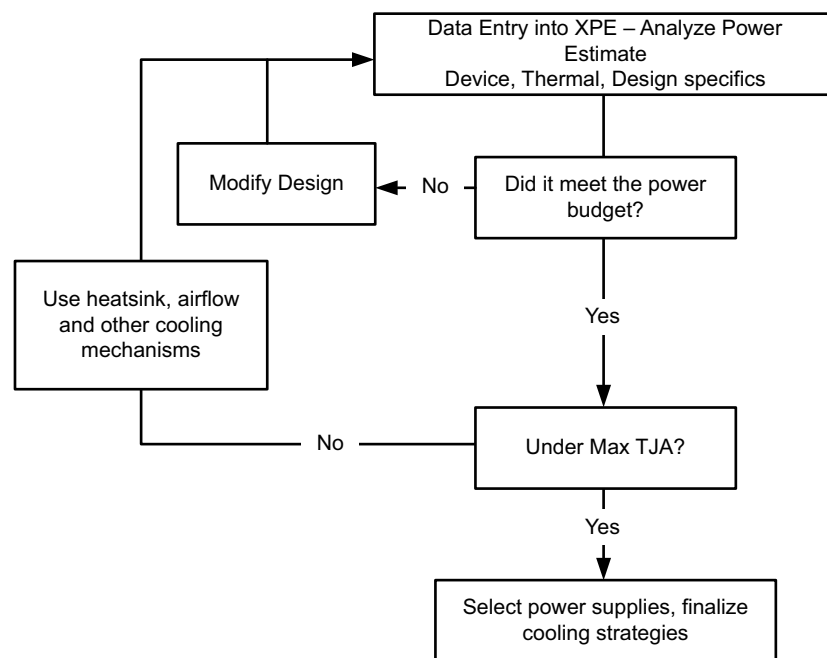
The Production device designation is released after enough production silicon of a particular device family member has been characterized to provide full power correlation over numerous production lots. Device models with Production characterization data are not expected to evolve further.

FPGA Device Power and the Overall System Design Process

From project conception to completion there are many different aspects to consider that influence power. Ignoring for a moment all other issues (including functionality, performance, cost, and time to market), power-related tasks can be sorted into the following separate classes:

- Physical domain, including enclosure, board shape, power delivery system, and thermal power dissipation system
- Functional domain, including area, performance, and I/O interfaces signal integrity

Typically, hardware selection and sizing occurs very early in the design flow to give time for prototype boards to be built. The effect of the FPGA device functionality on power consumption can be estimated early in the design flow, then refined as more and more of the design logic is completed. The following figure illustrates a typical system design process and highlights power-related decision points.



X13425

Figure 3-10: Managing Power Aspects in the PCB Planning Process

When you select your device and associated cooling parts, the FPGA logic is not yet available. A careful methodology to estimate the FPGA logic power requirements is needed. Xilinx recommends using logic data from an earlier design that is comparable to the current design. You can enter design data using your best estimate, then revise the data later. For information on importing data from an earlier design into XPE, see [Implementing the Design in Chapter 5](#).

Xilinx highly recommends a thermal simulation. Thermal models can be obtained from Xilinx. This allows a higher accuracy and it also provides ThetaJA, which can be specified in XPE.

- If a thermal simulation is not possible, provide a best guess for environment conditions, such as heatsink and airflow.
- If the power numbers provided by XPE exceed the power budget, the design may have to be optimized for power.
- If power estimates are within the budget, but TJA is higher than the maximum allowed, consider additional cooling techniques (such as heat sink or airflow) and power optimization techniques in an effort to both improve heat dissipation as well as reduce heat proliferation in the device. Re-run XPE with these new values for environmental conditions, since these also will impact power.

System Level Cooling Strategy

A cooling strategy ensures that the heat generated from the device is extracted and absorbed by the environment. The following cooling strategies are generally available at the beginning of the design, but become less feasible in the later stages. They significantly impact the device static power. These cooling strategies include increasing the airflow, lowering the ambient temperature, and using a heat sink (or a larger heat sink), or selecting a different regulator.

System Level Supply Strategy

Voltage has a large impact on both static and dynamic power. Active control of the voltage level ensures that the desired voltage is applied to the device.

- **Use switching regulators.**

Switching regulators are more power efficient than linear regulators, but at the expense of requiring a higher component count.

- **Use adjustable regulators.**

Sense voltage as close as possible to the FPGA device and to the highest consuming device if the same supply powers multiple FPGA devices.

- **Select regulators with tight tolerances.**

Regulators with tight tolerance ensure consistent voltage supply to the device.

Measuring Power and Temperature

This section briefly describes techniques for measuring FPGA device power consumption and heat dissipation. Some of these techniques use internal FPGA resources. Other techniques use board or external components. Some applications require power and temperature to be actively monitored and adjusted after deployment. Other applications use these measurement techniques in the lab during prototyping and validation phases.

Power Measurement Techniques

Power measurement techniques include:

- [Using a Current Sense Resistor](#)
- [Using Advanced Regulators and Digital Power Controllers](#)
- [Performing On-Board Monitoring](#)
- [Having Separate Voltage Rails](#)

Using a Current Sense Resistor

Inserting a Current Sense Resistor in series between the regulator output and the FPGA device creates a small voltage drop which, by Ohm's Law, is proportional to the flowing current. Measuring this voltage through an XADC gives you the current being supplied to the FPGA device. To understand the connections needed to obtain the desired accuracy of measurements, see the *7 Series FPGAs and Zynq-7000 All Programmable SoC XADC Dual 12-Bit 1 MSPS Analog-to-Digital Converter User Guide* (UG480) [[Ref 45](#)] (also known as the *XADC User Guide*).

Using Advanced Regulators and Digital Power Controllers

The latest evaluation kits include advanced regulator and digital power controllers that you can use to capture regulator output currents and voltages, then send this information to a monitoring computer over a USB interface. This is the simplest and most convenient way to monitor the power rails.

Most Xilinx development boards have integrated Texas Instrument UCD92xx controllers that can be accessed with the Fusions Digital Power Designer software on a PC using a PMBus (I2C) to USB interface module.

Performing On-Board Monitoring

Xilinx 7 series device families provide internal sensors and at least one analog-to-digital converter to measure supplied voltages and device temperature. The Vivado hardware manager provides real-time JTAG access to measure the different supply source voltages or device junction temperature before and after device configuration (see [Figure 3-11](#)). You can also instantiate a System Monitor or XADC component in your code to access these measurements from your FPGA application.

Having Separate Voltage Rails

When possible, have separate voltage rails for each of the supply voltages. If voltage rails are tied together, note it and account for it when power is measured across these rails.

Thermal Measurement Techniques

Thermal measurement techniques include:

- [Performing External Monitoring](#)
- [Performing On-Board Monitoring](#)

Performing External Monitoring

Because the device package prevents access to the silicon, junction temperature cannot be measured directly. Junction temperature can be approximated by measuring the temperature of the package, the heat sink, and other locations with a thermocouple.

Thermal cameras are also used to visualize the device temperature and thermal dissipation interactions with neighboring components and the larger environment.

Performing On-Board Monitoring

Thermal measurements are possible using the same techniques as power measurements. You can use the Vivado hardware manager before and after device configuration ([Figure 3-11](#)).

You can also use the System Monitor/XADC primitive within your design to read the device junction temperature.




Figure 3-11: Voltage and Junction Temperature Monitoring

Methodology for Power and Temperature Measurement

To evaluate the three factors contributing to the design total power, you must control the device junction temperature and let it stabilize before making measurements. This control and stabilization is required because the device and design static power is heavily dependent on the device junction temperature.

The three factors contributing to the design total power are:

- [Device Static](#)
- [Design Static](#)
- [Design Dynamic](#)

Device Static

Download a blank design to ensure that: (1) no input noise is captured; and (2) all internal logic and configuration circuits are in a known state.

Note: A blank design is a design with a single gate or flip-flop that never toggles, and in which all outputs are in a 3-state configuration.

Wait for the junction temperature to stabilize, then measure VCCINT, VCCAUX, and any other supply source of interest. With special equipment, a simple heat gun, or cold spray, you can force temperature changes to evaluate the influence of the environment on the device static power.

Design Static

Download the design onto the FPGA device and do not start any input or internal activity (input data and external and internal clock generation). Wait for the device temperature to stabilize, then measure power on all supply rails of interest.

Subtracting the device static measurement from these values gives you the additional static power from the specific logic resources and configuration used in your design (design static power).

Design Dynamic

Download the design onto the FPGA device and provide clocks and input stimulus representative of the design. Wait for the junction temperature to stabilize before measuring all supply sources of interest.

This power represents the instantaneous total power of the design. It will vary with the change in activity at each clock cycle.

Worst Case Power Analysis Using Xilinx Power Estimator (XPE)

The board should be designed for worst-case power. For details on power analysis using Xilinx Power Estimator (XPE), refer to *Xilinx Power Estimator User Guide* (UG440) [Ref 23].

Setting Expectations

Understanding the total power requirements will help you define your power delivery and cooling system specifications. You want to finalize on the number of:

- Voltage supplies
- Power drawn by each
- Absorbed energy that will generate heat

XPE can answer these questions. It helps you develop in parallel the FPGA logic and the printed circuit board on which the device will be soldered. This exercise will also help you understand the margin you can expect to have and therefore gain confidence that your system will work within budget once implemented. The following figure shows a sample of Xilinx Power Estimator interface.

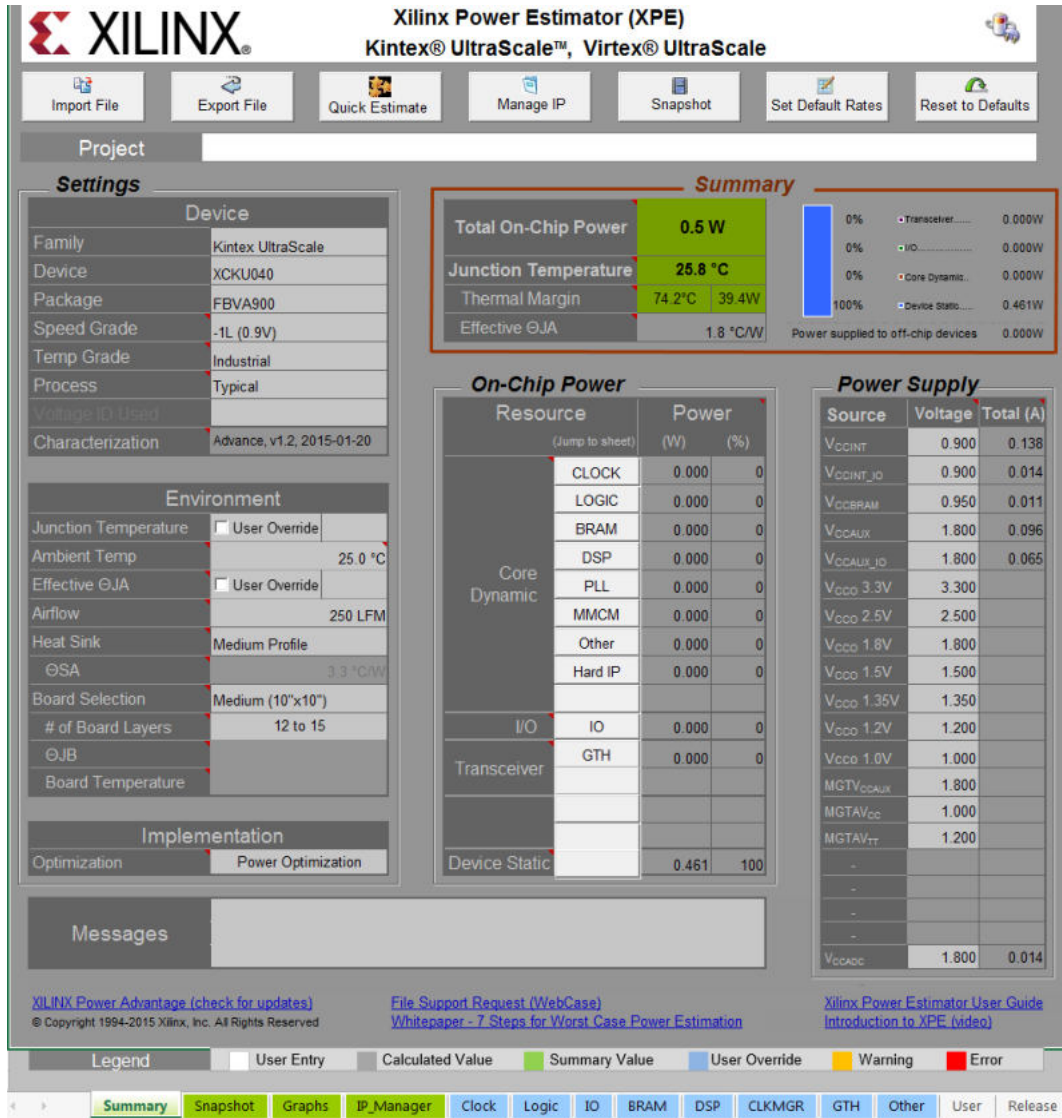


Figure 3-12: Xilinx Power Estimator (XPE) Presentation of Power Information

Estimating Power in the XPE

Power and cooling specifications must be properly set in order to create a functioning and reliable system. In most cases, these thermal and power specifications must be set before PCB design. Due to the flexibility of FPGA devices, the FPGA design often is not completed (or sometimes even started) before system design or PCB fabrication.

This sequence creates a significant challenge for FPGA designers, since thermal and power characteristics can vary dramatically depending on the bitstream (design), clocking, and data put into the chip.

Underdesigning the power or thermal system can make the FPGA device operate out of specification. This can result in the device not operating at the expected performance, and can have other potentially more serious consequences.

Overdesigning the power system is generally less serious, but is still not desirable because it can add unnecessary cost and complexity to the overall system.

There are several techniques for power optimization that can be explored and applied during the analysis and can result in significant power savings. These techniques are discussed in [Chapter 4, Design Creation](#), and [Chapter 5, Implementation](#).

Use Xilinx Power Estimator (XPE) for performing power analysis/estimation at the board-design stage. Refer to this [link](#) in the *Vivado Design Suite User Guide: Power Analysis and Optimization* (UG907) [Ref 22] for more information on obtaining Power Estimates through XPE. Some of the major highlights of the methodology involve:

Use the latest version of the Xilinx Power Estimator (XPE) tool. Power information is updated periodically to reflect the latest power modeling and characterization data. You can get the latest version of XPE from the [Power Efficiency](#) page on the Xilinx website.

Provide accurate information about your specific device selection. Device settings can significantly impact the static and clocking power calculations.

Set the environment conditions in which your device is expected to operate. These conditions impact the heat-dissipation ability and, therefore, the temperature of the device. The temperature, in turn, impacts power.

Increase each voltage rail for your device to the highest voltage value seen at the FPGA device based on tolerances from the supplies or regulators to each rail.

Import the XPower Export File (.xpe) from a design into XPE to help fill out the resource information if the design has already been run in the Vivado tools or if a previous revision of the design has been run, and that revision can be used as a good starting point for the analysis.



TIP: *After importing a Vivado Design Suite XPE file, check that the data is correct and relevant. Consider this information to be a good starting point, but not a complete solution.*

If the .xpe file for a comparable design is not available, examine and (if necessary) fill out the resources expected to be used in the design for each of the resource types, namely:

- Clock Tree Power
- Logic Power
- I/O Power
- Block RAM Power
- DSP Power
- Clock Manager (CLKMGR)
- GT

Review the set value for each tab containing a Toggle Rate, Average Fanout, or Enable Rate, and adjust if needed.

For example:

- If a memory interface has a training pattern routine that exercises a sustained high toggle rate on that interface, *raise* the toggle rate to reflect this additional activity.
- If a portion of a circuit is clock enabled in a way that reduces the overall activity of the circuit, *reduce* the toggle rate.

For more information on how to determine toggle rate, see the *Xilinx Power Estimator User Guide* (UG440) [Ref 23].

For logic fanout, the nature of the data and control paths must be thought out. For example:

- In designs with well-structured sequential data paths (such as DSP designs), fanouts generally tend to be lower than the set default.
- In designs with many data execution paths (such as some embedded designs), higher fanouts may be seen.

Before you review the results, iterate through the above sequence if necessary. After completing these steps, analyze the results.

Be sure that the junction temperature is not exceeded, and that the power drawn is within the desired budget for the project. If the thermal dissipation or power characteristics are not within targets:

- Adjust the environmental characteristics (for example, increase airflow or add a heatsink), or
- Adjust the resource and power characteristics of the design until an acceptable result is reached.

Many times, trade-offs can be made to derive the desired functionality with a tighter power budget. The best time to explore these options is early in the design process. Once the data is completely entered, and the part is operating within the thermal limits of the selected grade, the power reported by XPE can be used to specify the rails for the design.

If your confidence in the data entered is not very high, you can pad the numbers to circumvent the possibility of underdesigning the power system for the device.

As the design matures, continue to review and update the information in the spreadsheet to reflect the latest requirements and implementation details. This will present the most current picture of the power used in the design and can potentially allow early identification of adjustments to the power budgeting up or down depending on the current power trends of the design.

Configuration

Configuration is the process of loading application-specific data into the internal memory of the FPGA device.

Because Xilinx FPGA configuration data is stored in CMOS configuration latches (CCLs), the configuration data is volatile and must be reloaded each time the FPGA device is powered up.

Xilinx FPGA devices can load themselves through configuration pins from an external nonvolatile memory device. They can also be configured by an external smart source, such as a:

- Microprocessor
- DSP processor
- Microcontroller
- Personal Computer (PC)
- Board tester

Board Planning should consider configuration aspects up front, which makes it easier to configure as well as debug.

Each FPGA device family has a *Configuration User Guide* [Ref 37] that is the primary resource for detailed information about each of the supported configuration modes and their trade-offs on pin count, performance, and cost.

Board Design Tips

When designing a board, it is important to consider which interfaces and pins will assist with debug capability beyond configuration. For example, Xilinx recommends that you ensure the JTAG interface is accessible even when the interface is not the primary configuration mode. The JTAG interface allows you to check the device ID and device DNA information, and you can use the interface to enable indirect flash programming solutions during prototyping.

In addition, signals such as the INIT_B and DONE are critical for FPGA configuration debug. The INIT_B signal has multiple functions. It indicates completion of initialization at power-up and can indicate when a CRC error is encountered. Xilinx recommends that you connect the INIT_B and DONE signals to LEDs using LED drivers and pull-ups. See the FPGA family configuration user guide for recommended pull-up values.

The schematic checklists include these recommendations along with other key suggestions. Use these checklists to identify and check recommended board-level pin connections:

- 7 Series Schematic Review Recommendations (XMP277) [\[Ref 50\]](#)
- UltraScale Architecture Schematic Review Checklist (XTP344) [\[Ref 51\]](#)

Design Creation

Overview of Design Creation

You have planned your device I/O, and you have planned on how to lay out your PCB. You have also decided on your use model for the Vivado® Design Suite. You can now begin creating your design.

The main points to consider include:

- Achieving the desired functionality.
- Operating at the desired frequency.
- Operating with the desired degree of reliability.
- Fitting within the silicon resource and power budget.

Design creation (in order to achieve the above goals) involves:

- Planning the hierarchy of your design.
- Identifying the IP cores to use and customize in your design.
- Creating the custom RTL for interconnect logic and functionality for which a suitable IP was not found.
- Creating timing and physical constraints.
- Specifying additional constraints, attributes, and other elements used during synthesis and implementation.

Any decision here has a wide and deep impact on the end product. A wrong decision at this stage can result in problems at a later stage, causing iterations through the entire design cycle. Spending the time early in the process to create a carefully planned design is worth the effort. This will help achieve the desired design goals and minimize debug time in lab.

Defining a Good Design Hierarchy

The first step in design creation is to decide how to partition the design logically. The main factor when considering hierarchy is to partition a part of the design that contains a specific function. This allows a specific designer to design a piece of IP in isolation as well as isolating a piece of code for reuse.

However, defining a hierarchy based on functionality only, does not take into account how to optimize for timing closure, runtime, and debugging. The following additional considerations made during hierarchy planning also help in timing closure.

Infer I/O Components Near the Top Level

Where possible, infer I/O components near the top level for design readability. Components that can be inferred are simple single-ended I/O (IBUF, OBUF, OBUFT and IOBUF) and single data rate registers in the I/O. I/O components that need to be instantiated, such as differential I/O (IBUFDS, OBUFDS) and double data-rate registers (IDDR, ODDR, ISERDES, OSERDES), should also be instantiated near the top level.

Place Clocking Elements Towards the Top Level

Placing the clocking elements towards the top level allows for easier clock sharing between modules. This sharing may result in fewer clocking resources needed, which helps in resource utilization, improved performance, and power.

Aside from the module the clocks are created in, clock paths should only drive down into modules. Any paths that go through (down from top and then back to top) can create a delta cycle problem in VHDL simulation that is difficult and time consuming to debug.

Register Data Paths at Logical Boundaries

Register the outputs of hierarchical boundaries to contain critical paths within a single module or boundary. Consider registering the inputs also at the hierarchical boundaries. It is always easier to analyze and repair timing paths which lie within a module, rather than a path spanning multiple modules. Any paths that are not registered at hierarchy boundaries should be synthesized with hierarchy rebuilt or flat to allow cross hierarchy optimization. Registering the datapaths at logical boundaries helps to retain traceability (for debug) through the design process because cross hierarchical optimizations are kept to a minimum and logic does not move across modules.

Address Floorplanning Considerations

A floorplan ensures that cells belonging to a specific portion in the design netlist are placed at particular locations on the device. You can use manual floorplanning:

- To partition logic to a particular SLR when using SSI technology devices. This confines the launch, destination registers, or both.
- To close timing on a design when timing is not met using standard flows.
- When using hierarchical design flows such as partial configuration.

If the cells are not contained within a level of hierarchy, all objects must be included individually in the floorplan constraint. If synthesis changes the names of these objects, you must update the constraints. A good floorplan is contained at the hierarchy level, since this requires only a one line constraint.

Floorplanning is not always required. Floorplan only when necessary.

For more information on floorplanning, see this [link](#) in the *Vivado Design Suite User Guide: Design Analysis and Closure Techniques* (UG906) [Ref 21].



RECOMMENDED: *While the Vivado tools allow cross hierarchy floorplans, these require more maintenance. Avoid cross hierarchy floorplans where possible.*

Optimize Hierarchy for Functional and Timing Debug

As discussed earlier in this section, keeping the critical path within the same hierarchical boundary is helpful in debugging and repairing timing. Similarly, for functional debug (and modification) purposes, signals that are related should be kept in the same hierarchy. This allows the related signals to be probed and modified with relative ease.

Apply Attributes at the Module Level

Applying attributes at the module level can keep code tidier and more scalable. Instead of having to apply an attribute at the signal level, you can apply the attribute at the module level and have the attribute propagated to all signals declared in that region. Applying attributes at the module level also allows you to override global synthesis options. For this reason, it is sometimes advantageous to add a level of hierarchy in order to apply module level constraints in the RTL.



CAUTION! *Some attributes (e.g., DONT_TOUCH) do not propagate from a module to all the signals inside the module.*

Optimize Hierarchy for Advanced Design Techniques

Advanced design techniques such as bottom-up synthesis, partial reconfiguration, and out-of-context design require planning at the hierarchical level. The design must choose the appropriate level of hierarchy for the technique being used. These techniques are not covered in this version of the document. For more information, see this [link](#) in the *Vivado Design Suite User Guide: Hierarchical Design* (UG905) [Ref 20].

Example of Upfront Hierarchical Planning for High Speed DSP Designs

The following example is not applicable to all designs, but demonstrates what can be done with hierarchy. DSP designs generally allow latency to be added to the design. This allows registers to be added to them to be optimized for performance. In addition, registers can be used to allow for placement flexibility. This is important because at high speed, you cannot traverse the die in one clock cycle. Adding registers can allow hard-to-reach areas to be used. The following figure shows how effective hierarchy planning results in faster timing closure.

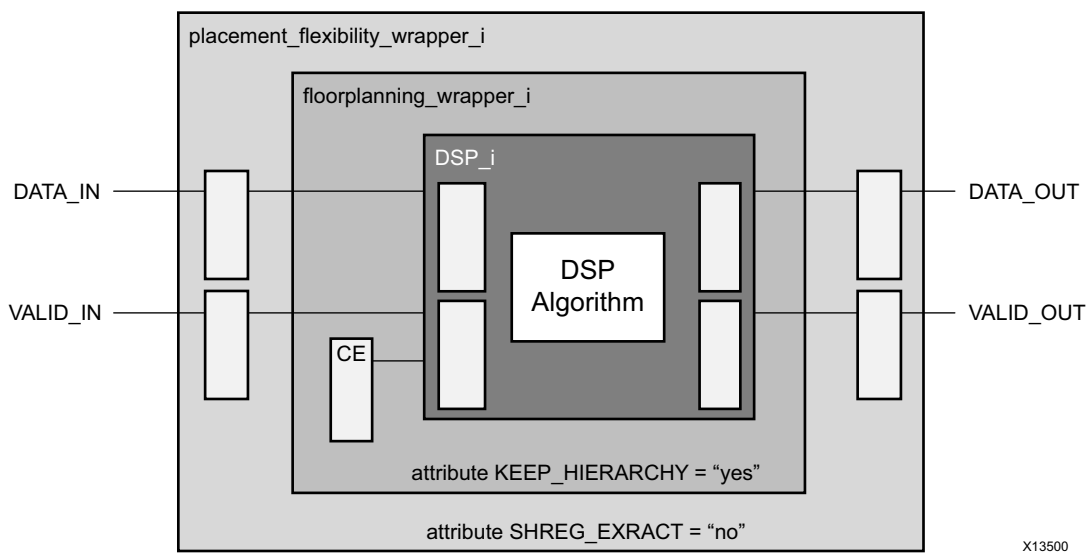


Figure 4-1: Effective Hierarchy Planning Example

There are three levels of hierarchy in this part of the design:

- `DSP_i`

In the `DSP_i` algorithm block, both the inputs and outputs are registered. Because registers are plentiful in an FPGA device, it is preferable to use this method to improve the timing budget.

- `floorplanning_wrapper_i`

In `floorplanning_wrapper_i`, there is a CE signal. CE signals are typically heavily-loaded and can present a timing challenge. They should be included in a floorplan. By creating a floorplanning wrapper, this module can be manually floorplanned later if needed.

In addition, `KEEP_HIERARCHY` has been added at the module level to ensure that hierarchy is preserved for floorplanning regardless of any other global synthesis options.

- `placement_flexibility_wrapper_i`

In `placement_flexibility_wrapper_i`, the `DATA_IN`, `VALID_IN`, `DATA_OUT` and `VALID_OUT` signals are registered. Because these signals are not intended to be part of the floorplan, they are outside `floorplanning_wrapper_i`. If they were in the floorplan, they would not be able to fulfill the requirement for placement flexibility.

In addition, more registers can be added later as long as both `DATA_IN + VALID_IN` or `DATA_OUT` and `VALID_OUT` are treated as pairs. If more registers are added, the synthesis tool may infer SRLs which will force all registers into one component and not help placement flexibility. To prevent this, `SHREG_EXTRACT` has been added at the module level and set to `NO`.

RTL Coding Guidelines

You might have to write your custom RTL to implement either the glue logic functionality, or for some function for which a suitable IP was not found.

Basic Functionality

The RTL must be coded in such a way as to be safely implementable. Otherwise, design functionality may differ from RTL simulation results. The RTL must be free of race conditions and of common coding pitfalls that can result in simulation-synthesis mismatches.

There are some basic guidelines on synthesizable RTL code. Most readily-available literature on synthesizable RTL code discusses these guidelines. Some of the more common guidelines are mentioned below, but the list is not exhaustive. You can also run a set of RTL DRCs as explained in [Check Your HDL Code in Chapter 5](#).

Blocking Statements vs. Non-Blocking Statements

Improper use of Verilog blocking statements or non-blocking statements might result in race conditions, causing RTL simulations to not match netlist simulation. As a general rule, Xilinx recommends using non-blocking statements for all sequential elements, and blocking statements for all combinational elements. This makes the simulation event ordering much more predictable, and less likely to cause race conditions.

Incomplete Sensitivity List

A sensitivity list in a process statement (VHDL) or always block (Verilog) is a list of signals to which the process statement (VHDL) or always block (Verilog) is sensitive. When a listed signal changes its value, the process statement (VHDL) or always block (Verilog) triggers and executes its statements.

In case of an incomplete sensitivity list, while you may get the hardware you intended, the RTL and post-synthesis simulation will differ. In this case, some synthesis tools may issue a message warning of an incomplete sensitivity list. In that event, check the synthesis log file and, if necessary, fix the RTL code.

The following example describes a simple AND function using a process and always block. The sensitivity list is complete and a single LUT is generated.

VHDL process coding example one:

```
process (a,b) begin
  c <= a and b;
end process;
```

Verilog always block coding example one

```
always @(a or b)
  c <= a & b;
```

However, if signal b is omitted from the sensitivity list, the synthesis tool still generates the combinational logic (AND function), though RTL simulation might not trigger the evaluation of c based on changes to b. This results in RTL simulation behaving one way, while the actual circuit behaves differently. Following is an example warning message:

```
WARNING: [Synth 8-567] referenced signal <signal name> should be on the sensitivity
list [<file name>:<line number>]
```

In Verilog, when defining a combinational always block, use an asterisk sensitivity list:

```
always @(*)
```

This automatically uses a fully specified sensitivity list.

Note: Alternatively, when using SystemVerilog, use of the `always_comb` also eliminates the need to specify a sensitivity list as well as documenting the intent of the block.

Delays in RTL Code

Avoid using any kind of delay in your RTL code, either through use of `wait` or `AFTER` (VHDL) or `#delay` (Verilog). Delays do not synthesize to a component. In designs that include explicit delay assignments, the functionality of the *simulated* design does not always match the functionality of the *synthesized* design.

Latch Inference

Synthesizers infer latches from incomplete conditional expressions in combinational, non-sequential logic, such as:

- An `if` statement without an `else` clause
- An intended register without a rising edge or falling edge construct

If statement without an else clause VHDL coding example

```
process (G, D) begin
  if (G='1') then
    Q <= D;
  end if;
end process;
```

If statement without an else clause Verilog coding example

```
always @(G or D)
  if (G)
    Q = D;
```

Many times a branch or edge is missing by mistake. Check your synthesis log to see the latches being inferred. Confirm that any latches inferred are intentional, rather than an oversight.

Xilinx recommends that you avoid using latches in FPGA designs, due to the more difficult timing analyses that occur when latches are used, even if the simulations pass.

Follow the recommended coding styles in the synthesis tool documentation to avoid inferring latches.

Note: When using SystemVerilog, if a latch is needed, Xilinx recommends the `always_latch` construct to document that a latch is the intended result. This can help during debug and code reviews.

Incomplete Reset Specification

In the following example snippet, only `reg1` is assigned within the reset branch, and `reg2` is missed. Synthesis will assume that `reg2` has to hold its value when reset is asserted. Thus, the reset signal will get hooked to the CE pin, thereby creating another unique control set. See [Control Signals and Control Sets](#).

```
always @(posedge clk)
    if (rst)
        reg1<= 1'b0;
    else
    begin
        reg1 <= din1;
        reg2 <= din2;
    end
```



TIP: *If a reset is being used, make sure that the registers have not been missed by mistake in the reset branch.*

Using Vivado Design Suite HDL Templates



RECOMMENDED: *Use the Vivado Design Suite language templates when creating RTL or instantiating Xilinx® primitives. The language templates include recommended coding constructs for proper inference to the Xilinx FPGA device architecture. Using the templates should both ease design and lead to improved results in many cases.*

To access the templates from the Vivado Design Suite IDE:

1. Go to **Windows > Language Template**.
2. Choose the desired template.

HDL Coding for Efficiency

Use of Loops in Code

Loops in HDL are often used to minimize coding effort. When inferring hardware, loop un-rolling may lead to inefficient structures thereby degrading performance (both area as well as timing). Mapping the un-rolled logic to available resources causes possibilities of sub-optimal implementation. Xilinx recommends representing the same functionality using constructs that are easier for the tool to interpret.

Consider a case of priority MUX code using a `for` loop:

```
always@(posedge clk) begin
  for(i=0;i<=3;i=i+1) begin
    if(en[i]) dout[i] <= i;
  end
end
```

The same functionality can be coded using `case/if-else`, and is easier for the tool to interpret to generate efficient hardware.

Sometimes though, a `for` loop might provide the required conciseness, without impacting the quality of results (for example a bus-reversal code).

Example of using of loops in code

```
reg [3:0] dout;
integer i;

always@(posedge clk)
  for(i=0;i<=3;i=i+1)
    dout[3-i] <= din[i];
```



TIP: *It is acceptable to infer loops for basic connectivity. However, when the code infers hardware resources (other than just wires/interconnects), it is better to avoid loops.*

State-Machine Guidance

There are several methods to code state machines. Following certain coding styles ensures that the synthesis tool FSM (Finite State Machine) extraction algorithms properly identify and optimize the state machines as well as possibly improving the simulation, timing and debug of the circuit. The choice of state machines depends on the target architecture and specifics of the state machines size and behavior. Some of the basic trade-offs for different implementation are explained below.

- **Mealy vs. Moore Styles**

There are two well-known implementation styles for state machines, Mealy and Moore. The main difference between Mealy and Moore is that a Mealy state machine determines the output values based on both the current state as well as the inputs to the state machines, whereas a Moore state machine determines its outputs solely on the state.

In general, Moore state machines implement best in FPGA devices due to the fact that most often one-hot state machines is the chosen encoding method, and there is little decode logic necessary for output values.

For a binary encoding, sometimes a more compact or faster state machines can be built using the Mealy machine. However, this is not easy to determine without knowing more specifics of the state machines.

- **One-Hot vs. Binary Encoding**

There are several encoding methods for state machines design. The two most popular for FPGA designs are binary and one-hot. Most modern synthesis tools contain FSM extraction algorithms that can identify state machines code and choose the best encoding method. Sometimes it can be more advantageous to manually code the encoding scheme for the design to allow better control, and possibly to ease debug of the implemented design. See your synthesis tool documentation for details about the state machines extraction capabilities.

- **Safe vs. Fast**

When coding state machines, there are two generally conflicting goals that must be understood: safe vs. fast. A safe state machine implementation is one in which, if a state machines gets an unknown input, or goes into an unknown state, it can recover into a known state (in the next cycle) and resume from that recovery state. On the other hand, if this requirement is discarded (no recovery state), many times the state machines can be implemented with less logic and more speed. Designing a safe-state involves coding in a default state into the state machines case clause and/or specifying to the synthesis tool to implement the state machines encoding in a "safe" mode. If a safe-state capability is desired, usually binary encoding works best as there are fewer unassigned states with binary encoding. Consult your synthesis tool documentation for details about implementing a safe state machines.

- **Enumerated Type**

SystemVerilog adds a new data type `enum` (short for enumerated), which in many cases is beneficial for state machines creation. The `enum` data type allows for named states without implicit mapping to a register encoding. The benefit this provides to synthesis is flexibility in state machine encoding techniques and for simulation, the ability to display and query specific states by name to improve overall debugging. For these reasons, Xilinx recommends using enum types when SystemVerilog, or VHDL (which always had this capability) is the chosen design language.

See *Vivado Design Suite User Guide: Synthesis* (UG901) [Ref 16].

Avoid Preserving Hierarchical Boundaries

Preserving hierarchical boundaries may lead to hardened boundaries, which can negatively impact cross-boundary optimizations.

Consider the following example code snippet:

```

assign ored_signal = din[3]|din[2];
sub sub_inst (.clk (clk),
.din0 (ored_signal),
.din1 (din[1:0]),
.dout (dout));
endmodule

module sub ....
assign din_tmp = |din1 || din0;
endmodule
    
```

There are two ORs:

- ored_signal in the top level
- din_tmp in sub

If the hierarchical boundary between them is preserved using synthesis attributes or constraints, these two ORs cannot be combined, impacting the area and timing of the design. For more information, see [Defining a Good Design Hierarchy](#).

Avoid Mixing Edges of a Flip Flop

If you use both positive and negative edges of clocks for triggering sequential elements, then the path between the elements being triggered by the two different polarities will get only half a clock cycle. This makes the timing more stringent.



TIP: *If both clock edges are being used to capture or provide external DDR type data, use Xilinx IDDR/ODDR primitives.*

Use of Debug Logic

Coding efficiency leads to efficiency in design implementation. Unnecessary constructs often lead to unnecessary logic. Keep this in mind when designing debug signals or logic that is not necessary for the design function, but which is useful in the design analysis. Many times, such debug code serves a valuable purpose during the design phase, but becomes unwanted surplus as the design matures. You should design such logic so that it can still serve its debug purpose, yet not remain in the final design.

Several methods can assist in this objective:

- Guard the logic with a ``ifdef`, parameter, or generic that can be set to disable or enable these sections of code.
- Code the logic in a way to more easily facilitate commenting it out for the future.
- Have a separate debug version of a module or entity to interchange for this purpose.

Regardless of the method chosen, the idea remains the same: it is important not only to have a good methodology for debugging the design code and the implemented hardware, but it is also important to have a good way to remove that logic when it is no longer necessary.

For the details of debug method, see [Chapter 6, Configuration and Debug](#).

Arrays in Port Declarations

Although VHDL allows you to declare a port as an array type, Xilinx recommends that you not do so, for the following reasons:

- Incompatibility with Verilog

There is no equivalent way to declare a port as an array type in Verilog. This limits portability across languages. It also limits as the ability to use the code for mixed-language projects.

- Inability to Store and Re-Create Original Array Declaration

When you declare a port as an array type in VHDL, the original array declaration cannot be stored and re-created. The Electronic Data Interchange Format (EDIF) netlist format, as well as the Xilinx database, are unable to store the original type declaration for the array. As a result, when a simulation netlist is generated, there is no information as to how the port was originally declared.

The resulting netlist generally has mismatched port declarations and resulting signal names. This is true not only for the top-level port declarations, but also for the lower-level port declarations of a hierarchical design since KEEP_HIERARCHY can be used to attempt to preserve those net names.

- Miscalculation of Software Pin Names

Array port declarations can cause a miscalculation of the software pin names from the original source code. Since the tool must treat each I/O as a separate label, the corresponding name for the broken-out port may not match your expectation. This makes design constraint passing, design analysis, and design reporting more difficult to analyze.

Control Signals and Control Sets

A control set is the grouping of control signals (set/reset, clock enable and clock) that drives any given SRL, LUTRAM, or register. For any unique combination of control signals, a unique control set is formed. The reason this is an important concept is registers within a 7 series slice all share common control signals and thus only registers with a common control set may be packed into the same slice.

Designs with several unique control sets may have a lot of wasted resources, as well as fewer options for placement resulting in higher power and lower performance. Designs with fewer control sets have more options and flexibility in terms of placement, generally resulting in improved results. In UltraScale™ devices, there is more flexibility in control set mapping within a CLB. However, it remains a good practice to limit unique control sets to give maximum flexibility in placement of a group of logic.

The following table provides a guideline on the number of control sets that might be acceptable for designs with Xilinx 7 series FPGAs.

Table 4-1: Control Set Guidelines for 7 series FPGAs

Condition	Typically Acceptable	Analysis Required	Recommended Design Change
Number of Unique Control Sets^a	< 7.5% of total slices ^b	>15% of total slices ^{a, b}	>25% of total slices ^b Reducing the number of control sets increases utilization and performance.
Number of registers lost to control set restriction^a	Slice utilization < 75% and registers lost < 4% ^c	Slice utilization > 85% and registers lost > 2% ^c	Slice utilization > 90% and registers lost > 1% ^c

- a. Run `report_utilization` and `report_control_sets -verbose`.
- b. Slices can be found in the device product table. (For example, XC7VX690T contains 108,300 slices.)
Acceptable: 108,300 slices x 7.5% = 8122 control sets
Analysis required: 108,300 slices x 15% = 16,245 control sets
- c. Total available registers



TIP: UltraScale devices are more flexible with the mapping and use of control sets. Therefore, the guidelines for control set usage is less stringent. However, Xilinx still recommends following the 7 series guidelines to allow for best design optimization, placement, and portability.

Resets

Resets are one of the more common and important control signals to take into account and limit in your design. Resets can significantly impact your design’s performance, area, and power.

Inferred synchronous code may result in resources such as:

- LUTs
- Registers
- Shift Register LUTs (SRLs)
- Block or LUT Memory
- DSP48 registers

The choice and use of resets can affect the selection of these components, resulting in less optimal resources for a given design. A misplaced reset on an array can mean the difference between inferring one block RAM, or inferring several thousand registers.

Asynchronous resets described at the input or output of a multiplier may result in registers placed in the slice(s) rather than the DSP block. In these and other situations, the amount of resources is obviously impacted. However, overall power and performance can also be significantly impacted.

When and Where to Use a Reset

FPGA devices have dedicated global set/reset signals (GSR). These signals initialize all registers to the initial value specified state in the HDL code at the end of device configuration.

If an initial state is not specified, it defaults to a logic zero. Accordingly, every register is at a known state at the end of configuration, regardless of the reset topology specified in the HDL code. It is not necessary to code a global reset for the sole purpose of initializing the device on power-up.

Xilinx highly recommends that you take special care when deciding when the design requires a reset, and when it does not. In many situations, resets might be required on the control path logic for proper operation. However, resets are generally less necessary on the data path logic. Limiting the use of resets:

- Limits the overall fanout of the reset net.
- Reduces the amount of interconnect necessary to route the reset.
- Simplifies the timing of the reset paths.
- Results in many cases in overall improvement in performance, area, and power.



RECOMMENDED: *Evaluate each synchronous block, and attempt to determine whether a reset is required for proper operation. Do not code the reset by default without ascertaining its real need.*

Functional simulation should easily identify whether a reset is needed or not.

For logic in which no reset is coded, there is much greater flexibility in selecting the FPGA resources to map the logic.

The synthesis tool can then pick the best resource for that code in order to arrive at a potentially superior result by considering, for example:

- Requested functionality
- Performance requirements
- Available device resources
- Power

Synchronous Reset vs. Asynchronous Reset

If a reset is needed, Xilinx recommends code synchronous resets. Synchronous resets have many advantages over asynchronous resets.


- Synchronous resets can directly map to more resource elements in the FPGA device architecture.
- Some resources such as the DSP48 and block RAM have only synchronous resets for the register elements within the block. When asynchronous resets are used on register elements associated with these elements, those registers may not be inferred directly into those blocks without impacting functionality.
- Asynchronous resets also impact the performance of the general logic structures. As all Xilinx FPGA general-purpose registers can program the set/reset as either asynchronous or synchronous, it can be perceived that there is no penalty in using asynchronous resets. That assumption is often wrong. If a global asynchronous reset is used, it does not increase the control sets. However, the need to route this reset signal to all register elements increases timing complexity. For more information, see [Use of Untimed Resets](#).
- If using asynchronous reset, remember to synchronize the deassertion of the asynchronous reset. For more information, see [Controlling and Synchronizing Device Startup](#).
- Synchronous resets give more flexibility for control set remapping when higher density or fine tuned placement is needed. A synchronous reset may be remapped to the data path of the register if an incompatible reset is found in the more optimally placed Slice. This can reduce wire length and increase density where needed to allow proper fitting and improved performance.
- Asynchronous resets might require multi-cycle assertion to ensure a circuit is properly reset and stable. When properly timed, synchronous resets do not have this requirement.
- Use synchronous resets if asynchronous resets have a greater probability of upsetting memory contents to BRAMs, LUTRAMs, and SRLs during reset assertion.

Control Signal Polarity (Active-High vs. Active-Low)

For high-fanout control signals like clock enables or resets, it is best to use active high in the entire design. If a block operates with active low resets or clock enables, inverters get added to the design and there is an associated timing penalty. It can restrict synthesis options to flat or rebuilt to optimize the inverters or require the implementation of a custom solution.

The Slice and internal logic of the Xilinx FPGA clock enables and resets are inherently active-High. Describing active-Low resets or clock enables may result in additional LUTs used as simple inverters for those routes.

For UltraScale devices, a programmable inversion is available on the reset. Therefore, reset polarity is more flexible. However, Xilinx still recommends keeping the reset polarity coding consistent (all active-High or all active-Low) to allow for maximum flexibility for packing logic. The enable does not have an inversion so Xilinx recommends always describing active-High enables.



X13426

Figure 4-2: Extra Inverters Due to Active Low Control Signals

Reset Coding Example One

The following coding example constructs a highly pipelined multiply and parity generation function:

```
// Reset synchronization
always @(posedge CLK) begin
    reset_sync <= SYS_RST;
    reset_reg <= reset_sync;
end
// Uses active-Low, async reset
// Also using an active-Low CE
always @(posedge CLK, negedge reset_reg)
if (!reset_reg) begin
    data1_reg <= 16'h0000;
    data2_reg <= 16'h0000;
    DATA_VALID <= 1'b0;
end else if (!NEW_DATA) begin
    data1_reg <= DATA1;
    data2_reg <= DATA2;
    DATA_VALID <= data_valid_delay[3];
end
// Uses an async reset when a reset is not necessary
always @(posedge CLK, negedge reset_reg)
if (!reset_reg) begin
    parity <= 4'h0;
    data1_pipe <= 32'h00000000;
    data2_pipe <= 32'h00000000;
    mult_data_reg <= 32'h00000000;
    mult_pipe <= 32'h00000000;
    mult_pipe2 <= 32'h00000000;
    mult_par_reg <= 36'h00000000;
    mult_par_pipe <= 36'h00000000;
    data_valid_delay <= 4'h0;
    DATA_OUT <= 36'h00000000;
end else begin
    data1_pipe <= data1_reg;
    data2_pipe <= data2_reg;
    mult_data_reg <= data1_pipe * data2_pipe;
    mult_pipe <= mult_data_reg;
    parity <= {^mult_pipe[31:24], ^mult_pipe[23:16],
    ^mult_pipe[15:8], ^mult_pipe[7:0]};
    mult_pipe2 <= mult_pipe;
    mult_par_reg <= {parity[3], mult_pipe2[31:24],
    parity[2], mult_pipe2[23:16],
    parity[1], mult_pipe2[15:8],
    parity[0], mult_pipe2[7:0]};
    data_valid_delay <= {data_valid_delay[2:0], NEW_DATA};
    mult_par_pipe <= mult_par_reg;
    DATA_OUT <= mult_par_pipe;
end
end
```

Reset Coding Example Two

The above code can be rewritten to:

- Remove unnecessary resets
- Change Async resets to Sync
- Change active-Low reset to active-High

```
// Reset synchronization, inversion moved here
always @(posedge CLK) begin
    reset_sync <= SYS_RST;
    reset_reg <= ~reset_sync;
end
// Notice the inversion above
// sync reset has become active High, though:
// from the top level port (SYS_RST) perspective,
// it is still active Low.
// Also changed to active-High CE
always @(posedge CLK)
if (reset_reg) begin
    data1_reg <= 16'h0000;
    data2_reg <= 16'h0000;
    DATA_VALID <= 1'b0;
end else if (NEW_DATA) begin
    data1_reg <= DATA1;
    data2_reg <= DATA2;
    DATA_VALID <= data_valid_delay[3];
end

// Removed unnecessary reset on datapath
always @(posedge CLK) begin
    data1_pipe <= data1_reg;
    data2_pipe <= data2_reg;
    mult_data_reg <= data1_pipe * data2_pipe;
    mult_pipe <= mult_data_reg;
    parity <= {^mult_pipe[31:24], ^mult_pipe[23:16],
^mult_pipe[15:8], ^mult_pipe[7:0]};
    mult_pipe2 <= mult_pipe;
    mult_par_reg <= {parity[3], mult_pipe2[31:24],
parity[2], mult_pipe2[23:16],
parity[1], mult_pipe2[15:8],
parity[0], mult_pipe2[7:0]};
    data_valid_delay <= {data_valid_delay[2:0], NEW_DATA};
    mult_par_pipe <= mult_par_reg;
    DATA_OUT <= mult_par_pipe;
end
```

The implementation of the second coding example, compared to the first coding example, is shown in the following table.

Table 4-2: Comparison of Coding Examples Targeting a 7 Series Device

Parameter	Result
Resources	33% to 75% less depending on specific resource type
Performance	36% better
Number of Timing End Points	40% less
Dynamic Power @220 MHz	40% less

In addition, the second coding example is more concise.

Reset Coding Example Three

Sometimes, a design might have an active low reset (for example, AXI standard dictates the reset to be active-Low). Since asynchronous resets have synchronizing circuits to ensure deassertion being timed, it is possible to make minor modifications to the synchronizing circuit, so that some of the LUT counts may be reduced. For an example schematic for synchronizing of asynchronous reset, see [Controlling and Synchronizing Device Startup](#).

Original HDL code:

```

always @ (posedge clk or negedge rst_n) //async. negedge reset
begin
    if (!rst_n)
        synchronizer_ckt <= 4'b0; // 4 stage reset syncornization
    else
        synchronizer_ckt <= {synchronizer_ckt[2:0], 1'b1};
end

assign synchronized_rst_n = synchronizer_ckt[3]; // the final reset signal which is
used to reset the actual flops in the design
    
```

There is a LUT in the reset path, as shown in the black circle in the following figure. Because reset signals feed many flops, saving on the delay for this LUT might impact many paths.

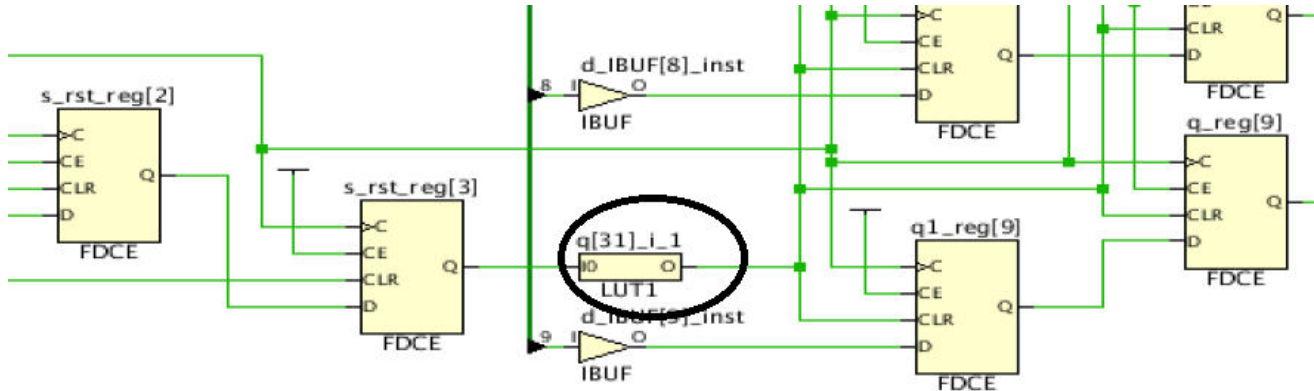


Figure 4-3: LUT on Reset Path

Modified HDL code:

```
always @ (posedge clk or negedge rst_n) //async. negedge reset
begin
  if (!rst_n)
    synchronizer_ckt <= 4'hf // 4 stage reset syncornization
  else
    synchronizer_ckt <= {synchronizer_ckt[2:0], 1'b0};
end

assign synchronized_rst_n = ~synchronizer_ckt[3]; // the final reset signal which is
used to reset the actual flops in the design
```

The synchronizer_ckt has been given an inverted logic, and another inversion has been added to the final synchronized_rst_n, in order to restore the polarity back. This slight modification to the synchronizing circuit can get rid of LUTs that exist between synchronizing circuit and the actual signal going into the flops of the design, as shown in the following figure.

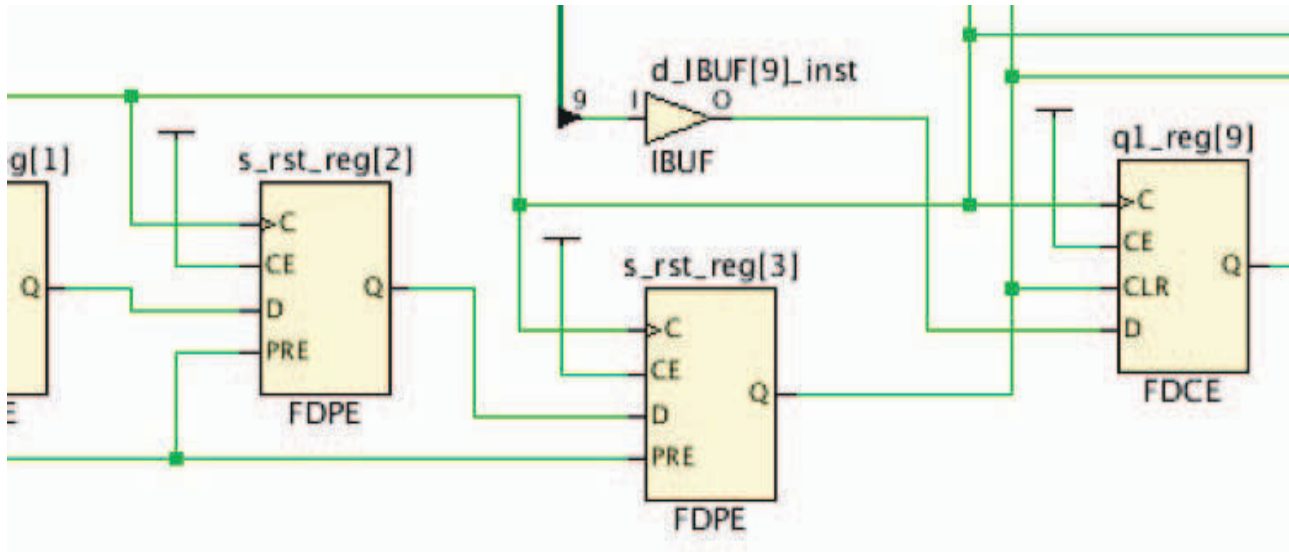


Figure 4-4: Modified Reset Circuit

Do Not Code Both a Set and a Reset in the Same Process or Always Block

Xilinx FPGA device registers have a built-in set or reset capability, but cannot natively do both at the same time. Accordingly, when using both a synchronous set and reset, an additional signal is added to the datapath. This might affect area and timing depending on placement, fanout, and timing. For this reason, Xilinx recommends coding both set and reset in the same sequential block only if absolutely required.

For an asynchronous set and reset, the effect on resource utilization and timing is more significant and should be avoided. Registers that contain both asynchronous reset and asynchronous set signals and/or with an asynchronous control signal with a dynamic value result in a circuit that cannot be timed correctly. For this reason, Xilinx requires that when an asynchronous set and reset condition is desired that at least one of set or reset condition be changed to a synchronous event.

Clock Enables

When used wisely, clock enables can significantly reduce design power with little impact on area or performance. However, when clock enables are used improperly, they can lead to:

- Increased area
- Decreased density
- Increased power
- Reduced performance

In many designs with a large number of control sets, low fanout clock enables might be the main contributor to the number of control sets.

Creating Clock Enables

Clock enables are created when an incomplete conditional statement is coded into a synchronous block. A clock enable is inferred to retain the last value when the prior conditions are not met. When this is the desired functionality, it is valid to code in this manner. However, in some cases when the prior conditional values are not met, the output is a don't care. In that case, Xilinx recommends closing off the conditional (that is, use an else clause), with a defined constant (that is, assign the signal to a one or a zero).

In most implementations, this does not result in added logic, and avoids the need for a clock enable. The exception to this rule is in the case of a large bus when inferring a clock enable in which the value is held can help in power reduction. The basic premise is that when small numbers of registers are inferred, a clock enable can be detrimental because it increases control set count. However, in larger groups, it can become more beneficial and is recommended.

Reset and Clock Enable Precedence

In Xilinx FPGA devices, all registers are built to have set/reset take precedence over clock enable, whether an asynchronous or synchronous set/reset is described. In order to obtain the most optimal result, Xilinx recommends that you always code the set/reset before the enable (if deemed necessary) in the if/else constructs within a synchronous block. Coding a clock enable first forces the reset into the data path if synchronous and if asynchronous, creates additional logic.

For information on clocking, see [Clocking Guidelines](#).

Tips for Control Signals

- Check whether a global reset is really needed.
- Avoid asynchronous control signals.
- Keep clock, enable, and reset polarities consistent.
- Do not code a set and reset into the same register element.
- If an asynchronous reset is absolutely needed, remember to synchronize its deassertion.

Know What You Infer

Your code finally has to map onto the resources present on the device. Make an effort to understand the key arithmetic, storage, and logic elements in the architecture you are targeting. Then, as you code the functionality of the design, anticipate the hardware resources to which the code will map. Understanding this mapping gives you an early insight into any potential problem.

The following examples demonstrate how understanding the hardware resources and mapping can help make certain design decisions:

- For larger than 4-bit addition, subtraction and add-sub, a carry chain is generally used and one LUT per 2-bit addition is used (that is, an 8-bit by 8-bit adder uses 8 LUTs and the associated carry chain). For ternary addition or in the case where the result of an adder is added to another value without the use of a register in between, one LUT per 3-bit addition is used (that is, an 8-bit by 8-bit by 8-bit addition also uses 8 LUTs and the associated carry chain).

If more than one addition is needed, it may be advantageous to specify registers after every two levels of addition to cut device utilization in half by allowing a ternary implementation to be generated.

- In general, multiplication is targeted to DSP blocks. Signed bit widths less than 18x25 (18x27 in UltraScale devices) map into a single DSP Block. Multiplication requiring larger products may map into more than one DSP block. DSP blocks have pipelining resources inside them.

Pipelining properly for logic inferred into the DSP block can greatly improve performance and power. When a multiplication is described, three levels of pipelining around it generates best setup, clock-to-out, and power characteristics. Extremely light pipelining (one-level or none) may lead to timing issues and increased power for those blocks, while the pipelining registers within the DSP lie unused.

- Two SRLs with 16 bits or less depth can be mapped into a single LUT, and single SRLs up to 32 bits can also be mapped into a single LUT.
- For conditional code resulting in standard MUX components:
 - A 4-to-1 MUX can be implemented into a single LUT, resulting in one logic level.
 - An 8-to-1 MUX can be implemented into two LUTs and a MUXF7 component, still resulting in effectively one logic (LUT) level.
 - A 16-to-1 MUX can be implemented into four LUTs and a combination of MUXF7 and MUXF8 resources, still resulting in effectively one logic (LUT) level.

A combination of LUTs, MUXF7, and MUXF8 within the same CLB/slice structure results in very small combinational delay. Hence, these combinations are considered as equivalent to only one logic level. Understanding this code can lead to better resource management, and can help in better appreciating and controlling logic levels for the data paths.

For general logic, the rule of thumb is to take into account the number of unique inputs for a given register. From that number, an estimation of LUTs and logic levels can be achieved. In general, six inputs or fewer always results in a single logic level. Theoretically, two levels of logic can manage up to thirty-six inputs. However, for all practical purposes, you should assume that approximately twenty inputs is the maximum that can be managed with two levels of logic. In general, the larger the number of inputs and the more complex the logic equation, the more LUTs and logic levels are required.



IMPORTANT: *Appreciating the hardware resources availability and how efficiently they are being utilized or wasted early in the design can allow for much easier modifications for better results than late in the design process during timing closure.*

Inferring RAM and ROM

RAM and ROM may be specified in multiple ways. Each has its advantages and disadvantages.

- Inference

Advantages:

- Highly portable
- Easy to read and understand
- Self-documenting
- Fast simulation

Disadvantages:

- May not have access to all RAM configurations available
- May produce less optimal results

Because inference usually gives good results, it is the recommended method, unless a given use is not supported, or it is not producing adequate results in performance, area, or power. In that case, explore other methods.

When inferring RAM, Xilinx highly recommends that you use the HDL Templates provided in the Vivado tools. As mentioned earlier, using asynchronous reset impacts RAM inference, and should be avoided. See [Using Vivado Design Suite HDL Templates](#).

- Direct Instantiation of RAM Primitives

Advantages:

- Highest level control over implementation
- Access to all capabilities of the block

Disadvantages:

- Less portable code
 - Wordier and more difficult to understand functionality and intent
- Use of a Core from IP Catalog

Advantages:

- Generally more optimized result when using multiple components
- Simple to specify and configure

Disadvantages:

- Less portable code
- Core management

Performance Considerations When Implementing RAM

In order to efficiently infer memory elements, consider these factors affecting performance:

- Using the Output Pipeline Register

Using an output register is required for high performance designs, and is recommended for all designs. This improves the clock to output timing of the block RAM. Additionally, a second output register is beneficial, as slice output registers have faster clock to out timing than a block RAM register. Having both registers has a total read latency of 3. When inferring these registers, they should be in the same level of hierarchy as the RAM array. This allows the tools to merge the block RAM output register into the primitive.



RECOMMENDED: *Determine early whether an extra clock cycle of latency during reads is tolerable. If it is, code in an extra stage of registers to the output of the memory array in order to use this dedicated resource to improve the overall timing of these paths.*

- Using Dedicated Blocks or Distributed RAMs

RAMs may be implemented in either: (1) the dedicated block RAM; or (2) within LUTs using distributed RAM. The choice not only impacts resource selection, but may also significantly impact performance and power.

In general, the required depth of the RAM is the first criterion. Memory arrays described up to 64-bits deep are generally implemented in LUTRAMs where depths 32-bits and less are mapped - two bits per LUT and depths up to 64-bits can be mapped one bit per LUT. Deeper RAMs may also be implemented in LUTRAM depending on available resources and synthesis tool assignment.

Memory arrays deeper than 256 are generally implemented in Block memory. Xilinx FPGA devices have the flexibility to map such structures in different width and depth combinations. You should be familiar with these configurations in order to understand the number and structure of block RAMs used for larger memory array declarations in the code.



IMPORTANT: *Slight deviations in coding styles for these blocks may result in sub-optimal utilization of resources. For example, an asynchronous read of memory infers LUTRAM instead of block RAM. However, adding reset of the memory causes this to be implemented in an array of registers rather than LUTRAM.*

Selecting the Proper Block RAM Write Mode

Xilinx block RAMs have the ability to change the write mode. This can impact functionality, behavior, and power. Xilinx recommends the following guidelines for selecting the best write mode for a particular operation:

- Consider Functionality First

When selecting a write mode, consider functionality first. When writing to a particular port of the block RAM, do you need the output read data to be a particular value? If you must see the prior value in the block RAM during write, select READ_FIRST. If you want to read the new data being written to the block RAM use WRITE_FIRST. If you do not care about the data read during writes, then the next selection criteria has to do with memory collisions.

- Use READ_FIRST Mode

If you are implementing a dual-port memory and connecting the same clock to the block RAM and cannot guarantee that a memory collision will not occur, select READ_FIRST. The READ_FIRST mode ensures that no memory collisions occur when the same clock is connected to both block RAM ports. For more information, see the *7 Series FPGA Devices Memory Resources User Guide (UG473)* [Ref 42].

- Use WRITE_FIRST Mode

When using 7 series block RAM in the wide SDP mode (RAMB36 in 72-bit wide or RAMB18 in 36-bit wide mode), where different clocks are on both ports or read/write collisions can be avoided, use WRITE_FIRST mode. In the wide SDP mode, this mode is the lowest power operation and the suggested mode.

- Use NO_CHANGE Mode

In all other cases, Xilinx recommends NO_CHANGE mode. NO_CHANGE has the best power characteristics. If read during write functionality (and if collisions are not a concern), then NO_CHANGE mode results in lower dissipated power in the block RAM and associated interconnect.

FIFO Creation

First-In, First-Out (FIFO) buffers are one of the most common uses for memory in FPGA designs.

Note: Asynchronous First-In-First-Out (FIFO) buffers are also known as *async FIFO* or *multi-rate FIFO*.

FIFO buffers are commonly used to transfer data from one clock domain to another. There are multiple methods and trade-offs to evaluate when creating and using FIFOs.

Selecting the Proper Entry Method and Resources for Your FIFO

Xilinx FPGA devices contain block RAM that possess dedicated FIFO circuitry. In general, Xilinx recommends using it in order to obtain the best area, power, performance and MTBF characteristics.

Using the hard FIFO also eases design by not requiring additional timing constraints or memory collision considerations. If, however, this circuit does not meet your needs, other soft implementations can be created resulting in an almost infinite amount of behaviors and characteristics.

If a soft FIFO is needed, it is better to be created from the IP Catalog. This eases implementation by not only creating the proper logic for most common FIFO implementations, but also creates the appropriate timing constraints and attributes for proper implementation and analysis. If ultimate customization is required, one can be inferred as well.

Design Challenges in Using a Soft Implementation for Asynchronous FIFO Buffer

If a soft FIFO is desired, here are some considerations to take into account. In order to determine the status of the FIFO and safely transfer the data, the design must monitor and react to status flags (empty and full signals).

Since these flags are based on two clock domains that do not have related phases or periods, the timing and predictability of the flags cannot always be readily determined. For this reason, you must take special precautions when using an asynchronous FIFO.

Flag assertion and de-assertion for most asynchronous FIFO implementations is not inherently cycle deterministic. A functional or timing simulation may show the status flag changing on one clock cycle, while on the FPGA device itself, the status flag may change in the previous or next clock cycle. This may occur when the timing and order of events in the simulator differs from the timing and order of events in the FPGA device.

The end timing of the FPGA device is determined by process, voltage, and temperature (PVT). It is therefore possible to have cycle differences on different chips, as well as under different environmental conditions on the same chip. You must be sure to take these differences into account when designing your circuits.

You may encounter problems if you expect data to be valid after or during a certain number of clock cycles, and you do not monitor the empty and full flags directly. In most FIFO implementations, even if there is memory space, reading from a FIFO that has its empty flag asserted, or writing to a FIFO that has its full flag asserted, gives an invalid read or write condition. This can lead to unexpected results, and can create a serious debugging problem. Xilinx strongly recommends that you always monitor the status flags, regardless of whether the asynchronous FIFO implementation passes simulation.

In most asynchronous FIFO implementations, empty and full flags default to a safe condition when a read and a write is performed at or near the same time at status flag boundaries. A full flag may assert even if the FIFO is not actually full. An empty flag may assert even if the FIFO is not actually empty. This provides a slight degree of safety, rather than taking the risk of flags not being asserted.

Various synthesis and simulation directives can allow the asynchronous FIFO to behave in a known manner when testing asynchronous conditions.

In many cases, a timing violation cannot be avoided when designing FIFO flag logic. If a timing violation occurs during timing simulation, the simulator produces an unknown (X) output to indicate the unknown state. For this reason, if logic is being driven from a known asynchronous source, and the proper design precautions were made to ensure proper operation regardless of the violation, Xilinx recommends adding the `ASYNC_REG=TRUE` attribute to the associated synchronizers in the FIFO flag logic. This indicates that the register can safely receive asynchronous input. Timing violations on the register no longer result in an X, but instead maintain its previous value. This also prevents the tool from replicating the register, or performing other optimizations that can have a negative affect on the register operation.

A memory collision can occur when a read occurs at the same time as a write to the same memory location. Avoid memory collisions when possible through effective use of flags (full and empty). Otherwise, the read data may be corrupted. If you have guarded your design well against reading corrupted data due to collisions, you can disable collision checking with the `SIM_COLLISION_CHECK` attribute on the RAM model. Doing so might accelerate the simulation model, but Xilinx recommends this method only if memory collisions are guaranteed not to occur.



TIPS:

- Use the HDL Templates within the Vivado tools.
- Determine whether Block memory or distributed memory is better suited for your memory function.
- Use output registers whenever possible.
- Avoid asynchronous resets around memory structures.
- Consider the best write mode depending on circuit requirements.
- For FIFO implementation, consider the dedicated hard FIFO first.

Coding for Proper DSP and Arithmetic Inference

The DSP blocks within the Xilinx FPGA devices can perform many different functions, including:

- Multiplication
- Addition and subtraction
- Comparators
- Counters
- General logic

The DSP blocks are highly pipelined blocks with multiple register stages allowing for high-speed operation while reducing the overall power footprint of the resource. Xilinx recommends that you fully pipeline the code intended to map into the DSP48, so that all pipeline stages are utilized. To allow the flexibility of use of this additional resource, a set condition cannot exist in the function for it to properly map to this resource.

DSP48 slice registers within Xilinx devices contain only resets, and not sets. Accordingly, unless necessary, do not code a set (value equals logic 1 upon an applied signal) around multipliers, adders, counters, or other logic that can be implemented within a DSP48 slice. Additionally, avoid asynchronous resets, since the DSP slice only supports synchronous reset operations. Code resulting in sets or asynchronous resets may produce sub-optimal results in terms of area, performance, or power.

Many DSP designs are well-suited for the Xilinx architecture. To obtain best use of the architecture, you must be familiar with the underlying features and capabilities so that design entry code can take advantage of these resources.

The DSP48 blocks use a signed arithmetic implementation. Xilinx recommends code using signed values in the HDL source to best match the resource capabilities and, in general, obtain the most efficient mapping. If unsigned bus values are used in the code, the synthesis tools may still be able to use this resource, but might not obtain the full bit precision of the component due to the unsigned-to-signed conversion.

The multiplier within the Xilinx 7 series DSP48E1 slice has an input bit precision of 18 bits by 25 bits signed data. Therefore, the bit precision for unsigned data is 17 bits by 24 bits. For

UltraScale devices, the multiplier within the DSP48E2 slice has an input bit precision of 18 bits by 27 bits signed data. Therefore, the bit precision for unsigned data is 17 bits by 26 bits. For Verilog code, data is considered unsigned unless otherwise declared in the code. If the target design is expected to contain a large number of adders, Xilinx recommends that you evaluate the design to make greater use of the DSP48 slice pre-adders and post-adders. For example, with FIR filters, the adder cascade can be used to build a systolic filter rather than using multiple successive add functions (adder trees). If the filter is symmetric, you can evaluate using the dedicated pre-adder to further consolidate the function into both fewer LUTs and flip-flops and also fewer DSP slices as well (in most cases, half the resources).

If adder trees are necessary, the 6-input LUT architecture can efficiently create ternary addition ($A + B + C = D$) using the same amount of resources as a simple 2-input addition. This can help save and conserve carry logic resources. In many cases, there is no need to use these techniques.

By knowing these capabilities, the proper trade-offs can be acknowledged up front and accounted for in the RTL code to allow for a smoother and more efficient implementation from the start. In most cases, DSP resources should be inferred.

For more information about the features and capabilities of the DSP48 slice, and how to best leverage this resource for your design needs, see the *7 Series DSP48E1 Slice User Guide* (UG479) [Ref 43] and *UltraScale Architecture DSP Slice User Guide* (UG579) [Ref 44].

Coding Shift Registers and Delay Lines

In general, a shift register is characterized by some or all of the following control and data signals:

- Clock
- Serial input
- Asynchronous set/reset
- Synchronous set/reset
- Synchronous/asynchronous parallel load
- Clock enable
- Serial or parallel output

Xilinx FPGA devices contain dedicated SRL16 and SRL32 resources (integrated in LUTs). These allow efficiently implemented shift registers without using flip-flop resources.

However, these elements support only LEFT shift operations, and have a limited number of I/O signals:

- Clock
- Clock Enable
- Serial Data In
- Serial Data Out

In addition, SRLs have address inputs (LUT A3, A2, A1, A0 inputs for SRL16) determining the length of the shift register. The shift register may be of a fixed static length, or it may be dynamically adjusted.

In dynamic mode each time a new address is applied to the address pins, the new bit position value is available on the Q output after the time delay to access the LUT. Synchronous and Asynchronous set/reset control signals are not available in the SRL primitives. However, if your RTL code includes a reset, the Xilinx synthesis tool infers additional logic around the SRL to provide the reset functionality.

To obtain the best performance when using SRLs, Xilinx recommends that you implement the last stage of the shift register in the dedicated Slice register. The Slice registers have a better clock-to-out time than SRLs. This allows some additional slack for the paths sourced by the shift register logic. Because synthesis tools often automatically infer this register for properly coded shift register inference code, it is not necessary to do additional work unless this resource is instantiated or the synthesis tool is prevented from inferring such a register.

Xilinx recommends that you use the HDL coding styles represented in the Vivado Design Suite HDL Templates.

When using registers to obtain placement flexibility in the chip, turn off SRL inference using the attribute:

```
SHREG_EXTRACT = "no"
```

For more information about synthesis attributes and how to specify those attributes in the HDL code. see *Vivado Design Suite User Guide: Synthesis* (UG901) [Ref 16].

Initialization of All Inferred Registers, SRLs, and Memories

The GSR net initializes all registers to the specified initial value in the HDL code. If no initial value is supplied, the synthesis tool is at liberty to assign the initial state to either zero or one. Vivado synthesis generally defaults to zero with a few exceptions such as one-hot state machines encoding.

Any inferred SRL, memory, or other synchronous element may also have an initial state defined that will be programmed into the associated element upon configuration.

Xilinx highly recommends that you initialize all synchronous elements accordingly. Initialization of registers is completely inferable by all major FPGA synthesis tools. This lessens the need to add a reset for the sole purpose of initialization, and makes the RTL code more closely match the implemented design in functional simulation, as all synchronous elements start with a known value in the FPGA device after configuration.

Initial state of the registers and latches VHDL coding example one:

```
signal reg1 : std_logic := '0'; -- specifying register1 to start as a zero
signal reg2 : std_logic := '1'; -- specifying register2 to start as a one
signal reg3 : std_logic_vector(3 downto 0):="1011"; -- specifying INIT value for
4-bit register
```

Initial state of the registers and latches Verilog coding example one

```
reg register1 = 1'b0; // specifying register1 to start as a zero
reg register2 = 1'b1; // specifying register2 to start as a one
reg [3:0] register3 = 4'b1011; //specifying INIT value for 4-bit register
```

Initial state of the registers and latches Verilog coding example two

Another possibility in Verilog is to use an initial statement:

```
reg [3:0] register3;
initial begin
    register3= 4'b1011;
end
```

To ensure that all the sequential elements come out of the reset at the same time, see [Controlling and Synchronizing Device Startup](#).

Deciding When to Instantiate or Infer

Xilinx recommends that you have an RTL description of your design; and that you let the synthesis tool do the mapping of the code into the resources available in the FPGA device. In addition to making the code more portable, all inferred logic is visible to the synthesis tool, allowing the tool to perform optimizations between functions. These optimizations include logic replications; restructuring and merging; and retiming to balance logic delay between registers.

Synthesis Tool Optimization

When device library cells are instantiated, synthesis tools do not optimize them by default. Even when instructed to optimize the device library cells, synthesis tools generally cannot perform the same level of optimization as with the RTL. Therefore, synthesis tools typically only perform optimizations on the paths to and from these cells but not through the cells.

For example, if an SRL is instantiated and is part of a long path, this path might become a bottleneck. The SRL has a longer clock-to-out delay than a regular register. To preserve the area reduction provided by the SRL while improving its clock-to-out performance, an SRL of

one delay less than the actual desired delay is created, with the last stage implemented in a regular flip-flop.

When Instantiation Is Desirable

Instantiation may be desirable when the synthesis tool mapping does not meet the timing, power, or area constraints; or when a particular feature within an FPGA device cannot be inferred.

With instantiation, you have total control over the synthesis tool. For example, to achieve better performance, you can implement a comparator using only LUTs, instead of the combination of LUT and carry chain elements usually chosen by the synthesis tool.

Sometimes instantiation may be the only way to make use of the complex resources available in the device. This can be due to:

- HDL Language Restrictions

For example, it is not possible to describe double data rate (DDR) outputs in VHDL because it requires two separate processes to drive the same signal.

- Hardware Complexity

It is easier to instantiate the I/O SerDes elements than to create synthesizable description.

- Synthesis Tools Inference Limitations

For example, synthesis tools currently do not have the capability to infer the hard FIFOs or the DSP48 symmetric rounding and saturation from RTL descriptions. Therefore, you must instantiate it.

If you decide to instantiate a Xilinx primitive, see the appropriate User Guide and Libraries Guide for the target architecture to fully understand the component functionality, configuration, and connectivity.

In case of both inference as well as instantiation, Xilinx recommends that you use the instantiation and language templates from the Vivado Design Suite language templates.



TIPS:

- *Infer functionality whenever possible.*
 - *When synthesized RTL code does not meet requirements, review the requirements before replacing the code with device library component instantiations.*
 - *Consider the Vivado Design Suite language templates when writing common Verilog and VHDL behavioral constructs or if necessary instantiating the desired primitives.*
-

Coding Styles for Higher Reliability

Some specific design situations require specific considerations in order to achieve higher reliability.

Clock Domain Crossings

Whenever a data or control signal transfers from one clock domain to another, you must understand the nature of the crossing. Clock crossings may be categorized into:

- **Synchronous crossings**

Synchronous crossings are crossings in which there is a known and predictable phase relationship between domains.

- **Asynchronous crossings**

Asynchronous crossings are crossings in which phase relationship may not be determined predictably.

Within synchronous crossings, there can be situations in which clock skew is very high. Situations with very high skew can make it much more difficult to meet timing.



IMPORTANT: *If skew is in a direction that helps to meet **setup**, that makes it difficult to meet **hold**. Conversely, if skew is in a direction that helps to meet **hold**, that makes it difficult to meet **setup**.*

Synchronous Domain Crossing

- Going from one BUFG network to another driven from the same MMCM, PLL, or device pin, where both BUFGs exist in the same half of the chip (top half or bottom half).
- For 7 series devices, going from a BUFH network to another BUFH network that is placed horizontally adjacent (assuming, these 2 BUFH themselves are driven by the same clock source).
- For 7 series devices, going from a BUFR to another similarly configured BUFR both driven by the same BUFMR.

Note: If the BUFRs are not in BYPASS mode, it must also be phase aligned by synchronizing the reset on all associated BUFRs.

- For 7 series devices, going to or from a BUFIO to or from a BUFR in the same clock region from the same clock source.

Synchronous Domain Crossing, But Potentially Very High Skew

- Going to or from a clock network using an MMCM or PLL to or from a network not using one, even if generated from the same source clock.
- Going to or from a BUFH to any other network, except the BUFH network horizontally adjacent to it.
- Going to or from a clock network that is not directly driven by a dedicated clocking resource (such as external clock pin, MMCM, or PLL).
- Going to or from a BUFG located in the top half of the device to one located in the bottom half of the device.

Asynchronous Domain Crossing

- Going from a clock network to another clock network that have no phase relationship.
- Going to or from domains generated from the same MMCM or PLL if the clock periods do not have a simple ratio and their edges never align within a small number of cycles. This is also referred to as clocks with no common period or unexpandable.
- Going to or from any GT TXCLKOUT or RXCLKOUT outputs on any Xilinx FPGA device to another domain. This also includes paths amongst TXCLKOUT and RXCLKOUT.

In short, synchronous domains are domains in which there is a known and predictable phase relationship between clocks. This generally occurs when the domains: (1) are derivatives of each other, or; (2) are provided from the same internal or external source. In these cases, timing may be analyzed and (with some considerations), be safely transferred from one domain to another.

Depending on the distance and nature of clocking resources traveled after the common node to reach source and destination, clock skew can be significant enough to not ignore. For small data paths such as register to register paths, the clock skew may be longer than the data delay that can result in a hold violation. If there are several logic levels, then the additional skew can make timing very difficult. Xilinx recommends that you closely monitor logic levels in such cross clocking and take into account the effects of too few or too many logic levels.

For asynchronous domain crossing, special steps must be taken to mitigate improper bus capture, metastability, and other occurrences that can affect the data integrity in such paths.

In general, there are two popular methods to allow data to cross asynchronous clock domains safely. If only a single bit is needed or if methods such as grey-coding are used to transfer more than one bit of related data, register synchronizers can be inserted to reduce the Mean Time Before Failure (MTBF) of the circuit. For multiple bits of data (that is, a bus), the generally recommended practice is to use an independent clock (asynchronous) FIFO to safely transfer data from one domain to another. Such a FIFO can be inferred if built from soft logic. However, if the use of the dedicated hard-FIFO is desired (or if pre-characterized

and predefined FIFO Logic makes the task easier), the FIFO can be directly instantiated. FIFO primitives or the FIFO Generator can be used to construct the FIFO.

Use the ASYNC_REG attribute in your HDL code to identify all synchronizing registers. By doing so, the Vivado Design Suite design tools can better understand and use special algorithms to improve synthesis, simulation, placement, and routing to improve MTBF, by reducing metastability occurrences.

ASYNC_REG example:

```

module synchronizer #(
    parameter SYNC_STAGES = 2
) (
    input ASYNC_IN,
    input CLK,
    output SYNC_OUT
);
(* ASYNC_REG = "TRUE" *) reg [SYNC_STAGES-1:0] sync_regs = {SYNC_STAGES{1'b1}};

always @(posedge CLK)
    sync_regs <= {sync_regs[SYNC_STAGES-2:0], ASYNC_IN};

assign SYNC_OUT = sync_regs[SYNC_STAGES-1];

endmodule

```



TIP: Consider running static checkers to identify clock domain crossings and confirm appropriate synchronization.

Controlling and Synchronizing Device Startup

Once the FPGA device completes configuration, a sequence of events occurs in which the device comes out of the configuration state and into general operation. In most configuration sequences, one of the last steps is the deassertion of the Global Set Reset (GSR), followed by the deassertion of the Global Enable (GWE) signal. When this happens, the design is in a known initial state, and is then released for operation.

If this release point is not synchronized to the given clock domain, or if the clock is operating at a faster time than the GWE can safely be released, portions of the design can go into an unknown state. For some designs, this is inconsequential. In other designs, this can cause the design to become unstable, or to misprocess the initial data set. If the design must come up in a known state, Xilinx recommends that you take action to control the start-up synchronization process. This can be done in several ways.

One method is to delay all design clocks until a period of time after GWE is asserted. To do this, Xilinx recommends:

- Use instantiated BUFGCE, BUFHCE, or BUFR components.
- Use the enables of those components to delay clocking a few clock cycles post-configuration.
- When using an MMCM, do so on the output clocks, but not on the feedback clock, by selecting the **Safe Clock Startup** option from the Clocking Wizard.

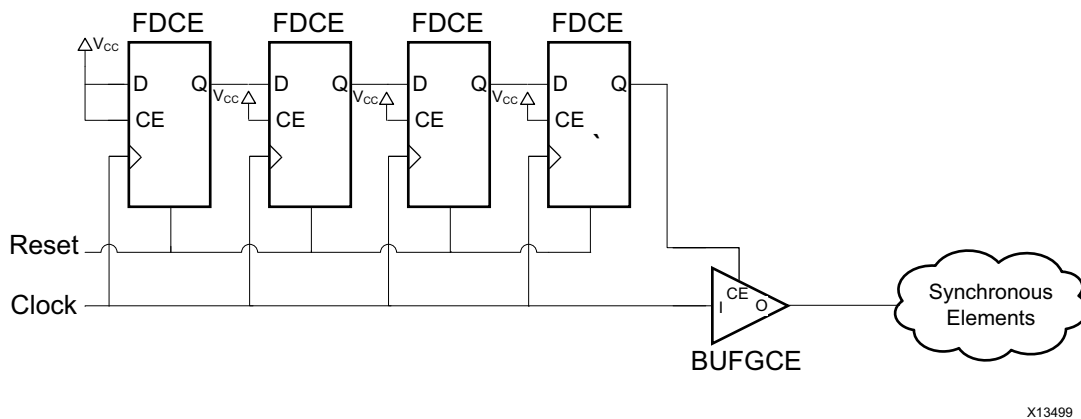


Figure 4-5: Clock Startup

You can also use clock enables, local reset (synchronized), or both on critical parts of the design (such as a state machine) to ensure that the startup of those portions of the designs are controlled and known.

Use of Untimed Resets

Resets (especially global resets) may have a very high fanout spanning a large portion of the FPGA array. In these cases, the reset timing may be difficult to meet regardless of the clock frequency or timing requirements. This can be especially challenging in high-speed designs.

Missing timing on a reset path can cause indeterminate behavior such as initial or intermittent data corruption; and full design lock up or catastrophic malfunction in extreme cases.

The issue for most designs is during deassertion, rather than the assertion of the reset (although in some cases assertion may pose an issue as well). If reset is deasserted at the same time the clock asserts at a particular register, the register output may be indeterminate (recovery/removal violation). Or if the skew of the reset signal is such that parts of the design may be released from reset at one clock cycle and other parts in another, it may cause unknown circuit behavior. [Controlling and Synchronizing Device Startup](#), shows an example method for synchronizing the deassertion of reset.

Xilinx recommends that besides adding a synchronizing circuit, you place proper timing constraints on the reset path; and, during design entry, minimize the impact for timing closure on the reset. One way to do so is to limit the number of loads a reset signal drives by either removing the reset or replicating the driver. Use of ASYNC_REG attribute on the flops used for synchronizing the resets will prevent the driver to be replicated. Hence, if you want replication (i.e. reset is driving a high fanout), put an additional flop at the end of the synchronizer, but exclude ASYNC_REG property on this last flop. This last flop can now be replicated and is not really a part of the synchronization chain.

If this is not adequate to close timing, you can stop the clock internally using the BUFGCE capability, or externally during reset allowing for a multi-cycle deassertion period for the reset signal.



IMPORTANT: *The reset deassertion must be timed, so that the entire design comes out of reset in the same cycle.*

UltraScale Device IDDRE1 Reset Considerations

The recovery time to the RST pin of the IDDRE1 in UltraScale devices has a tight timing relationship with the falling edge of the CLK pin. The tight timing need could potentially cause timing violations. If it becomes difficult to meet recovery timing, Xilinx has the following design recommendations:

Wait several cycles after the deassertion of the IDDRE1 RST and apply a `set_multicycle_path` constraint to the destination IDDRE1 RST pins. Examples include:

- Hold the input data on the D-pin stable and wait 3 clock cycles after the deassertion of RST before using the data presented on the Q-pins of the IDDRE1
- Ignore the data coming from the Q-pins of the IDDRE1 for 3 clock cycles after the deassertion of RST (e.g., by disabling the downstream registers from sampling the data)
- Since the IDDRE1 was under reset, hence, its outputs were at "0". Hold the D input of IDDRE1 also at "0" for 3 cycles.

If the design requires that the data from the IDDRE1 Q-pins must be sampled immediately after the deassertion of RST, Xilinx recommends to replicate the registers driving the RST pin of the IDDRE1. In some cases it might be required to have one register per IDDRE1. Position the registers driving the RST pin of the IDDRE1 as close as possible to the IDDRE1 by using LOC constraints or PBlocks, so that the RST signal can reach in minimum possible time.

Avoid Combinational Loops

Do not use combinational feedback paths in FPGA designs. The timing ramifications are difficult to simulate, analyze, and fully take into account under all operating conditions. Results can be unpredictable.

For information on using the `check_timing` command to check for any inadvertent combinational loops, see [Chapter 5, Implementation](#).

Coding Styles to Improve Performance

Violating the coding techniques discussed in the previous section ([Coding Styles for Higher Reliability](#)) generally has a detrimental impact on performance. For high performance designs, the coding techniques discussed in this section (Coding Styles to Improve Performance) can mitigate possible timing hazards.

High Fanouts in Critical Paths

High fanout nets are much easier to deal with early in the design process. What constitutes too high of a fanout is often dictated by performance requirements and the construction of the paths.



RECOMMENDED: *Examine nets with thousands of loads early to assess their impact on the overall design.*

If you identify a high fanout net, mitigation techniques include:

- [Reduce Loads to Portions of the Design That Do Not Require It](#)
- [Use Register Replication](#)

Reduce Loads to Portions of the Design That Do Not Require It

For high fanout control signals, evaluate whether all coded portions of the design require that net. Reducing the load demand can greatly ameliorate timing problems. For data paths, determine whether there is any restricting of the logic that might result in fanout reduction.

Use Register Replication

Register replication can increase the speed of critical paths by making copies of registers to reduce the fanout of a given signal. This gives the implementation tools more flexibility in placing and routing the different loads and associated logic. Synthesis tools use this technique extensively.

If high fanout nets with long route delays are reported in the timing report as critical paths, consider replication constraints on the synthesis tool; and manual replication of registers. Often, you must add an additional synthesis constraint to ensure that a manually duplicated register is not optimized away by the synthesis tool. Most synthesis tools use a fanout threshold limit to automatically determine whether or not to duplicate a register.

Adjusting this global threshold allows for automatic duplication of high fanout nets, but it does not offer a finer level of user control as to which specific registers can be duplicated. A better method is to apply attributes on specific registers or levels of hierarchy to specify

which registers can or cannot be replicated. If a LUT1 (rather than a register) is being used for replication, it indicates that an attribute or constraint is being applied incorrectly.

Do not replicate registers used for synchronizing signals that cross clock domains. The presence of `ASYNC_REG` attribute on these registers prevents the tool from replicating these registers. If the synchronizing chain has a very high fanout and there is a need for replication in order to meet the timing, the last flop might be replicated by removing the `ASYNC_REG` attribute on it. However, this register is then no longer a part of the synchronization chain.

The following table gives an indicative guideline on the number of fanouts that might be acceptable for your design.

Table 4-3: Fanout Guidelines

Condition	Fanout < 5000	Fanout < 200	Fanout < 100
Low Frequency 1 to 125 MHz	Few logic levels between synchronous logic <13 levels of logic at maximum frequency		
Medium Frequency 125 to 250 MHz	Results dependent. Might need to reduce fanout and/or logic levels to achieve.	<6 levels of logic at maximum frequency. (Driver and load types impact performance.)	
High Frequency > 250 MHz	Not recommended for most designs.	Small number of logic levels is typically necessary for higher speeds.	Advance pipelining methods required. Careful logic replication. Compact functions. Low logic levels required. (Driver and load types impact performance.)



TIP: If the timing reports indicate that high-fanout signals are limiting the design performance, consider replicating them. The `phys_opt_design` command might do a much better job of replicating registers. For more information, see [MAX_FANOUT in Chapter 5](#).



TIP: When replicating registers, consider naming them using a convention, such as `<original_name>_a`, `<original_name>_b`, etc. to make it easier to understand intent of the replication, thereby making future maintenance of the RTL code easier.

Pipelining Considerations

Another way to increase performance is to restructure long datapaths with several levels of logic and distribute them over multiple clock cycles. This method allows for a faster clock cycle and increased data throughput at the expense of latency and pipeline overhead logic management.

Because FPGA devices contain many registers, the additional registers and overhead logic are usually not an issue. However, the datapath spans multiple cycles, and you must make special considerations for the rest of the design to account for the added path latency.

Consider Pipelining for SSI Devices

When designing high performance register-to-register connections for SLR boundary crossings, the appropriate pipelining must be described in the HDL code and controlled at synthesis. This ensures that the Shift Register LUT (SRL) inference and other optimizations do not occur in the logic path that must cross an SLR boundary. Modifying the code in this manner along with appropriate use of Pblocks defines where the SLR boundary crossing occurs.

Consider Pipelining Up Front

Considering pipelining up front rather than later on can improve timing closure. Adding pipelining at a later stage to certain paths often propagates latency differences across the circuit. This can make one seemingly small change require a major redesign of portions of the code.

Identifying pipelining opportunities early in the design can often significantly improve timing closure, implementation runtime (due to easier-to-solve timing problems), and device power (due to reduced switching of logic).

Check Inferred Logic

As you code your design, be aware of the logic being inferred. Monitor the following conditions for additional pipelining considerations:

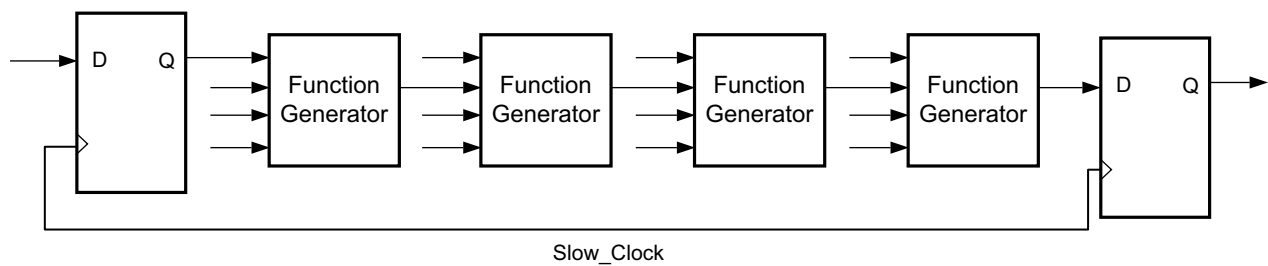
- Cones of logic with large fanin
For example, code that requires large buses or several combinational signals to compute an output
- Blocks with restricted placement or slow clock-to-out or large setup requirements
For example, block RAMs without output registers or arithmetic code that is not appropriately pipelined

- Forced placement that causes long routes

For example, a pinout that forces a route across the chip might require pipelining to allow for high-speed operation

In the following figure the clock speed is limited by:


- Clock-to out-time of the source flip-flop
- Logic delay through four levels of logic
- Routing associated with the four function generators
- Setup time of the destination register



X13429

Figure 4-6: Before Pipelining Diagram

The following figure is an example of the same data path shown in Figure 4-6. Because the flip-flop is contained in the same slice 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 after pipelining than before pipelining.



X13430

Figure 4-7: After Pipelining Diagram

Determine Whether Pipelining is Needed

A commonly used pipelining technique is to identify a large combinatorial logic path, break it into smaller paths, and introduce a register stage between these paths, ideally balancing each pipeline stage.

To determine whether a design requires pipelining, identify the frequency of the clocks and the amount of logic distributed across each of the clock groups. You can use the `report_design_analysis` Tcl command to determine the logic-level distribution for each of the clock groups.



TIP: *The design analysis report also highlights the number of paths with zero logic levels, which you can use to determine where to make modifications in your code.*

Balance Latency

To balance the latency by adding pipeline stages, add the stage to the control path not the data path. The data path includes wider buses, which increases the number of flip-flop and register resources used.

For example, if you have a 128-bit data path, 2 stages of registers, and a requirement of 5 cycles of latency, you can add 3 stages of registers: $3 \times 128 = 384$ flip-flops. Alternatively, you can use control logic to enable the data path. Use 5 stages of the single-bit pipeline, and enable the data path flip-flops.

Note: This example is only possible for certain designs. For example, in cases where there is a fanout from the intermediate data path flip-flops, having only 2 stages does not work.



RECOMMENDED: *Xilinx recommends keeping the flip-flop to LUT ratio below 1:5. A higher flip-flop to LUT ratio increases the amount of unrelated logic packing into slices, which increases routing complexity and can degrade QoR.*

Infer SRLs

To add a deeper pipeline stage, map as much as possible into the SRLs. This approach preserves the flip-flops/registers on the device. For example, a 9-deep pipeline stage (for a data width of 32) results in 9 flip-flops/registers for each bit, which uses $32 \times 9 = 288$ flip-flops. Mapping the same structure to SRLs uses 32 SRLs. Each SRL has address pins (A4A3A2A1A0) connected to 5'b01000.

There are multiple ways to infer SRLs during synthesis, including the following:

- SRL
- REG -> SRL
- SRL -> REG
- REG -> SRL -> REG

You can create these structures using the `srl_style` attribute in the RTL code as follows:

- `(* srl_style = "srl" *)`
- `(* srl_style = "reg_srl" *)`
- `(* srl_style = "srl_reg" *)`
- `(* srl_style = "reg_srl_reg" *)`

A common mistake is to use different enable/reset control signals in deeper pipeline stages. Following is an example of a reset used in a 9-deep pipeline stage with the reset connected to the third, fifth, and eighth pipeline stages:

```
FF->FF->FF(reset) -> FF->FF(reset)->FF->FF->FF(reset)->FF
```

To take advantage of SRL inference:

- Ensure there are no resets for the pipeline stages.
- Analyze whether the reset is really required.
- Use the reset on one flip-flop (for example, on the first and last stage of the pipeline).

Avoid Unnecessary Pipelining

For highly utilized designs, too much pipelining can lead to sub-optimal results. For example, unnecessary pipeline stages increase the number of flip-flops and routing resources, which might limit the place and route tool if the utilization is high.

Note: If there are many paths with 0/1 levels of logic, check to make sure this is intentional.

Consider Pipelining Dedicated Blocks

Based on the target architecture, dedicated primitives such as block RAMs and DSPs can work at over 500 MHz if enough pipelining is used. For high frequency designs, Xilinx recommends using all of the pipelines within these blocks.

Managing Wide Buses

The need for high throughput gives rise to wider bus functions at high frequencies. For example, a 200 Gb/s throughput data transfer needs a 1024-bit wide bus transferring data at 200 MHz.

Wider functions and wider memories are dominating the new era of FPGA designs. Xilinx provides various route resources on silicon and advanced placer and router algorithms to take care of these situations.

You should take into account the available resources and use them to achieve better performance. Design techniques that can be adopted to assist the flow include:

- Memory Organization
 - Register address and data-out of memories.
 - RTL coding should split the memories width-wise to achieve high performance.
- Wide Functions
 - Have sufficient pipeline stages when implementing wide arithmetic functions and reduction operators.
 - Ensure minimal SLR crossings when using SSI technology targets and performing manual design partition.
 - Use Xilinx IP to implement wide-bus complex arithmetic operations. The IP cores take care of re-timing and pipelining requirements for high-performance.
- I/Os
 - Use Xilinx serial IP cores that support wide-bus and high-throughput reliable chip-to-chip data transfer.
 - Allocate primary I/Os of the same bank or adjacent banks in order to minimize skew effect between individual bits of an in-coming or out-going bus interface.

Coding Styles to Improve Power

Coding styles to improve power include:

Gate Clock or Data Paths

Gating the clock or data paths is common technique to stop transition when the results of these paths are not used. Gating a clock stops all driven synchronous loads; and prevents data path signal switching and glitches from continuing to propagate.

The tools analyze the description and netlist to detect unwanted conditions. However, there are things you know about the application, data flow, and dependencies that are not available to the tool, and that only you can specify.

Maximize Gating Elements

Maximize the number of elements affected by the gating signal. For example, it is more power efficient to gate a clock domain at its driving source than to gate each load with a clock enable signal.

Use Clock Enable Ports of Dedicated Clock Buffers

When gating or multiplexing clocks to minimize activity or clock tree usage, use the clock enable ports of dedicated clock buffers. Inserting LUTs or using other methods to gate-off clock signals is not efficient for power and timing.

Keep an Eye on Control Sets

As discussed earlier, the number of control sets should be minimized. Xilinx recommends clock gating only if the gated clock drives a high number of synchronous elements. Otherwise, there is a risk of wasted flops. Adding gating signals to stop the data or clock path can require additional logic and routing (and, thus power). Minimize the number of additional structures to avoid defeating the original purpose.



RECOMMENDED: *Do not use too fine-grained clock gating. Each gated clock should impact a large number of synchronous elements.*

Use Case Block When Priority Encoder Not Needed

When a priority encoding is not needed, use a case block instead of an if-then-else block or ternary operator.

Inefficient coding example

```
if (reg1)
    val = reg_in1;
else if (reg2)
    val = reg_in2;
else if (reg3)
    val = reg_in3;
else val = reg_in4;
```

Correct coding example

```
(* parallel_case *) casex  ({reg1, reg2, reg3})
1xx: val = reg_in1 ;
01x: val = reg_in2 ;
001: val = reg_in3 ;
default: val = reg_in4 ;
endcase
```


Best Practices for Block RAM in Design

The amount of power the block RAM consumes is directly proportional to the amount of time that the block RAM is enabled. To save power, the block RAM enable can be driven Low on clock cycles when the block RAM is not used in the design. Both the block RAM enable rate and the clock rate are important parameters that must be considered for power optimization.

The mode settings for block RAM (such as NO_CHANGE, READ_FIRST, and WRITE_FIRST) were explained earlier in [Selecting the Proper Block RAM Write Mode](#).

Deeper and wider memories must follow the depth wise splitting mechanism to save dynamic power. During IP customization, if you choose power-efficient realization, Vivado IDE creates the depth wise splitting.

The following figure shows an example of the above statement with memory configuration 4Kx1024 bit split width wise.




X13432

Figure 4-8: RTL Representation of 4K x 1024 Using 4K x 9

In this implementation, all block RAMs are always enabled (for each read or write) and consume more power.

The following figure shows an example of depth wise splitting.



X13433

Figure 4-9: RTL Representation of 4K x 1024 Using 2K x 18

In this implementation, because one block RAM at a time is selected (from each unit), the dynamic power contribution is almost half. UltraScale device block RAMs have a dedicated cascade MUX and routing structure that allows the construction of wide, deep memories requiring more than one block RAM primitive to be built in a very power efficient configuration with a limited impact on performance. Both synthesis inference and the Block Memory Generator can automatically take advantage of this circuitry based on the specified RAM configuration.

The following figure shows a high-level, conceptual view of four cascaded block RAMs.

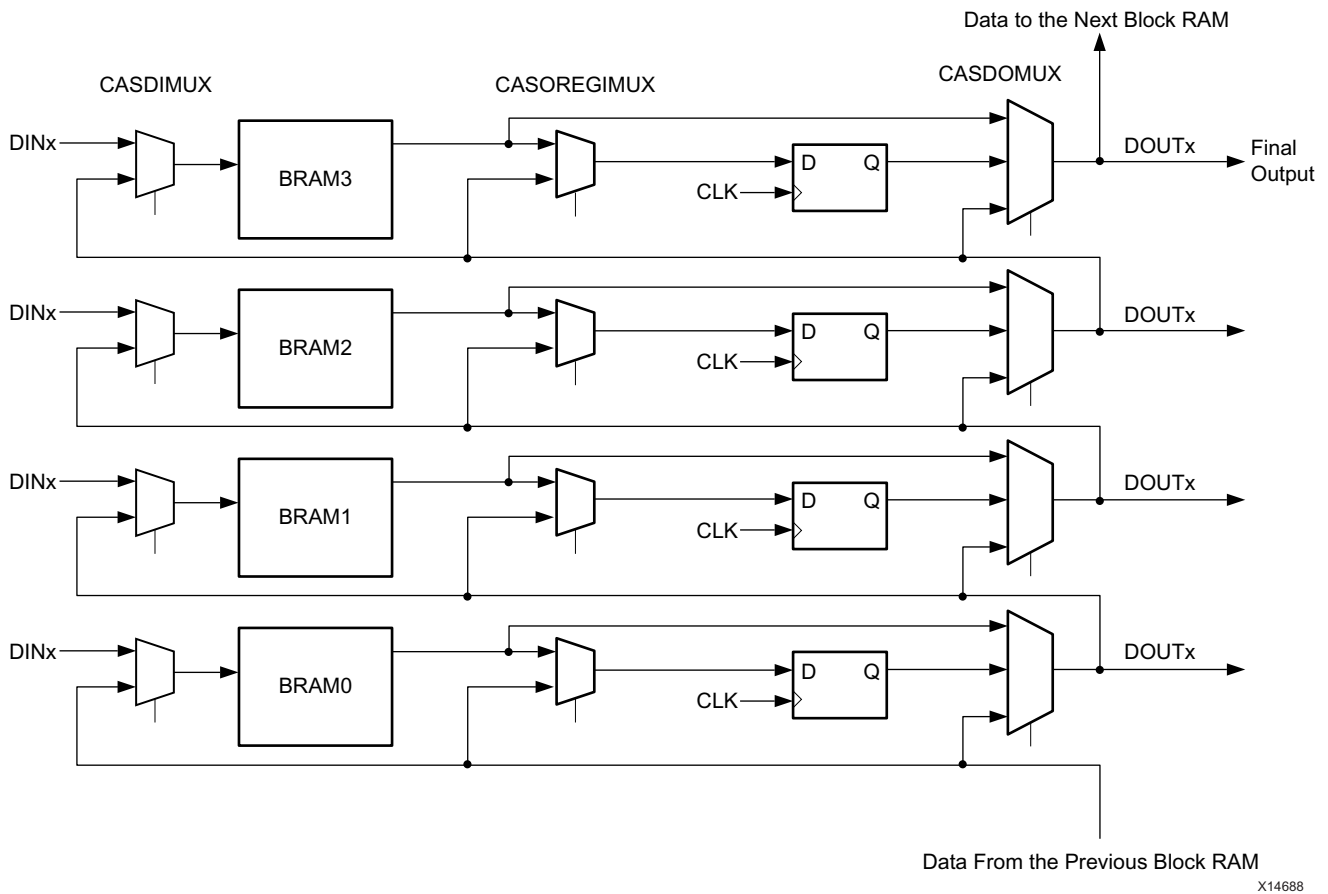


Figure 4-10: High-level View of the Block RAM Cascade Architecture

Parameters, Attributes, and Constraints

Depending on the context, terminologies might be used interchangeably among parameters, attributes, or constraints. This section explains these concepts. You should understand the concepts, so that even if another literature uses some other terminology, you are still able to appreciate the underlying message in that literature.

- [Parameters](#)
- [Constraints and Attributes](#)

Parameters

A parameter (generic in VHDL) is a property associated with a device architecture primitive component that affects an instantiated component's functionality or implementation. These properties are passed by means of generic maps (VHDL) or inline parameter passing (Verilog). These properties are called a generic or parameter in both VHDL and Verilog.

Note: Although defparams can also be used to modify parameters, Xilinx does not recommend that you use them for this purpose.

Examples of parameters include:

- The INIT property on a LUT6 component
- The DIVCLK_DIVIDE property on a MMCM

All parameters are described in the Xilinx *Libraries Guides* as a part of the primitive component description. You can use parameters to customize specific behavior of Xilinx primitives.

In the context of *primitives*, when explaining the properties that can be modified during instantiation, some literature uses the term *attributes* for what is defined here as *parameters* or *generics*.

VHDL Primitive Parameter (Generic) Coding Example:

The following VHDL coding example shows an example of setting the INIT generic for an instantiated RAM32X1S primitive that specifies the initial contents of this RAM symbol to the hexadecimal value of A1B2C3D4.

```
small_ram_inst : RAM32X1S
generic map (
  INIT => X"A1B2C3D4")
port map (
  O => ram_out, -- RAM output
  A0 => addr(0), -- RAM address[0] input
  A1 => addr(1), -- RAM address[1] input
  A2 => addr(2), -- RAM address[2] input
  A3 => addr(3), -- RAM address[3] input
  A4 => addr(4), -- RAM address[4] input
  D => data_in, -- RAM data input
  WCLK => clock, -- Write clock input
  WE => we -- Write enable input
);
```

Verilog Primitive Parameter Coding Example:

The following Verilog coding example shows an instantiated IBUFDS symbol in which DIFF_TERM and IOSTANDARD are specified as FALSE and LVDS_25 respectively.

```
IBUFDS #(
  .DIFF_TERM("FALSE"), // Differential Termination
  .IOSTANDARD("DEFAULT") // Specify the input I/O standard
) IBUFDS_inst (
  .O(O), // Buffer output
  .I(I), // Diff_p buffer input (connect directly to top-level port)
  .IB(IB) // Diff_n buffer input (connect directly to top-level port)
);
```

Constraints and Attributes

Constraints and attributes are often used interchangeably. Strictly speaking, *attributes* are directives that are provided in the HDL code itself, while *constraints* are provided in a constraints file (XDC). Both attributes and constraints provide guidance to specific tools on how to interpret and implement certain signals or instances.

Several properties can be provided as an *attribute* in the HDL or as a *constraint* in the XDC. For this reason, the specific property is considered both an attribute as well as a constraint. Accordingly, in the context of those properties, attributes and constraints are used interchangeably.

Constraints fall into three categories:

- Synthesis Constraints
- Timing Constraints
- Physical Constraints

Synthesis Constraints and Attributes

Synthesis constraints direct the synthesis tool's optimization techniques for a particular design or piece of HDL code. They are either embedded (also known as attribute) in the VHDL or Verilog code, or in a separate synthesis constraints file. Examples of synthesis attributes include USE_DSP48 and RAM_STYLE.

Xilinx recommends the following:

- Embed directives that impact functionality as an attribute in the HDL code. The HDL code always accompanies the associated attributes.
- Put temporary constraints (such as those required for debugging) in a separate constraints file. These constraints can then be easily dropped or added without modifying the actual HDL code.

Synthesis attributes, constraints, and directives are often embedded in the code or synthesis constraints file in an earlier implementation or architecture. Xilinx recommends that you comment out or remove these elements. They might lead to an inferior result, and not be the best choice in future implementations.



TIP: Remove any LOC, RLOC, or BEL constraints, or other physical constraints, embedded in the code or netlist of an existing design before retargeting to a new design or device.

An optimal placement for an older architecture is likely not optimal for new design or architecture. In some cases, certain constraints (for example related to location) may not even be valid for the new architecture.

The following examples illustrates passing attributes through HDL code only.

Attribute Declaration Example

```
attribute attribute_name : attribute_type;
```

Attribute Use on a Port or Signal Example

```
attribute attribute_name of object_name : signal is attribute_value
```

See the following example:

```
library IEEE;
use IEEE.std_logic_1164.all;
entity d_reg is
port (
CLK, DATA: in STD_LOGIC;
Q: out STD_LOGIC
);
attribute KEEP_HIERARCHY : string;
attribute KEEP_HIERARCHY of d_reg : entity is "true";
end d_reg;
```

Attribute Use on an Instance Example

```
attribute attribute_name of object_name : label is attribute_value
```

See the following example:

```
architecture struct of spblkrams is
attribute LOC: string;
attribute LOC of SDRAM_CLK_IBUFG: label is "AA27";
begin
-- IBUFG: Single-ended global clock input buffer
-- All FPGA
-- Xilinx HDL Language Template
SDRAM_CLK_IBUFG : IBUFG
generic map (
IOSTANDARD => "DEFAULT")
port map (
O => SDRAM_CLK_o, -- Clock buffer output
I => SDRAM_CLK_i -- Clock buffer input
);
-- End of IBUFG_inst instantiation
```

Attribute Use on a Component Example

```
attribute attribute_name of object_name : component
is attribute_value
```

See the following example:

```
architecture xilinx of tenths_ex is
attribute black_box : boolean;
component tenths
port (
CLOCK : in STD_LOGIC;
CLK_EN : in STD_LOGIC;
Q_OUT : out STD_LOGIC_VECTOR(9 downto 0)
);
end component;
attribute black_box of tenths : component is true;
begin
```

Historically, Verilog did not have a concept similar to VHDL *attribute*. For this reason, most tools had their own *pragmas*, for Verilog. Verilog 2001 provides a uniform syntax for passing VHDL-like attributes. Since the attribute is declared immediately before the object is declared, the object name is not mentioned during the attribute declaration.

```
(* (attribute_name = "attribute_value" *)
Verilog_object;
```

See the following example:

```
(* (RLOC = "R1C0.S0" *) FDCE #(
.INIT(1'b0) // Initial value of register (1'b0 or 1'b1)
) U2 (
.Q(q1), // Data output
.C(clk), // Clock input
.CE(ce), // Clock enable input
.CLR(rst), // Asynchronous clear input
.D(q0) // Data input
);
```

Running RTL DRCs

A set of RTL DRC rules identify potential coding issues with your HDL. These DRC checks may be made through **Elaborated Design > Report DRC** in the Flow Navigator or by executing `report_drc -ruledck methodology_checks` at the Tcl command prompt. You can perform these checks on the elaborated views, which you can open by clicking **Open Elaborated Design** in the Flow Navigator.

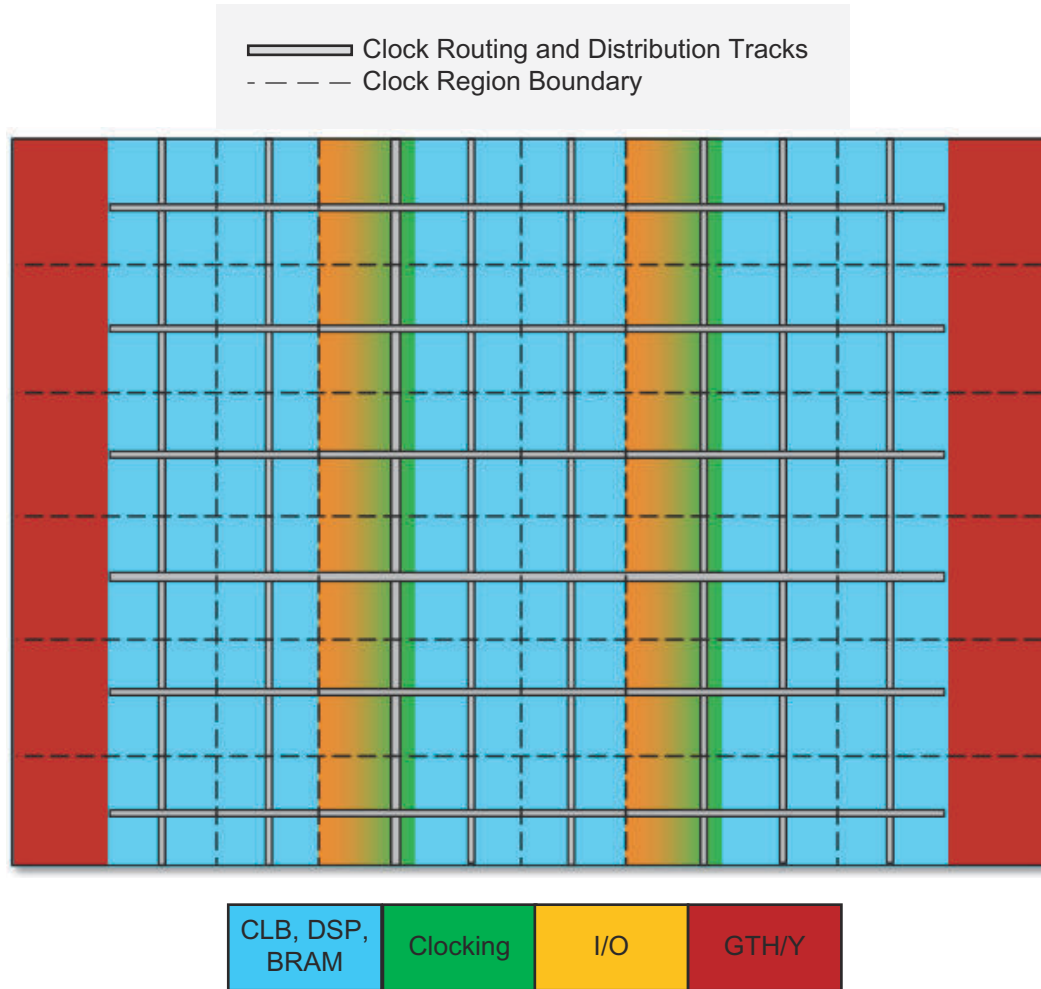
Clocking Guidelines

Each FPGA architecture has some dedicated resources for clocking. Understanding the clocking resources for your FPGA architecture can allow you to plan your clocking to best utilize those resources. Most designs might not need you to be aware of these details. However, if you can control the placement and have a good idea of the fanout on each of the clocking domains, you can explore alternatives based on the following clocking details. If you decide to exploit any of these clocking resources, you need to explicitly instantiate the corresponding clocking element.

UltraScale Device Clocking

UltraScale devices have a different clocking structure from previous device architectures, which blurs the line between global versus regional clocking. UltraScale devices do not have regional clock buffers like 7 series devices and instead use a common buffer and clock routing structure whether the loads are local/regional or global.

UltraScale devices feature smaller clock regions of a fixed size across devices, and the clock regions no longer span half of the device width in the horizontal direction. The number of clock regions per row varies per UltraScale device. Each clock region contains a clock network routing that is divided into 24 vertical and horizontal routing tracks and 24 vertical and horizontal distribution tracks. The following figure shows a device with 36 clock regions (6 columns x 6 rows). The equivalent 7 Series device has 12 clock regions (2 columns x 6 rows).



X15241-110415

Figure 4-11: UltraScale Device Clock Region Tiles

The clocking architecture is designed so that only the clock resources necessary to connect clock buffers and loads for a given placement are used, and no resource is wasted in clock regions with no loads. The efficient clock resource utilization enables support for more design clocks in the architecture while improving clock characteristics for performance and power. Following are the main categories of clock types and associated clock structures grouped by their driver and use:

- High-Speed I/O Clocks

These clocks are associated with the high-speed SelectIO™ interface bit slice logic, generated by the PLL, and routed via dedicated, low-jitter resources to the bit slice logic for high-speed I/O interfaces. In general, this clocking structure is created and controlled by Xilinx IP, such as memory IP or the High Speed SelectIO Wizard, and is not user specified.

- General Clocks

These clocks are used in most clock tree structures and can be sourced by a GCIO package pin, an MMCM/PLL, or fabric logic cells (not generally suggested). The general clocking network must be driven by BUFGCE/BUFGCE_DIV/BUFGCTRL buffers, which are available in any clock region that contains an I/O column. Any given clock region can support up to 24 unique clocks, and most UltraScale devices can support over 100 clock trees depending on their topology, fanout, and load placement.

- Gigabit Transceiver (GT) Clocks

Transmit, receive, and reference clocks of gigabit transceivers (GTH or GTY) use dedicated clocking in the clock regions that include the GTs. You can use GT clocks to achieve the following:

- Drive the general clocking network using the BUFG_GT buffers to connect any loads in the fabric
- Share clocks across several transceivers in the same or different Quad

Clock Primitives

Most clocks enter the device through a global clock-capable I/O (GCIO) pin. These clocks directly drive the clock network via a clock buffer or are transformed by a PLL or MMCM located in the clock management tile (CMT) adjacent to the I/O column.

The CMT contains the following clocking resources:

- Clock generation blocks
 - 2 PLLs
 - 1 MMCM
- Global clock buffers
 - 24 BUFGCEs
 - 8 BUFGCTRLs
 - 4 BUFGCE_DIVs

Note: Clocking resources in CMTs that are adjacent to I/O columns with unbonded I/Os are available for use.

The GT user clocks drive the global clock network via BUFG_GT buffers. There are 24 BUFG_GT buffers per clock region adjacent to the GTH/GTY columns.

Following is summary information for each of the UltraScale device clock buffers:

- BUFGCE

The most commonly used buffer is the BUFGCE. This is a general clock buffer with a clock enable/disable feature equivalent to the 7 series BUFHCE.

- BUFGCE_DIV

The BUFGCE_DIV is useful when a simple division of the clock is required. It is considered easier to use and more power efficient than using an MMCM or PLL for simple clock division. When used properly, it can also show less skew between clock domains as compared to an MMCM or PLL when crossing clock domains. The BUFGCE_DIV is often used as replacement for the BUFR function in 7 series devices. However, because the BUFGCE_DIV can drive the global clock network, it is considered more capable than the BUFR component.

- BUFGCTRL (also BUFGMUX)

The BUFGCTRL can be instantiated as a BUFGMUX and is generally used when multiplexing two or more clock sources to a single clock network. As with the BUFGCE and BUFGCE_DIV, it can drive the clock network for either regional or global clocking.

- BUFG_GT

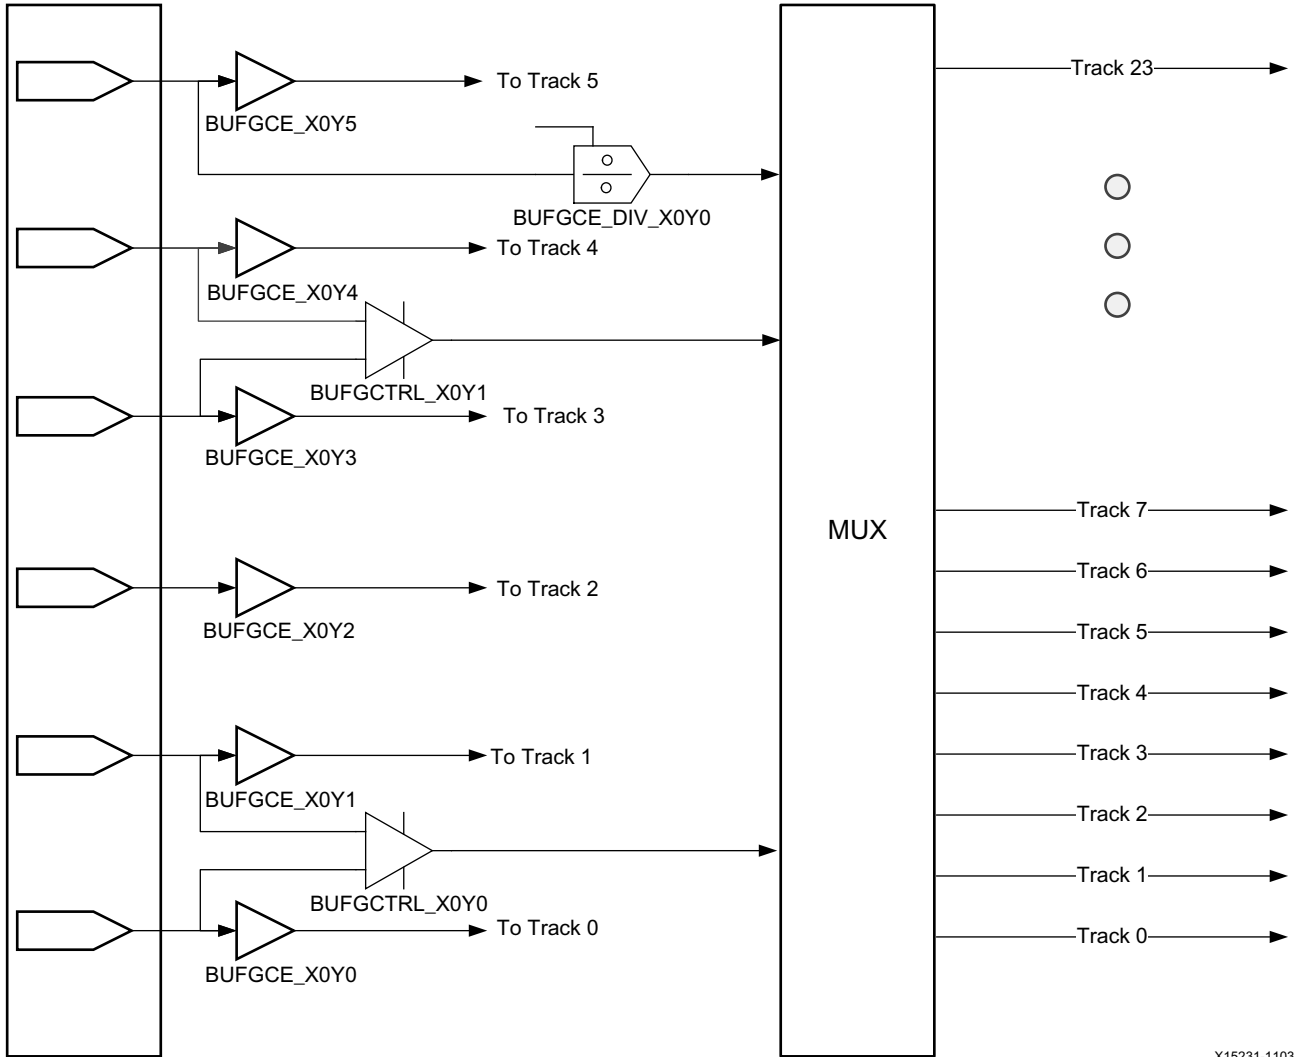
When using clocks generated by GTs, the BUFG_GT clock buffer allows connectivity to the global clock network. In most cases, the BUFG_GT is used as a regional buffer with its loads placed in one or two adjacent clock regions. The BUFG_GT has built-in dynamic clock division capability that you can use in place of an MMCM for clock rate changes.

For more information on the BUFGCE, BUFGCE_DIV, and BUFGCTRL buffers, see the *UltraScale Architecture Clocking Resources User Guide* (UG572) [Ref 39]. For details on connectivity and use of the BUFG_GT buffer, see the appropriate *UltraScale Architecture Transceiver User Guide* [Ref 40].

Global Clock Buffer Connectivity and Routing Tracks

Each of the 24 BUFGCE buffers in a clock region can only drive a specific clock routing track. However, the BUFGCTRL and BUFGCE_DIV outputs can use any of the 24 tracks by going through a MUX structure. Each BUFGCE_DIV shares the input connectivity with a specific BUFGCE site, and each BUFGCTRL shares input connectivity with two specific BUFGCE sites. Consequently, when BUFGCE_DIV or BUFGCTRL buffers are used in the clock region, use of the BUFGCE buffers is limited. The following figure shows the bottom 6 BUFGCE in a clock region, which are replicated 4 times within a clock region.

Note: A global clock net is assigned to a specific track ID in the device for all the vertical, horizontal routing, and distribution resources the clock uses. A clock *cannot* change track IDs unless the clock goes through another clock buffer.



X15231-110315

Figure 4-12: BUFCE, BUFCE_DIV, and BUFCTRL Shared Inputs and Output Muxing

Clock Routing, Root, and Distribution


To properly understand the clocking capacity of an UltraScale device and the clocking utilization of a design, it is important to know how the clock routes use the dedicated routing resources:

- From the clock buffer to the clock root, the clock signal goes through one or several segments of vertical and horizontal routing. Each segment must use the same track ID (between 0 and 23).
- At the clock root, the clock signal transitions from the routing track to the distribution track with the same track ID. To get the most favorable skew, the clock root is usually located in the middle of a clock region, which is located in the center of the clock window. The clock window is the rectangular area that includes all the clock regions where the clock net loads are placed. For skew optimization reasons, the Vivado IDE might move the clock root to off center.
- From the clock root to the CLB columns where the loads are located, the clock signal travels on the vertical distribution (both up and down the device as needed) and then onto the horizontal distribution (both to the left and right as needed).
- The CLB columns are split into two halves, which are located above and below the horizontal distribution resources. Each half of the CLB column contains several leaf clock routing resources that can be reached by any of the horizontal distribution tracks.

In some cases, a clock buffer can directly drive onto the clock distribution track. This usually happens when the clock root is located in the same clock region as the clock buffer or when the clock buffer only drives non-clock pins (for example, high fanout nets).

Because clock routing resources are segmented, only the routing and distribution segments used to traverse a clock region or to reach a load in a clock region are consumed.

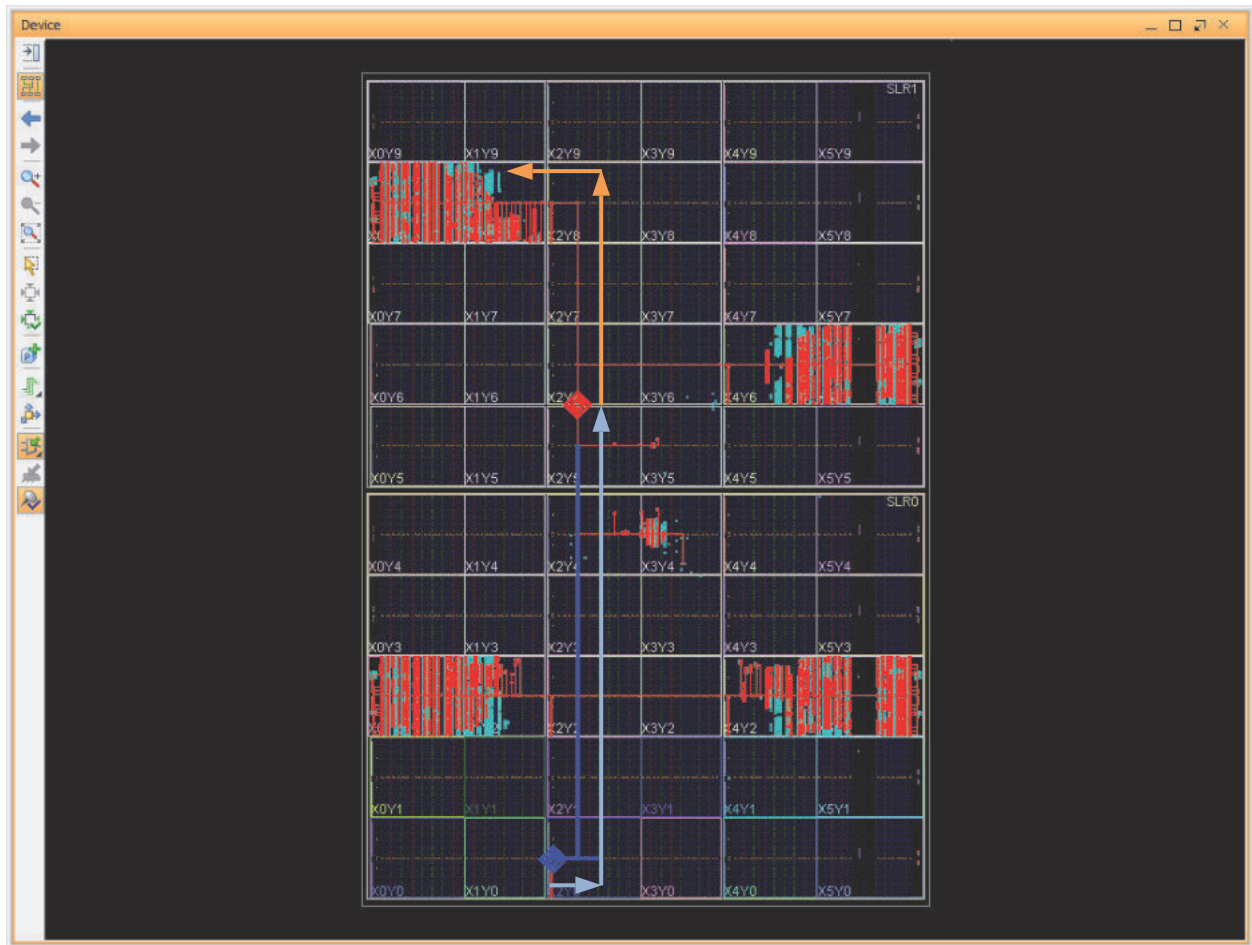
The following figure shows how a clock buffer located in clock region X2Y1 reaches its loads placed inside the clock window, which is formed by a rectangle of clock regions from X1Y3 to X5Y5.



X15389-111015

Figure 4-13: UltraScale Device Clock Routing from Driver to Loads

In the following figure, a routed device view shows an example of a global clock that spans most of the device. The clock buffer driving the network is highlighted in blue in clock region X2Y0 and drives onto the horizontal routing in that clock region. The net then transitions from the horizontal routing onto the vertical routing in clock region X2Y0 reaching the clock root in clock region X2Y5. All clock routing is highlighted in blue. The clock root is highlighted in red in the clock region X2Y5. From the clock root in X2Y5, the net transitions onto the vertical distribution and then the horizontal distribution to the clock leaf pins. The distribution layer and the leaf clock routing resources in the CLB columns are highlighted in red.



X15233-110315

Figure 4-14: Routed Device View of a Routed Clock Network

Clock Tree Placement and Routing

During the following phases, the Vivado placer determines the placement of MMCM/PLLs, global clock buffers, and the clock root while honoring the physical XDC constraints:

1. I/O and clock placement

The placer places I/O buffers and MMCM/PLLs based on connectivity rules and user constraints. The placer assigns clock buffers to clock regions but not to individual sites unless constrained using the LOC property. For details, see [Table 4-4](#). Only the clock buffers that only drive non-clock loads can move to a different clock region later in the flow based on the placement of their driver and loads.

Any placer error at this phase is due to conflicting connectivity rules, user constraints, or both. The log file shows extensive information about the possible root cause of the error, which you must review in detail to make the appropriate design or constraint change.

2. SLR partitioning (SSI technology devices only) and global placement

The placer performs the initial clock tree implementation based on early driver and load placements. Each clock net is associated with a clock window. The excessive overlap of clock windows can lead to placer errors due to anticipated clock routing contention.

When a clock partitioning error occurs, the log file shows the last clock budgeting solution for each clock net as well as the number of unique clock nets present in each clock region. Review the log file in detail to determine which clocks to remove from the overutilized clock regions. You can remove clocks using the following methods:

- Reduce the number of clocks in the design by combining identical synchronous clocks, removing unnecessary MMCM feedback clocks, or consolidating lower fanout clocks with high fanout clocks.
- Move clock primitives to different clock regions, especially those without connectivity-based placement rules.
- Add floorplanning constraints on clock loads to keep clocks with smaller fanout closer to their driver or away from the highly utilized clock regions.

The placer refines the clock tree implementation several times to help improve timing QoR. For example, during the later placement optimization phases, the placer analyzes each challenging clock to determine a better clock root location.

3. Clock tree pre-routing

The placer guides the subsequent implementation steps and provides accurate delay estimates for post-place timing analysis.

After placement, the Vivado tools can modify the clock tree implementation as follows:

- The Vivado physical optimizer can replicate and move cells to clock regions without associated clocks.
- The Vivado router can make adjustments to improve timing QoR and legalize the clock routing. The Vivado router can also modify the clock root location to improve timing QoR when you use the `Explore` routing directive.

The following table summarizes the placement rules for the main clock topologies and how constraints affect these rules.

Table 4-4: Topologies with and without Placement Rules

Constrained Source	Unconstrained Destination	Behavior
GCIO	BUFGCE, BUFGCTRL, BUFGCE_DIV, PLL/MMCM	Automatically placed in same clock region.
PLL/MMCM	BUFGCE, BUFGCTRL, BUFGCE_DIV	Automatically placed in same clock region.
GT*_CHANNEL	BUFG_GT	Automatically placed in same clock region.
BUFGCTRL	BUFGCTRL	Automatically placed in same clock region. Note: You can override placement within same clock region using the <code>CLOCK_REGION</code> constraint.
BUFG*	BUFG*	Unpredictable placement of unconstrained destination BUFG. Recommend constraining destination BUFG* using the <code>CLOCK_REGION</code> constraint. Note: This excludes BUFGCTRL -> BUFGCTRL.
BUFG*	MMCM/PLL	Unpredictable placement of unconstrained destination MMCM/PLL. Recommend constraining MMCM/PLL using a <code>LOC</code> constraint. Recommend <code>CLOCK_DEDICATED_ROUTE</code> constraint when the route spans adjacent or multiple clock regions.

Clocking Capability

Clock planning must be based on the total number of high fanout clocks and low fanout clocks in the target device.

High Fanout Clocks

A high fanout clock spans almost an entire SLR of an SSI technology device or almost all clock regions of a monolithic device. The following figure shows a high fanout clock that spans almost an entire SLR with the BUFGCE driver shown in red.

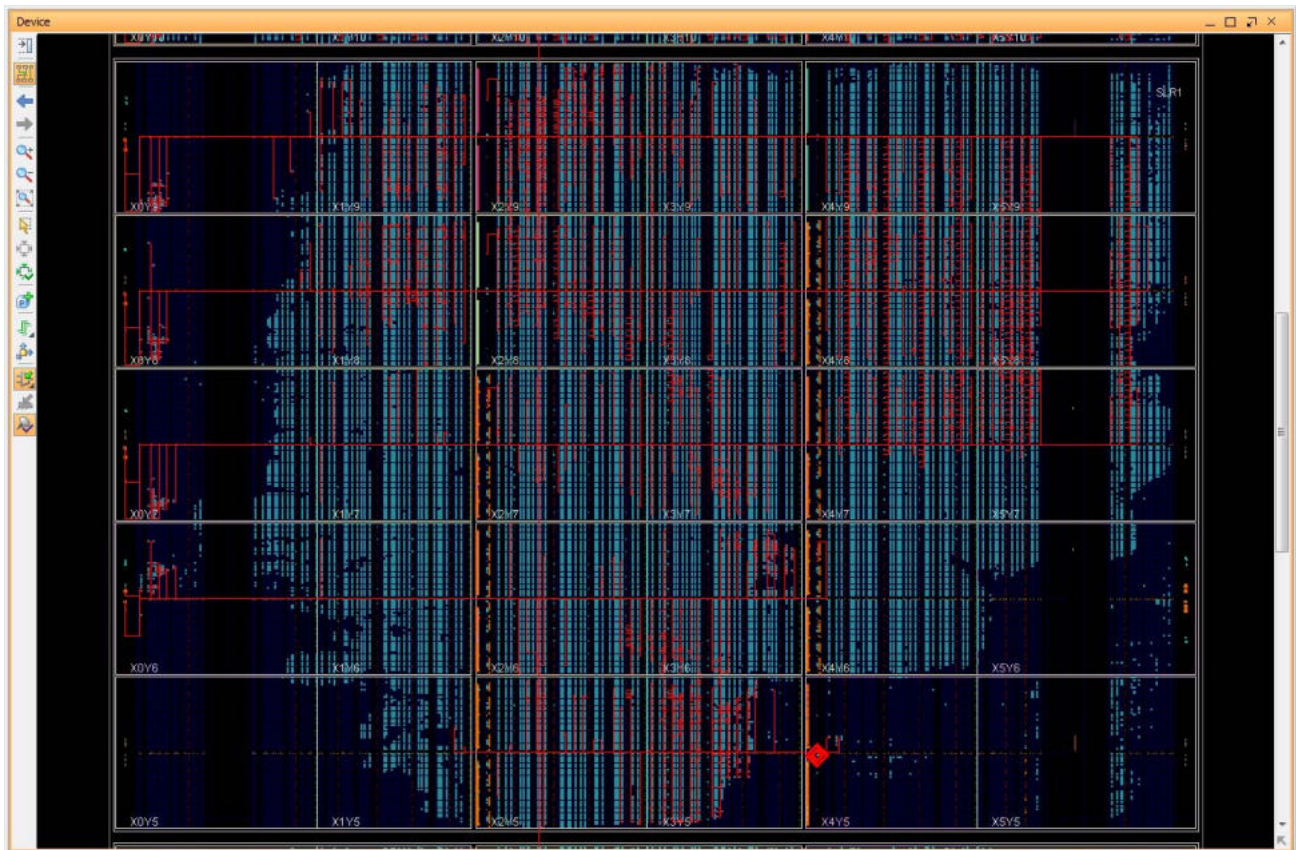


Figure 4-15: High Fanout Clock Spanning an SLR

Note: Using more than 24 clocks in a design might cause issues that require special design considerations or other up-front planning.



IMPORTANT: In *ZHOLD* and *BUF_IN* compensation modes, the MMCM feedback clock path matches the *CLKOUT0* clock path in terms of routing track, clock root location, and distribution tracks. Therefore, the feedback clock can be considered a high fanout clock when the clock buffer and clock root are far apart. For more information on MMCM compensation mechanism, see [I/O Timing with MMCM ZHOLD/BUF_IN Compensation](#).

Low Fanout Clocks

In most cases, a low fanout clock is a clock net that is connected to less than 5,000 clock pins, which are placed in 3 or fewer horizontally adjacent clock regions. The clock routing, clock root, and clock distribution are all contained within the localized area. In some cases, the placer is expected to identify a low fanout clock but fails. This can be caused by design size, device size, or physical XDC constraints, such as a LOC constraint or Pblock, which prevent the placer from placing the loads in a local area. To address this issue, you might need to guide the tool by manually creating a Pblock or modifying the existing physical constraints.

Clocks driven by BUFG_GT are an example of a low fanout clock. The Vivado placer automatically identifies these clock nets and contains the loads to the clock regions adjacent to the GT interface. The following figure shows a low fanout clock contained in two clock regions with the BUFG_GT driver shown in red.




Figure 4-16: Low Fanout Clock Contained in Two Clock Regions

Balanced Utilization of High and Low Fanout Clocks

UltraScale devices support more clocks than previous Xilinx FPGA families. This enables a wide range of clocking utilization scenarios, such as the following:

- 24 clocks or less

Unless conflicting user constraints exist, all clocks can be treated as high fanout clocks without risking placement or routing contention.

- Almost 300 clocks

For a design that targets a device with 6 clock region rows and includes only low fanout clocks with each clock included in 3 clock regions at most, the following clocks are required: 6 rows x 2 clock windows per row x 24 clocks per region = 288 clocks.

Low fanout clock windows do not have a fixed size but are usually between 1 and 3 clock regions. High fanout clocks rarely span the entire device or an entire SLR.

The following method shows how to balance high fanout clocks and low fanout clocks, assuming that a few low fanout clocks come from I/O interfaces and most from GT interfaces. You can apply the same method for each SSI technology device SLR.

- High fanout clocks
 - Up to 12 for monolithic devices
 - Up to 24 for SSI technology devices (assuming some high fanout clocks are only present in 1 SLR)
- Low fanout clocks
 - Up to 12 plus 8 per GT utilized Quad
 - Alternatively, up to 12 plus 6 per GT interface (group of GT channels that share the RXUSRCLK and TXUSRCLK)

Clock Constraints

Physical XDC constraints drive the implementation of clock trees and control the use of high fanout clocking resources. Because UltraScale device clocking is more flexible than clocking with previous architectures and includes additional architectural constraints, it is important to understand how to properly constrain your clocks for implementation.

Using LOC Constraints for IO/MMCM/PLL/GT

To constrain clocks, you can use a LOC constraint as follows:

- On a clock input at the I/O port

Assigning a PACKAGE_PIN constraint for a clock on a GCIO or an IOB LOC affects the clock network. The MMCM/PLL and clock buffers directly connected to the input port must be placed in the same clock region.

- On an MMCM or PLL

The clock buffers directly connected to the MMCM or PLL outputs and the input clock ports connected to the MMCM or PLL inputs are automatically placed in the same clock region. If an input clock port and an MMCM or PLL are directly connected and constrained to different clock regions, you must manually insert a clock buffer and set a CLOCK_DEDICATED_ROUTE constraint on the net connected to the MMCM or PLL.

- On a GT*_CHANNEL or IBUFDS_GTE3 cell

The BUFG_GTs driven by the cell are placed in the same clock region.



CAUTION! Xilinx does not recommend using LOC constraints on the clock buffer cells. This method forces the clock onto a specific track ID, which can result in placement that cannot be legally routed. Only use LOC constraints to place high fanout clock buffers in UltraScale devices when you understand the entire clock tree of the design and when placement is consistent in the design. Even after taking these precautions, collisions might occur during implementation due to design or constraint changes.

Using the CLOCK_REGION Property on Clock Buffers

You can use the CLOCK_REGION constraint to assign a clock buffer to a clock region without specifying a site. This gives the placer more flexibility when optimizing all the clock trees and when determining the appropriate buffer sites to successfully route all clocks.

You can also use a CLOCK_REGION constraint to provide guidance on the placement of cascaded clock buffers or clock buffers driven by non-clocking primitives, such as fabric logic.

In the following example, the XDC constraint assigns the clkgen/clkout2_buf clock buffer to the CLOCK_REGION X2Y2.

```
set_property CLOCK_REGION X2Y2 [get_cells clkgen/clkout2_buf]
```

Note: In most cases, the clock buffers are directly driven by input clock ports, MMCMs, PLLs, or GT*_CHANNELs that are already constrained to a clock region. If this is the case, the clock buffers are automatically placed in the same clock region, and you do not need to use the CLOCK_REGION constraint.

Using a Pblock to Restrict Clock Buffer Placement

When a clock buffer does not need to be placed in a specific clock region, you can use a Pblock to specify a range of clock regions. For example, use a Pblock when a BUFCTRL is needed to multiplex two clocks that are located in different areas. You can assign the BUFCTRL to a Pblock that includes the clock regions between the two clock drivers and let the placer identify a valid placement.

Note: Xilinx does *not* recommend using a Pblock for a single clock region.

Using the USER_CLOCK_ROOT Property on a Clock Net

You can use the USER_CLOCK_ROOT property to force the clock root location of a clock driven by a clock buffer. Specifying the USER_CLOCK_ROOT property influences the design placement, because it impacts both insertion delay and skew by modifying the clock routing. The USER_CLOCK_ROOT value corresponds to a clock region, and you must set the property on the net segment directly driven by the high fanout clock buffer. Following is an example:

```
set_property USER_CLOCK_ROOT X2Y3 [get_nets clkgen/wbClk_o]
```

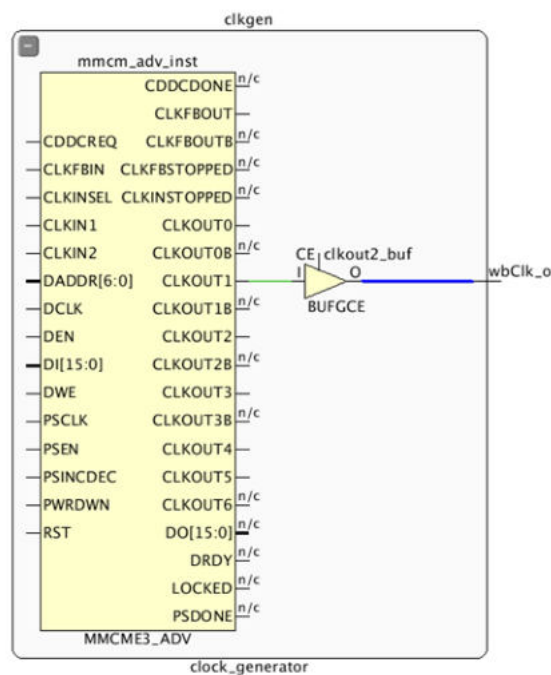


Figure 4-17: USER_CLOCK_ROOT Applied on the Net Segment Driven by the Clock Buffer

After placement, you can use the CLOCK_ROOT property to query the actual clock root as shown in the following example. The CLOCK_ROOT reports the assigned root whether it was user assigned or automatically assigned by the Vivado tools.

```
get_property CLOCK_ROOT [get_nets clkgen/wbClk_o]
=> X2Y3
```


Another way to review the clock root assignments of your implemented design is to use the `report_clock_utilization` Tcl command. For example:

```
report_clock_utilization [-clock_roots_only]
```

The following figure shows this report.

Index	Clock Net	Root Clock Region
1	clkgen/clkfbout_buf	X4Y1
2	clkgen/cpuClk_o	X4Y1
3	clkgen/fftClk_o	X3Y2
4	clkgen/phyClk0_o	X3Y3
5	clkgen/phyClk1_o	X3Y2
6	clkgen/usbClk_o	X3Y3

Figure 4-18: `report_clock_utilization` Clock Root Assignments

Using the `CLOCK_DELAY_GROUP` Constraint on Several Clock Nets

You can use the `CLOCK_DELAY_GROUP` constraint to match the insertion delay of multiple, related clock networks driven by different clock buffers. This constraint is commonly used to minimize skew on synchronous CDC timing paths between clocks originating from the same MMCM or PLL source. You must set the `CLOCK_DELAY_GROUP` constraint on the net segment directly connected to the clock buffer. Following is an example that shows the `clk1_net` and `clk2_net` clock nets, which are directly driven by the clock buffers:

```
set_property CLOCK_DELAY_GROUP grp12 [get_nets {clk1_net clk2_net}]
```

For more information on using this constraint on paths between clocks, see [Synchronous CDC](#).

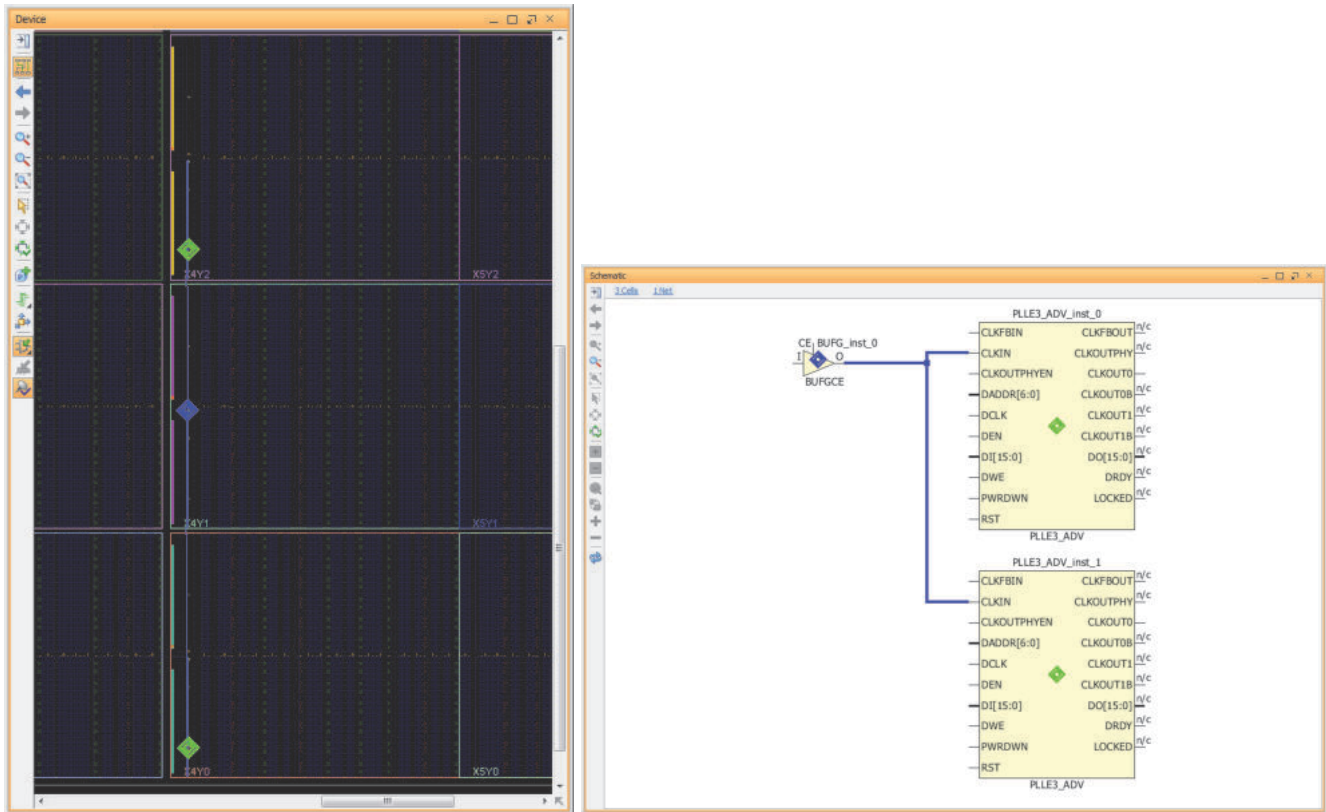
Using the `CLOCK_DEDICATED_ROUTE` Constraint

The `CLOCK_DEDICATED_ROUTE` constraint is typically used when driving from a clock buffer in one clock region to an MMCM or PLL in another clock region. By default, the `CLOCK_DEDICATED_ROUTE` constraint is set to `TRUE`, and the buffer/MMCM or PLL pair must be placed in the same clock region.

When driving from a clock buffer in one clock region to a MMCM or PLL in a vertically adjacent clock region, you must set the `CLOCK_DEDICATED_ROUTE` to `BACKBONE`. This prevents implementation errors and ensures that the clock is routed with global clock resources only. The following example shows a clock buffer driving two PLLs in vertically adjacent clock regions.

```
set_property CLOCK_DEDICATED_ROUTE BACKBONE [get_nets clk_buf]
```

The following figure shows the CLOCK_DEDICATED_ROUTE constraint set to BACKBONE to drive MMCMs or PLLs in vertically adjacent clock regions.

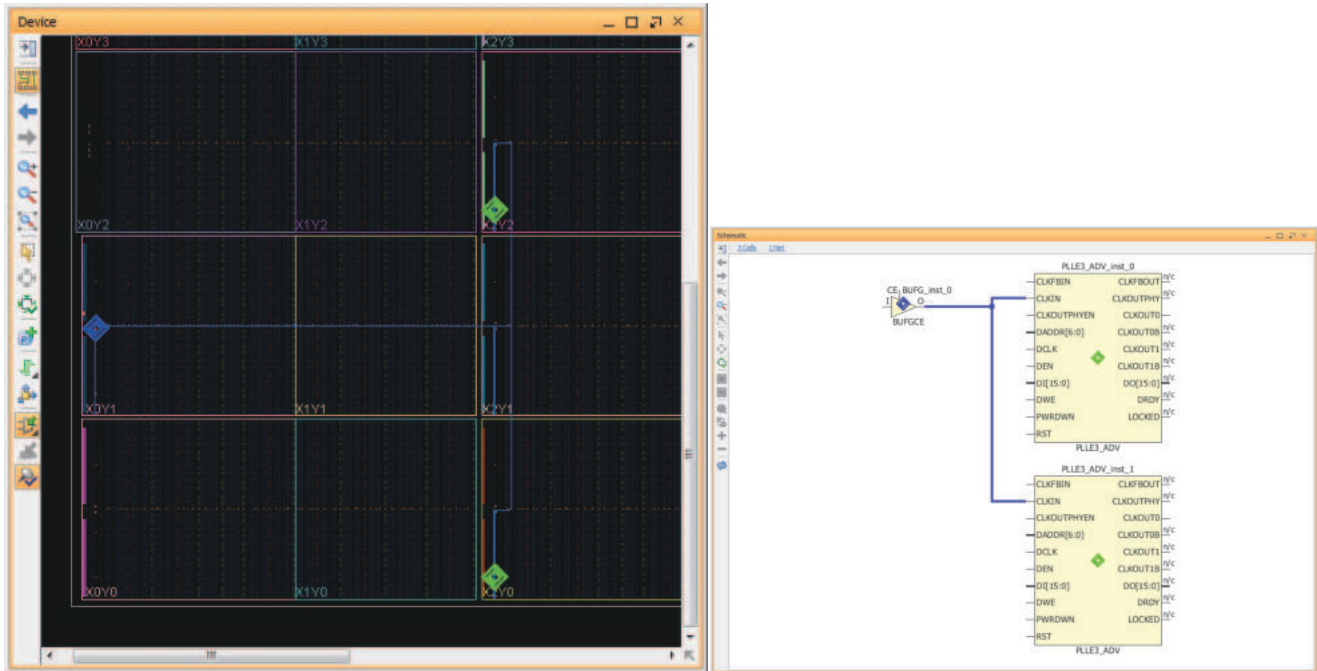


X15235-110415

Figure 4-19: CLOCK_DEDICATED_ROUTE Constraint Set to BACKBONE

When driving from a clock buffer to other clock regions that are not vertically adjacent, you must set the `CLOCK_DEDICATED_ROUTE` to `FALSE`. This prevents implementation errors and ensures that the clock is routed with global clock resources only. The following example and figure show a `BUFGCE` driving two PLLs that are not vertically adjacent.

```
set_property CLOCK_DEDICATED_ROUTE FALSE [get_nets clk_buf]
```



X15236-110415

Figure 4-20: `CLOCK_DEDICATED_ROUTE` Set to `FALSE`

Clocking Topology Recommendations

Xilinx recommends using simple clock tree topologies with the minimum number of clock buffers required for the design. Using extra clock buffers requires more routing tracks, which can lead to placement errors or routing conflicts in clock regions where the clock routing requirement is high and is close to the maximum capacity.

Following are clocking topology recommendations for `BUFGCE`/`BUFGCTRL`/`BUFGCE_DIV` connectivity.

Parallel Clock Buffers

Use parallel clock buffers to achieve the following:

- Ensure predictable placement across implementation runs

When the parallel clock buffers are directly driven by the same input clock port, MMCM, PLL, or GT*_CHANNEL, the buffers are always placed in the same clock region as their driver regardless of the netlist changes or logic placement variation.

- Match the insertion delays between parallel branches of the clock tree

Xilinx recommends parallel buffers over cascaded clock buffers, especially when there are synchronous paths between the branches. When using cascaded buffers, the clock insertion delay is *not* matched between the branches of the clock trees even when using the CLOCK_DELAY_GROUP or USER_CLOCK_ROOT constraints. This can result in high clock skew, which makes timing closure challenging if not impossible.

The following figure shows three parallel BUFGCE buffers driven by the MMCM CLKOUT0 port.

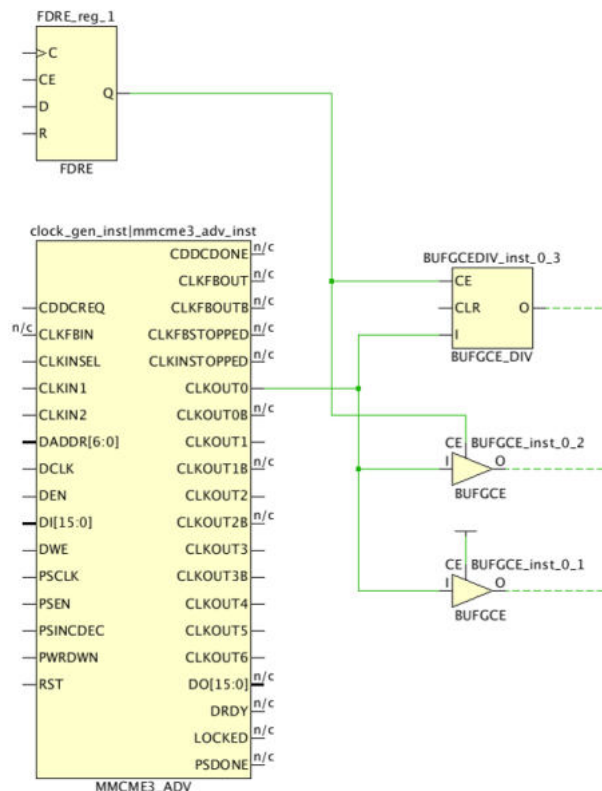


Figure 4-21: Parallel BUFGCE on MMCM Output

Cascaded Clock Buffers

In general, Xilinx does not recommend using cascaded buffers to artificially increase the delay and reduce the skew between unrelated clock trees branches. Unlike connections between BUFCTRLs, other clock buffer connections do not have a dedicated path in the architecture. Therefore, the relative placement of clock buffers is not predictable, and all placement rules take precedence over placing unconstrained cascaded buffers.

However, you can use cascaded clock buffers to achieve the following:

- Route the clock to another clock buffer located in a different clock region

This method is typical when using a clock multiplexer for clocks generated by MMCMs located in different clock regions. Although one of the MMCMs can directly drive the BUFCTRL (BUFGMUX), the other MMCM requires an intermediate clock buffer to route the clock signal to the other region. The following figure shows an example.

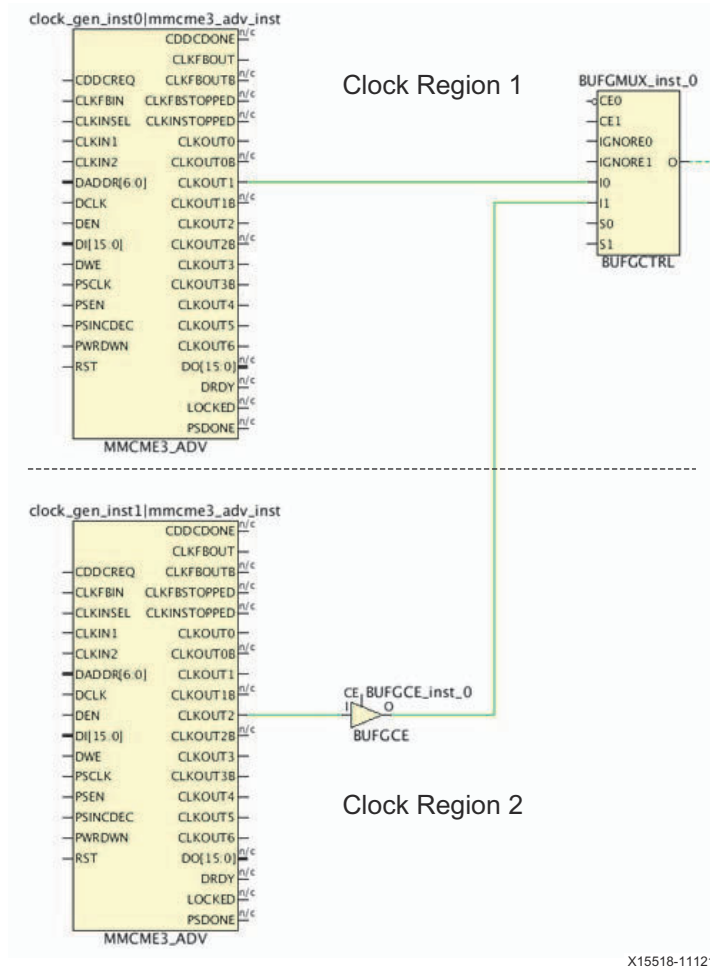


Figure 4-22: Routing the Clock to Another Clock Region

- Balance the number of clock buffer levels across the clock tree branches when there is a synchronous path between those branches

For example, consider an MMCM clock called clk0 that drives both group A (sequential cells driven via a BUFCTRL located in a different clock region) and group B (sequential cells). To better match the delay between the branches, insert a BUFCTRL for group B and place it in the same clock region as the BUFCTRL. This ensures that the synchronous paths between group A and group B have a controlled amount of skew. The following figure shows an example.

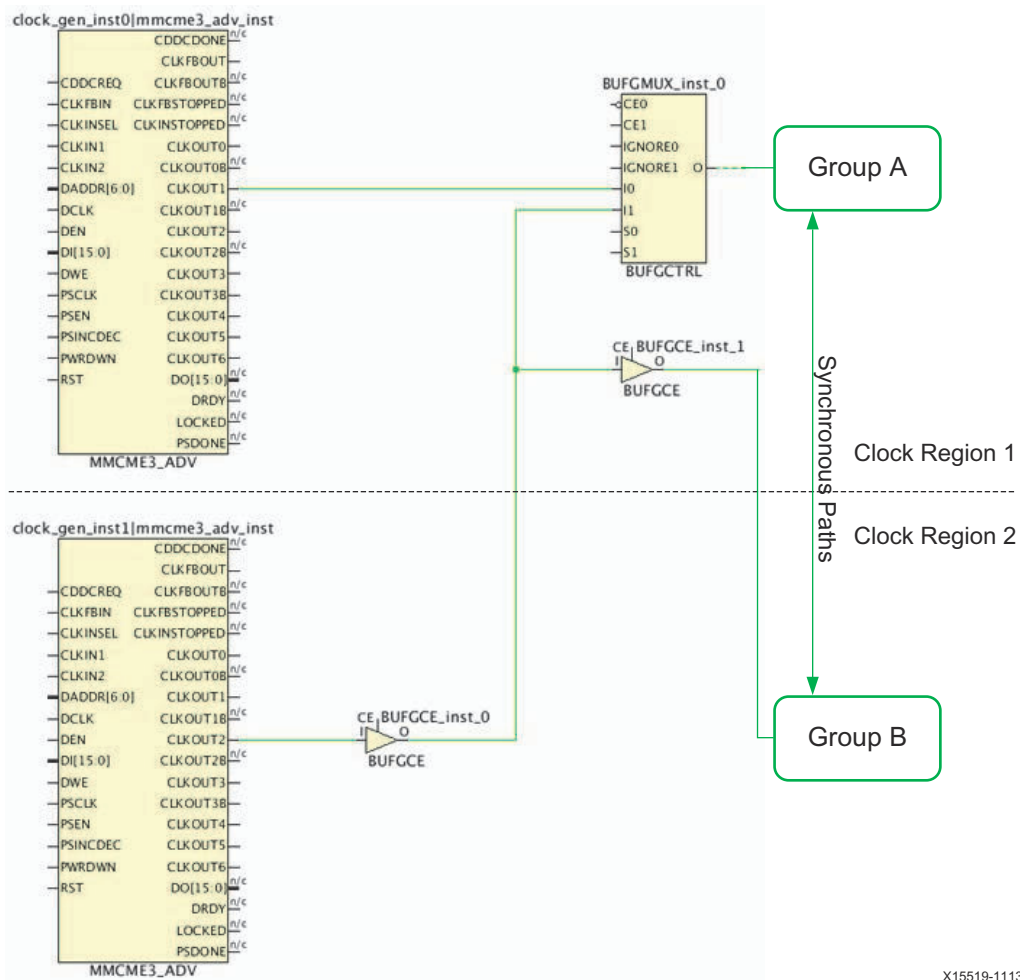


Figure 4-23: Balancing Clock Trees for Synchronous Paths Between Clock Regions

Note: If there are only asynchronous paths between the clock tree branches, the branches do not need to be balanced as long as there is proper synchronization circuitry on the receiving clock domain.

- Build clock multiplexers as described in [Clock Multiplexing](#).

To reduce the variation of insertion delays and skew, Xilinx recommends the following when using cascaded clock buffers:

- Keep the cascaded buffers in the same or adjacent clock regions.
- When clock tree branches are balanced, assign all the clock buffers of the same level to the same clock region.

Note: If absolutely required, Xilinx recommends using two cascaded BUFGCTRLs instead of cascaded BUFGCEs. Using dedicated routing, you can cascade two adjacent BUFGCTRLs with minimum delay when both BUFGCTRLs are placed inside the same clock region.

Clock Multiplexing

You can build a clock multiplexer using a combination of parallel and cascaded BUFGCTRLs. The placer finds the optimal placement based on the clock buffer site availability. If possible, the placer places BUFGCTRLs in the same clock region to take advantage of the dedicated cascade paths. If that is not possible, the placer places BUFGCTRLs from the same level in the adjacent clock regions.

The following figure shows a 4:1 MUX with balanced cascading. The first level of BUFGCTRL buffers are both placed in the directly adjacent clock regions (X0Y2, X0Y0) of the last BUFGCTRL (X0Y1). This configuration ensures a comparable insertion delay for all the clocks reaching the last BUFGCTRL. You can use an equivalent structure for a 3:1 MUX.

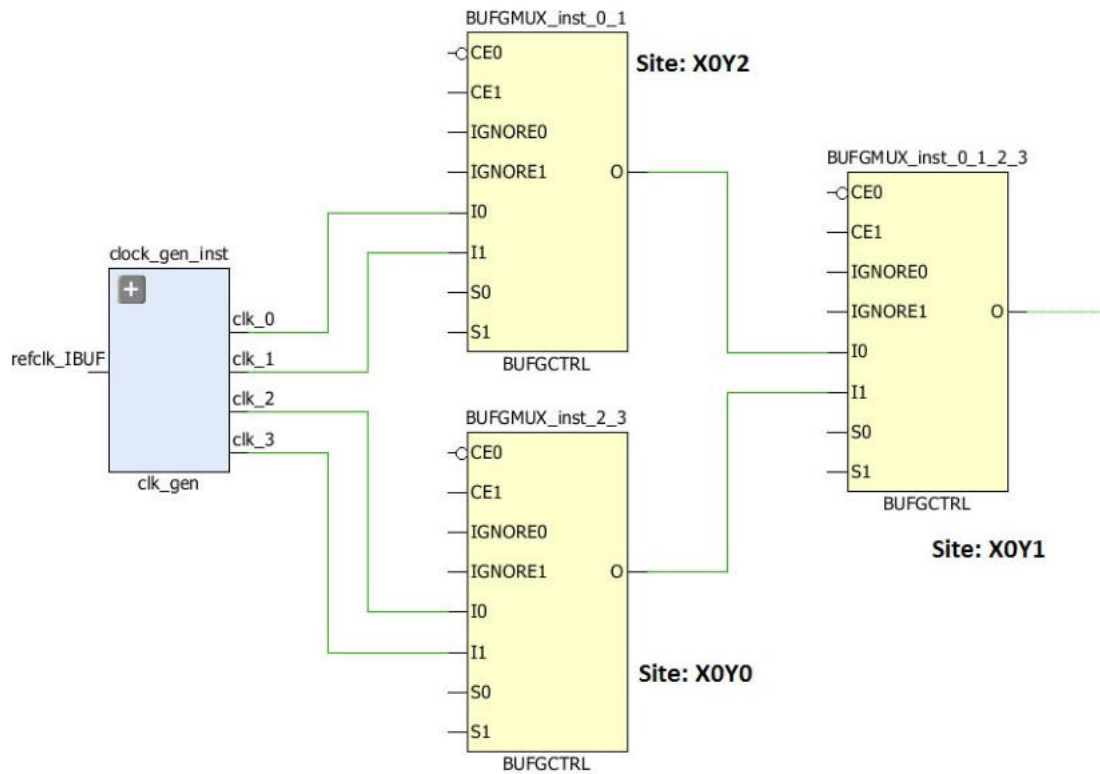


Figure 4-24: 4:1 MUX Using Parallel BUFCTRL

When creating a 5:1 or larger clock MUX structure, it is common to create a symmetrical clock structure as shown in the following figure. However, this is a sub-optimal solution, because each BUFCTRL only has one cascade path to the two adjacent BUFCTRLs, which does not provide minimal delay for all connections between the BUFCTRLs.

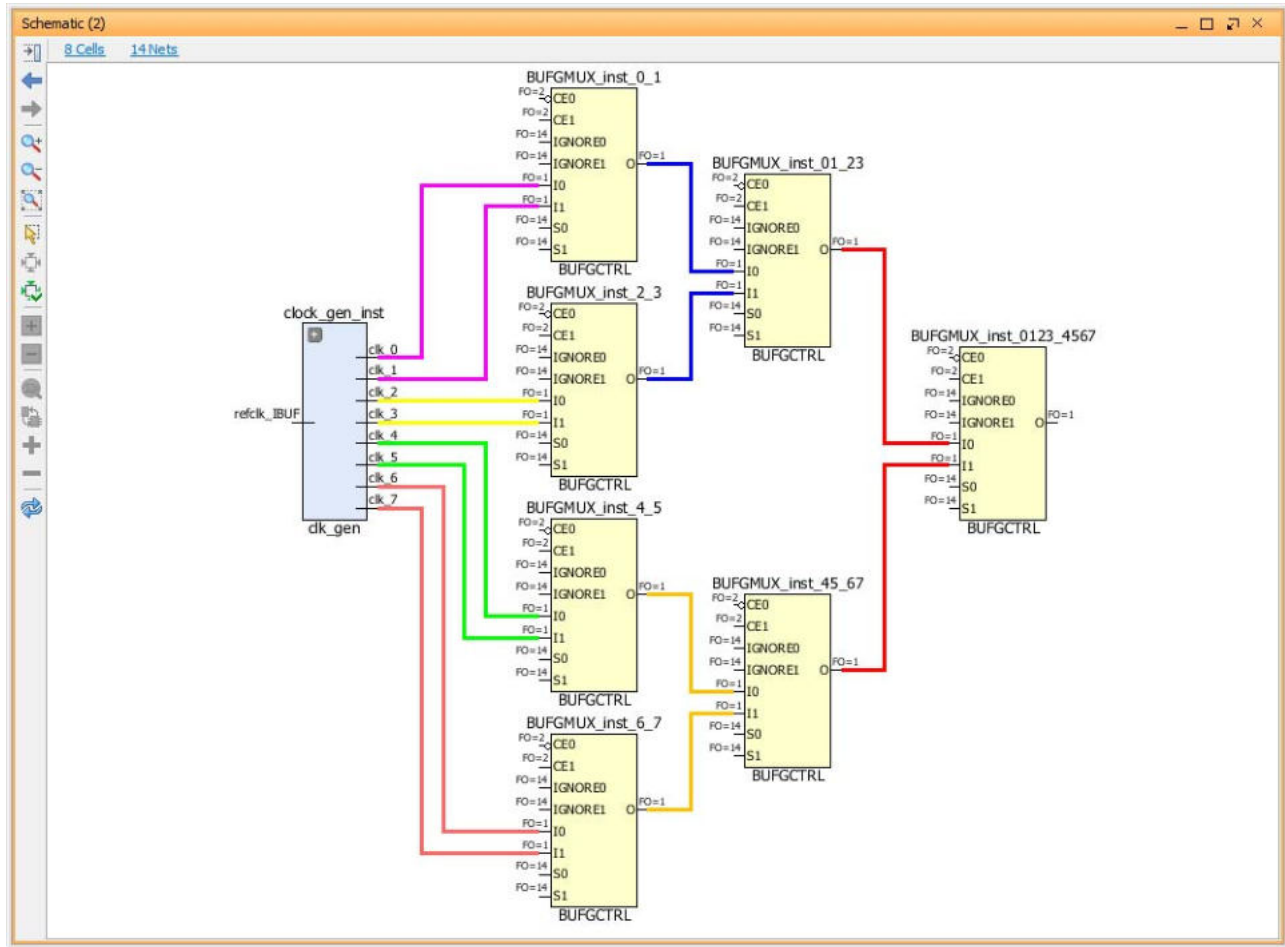


Figure 4-25: Non-Recommended 8:1 Balanced Clock MUX Structure

To support larger clock multiplexers (from 5:1 to 8:1 MUX), Xilinx recommends using cascaded BUFCTRL buffers as shown in the following figure. This figure shows an optimal 8:1 MUX that uses 7 BUFCTRL buffers.

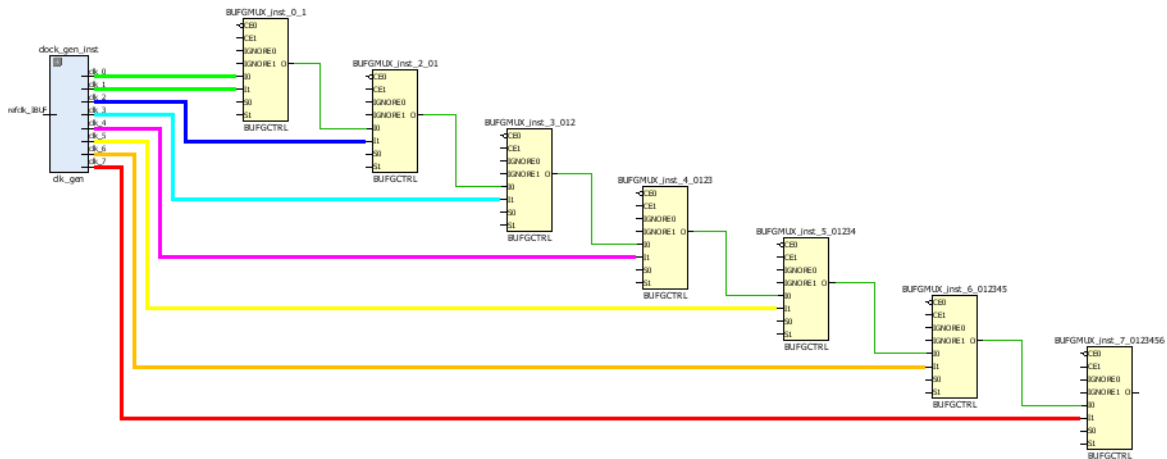


Figure 4-26: 8:1 MUX Using Cascaded BUFCTRL

Note: When using wide BUFCTRL-based clock multiplexers, the clock insertion delays cannot be balanced because some paths are longer than other paths in hardware. Therefore, this method is recommended for multiplexing asynchronous clocks only.

PLL/MMCM Feedback Path and Compensation Mode

PLLs do not support delay compensation and always operate in INTERNAL compensation mode, which means they do not need a feedback path. Similarly, MMCMs set to INTERNAL compensation mode do not need a feedback path. In both cases, the Vivado tools do not always automatically remove unnecessary feedback clock buffers. You must remove the clock buffers manually to reduce the amount of high fanout clock resource utilization. This is especially important for designs with high clocking usage where clock contention might occur.

When the MMCM compensation is set to ZHOLD or BUF_IN, the placer assigns the same clock root to the nets driven by the feedback buffer and by all buffers directly connected to the CLKOUT0 pin. This ensures that the insertion delays are matched so that the I/O ports and the sequential cells connected to CLKOUT0 are phase-aligned and hold time is met at the device interface. The Vivado tools consider all the loads of these nets to optimally define the clock root.

The Vivado tools do not automatically match the insertion delay with the other MMCM outputs. To match the insertion delay for the nets driven by other MMCM output buffers, use the following properties:

- `CLOCK_DELAY_GROUP`

Apply the same `CLOCK_DELAY_GROUP` property value to the nets directly driven by feedback clock buffer, the `CLKOUT0` buffers, and the other MMCM output buffers as needed. This is the preferred method.

- `USER_CLOCK_ROOT`

If you need to force a specific clock root, use the same `USER_CLOCK_ROOT` property value on the nets driven by the feedback clock buffer, the `CLKOUT0` buffers, and the other MMCM output buffers as needed.

BUFG_GT Divider

The `BUFG_GT` buffers can drive any loads in the fabric and include an optional divider you can use to divide the clock from the `GT*_CHANNEL`. This eliminates the need to use an extra MMCM or `BUFG_DIV` to divide the clock.

I/O Timing with MMCM ZHOLD/BUF_IN Compensation

Because the clock insertion delay varies with the clock root locations and the clock root placement depends on placement of the loads, there might be variability between runs. This variability affects the timing inside the FPGA as well as the I/O timing.

When dealing with high-frequency I/Os, you might want more control over the I/O timing and less variability between runs. One way to achieve this is to force the clock root placement. You can run the tool in automated mode and look at the clock root region. If the I/O timing is satisfactory, you can force the clock root placement on the buffer nets associated with I/O timing. To determine the placement of the clock roots, use the `report_clock_utilization [-clock_roots_only]` Tcl command.

In the following example, the I/O ports are located in the X0Y0 region. The Vivado placer determined the placement of the clock roots in X1Y2 based on the I/O placement as well as placement of other loads.

Index	Clock Net	Root Clock Region	Clock Root Node
1	mmcm/inst/clk0	X1Y2	RCLK_BRAM_L_X30Y209/CLK_VDISTR_B0T0
2	mmcm/inst/clkfbout_buf_mmcm_zhold	X1Y2	RCLK_CLEL_R_L_X25Y209/CLK_VDISTR_B0T

Figure 4-27: Clock Utilization Summary with Unconstrained Clock Root

The following summary shows the I/O timing when the clock root is unconstrained.

Setup	Hold
Worst Negative Slack (WNS): -0.279 ns	Worst Hold Slack (WHS): 0.036 ns
Total Negative Slack (TNS): -5.394 ns	Total Hold Slack (THS): 0.000 ns
Number of Failing Endpoints: 47	Number of Failing Endpoints: 0

Figure 4-28: Timing Summary with Unconstrained Clock Root

In the following example, the clock roots are moved next to the I/O registers in X0Y0, which reduces the clock insertion delays and timing pessimism and therefore, improves the I/O timing.

Index	Clock Net	Root Clock Region	Clock Root Node
1	mmcm/inst/clk0	X0Y0	XIPHY_L_X0Y60/CLK_VDISTR_B0T20
2	mmcm/inst/clkfbout_buf_mmcm_zhold	X0Y0	XIPHY_L_X0Y60/CLK_VDISTR_B0T14

Figure 4-29: Clock Utilization Summary with User Constrained Clock Root

The following summary shows the I/O timing when the clock root is moved.

Setup	Hold
Worst Negative Slack (WNS): -0.217 ns	Worst Hold Slack (WHS): 0.060 ns
Total Negative Slack (TNS): -1.488 ns	Total Hold Slack (THS): 0.000 ns
Number of Failing Endpoints: 16	Number of Failing Endpoints: 0

Figure 4-30: Timing Summary with User Constrained Clock Root

Synchronous CDC

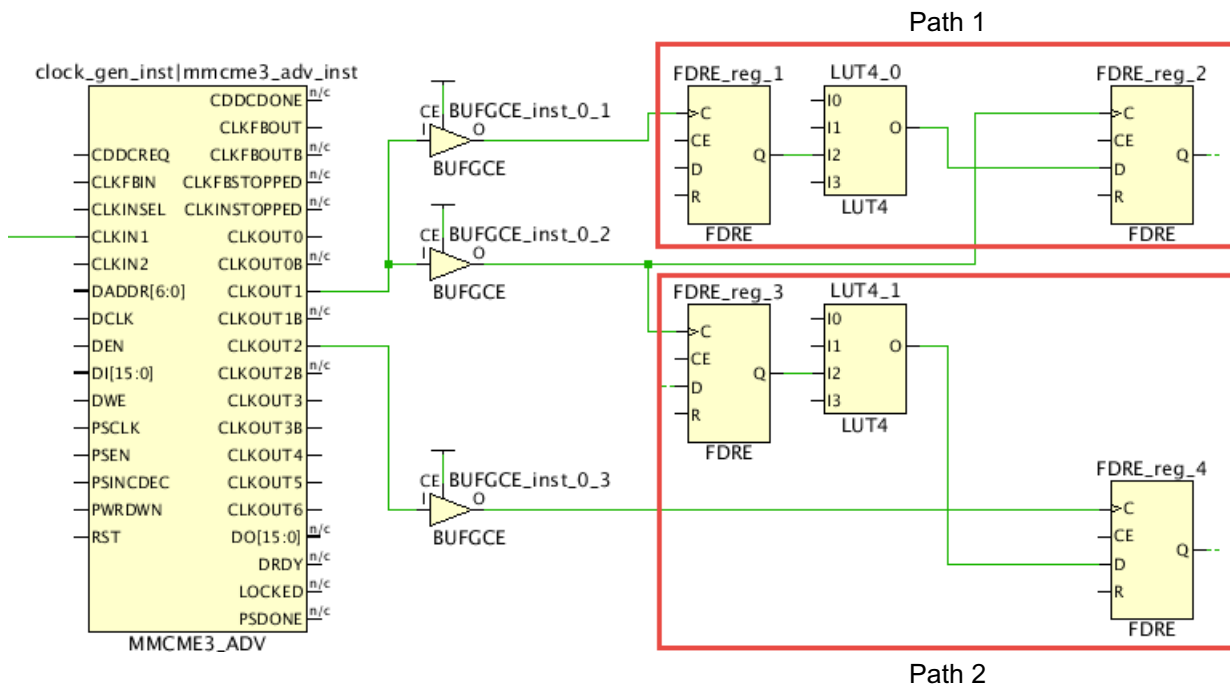
When the design includes synchronous CDC paths between clocks that originate from the same MMCM/PLL, you can use the following techniques to better control the clock insertion delays and skew and therefore, the slack on those paths.



IMPORTANT: *If the CDC paths are between clocks that originate from different MMCM/PLLs, the clock insertion delays across the MMCMs/PLLs are more difficult to control. In this case, Xilinx recommends that you treat these clock domain crossings as asynchronous and make design changes accordingly.*

When a path is timed between two clocks that originate from different output pins of the same MMCM/PLL, the MMCM/PLL phase error adds to the clock uncertainty for the path. For designs using high clock frequencies, the phase error can cause issues with timing closure both for setup and hold.

The following figure shows an example of paths both with and without the phase error. Path 1 is a CDC path clocked by two buffers connected to the same MMCM output and does not include the phase error. Path 2 is clocked by two clocks that originate from two different MMCM outputs and does include the phase error.



X15234-110315

Figure 4-31: MMCM and Phase Error

When two synchronous clocks from the same MMCM/PLL have a simple period ratio (2/4 /8), you can prevent the phase error between the two clock domains using a single MMCM/PLL output connected to two BUFGE_DIV buffers. The BUFGE_DIV buffer performs the clock division (1/2 /4 /8). Other ratios are possible (3/5 /7) but this requires modifying the clock duty cycle and making mixed edge timing paths more challenging.

Note: Because the BUFGE and BUFGE_DIV do not have the same cell delays, Xilinx recommends using the same clock buffer for both synchronous clocks (two BUFGE or two BUFGE_DIV buffers).

The following figure shows two BUFGCE_DIVs that divide the CLKOUT0 clock by 1 and by 2 respectively.

Note: Although the following figure only uses two BUFGCE_DIVs in parallel, you can use up to four BUFGCE_DIVs in a clock region.

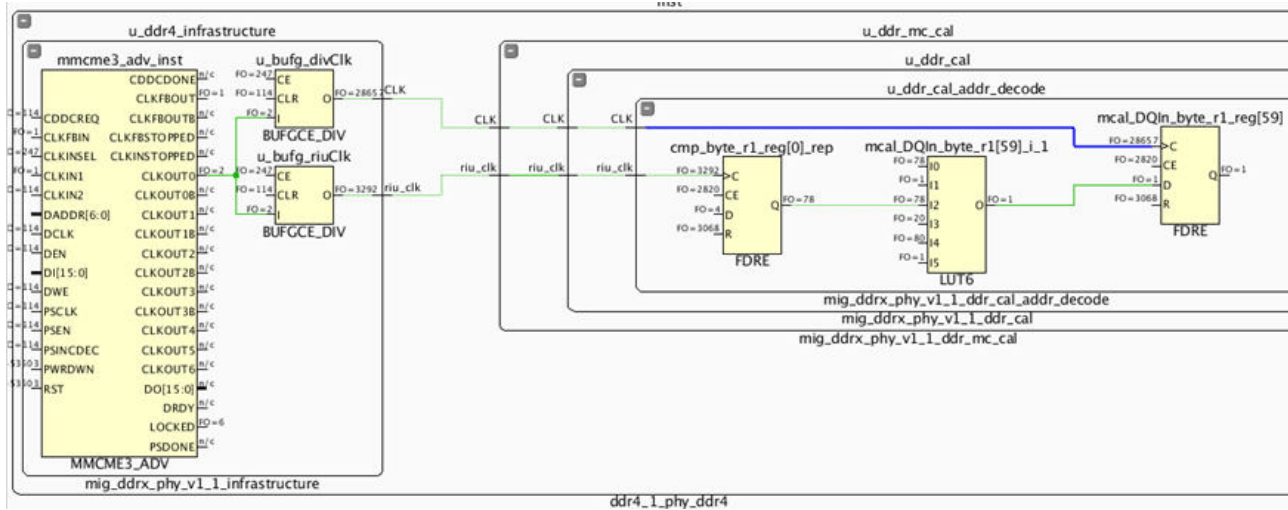


Figure 4-32: MMCM Synchronous CDC with BUFGCE_DIVs Connected to One MMCM Output

To automatically balance several clocks that originate from the same MMCM or PLL, set the same CLOCK_DELAY_GROUP property value on the nets driven by the clock buffers that need to be balanced. Following are additional recommendation:

- Avoid setting the CLOCK_DELAY_GROUP constraint on too many clocks, because this stresses the clock placer resulting in sub-optimal solutions or errors.
- Review the critical synchronous CDC paths in the Timing Summary Report to determine which clocks must be delay matched to meet timing.
- Limit the use of the CLOCK_DELAY_GROUP on groups of synchronous clocks with tight requirements and with identical clocking topologies.

GT Interface Clocking

Each GT interface requires several clocks, including some clocks that are shared across bonded GT*_CHANNEL cells located in one or several GT quads. UltraScale devices provide up to 128 GT*_CHANNEL sites, which can lead to the use of several hundreds of clocks in a design. Most GT clocks have a low fanout with loads placed locally in the clock region next to the associated GT*_CHANNEL. Some GT clocks drive loads across the entire device and require the utilization of clock routing resource in many clock regions. The UltraScale architecture includes the following enhancements to efficiently support the high number of GT clocks required.

BUFG_GT with Dynamic Divider

In UltraScale devices, the BUF_{GT} buffer simplifies GT clocking. Because the BUF_{GT} includes dynamic division capabilities, MMCMs are no longer required to perform simple integer divides on GT output clocks. This saves clocking resources and provides an improved low skew clock path when both a divided GT*_CHANNEL output clock and full-rate clock are required.

You can use the BUF_{GT} global clock buffer for GT interfaces where the user logic operates at half the clock frequency of the internal PCS logic or for PCIe® interfaces where the GT*_CHANNEL needs to generate multiple clock frequencies for user_clk, sys_clk, and pipe_clk. The following figure compares clocking requirements between 7 series and UltraScale devices for a single-lane GT interface where the frequency of TXUSRCLK2 is equal to half of the frequency of TXUSRCLK.

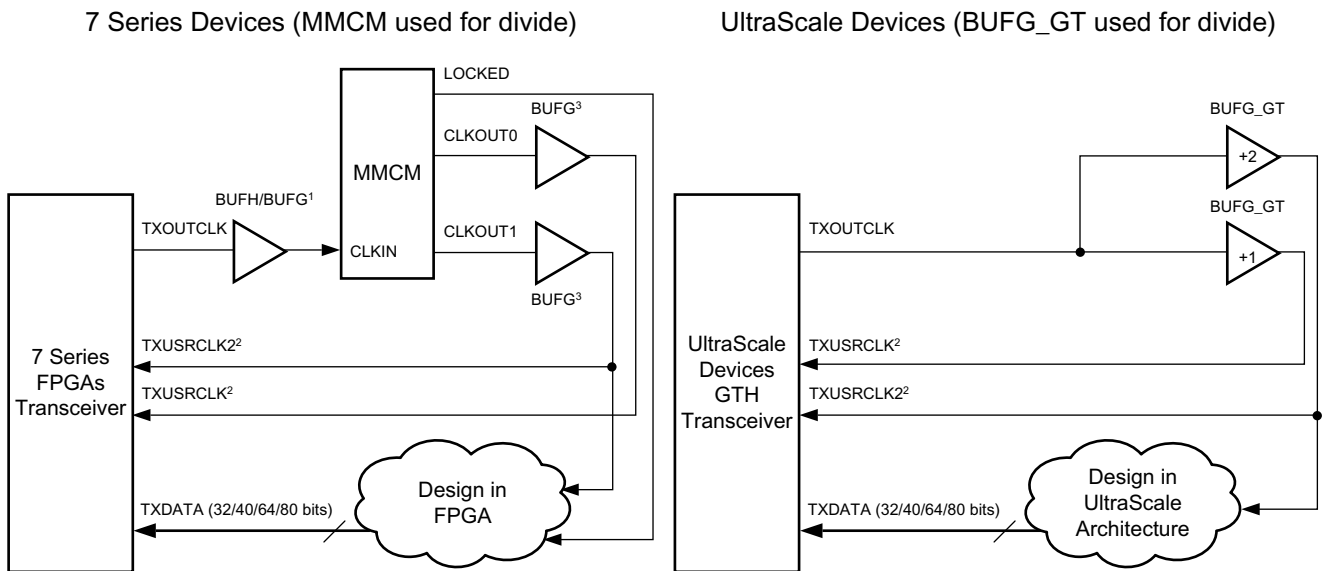


Figure 4-33: Clocking Requirements Comparison

You can use any output clock of the GT*_CHANNELS within a Quad or any reference clock generated by an IBUF_{DS}_GTE3/ODIV2 pin within a Quad to drive any of the 24 BUF_{GT} buffers located in the same clock region. A BUF_{GT}_SYNC is always required to synchronize reset and clear of BUF_{GT}s driven by a common clock source.

Note: The Vivado tools automatically insert the BUF_{GT}_SYNC primitive if it is not present in the design.

Some applications still require the use of an MMCM to generate complex non-integer clock division of the GT output clocks or the IBUF_{DS}_GTE3/ODIV2 reference clock. In these cases, a BUF_{GT} must directly drive the MMCM. By default, the placer tries to place the MMCM on the same clock region row as the BUF_{GT}. If other MMCMs try to use the same MMCM

site, you must verify that the automated MMCM placement is still as close as possible to the BUFG_GT to avoid wasting clocking resources due to long routes.

Single Quad vs. Multi-Quad Interface

In a multi-channel interface, a master channel can generate [RT]XUSRCLK[2] for all the GT*CHANNELs of the interface. If a multi-channel interface spans multiple quads, the maximum allowed distance for a GT*CHANNEL from the reference clock source is 2 clock regions above and below.

If the GT interface is contained within a single Quad, the placer treats the BUFG_GT clocks as local clocks. In this case, the placer attempts to place the BUFG_GT clock loads in the clock regions horizontally adjacent to the BUFG_GT, starting with the clock region that contains the BUFG_GT and potentially using up to half the width of the device.

[RT]XUSRCLK/[RT]XUSRCLK2 Skew Matching

When [RT]XUSRCLK2 operates at half the frequency of [RT]XUSRCLK (i.e., separate BUFG_GTs with divide by 1 and divide by 2), a tight skew requirement exists between the [RT]XUSRCLK/[RT]XUSRCLK2 pair at each GT*CHANNEL of a GT interface. To meet the skew requirement, GT*CHANNELs can be a maximum of 2 clock regions above or below the master channel that generates the [RT]XUSRCLK/[RT]XUSRCLK2 pair. In addition, the placer tightly controls skew as follows:

- Assigns the BUFG_GT pairs to the upper or lower 12 BUFG_GTs in a Quad
- Assigns the clock root for both clocks to the clock region containing the BUFG_GTs



RECOMMENDED: To avoid skew violations, Xilinx highly recommends following this clocking topology when [RT]XUSRCLK2 operates at half the frequency of [RT]XUSRCLK.

Integrated Block for PCI Express CORECLK/PIPECLK/USERCLK Skew Matching

The UltraScale Integrated Block for PCI Express® requires three clocks: CORECLK, USERCLK, and PIPECLK. The three clocks are sourced by BUFG_GTs driven by the TXOUTCLK pin of one of the GT*_CHANNELs of the physical interface. A tight skew requirement exists between the CORECLK and PIPECLK pins and the CORECLK and USERCLK pins. To meet the skew requirement, the placer tightly controls skew as follows:

- Assigns the BUFG_GTs that drive the three PCIe clocks in groups to the upper or lower 12 BUFG_GTs in a Quad
- Assigns the clock root for all three clocks to the same clock region

Note: For more information on PCIe clocking requirements, see the *UltraScale Architecture Gen3 Integrated Block for PCI Express LogiCORE IP Product Guide* (PG156) [Ref 41].

7 Series Device Clocking

This section uses Virtex®-7 clocking resources as an example. The clocking resources for Virtex-6 devices are similar. If you are using a different architecture, see the *Clocking Resources Guide* [Ref 39] for your device.

The Virtex-6 and Virtex-7 devices contain thirty-two global clock buffers known as BUFGs. BUFGs can serve most clocking needs for designs with less demanding needs in terms of:

- Number of clocks
- Design performance
- Low power demands
- Other clocking characteristics such as:
 - Clock gating
 - Multiplexing
 - Division
 - Other clocking control

They are inferred by synthesis, and have very few restrictions allowing for most general clocking.



RECOMMENDED: *If clocking demands exceed the number of BUFGs, or if better overall clocking characteristics are desired, analyze the clocking needs against the available clocking resources, and select the best resource for the task.*

Global Clocking Resources

This section discusses the following global clocking resources:

- BUFG

BUFG elements are commonly used for clocking. The global clocking buffers have additional functionality. However, these additional features can be accessed with some manual intervention to your design code or synthesis.

- BUFGCE

A synchronous, glitchless clock enable (gating) capability may be accessed without using any additional logic or resources by using the BUFGCE primitive. The BUFGCE may be used to stop the clock for a period of time or create lower skew and lower power clock division such as one-half ($\frac{1}{2}$) or one-fourth ($\frac{1}{4}$) frequency clocks from a higher frequency base clock especially when different frequencies may be desired at different times of circuit operation.

- BUFGMUX

A BUFGMUX can be used to safely change clocks without glitches or other timing hazards from one clock source to another. This can be used when two distinct frequencies of the clock are desired depending on time or operating conditions.

- BUFGCTRL

The BUFGCTRL gives access to all capabilities of the global clocking network allowing for asynchronous control of the clocking for more complicated clocking situations such as a lost or stopped clock switch-over circuit.

In most cases, the component must be instantiated in the code, and the proper connections made to obtain the desired clocking behavior.

In some situations, IP and synthesis may use these more advanced clocking features. For example, when using the Memory Interface Generator (MIG) special clocking buffers may be used for high-speed data transmit and capture at the I/Os. It is always a good idea to recognize the clock resources required and used for the individual IP, and account for it in your overall clocking architecture and planning.

For more information on using these components, see the *Clocking Resource User Guide* and *Libraries Guides* for the specific devices.

Regional Clocking Resources

In addition to global clocking resources, there are also regional clocking resources:

- Horizontal Clock Region Buffers (BUFH, BUFHCE)

Horizontal Clock Region Buffers (BUFH, BUFHCE) may be used standalone, or in conjunction with BUFGs. These buffers allow you to derive tighter control of the clocking and placement of the associated logic connected to the clock, and provide additional clocking resources for designs with a large number of clock domains.

The BUFH and BUFHCE resources allow the design to use the portions of the global clock network (BUFG) that connects to a given clock region. This allows access to a low skew resource from otherwise unused portions of the global clock network for smaller clock domains that can be located within a clock region. The BUFHCE has the same glitchless clock enable allowing for simple and safe clock gating of a particular clock domain.

When driven by a BUFG, the BUFHCE can be used as a medium-grained clock gating function. For portions of a clock domain ranging from a few hundred to a few thousand loads in which it is desired to stop clocking intermittently, the BUFHCE can be an effective clocking resource. A BUFG can drive multiple BUFHs in the same or different clock regions, allowing for several low skew clock domains in which the clocking can be individually controlled.




Figure 4-34: Horizontal Clock Region Buffers

When used independently, all loads connected to the BUFH must reside in the same clock region. This makes it well-suited for very high-speed, more fine-grained (fewer loads) clocking needs. BUFHCE may be used to achieve medium-grained clock-gating within the specific clock region. You must ensure that the resources driven by the BUFH do not exceed the available resources in the clock region, and that no other conflicts exist.



TIP: Keep loads small on these networks to avoid this problem.

The phase relationship may be different between the BUFH and clock domains driven by BUFGs, other BUFHs, or any other clocking resource. The single exception is when two BUFHs are driven to horizontally adjacent regions. In this case, the skew between left and right clock regions when both BUFHs driven by the same clock source should have a very controlled phase relationship in which data may safely cross the two BUFH clock domains. BUFHs can be used to gain access to MMCMs or PLLs in opposite regions to a clock input or GT. However, care must be taken in this approach to ensure that the MMCM or PLL is available.

- Regional Buffer (BUFR)

The Regional Clock Buffer (BUFR) is generally used as the slower speed I/O and fabric clock for capturing and providing higher-speed I/O data. The BUFR has the ability to enable and disable (gate) the clock as well as perform some common clock division. In Virtex-7 devices, the BUFR can drive only the clock region in which it exists. This makes the buffer better suited for slightly smaller clocking networks.

Because the performance of the BUFR is somewhat lower than the BUFG and BUFH, Xilinx does not recommend it for very high-speed clocking. However, it is well-suited for many medium to lower speed clocking needs. The added capability of built-in clock division makes it suitable for divided clock networks coming from an external clock source such as a high-speed I/O interface clock. It does not consume a global route, and is an alternative to using a BUFH.

- I/O Buffer (BUFIO)

The I/O Clock Buffer (BUFIO) is used exclusively to capture I/O data into the input logic, and provide an output clock to the output logic from the device. BUFIO is generally used to:

- Capture high-speed, source synchronous data within a bank
- Gear down the data (when used in conjunction with a BUFR and an ISERDES or OSERDES logic) to more manageable speeds within the device



IMPORTANT: A BUFIO may drive only the input and output components that exist in the ILOGIC and OLOGIC structures such as the IDDR, ODDR, ISERDES, OSERDES, or simple dedicated input or output registers.

When using the BUFIO, you must take into account the need to reliably transfer the data from the I/O logic to the fabric and vice-versa.

- Multi-Regional Clock Buffer (BUFMR)

The Multi-Regional Clock Buffer (BUFMR) allows a single clock pin (MRCC) to drive the BUFIO and BUFR within its bank, as well as the I/O banks above and below it (assuming they exist).

For more information on Clocking resources for Xilinx 7 series FPGA devices, see the *7 Series FPGAs Clocking Resources User Guide* (UG472) [Ref 39].

Additional Clocking Considerations for SSI Devices

In general, all clocking considerations mentioned above also apply to SSI technology devices. However, there are additional considerations when targeting these devices due to their construction. As mentioned in the prior section, regional clocking can be considered the same with the exception of when using a BUFMR, it cannot drive clocking resources across an SLR boundary. Accordingly, Xilinx recommends that you place the clocks driving BUFMRs into the bank or clocking region in the center clock region within an SLR. This gives access to all three clock regions on the left or right side of the SLR.

In terms of global clocking, for designs requiring sixteen or fewer global clocks (BUFGs), no additional considerations are necessary. The tools automatically assign BUFGs in a way to avoid any possible contention. When more than sixteen (but fewer than thirty-two) BUFGs are required, some consideration to pin selection and placement must be done in order to avoid any chance of contention of resources based on global clocking line contention and/or placement of clock loads.

As in all other Xilinx 7 series devices, Clock-Capable I/Os (CCIOs) and their associated Clock Management Tile (CMT) have restrictions on the BUFGs they can drive within the given SLR. CCIOs in the top or bottom half of the SLR can drive BUFGs only in the top or bottom half of the SLR (respectively). For this reason, pin and associated CMT selection should be done

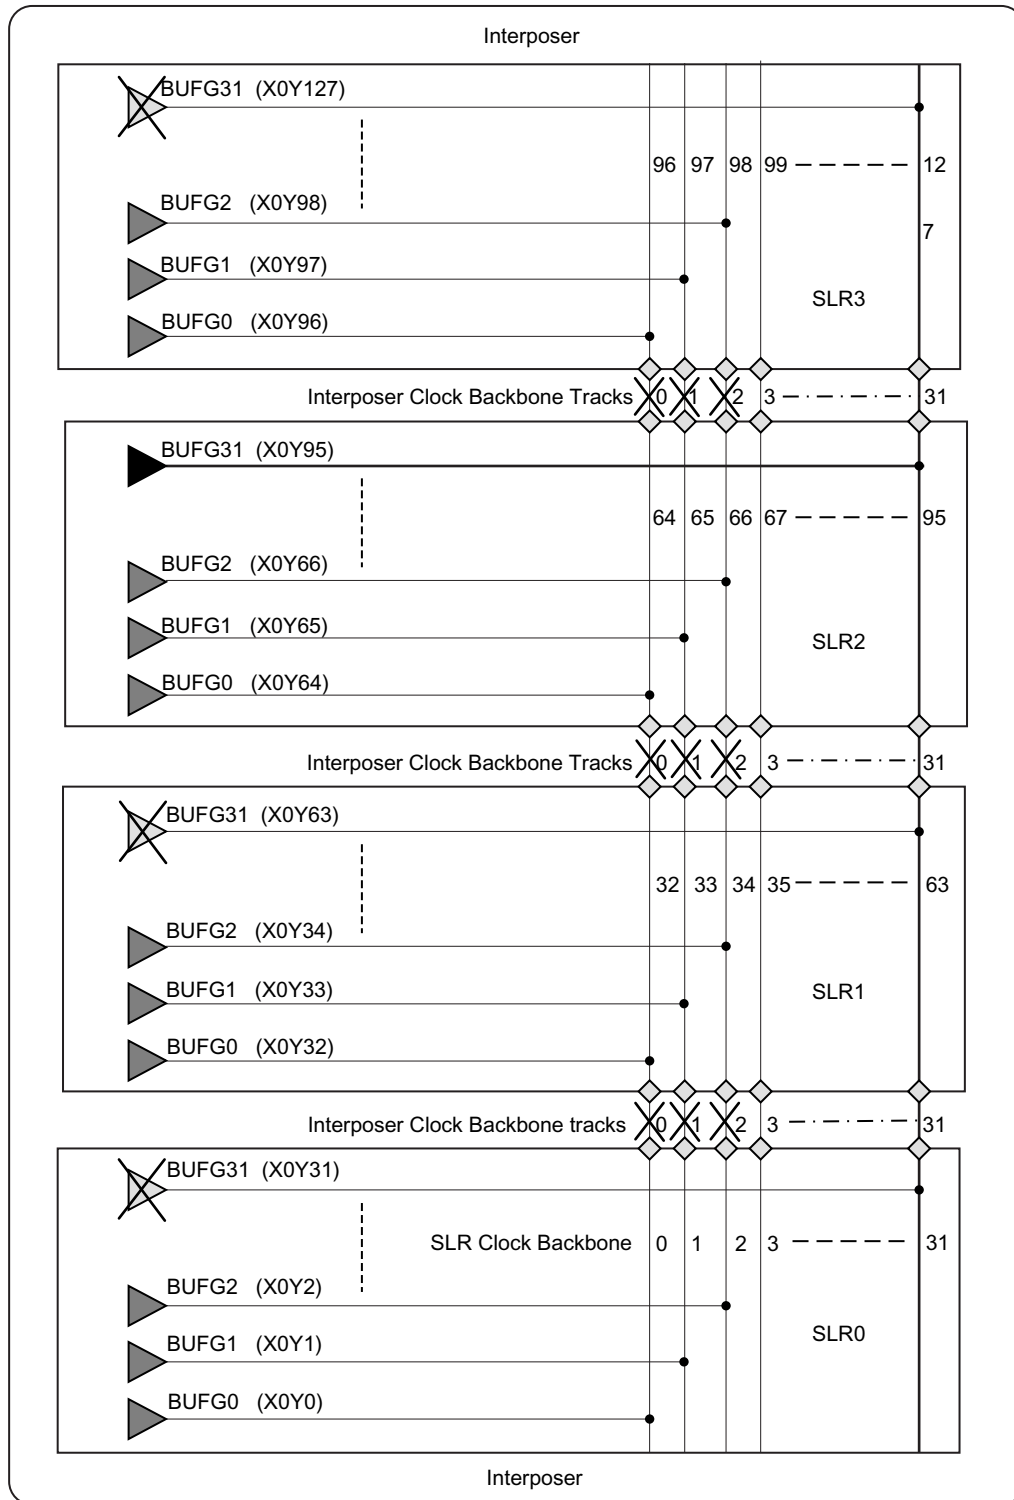
in a way in which no more than sixteen BUFMs are required in either the top or bottom half of all SLRs collectively. In doing so, the tools can automatically assign all BUFMs in a way to allow all clocks to be driven to all SLRs without contention.

For designs that require more than thirty-two global clocks, Xilinx recommends that you explore using BUFMs and BUFHs for smaller clock domains to reduce the number of needed global clock domains. BUFMs with the use of a BUFMR to drive resources within three clock regions that encompasses one-half of an SLR (approximately 250,000 logic cells in a Virtex-7 class SLR). Horizontally adjacent clock regions may have both left and right BUFH buffers driven in a low-skew manner enabling a clocking domain of one-third of an SLR (approximately 167,000 logic cells).

Using these resources when possible not only leads to fewer considerations for clocking resource contention, but many times improves overall placement, resulting in improved performance and power.

If more than thirty-two global clocks are needed that must drive more than half of an SLR or to multiple SLRs, it is possible to segment the BUFM global clocking spines. Isolation buffers exist on the vertical global clock lines at the periphery of the SLRs that allow use of two BUFMs in different SLRs that occupy the same vertical global clocking track without contention. To make use of this feature, more user control and intervention is required. In the figure below, BUF0 through BUF2 in the three SLRs have been isolated, and hence have independent clocks within their respective SLRs. On the other hand, the BUF31 line has not been isolated. Hence, the same BUF31 (located in SLR2 in the figure) drives the clock lines in all the 3 SLRs - and BUF31 located in other SLRs should be disabled.

Careful selection and manual placement (LOCs) must be used for the BUFMs. Additionally, all loads for each clock domain must be manually grouped and placed in the appropriate SLR to avoid clocking contention. If all global clocks are placed and all loads managed in a way to not create any clocking contention and allow the clock to reach all loads, this can allow greater use of the global clocking resources beyond thirty-two.



X14051

Figure 4-35: Optional Isolation on Clock Lines for SSI Devices

Clock Skew for Global Clocking Resources in SSI Devices

Clock skew in any large FPGA device may represent a significant portion of the overall timing budget for a given path. Too much clock skew may not only represent issues with maximum clock speed, but may also manifest itself into stringent hold time requirements. Having multiple die in a device worsens the process portion of the PVT equation, but is managed by the Xilinx assembly process in which only die of similar speed are packaged together.

Even with that extra action, the Xilinx timing tools accounts for these differences as a part of the timing report. During path analysis, these aspects are analyzed as a part of the setup and hold calculations, and are reported as a part of the path delay against the specified requirements. No additional user calculations or consideration are necessary for SSI technology devices, because the timing analysis tools consider these factors in their calculations.

Skew can increase if using the top or bottom SLR as the delay-differential is higher among points farther away from each other. For this reason, Xilinx recommends for global clocks that must drive more than one SLR to be placed into the center SLR. This allows a more even distribution of the overall clocking network across the part resulting in less overall clock skew.

When targeting UltraScale devices, there is less repercussion to clock placement. However, it is still highly suggested to place the clock source as close as possible to the central point of the clock loads to reduce clock insertion delay and improve clock power.

Designing the Clock Structure

Now that you understand the major considerations for clocking decisions, let us see how you can achieve the desired clocking for your design.

Inference

Without user intervention, Vivado synthesis automatically specifies a global buffer (BUFG) for all clock structures up to the maximum allowed in an architecture (unless otherwise specified or controlled by the synthesis tool). As discussed above, the BUFG provides a well-controlled, low-skew network suitable for most clocking needs. Nothing additional is required unless your design clocking exceeds the number or capabilities of BUFGs in the part.

Applying additional control of the clocking structure, however, may prove to show better characteristics in terms of jitter, skew, placement, power, performance, or other characteristics.

Synthesis Constraints and Attributes

A simple way to control clocking resources is to use the `CLOCK_BUFFER_TYPE` synthesis constraint or attribute. Synthesis constraints may be used to:

- Prevent BUFG inference.
- Replace a BUFG with an alternative clocking structure.
- Specify a clock buffer where one would not exist otherwise.

Using synthesis constraints allows this type of control without requiring any modification to the code.

Attributes can be placed in either of the following locations:

- Directly in the HDL code, which allows them to persist in the code
- As constraints in the XDC file, which allows this control without any changes needed to the source HDL code

For more information on constraints and attributes, see [Parameters, Attributes, and Constraints](#).

Use of IP

Certain IP assists in the creation of the clocking structures. Clocking Wizard and I/O Wizard specifically can assist in the selection and creation of the clocking resources and structure, including:

- BUFG
- BUFGCE
- BUFGCE_DIV (UltraScale devices)
- BUFGCTRL
- BUFIO (7 series devices)
- BUFR (7 series devices)
- Clock modifying blocks such as:
 - Mixed Mode Clocking Manager (MMCM)
 - Phase Lock Loop (PLL) components

More complex IP such as memory Interface Generator (MIG), PCIe, or Transceiver Wizard may also include clocking structures as part of the overall IP. This may provide additional clocking resources if properly taken into account. If not taken into account, it may limit some clocking options for the remainder of the design.

Xilinx highly recommends that, for any instantiated IP, the clocking requirements, capabilities and resources are well understood and leveraged where possible in other portions of the design.

For more information, see [Working With Intellectual Property \(IP\)](#).

Instantiation

The most low-level and direct method of controlling clocking structures is to instantiate the desired clocking resources into the HDL design. This allows you to access all possible capabilities of the device and exercise absolute control over them. When using BUFGCE, BUFGMUX, BUFHCE, or other clocking structure that requires extra logic and control, instantiation is generally the only option. However, even for simple buffers, sometimes the quickest way to obtain a desired result is to be direct and instantiate it into your design.

An effective style to manage clocking resources (especially when instantiating) is to contain the clocking resources in a separate entity or module instantiated at the top or near the top of the code. By having it at the top-level of code, it may more easily be distributed to multiple modules in your design.

Be aware of where clocking resources can and should be shared. Creating redundant clocking resources is not only a waste of resources, but generally consume more power, create more potential conflicts and placement decisions resulting in longer overall implementation tool runtimes and potentially more complex timing situations. This is another reason why having the clocking resources near the top module is important.



TIP: You can use Vivado HDL templates to instantiate specific clocking primitives. See [Using Vivado Design Suite HDL Templates](#).

Controlling the Phase, Frequency, Duty-Cycle, and Jitter of the Clock

This section explains some fine-grained tuning to the clock characteristics:

- [Using Clock Modifying Blocks \(MMCM and PLL\)](#)
- [Using IDELAYs on Clocks to Control Phase](#)
- [Using Gated Clocks](#)

Using Clock Modifying Blocks (MMCM and PLL)

You can use an MMCM or PLL to change the overall characteristics of an incoming clock.

An MMCM is most commonly used to remove the insertion delay of the clock (phase align the clock to the incoming system synchronous data).

The MMCM can also be used to:

- Create tighter control of phase.
- Filter jitter in the clock.
- Change the clock frequency.
- Correct or change the clock duty cycle, thus giving tight control over an important aspect of your design.

Using an MMCM or PLL is fairly common for conditioning and controlling the clock characteristics.

In order to use the MMCM or PLL, several attributes must be coordinated to ensure that the MMCM is operating within specifications, and delivering the desired clocking characteristics on its output. For this reason, Xilinx highly recommends that you use the Clocking Wizard to properly configure this resource.

The MMCM or PLL can also be directly instantiated, allowing even greater control. However, be sure to use the proper settings. Incorrect settings on the MMCM or PLL may:

- Increase clock uncertainty due to increased jitter.
- Build incorrect phase relationships.
- Make timing more difficult.



IMPORTANT: *When using the Clocking Wizard to configure the MMCM or PLL, the Clocking Wizard by default attempts to configure the MMCM for low output jitter using reasonable power characteristics.*

Depending on your goals, however, the settings in the Clocking Wizard may be changed to:

- Further minimize jitter, and thus improve timing at the cost of higher power, or
- Go the other way to reduce power but increase output jitter.

While using MMCM or PLL, pay attention to the following:

- Do not leave any inputs floating. Relying on synthesis or other optimization tools to tie off the floating values is not recommended, since the values that they tie to might be different from what you desire.
- RST should be connected to the user logic, so that it can be asserted as described in the *7 Series FPGAs Clocking Resources User Guide* (UG472) [Ref 39]. Grounding of RST can cause problems if the clock is interrupted.
- LOCKED output should be used in the implementation of reset, for example, synchronous logic clocked by the clock coming out of the PLL should be held in reset until LOCKED is asserted. The LOCKED signal would need to be synchronized before getting used in a synchronous portion of the design.
- The need for BUFG in the feedback path is important only if the PLL/MMCM output clock needs to be phase aligned with the input reference clock.
- Confirm the connectivity between CLKFBIN and CLKFBOUT.



RECOMMENDED: *Explore the different settings within the Clocking Wizard to ensure that the most desirable configuration is created based on your overall design goals.*

Using IDELAYS on Clocks to Control Phase

For 7 series devices, if only minor phase adjustments are necessary, you can use IDELAY or ODELAY (instead of MMCM or PLL) to add additional delay. This increases the phase offset of the clock in relation to any associated data. When using UltraScale devices, you cannot use an IDELAY on an input clock source. Therefore, if phase manipulation is necessary, Xilinx recommends using an MMCM.

Using Gated Clocks

Xilinx FPGA devices include dedicated clock networks that can provide a large-fanout, low-skew clocking resource. Fine-grained clock gating techniques implied in the HDL code can disrupt the functionality and mapping to this dedicated resource. Therefore, when coding to directly target an FPGA device, Xilinx does not recommend that you code clock gating constructs into the clock path. Instead, control clocking by using coding techniques to infer clock enables in order to stop portions of the design, either for functionality or power reasons.

If the code already contains clock gating constructs, or if it is intended for a different technology that requires such coding styles, Xilinx recommends that you use a synthesis tool that can remap gates placed within the clock path to clock enables in the data path. Doing so allows for a better mapping to the clocking resources; and simplifies the timing analysis of the circuit for data entering and exiting the gated domain.

When larger portions of the clock network can be shut down for periods of time, the clock network can be enabled or disabled by using a BUFGCE or BUFGCTRL. Alternatively, when targeting UltraScale devices, the BUFGCE_DIV, and BUFG_GT can also be gated. For 7 series devices, the BUFHCE, BUFR, and BUFMRCE can also be used to gate the clock. When a clock may be slowed down during periods of time, these buffers can also be used with additional logic to periodically enable the clock net. Alternatively, you can use a BUFGMUX to switch the clock source from a faster clock signal to a slower clock.

Any of these techniques can effectively reduce dynamic power. However, depending on the requirements and clock topology, one technique may prove more effective than another. For example:

- A BUFR may work best if it is an externally generated clock (under 450 MHz) that is only needed to source up to three clock regions.
- For Virtex-7 devices, a BUFMRCE may be needed in addition in order to use this technique with more than one clock region (but only up to three vertically adjacent regions).
- A BUFHCE is better-suited for higher speed clocks that can be contained in a single clock region. While a BUFGCE may span the device (and is the most flexible), it may not be the best choice for the greatest power savings.

Creating an Output Clock

An effective way to forward a clock out of an FPGA device for clocking devices external to the FPGA device, is to use an ODDR component. By tying one of the inputs high and the other low, you can easily create a well controlled clock in terms of phase relationship and duty cycle. (for example, by holding D1 to 0 and the D2 pin to 1, you can achieve a 180 degree phase shift). By utilizing the set/reset and clock enable, you also have control over stopping the clock and holding it at a certain polarity for sustained amounts of time.

If further phase control is necessary for an external clock, an MMCM or PLL can be used with external feedback compensation and/or coarse or fine grained, fixed or variable phase compensation. This allows great control over clock phase and propagation times to other devices simplifying external timing requirements from the device.

Clock Resource Selection Summary

BUFG

- Use when a high-fanout clock must be provided to several clock regions throughout the device. If you see cascaded BUFG, assure yourself of the need for it, or, did a BUFG come in unintentionally?
- Use when it is not desired to instantiate or have any manual control of clocking.
- Use for very high fanout non-clock nets such as a global reset for medium to slower speed clocks where mixed polarities does not exist. Xilinx recommends limit this use to only two in any design.
- For SSI technology devices where clocks that must span more than one SLR, locate them in one of the center SLRs. This more evenly distributes the clocking net across the entire device thus minimizing skew.

BUFGCE

- Use to stop a large-fanout several-region clock domain.

BUFGMUX/BUFGCTRL

- Use to change clock frequencies or clock sources during the operation of your design.

BUFGCE_DIV (UltraScale devices only)

- Use to generate simple clock division.

BUFG_GT (UltraScale devices only)

- Use when sourced from a GT or sourced from dedicated reference clock of a GT.

BUFH (7 series devices only)

- Use for smaller clock domains of logic that can be contained within a single clock region.
- Use for very high-speed clocking domains
- Use in clock domains that are less likely to compete for clocking resources with BUFGs
- For SSI technology devices, Xilinx recommends generally use in the upper or lower SLRs in order to lessen the chances of competition for resources with the BUFGs placed into the center SLRs

BUFHCE (7 series devices only)

- Use for clock-gating on a medium grained portion of the clock network which can be placed into a single clock region, The BUFHCE may be driven by BUFG.
- Use for high-fanout non-clock signals such as a reset that can be contained within a single clock region.

BUFR (7 series devices only)

- Use for small to medium sized clock networks that do not require performance higher than 450 MHz.
- Use for externally provided clocks that can be constrained within up to three vertically adjacent clock region that require clock division.
- For SSI technology devices, Xilinx recommends generally use in the upper or lower SLRs in order to lessen the chances of competition for resources with the BUFGs placed into the center SLRs.

BUFIO (7 series devices only)

- Use for externally provided high-speed I/O clocking generally in source synchronous data capture.

BUFMR (7 series devices only)

- Use when you need to use BUFRs or BUFIOs in more than one vertically adjacent clock regions for a single clock source.
- For SSI technology devices, Xilinx recommends locating the BUFMR and associated pin into the center clock region within an SLR. This allows access to all three clock regions from the BUFMR in case needed.

BUFMRCE (7 series devices only)

- Use when you need to use BUFRs or BUFIOs in more than one vertically adjacent clock regions for a single clock source where the clock is desired to be periodically stopped.
- If using more than one BUFR where clock division is used. The BUFMRCE can be used to ensure proper phase startup of all connected BUFRs

PLL and MMCM

- Use to remove the clock insertion delay (phase align the clock to the incoming data) for system synchronous inputs and outputs.
- Use for clock phase control to align source synchronous data to the clock for proper data capture.
- Use to change the clock frequency or duty cycle of an incoming clock more than a simple division.
- Use to filter clock jitter.

PLL provides a better control of jitter, while MMCM can provide a wider range of output frequencies. For tighter timing requirement, PLLs might be best, provided they can provide the frequency of interest.

IDELAY / IODELAY

- When targeting 7 series devices, use on an input clock to add small amounts of additional phase offset (delay).
- Use on input data to add additional delay to data thus effectively reducing clock phase offset in relation to the data.

ODDR

- Use to create an external forwarded clock from the device.

Special Clocking Considerations for SSI Devices

In addition to all the considerations previously mentioned in this section, you should consider the following for the design of the clocking structure for SSI technology:

- If clocks must span more than one SLR, locate BUFGs in one of the center SLRs to minimize skew.
- For 7 series devices, generally use BUFHs and BUFRs in the upper or lower SLRs to reduce the chance of resource competition with the BUFGs placed in the center SLRs.

For 7 series devices, locate the BUFMR and associated pin in the center clock region within an SLR. This placement allows access to all three clock regions from the BUFMR if needed.

Working With Intellectual Property (IP)

Pre-validated Intellectual Property (IP) cores significantly reduce design and validation efforts, and ensure a large advantage in time-to-market. See the following resources for more information on working with IP:

- *Vivado Design Suite User Guide: Designing with IP* (UG896) [Ref 9]
- [Vivado Design Suite QuickTake Video: Configuring and Managing Reusable IP in Vivado](#)

Planning IP Requirements

Planning IP requirements is one of the most important stages of any new project.

Evaluate the IP options available from Xilinx or third-party partners against required functionality and other design goals. Ask yourself:

- Is custom logic more desirable compared to an available IP core?
- Does it make sense to package a custom design for reuse in multiple projects in an industry standard format?

Consider the interfaces that are required such as, memory, network, and peripherals.

AMBA AXI

Xilinx has standardized IP interfaces on the open AMBA[®] 4 AXI4 interconnect protocol. This standardization eases integration of IP from Xilinx and third-party providers, and maximizes system performance. Xilinx worked with ARM to define the AXI4, AXI4-Lite, and AXI4-Stream specifications for efficient mapping into its FPGA device architectures.

AXI is targeted at high performance, high clock frequency system designs, and is suitable for high-speed interconnects. AXI4-Lite is a light-weight version of AXI4, and is used mostly for accessing control and status registers.

AXI-Stream is used for unidirectional streaming of data from Master to Slave. This is typically used for DSP, Video and Communications applications.

Vivado Design Suite IP Catalog

The IP Catalog is a single location for Xilinx-supplied IP. In the IP Catalog, you can find IP cores for embedded systems, DSP, communication, interfaces, and more.

From the IP Catalog, you can explore the available IP cores, and view the Product Guide, Change Log, Product Web page, and Answer Records for any IP.

You can access and customize the cores in the IP Catalog through the GUI or Tcl shell. You can also use Tcl scripts to automate the customization of IP cores.

Custom IP

Xilinx uses the industry standard IP-XACT format for delivery of IP, and provides tools (IP Packager) to package custom IP. Accordingly, you can also add your own customized IP to the catalog and create IP repositories that can be shared in a team or across a company. IP from third-party providers can also be added to this catalog.

Selecting IP from the IP Catalog

All Xilinx and third-party vendor IP is categorized based on applications such as communications and networking; video and image processing; and automotive and industrial. Use this categorization to browse the catalog to see which IP is available for your area of interest.



VIDEO: For more on customizing, adding, and instantiating IP into a project using the IP Catalog, see [Vivado Design Suite QuickTake Video: Customizing and Instantiating IP](#).

A majority of the IP in the IP Catalog is free. However, some high value IP has an associated cost and requires a license. The IP Catalog informs you about whether or not the IP requires purchase, as well as the status of the license. To select an IP from the catalog, consider the following key features, based on your design requirements, and what the specific IP offers:

- Silicon Resources required by this IP (found in the respective IP Product Guide)
- Is this IP supported in the device and speed grade being considered (the selection of the IP often drives the speed grade decision)? If supported, what is the max achievable throughput and Fmax?
- External interface standards, needed for your design to talk to its companion chip on board:
 - Industry-standard interfaces such as Ethernet, PCIe[®] interfaces, etc.
 - Memory interfaces - number of memory interfaces, including their size and performance.
 - Xilinx proprietary interfaces such as Aurora

Note: You can also choose to design your own custom interface.
- On-chip bus protocol supported by the IP (Application interface)

- On-chip bus protocol, needed for interaction with the rest of your design. Examples:
 - AXI4
 - AXI4-Lite
 - AXI4-Stream
- If multiple protocols are involved, bridging IP cores might have to be chosen using infrastructure IP from the IP Catalog. Examples:
 - AXI-AHB bridge
 - AXI-AXI interconnect
 - AXI-PCIe bridge
 - AXI-PLB bridge

IP and I/O

IP that interacts with the external world must be associated with I/O pins. For this reason, Xilinx recommends that you consider the I/O assignments while choosing IP. These include:

- [Parallel Interface](#)
- [Serial Interface](#)
- [I/O Voltages and I/O Standards](#)

Parallel Interface

The number of available I/Os in the I/O bank determine which I/O bank to choose.

Serial Interface

- Low-Speed Serial Interface: The ISERDES/OSERDES that are part of the General IOB can be used.
- High-Speed Serial Interface: The low-power Gigabit transceivers (GTs) can be used.

I/O Voltages and I/O Standards

- If the I/O voltage is 1.8V, choose the I/O bank that supports Vccio of 1.8V.
- For low data rates, use single ended I/O standard such as LVCMOS.
- For high data rates, use differential I/O standard, such as:
 - LVDS
 - DIFF-SSTL
 - DIFF-HSTL

Example Decision Process for IP Selection and Customization

Consider a communication and networking system with the following requirements:

- [10-Port 10G Ethernet MAC Aggregation](#)
- [PCIe Interface for System Configuration](#)
- [External Memory Storage](#)

Based on the requirements and the available IP, you must now check for the key functional features of each IP to decide its suitability and customization for your design. This process allows you to select the right IP needed for your purposes.

10-Port 10G Ethernet MAC Aggregation

The IP supports an optional XGMII interface. If the system needs an XAUI or a 10G PCS/PMA as its external interface, you must choose the XGMII option. The XAUI or the 10G PCS/PMA IP from Xilinx supports the XGMII interface.

Because data transfer from the MAC happens through the AXI4-Stream interface, the system must be able to consume the data from the MAC. It must be then connected either to an AXI interconnect to talk to other IP cores, or terminate in a wrapper with a proprietary protocol.

The core can be configured either through an optional AXI-Lite interface or a simple read/write interface. If AXI-Lite is chosen, the system should have an AXI-Lite support internally.

PCIe Interface for System Configuration

In addition to the considerations mentioned above, you must be aware of the data rates requirements of the interface. For assistance in making the selection, see the following table.

Table 4-5: Data Rate Requirements by Device

	Artix®-7	Kintex®-7	Virtex®-7T	Virtex-7 XT	Virtex-7
GEN (integrated block)	Gen2	Gen2	Gen2	Gen3	Gen3
Width	X4	X8	X8	X8	X8
Number of Blocks	1	1	3-4	2-4	1-3
Serial Data Rate (Gb/s)	5	5	8	8	8

External Memory Storage

You need to be aware of the number of DDR3 memories to be supported in the system. For this design, the total storage data rate requirement is about 80 Gb/s effective bandwidth or 100 Gb/s raw bandwidth when taking the MC efficiency into account. This can be achieved by multiple ways as illustrated below:

- Single controller with 64-bit DDR3 @ 1600 Mbps
- Four controllers with 16-bit DDR3 @ 1600 Mbps

Because data transfer from the Memory controller happens through AXI4 interface, the system should be able to consume the data from the MC. You can use $\frac{1}{2}$ rate or $\frac{1}{4}$ rate interfaces. When using $\frac{1}{4}$ rate, the application data width of the MC is $8 \times \text{DDR3-width}$. For example, for 16-bit DDR3, the AXI-Stream has 128-bits data width. The AXI Interconnect IP may be useful to connect the slave memory controller IP to the master peripherals which are accessing the memory. In this case, data from the IP cores comes through AXI4-Stream interfaces, a DMA needs to be added to convert the AXI4-Stream to AXI4 for the controller.

Xilinx recommends using MIG (Memory Interface Generators) to generate your memory controllers, which also guide the selection of I/O banks.

Customizing IP

IP can be customized through the GUI or through Tcl scripts.

- [Using the Customization GUI](#)
- [Using a Tcl Script](#)

Using the Customization GUI

Using the graphical interface is the easiest way to find, research, and customize IP. Each IP is customized with its own set of tabs or pages. Related configuration options are grouped together. An example of a customization window is shown in the following figure. A unique customization of an IP can be created, which is represented in an XCI file. From this, the various output products of an IP can be created.

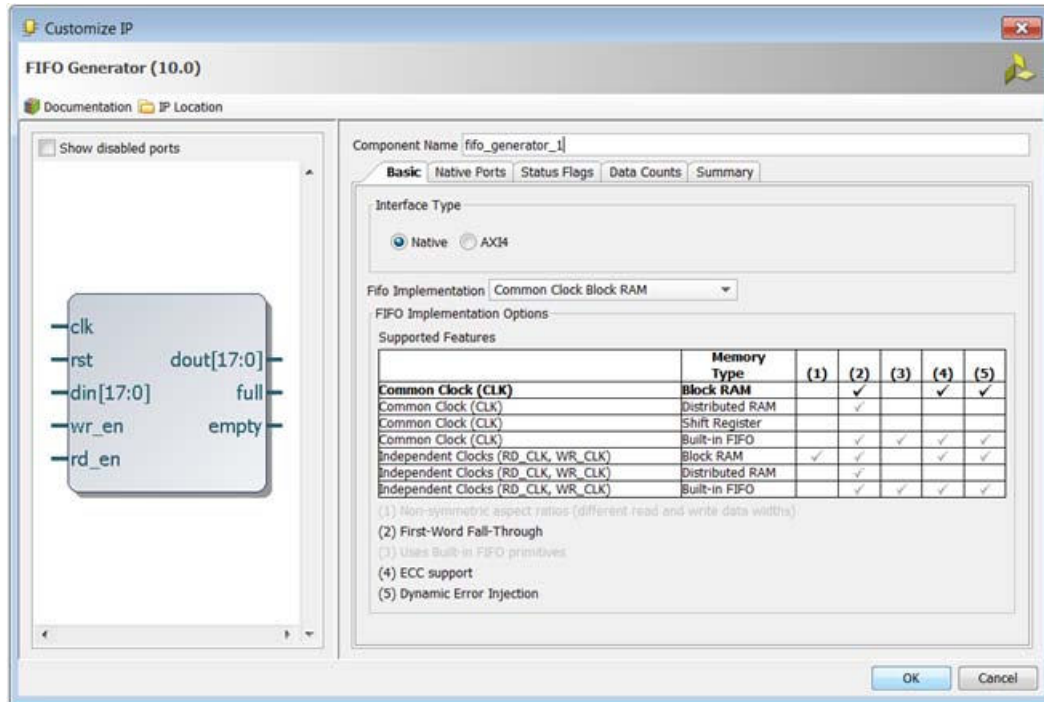


Figure 4-36: Customization Window for an IP

Using a Tcl Script

Almost every GUI action results in the issuance of a Tcl command. The creation of an IP including the setting of all the customization options can be performed in a Tcl script without user interaction.

You would need to know the names of the configuration options, and the values to which they can be set. Typically, you first perform the customization through the GUI, and then create the script from that. Once you see the resulting Tcl script, you can easily modify the script for your needs, such as changing data sizes.

Tcl script based IP creation is useful for automation, for example working with version control system. See [Source Management and Revision Control Recommendations in Chapter 2](#).

IP Output Products

When IP is customized, the tool creates an XCI file containing all the selected parameterization values. Each Vivado IDE version supports only one version of an IP. Xilinx recommends that you use this latest IP version. If you use an older IP version, you should have saved all the output products for the older version. For more information, see [Managing IP in Chapter 2](#).



TIP: For MIG, instead of an XCI file, a `.prj` file is created. All future references to XCI in the context of IP also means `.prj` for MIG.

Generated Output Products

Depending on the settings in the Vivado tools, by default some output products might be created automatically during IP customization. You can change this from the Project Settings. Using the XCI file created during customization, you can create any of the associated Output Products that an IP provides, including:

- **Instantiation template**

Shows how to instantiate the IP.

- **Synthesis**

Creates the script for running synthesis on this specific customized view of the IP.

- **DCP (out-of-context synthesis)**

Generates the synthesized netlist, which can be used directly in a higher level design, without having to resynthesize the IP part of the design.

- **Test bench**

Generates a test bench to be used for a simulation environment that can be used to verify the functionality of the IP.

- **Simulation**

Generates the script for running simulation on this specific customized view of the IP, using the test bench mentioned above.

The XCI file and these output products contain all that is required (for example HDL files and constraints files) to correctly simulate and synthesize the IP in a design.

Working with Constraints

Organizing the Design Constraints

Design constraints define the requirements that must be met by the compilation flow in order for the design to be functional in hardware. For more complex designs, they also define guidance for the tools to help with convergence and closure. Not all constraints are used by all steps in the compilation flow. For example, physical constraints are used only during the implementation steps (that is, by the placer and the router).

Because synthesis and implementation algorithms are timing-driven, creating proper timing constraints is essential. Over-constraining or under-constraining your design makes timing closure difficult. You must use reasonable constraints that correspond to your application requirements. For more information on constraints, see the following resources:

- *Vivado Design Suite User Guide: Design Analysis and Closure Techniques* (UG906) [Ref 21]
- Applying Design Constraints video tutorials available from the [Vivado Design Suite Video Tutorials](#) page on the Xilinx website

The constraints are usually organized by category, by design module, or both, in one or many files. Regardless of how you organize them, you must understand their overall dependencies and review their final sequence once loaded in memory. For example, because timing clocks must be defined before they can be used by any other constraints, you must make sure that their definition is located at the beginning of your constraint file, in the first set of constraint files loaded in memory, or both.

Recommended Constraint Files

There are many ways to organize your constraints depending on the size and complexity of your project. Following are a few suggestions.

Simple Design

For a simple design with a small team of designers:

- 1 file for all constraints
- 1 file for physical + 1 file for timing
- 1 file for physical + 1 file for timing (synthesis) + 1 file for timing (implementation)

Complex Design

For a complex design with IP cores or several designer teams:

- 1 file for top-level timing + 1 file for top-level physical + 1 file per IP/major block

Validating the Read Sequence

Once you have settled on the organization of your project constraint files, you must validate the read sequence of the files depending on the content of the files. In Project Mode, you can modify the constraint file sequence in the Vivado IDE or by using the `reorder_files` Tcl command. In Non-Project Mode, the sequence is directly defined by the `read_xdc` (for XDC files) and `source` (for constraints generated by Tcl scripts) commands in your compilation flow Tcl script.

Recommended Constraints Sequence

The constraints language (XDC) is based on Tcl syntax and interpretation rules. Like Tcl, XDC is a sequential language:

- Variables must be defined before they can be used. Similarly, timing clocks must be defined before they can be used in other constraints.
- For equivalent constraints that cover the same paths and have the same precedence, the last one applies.

When considering the priority rules above, the timing constraints should overall use the following sequence:

```
## Timing Assertions Section
# Primary clocks
# Virtual clocks
# Generated clocks
# Delay for external MMCM/PLL feedback loop
# Clock Uncertainty and Jitter
# Input and output delay constraints
# Clock Groups and Clock False Paths

## Timing Exceptions Section
# False Paths
# Max Delay / Min Delay
# Multicycle Paths
# Case Analysis
# Disable Timing
```

When multiple XDC files are used, you must pay particular attention to the clock definitions and validate that the dependencies are ordered correctly.

The physical constraints can be located anywhere in any constraint file.

Creating Synthesis Constraints

Synthesis takes the RTL description of the design and transforms it into an optimized technology mapped netlist by using timing-driven algorithms. The quality of the results is affected by the quality of the RTL code and the constraints provided. At this point of the compilation flow, the net delay modeling is approximate and does not reflect placement constraints or complex effects such as congestion. The main objective is to obtain a netlist which meets timing, or fails by a small amount, with realistic and simple constraints.

The synthesis engine accepts all XDC commands, but only some have a real effect:

- Timing constraints related to setup/recovery analysis influence the QoR:
 - `create_clock`
 - `create_generated_clock`
 - `set_input_delay` and `set_output_delay`
 - `set_clock_groups`
 - `set_false_path`
 - `set_max_delay`
 - `set_multicycle_path`
- Timing constraints related to hold and removal analysis are ignored during synthesis:
 - `set_false_path -hold`
 - `set_min_delay`
 - `set_multicycle_path -hold`
- RTL attributes forces decisions made by the mapping and optimization algorithms. Following are a few examples:
 - `DONT_TOUCH` / `KEEP` / `KEEP_HIERARCHY` / `MARK_DEBUG`
 - `MAX_FANOUT`
 - `RAM_STYLE` / `ROM_STYLE` / `USE_DSP48` / `SHREG_EXTRACT`
 - `FULL_CASE` / `PARALLEL_CASE` (Verilog RTL only)

The same attribute can also be set as a property from an XDC file. Using XDC-based constraints is convenient for influencing the synthesis results only in some cases without changing the RTL.

- Physical constraints are ignored (LOC, BEL, Pblocks)

Synthesis constraints must use names from the elaborated netlist, preferably ports and sequential cells. During elaboration, some RTL signals can disappear and it is not possible to attach XDC constraints to them. In addition, due to the various optimizations after elaboration, nets or logical cells are merged into the various technology primitives such as LUTs or DSP blocks. To know the elaborated names of your design objects, refer to [Using the Elaborated Design in Chapter 5](#).

Some registers are absorbed into RAM blocks and some levels of the hierarchy can disappear to allow cross-boundary optimizations.

Any elaborated netlist object or level of hierarchy can be preserved by using a DONT_TOUCH, KEEP, KEEP_HIERARCHY, or MARK_DEBUG constraint, at the risk of degrading timing or area QoR.

Finally, some constraints can conflict and cannot be respected by synthesis. For example, if a MAX_FANOUT attribute is set on a net that crosses multiple levels of hierarchy, and some hierarchies are preserved with DONT_TOUCH, the fanout optimization will be limited or fully prevented.



IMPORTANT: *Unlike during implementation, netlist objects that are used for defining timing constraints can be optimized away by synthesis to allow better QoR. This is usually not a problem as long as the constraints are updated and validated for implementation. But if needed, you can preserve any object by using the DONT_TOUCH constraint so that the constraints will apply during both synthesis and implementation.*

Once synthesis has completed, Xilinx recommends that you review the timing and utilization reports to validate that the netlist quality meets the application requirements and can be used for implementation.

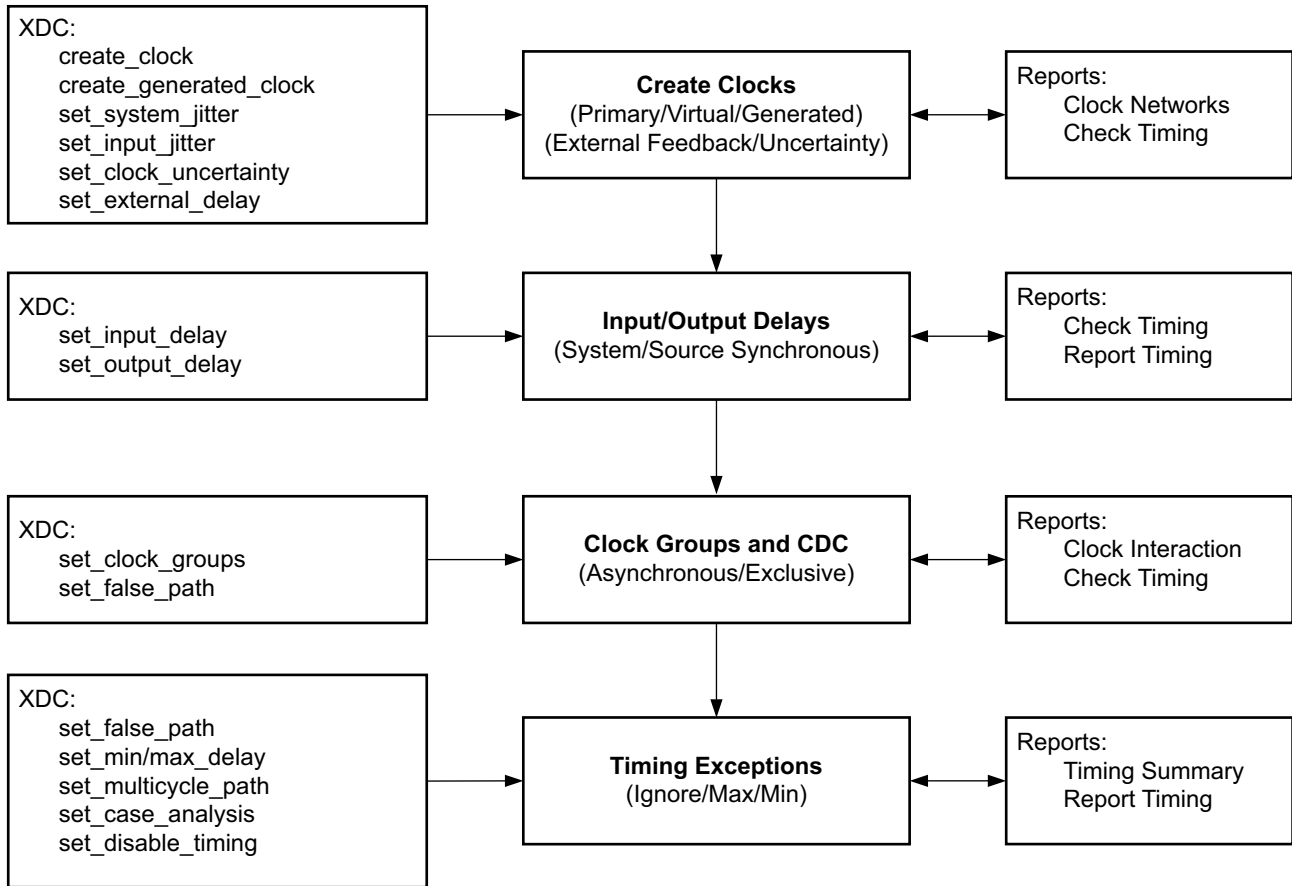
Creating Implementation Constraints

The implementation constraints must accurately reflect the requirements of the final application. Physical constraints such as I/O location and I/O standard are dictated by the board design, including the board trace delays, as well as the design internal requirements derived from the overall system requirements. Before you proceed to implementation, Xilinx highly recommends that you validate the correctness and accuracy of all your constraints. An improper constraint will likely contribute to degradation of the implementation QoR and can lower the confidence level in the timing signoff quality.

In many cases, the same constraints can be used during synthesis and implementation. However, because the design objects can disappear or have their name changed during synthesis, you must verify that all synthesis constraints still apply properly with the implementation netlist. If this is not the case, you must create an additional XDC file containing the constraints that are valid for implementation only.

Defining Timing Constraints in Four Steps

The process of defining good constraints is broken into the four major steps shown in the following figure. The steps follow the timing constraints precedence and dependency rules, as well as the logical way of providing information to the timing engine to perform the analysis.



X13445

Figure 4-37: Steps for Developing Timing Constraints

- The first two steps refer to the timing assertions where the default timing path requirements are derived from the clock waveforms and I/O delay constraints.
- During the third step, relationships between the asynchronous/exclusive clock domains that share at least one logical path are reviewed. Based on the nature of the relationships, clock groups or false path constraints are entered to ignore the timing analysis on these paths.
- The last step corresponds to the timing exceptions, where the designer can decide to alter the default timing path requirements by ignoring, relaxing or tightening them with specific constraints.

Constraints creation is associated with constraints identification and constraints validation tasks that are only possible with the various reports generated by the timing engine. The timing engine only works with a fully mapped netlist, for example, after synthesis. While it is possible to enter constraints with an elaborated netlist, it is recommended to create the first set of constraints with the post-synthesis netlist so that analysis and reports on the constraints can be performed interactively.

When creating timing constraints for a new design or completing existing constraints, Xilinx recommends using the Timing Constraints Wizard to quickly identify missing constraints for the first three steps in [Figure 4-37](#). The Timing Constraints Wizard follows the methodology described in this section to ensure the design constraints are safe and reliable for proper timing closure. You can find more information on the Timing Constraints Wizard in *Vivado Design Suite User Guide: Using Constraints* (UG903) [\[Ref 18\]](#).

The following sections describe in detail the four steps described above:

- [Defining Clock Constraints](#)
- [Constraining Input and Output Ports](#)
- [Defining Clock Groups and CDC Constraints](#)
- [Specifying Timing Exceptions](#)

Refer to each section for a detailed methodology and use case when you are at the appropriate step in the constraint creation process.

Defining Clock Constraints

Clocks must be defined first so that they can be used by other constraints. The first step of the timing constraint creation flow is to identify where the clocks must be defined and whether they must be defined as *primary clock* or a *generated clock*.



IMPORTANT: *When defining a clock with a specific name (-name option), you must verify that the clock name is not already used by another clock constraint or an existing auto-generated clock. The Vivado Design Suite timing engine issues a message when a clock name is used in several clock constraints to warn you that the first clock definition is overridden. When the same clock name is used twice, the first clock definition is lost as well as all constraints referring to that name and entered between the two clock definitions. Xilinx recommends that you avoid overriding clock definitions unless no other constraints are impacted and all timing paths remain constrained.*

Identifying Clock Sources

The unconstrained clock roots can be identified in the design by the following two reports:

- [Clock Networks Report](#)
- [Check Timing Report](#)

Clock Networks Report

Both constrained and unconstrained clock source points are listed in two separate categories. For each unconstrained source point, you must identify whether a primary clock or a generated clock must be defined.

```
% report_clock_networks

Unconstrained Clocks
Clock sysClk (endpoints: 15633 clock, 0 nonclock)
Port sysClk

Clock TXOUTCLK (endpoints: 148 clock, 0 nonclock)
GTXE2_CHANNEL/TXOUTCLK
(mgtEngine/ROCKETIO_WRAPPER_TILE_i/gt0_ROCKETIO_WRAPPER_TILE_i/gtxe2_i)

Clock Q (endpoints: 8 clock, 0 nonclock)
FDRE/Q (usbClkDiv2_reg)
```

Check Timing Report

The `no_clock` check reports the groups of active leaf clock pins with no clock definition. Each group is associated with a clock source point where a clock must be defined in order to clear the issue.

```
% check_timing -override_defaults no_clock

1. checking no_clock
-----
There are 15633 register/latch pins with no clock driven by root clock pin: sysClk
(HIGH)

There are 148 register/latch pins with no clock driven by root clock pin:
mgtEngine/ROCKETIO_WRAPPER_TILE_i/gt0_ROCKETIO_WRAPPER_TILE_i/gtxe2_i/TXOUTCLK
(HIGH)

There are 8 register/latch pins with no clock driven by root clock pin:
usbClkDiv2_reg/C (HIGH)
```

With `check_timing`, the same clock source pin or port can appear in several groups depending on the topology of the entire clock tree. In such case, creating a clock on the recommended source pin or port will resolve the missing clock definition for all the associated groups.

For more information, see [Checking That Your Design is Properly Constrained in Chapter 5](#).

Creating Primary Clocks

A primary clock is a clock that defines a timing reference for your design and that is used by the timing engine to derive the timing path requirements and the phase relationship with other clocks. Their insertion delay is calculated from the clock source point (driver pin/port where the clock is defined) to the clock pins of the sequential cells to which it fans out.

For this reason, it is important to define the primary clocks on objects that correspond to the boundary of the design, so that their delay, and indirectly their skew, can be accurately computed.

Typical primary clock roots are:

- [Input Ports](#)
- [Gigabit Transceiver Output Pins in 7 Series Devices](#)
- [Certain Hardware Primitive Output Pins](#)

Input Ports




Figure 4-38: **create_clock** for Input Ports

Constraint example:

```
create_clock -name SysClk -period 10 -waveform {0 5} [get_ports sysclk]
```

In this example, the waveform is defined to have a 50% duty cycle. The `-waveform` argument is shown above to illustrate its usage and is only necessary to define a clock with a duty cycle other than 50%. For a differential clock input buffer, the primary clock only needs to be defined on the P-side of the pair.

Gigabit Transceiver Output Pins in 7 Series Devices

Gigabit transceiver output pin, for example, a recovered clock:

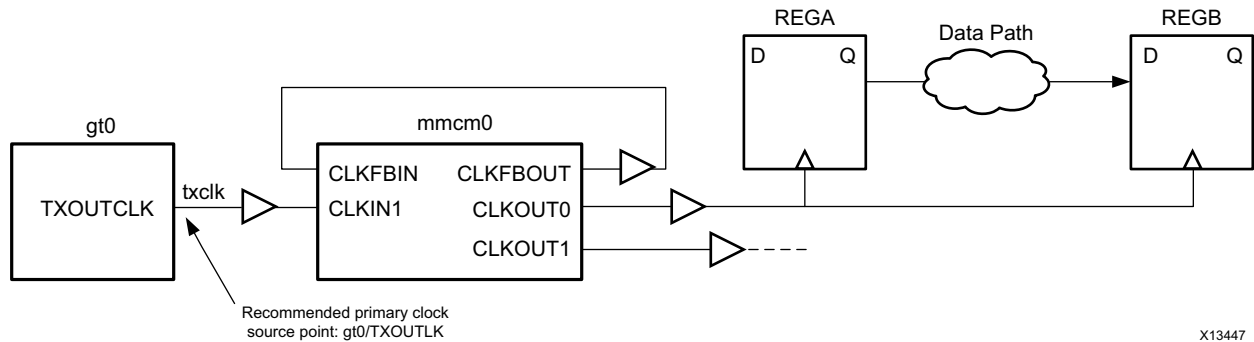


Figure 4-39: create_clock on a Primitive Pin

Constraint example:

```
create_clock -name txclk -period 6.667 [get_pins gt0/TXOUTCLK]
```

Note: In UltraScale devices, the Gigabit Transceiver clocks are automatically derived as long as the related board input clocks are defined. For this reason, Xilinx does not recommend defining a primary clock on the output of Gigabit Transceivers in designs that target UltraScale devices.

Certain Hardware Primitive Output Pins

The output pin of certain hardware primitives (for example, BSCANE2) which does not have a timing arc from an input pin of the same primitive.

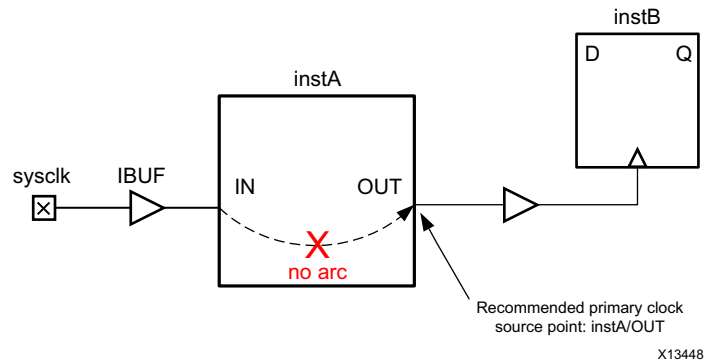


Figure 4-40: Clock Path Broken Due to a Missing Timing Arc



IMPORTANT: No primary clock should be defined in the transitive fanout of another primary clock because this situation does not correspond to any hardware reality. It will also prevent proper timing analysis by preventing the complete clock insertion delay calculation. Any time this situation occurs, the constraints must be revisited and corrected.

The following figure shows an example in which the clock `clk1` is defined in the transitive fanout of the clock `clk0`: `clk1` overrides `clk0` starting at the output of `BUFG1`, where it is defined. As a consequence, the timing analysis between `REGA` and `REGB` will not be accurate because of the invalid skew computation between `clk0` and `clk1`.

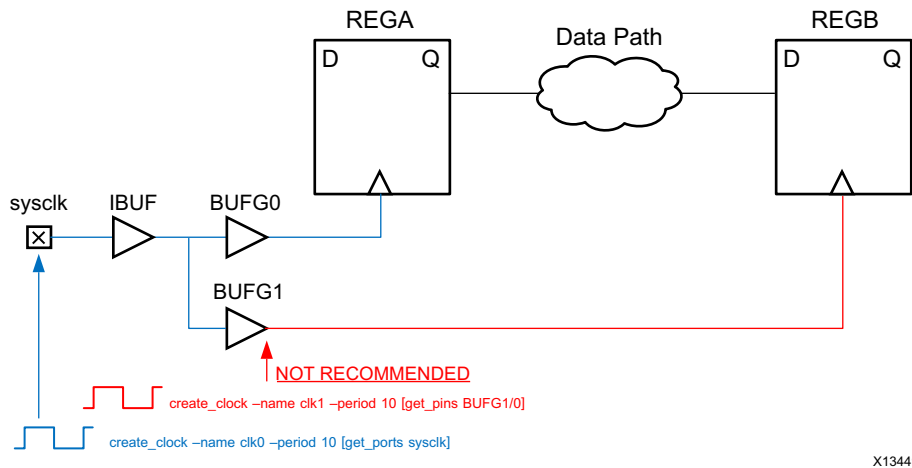


Figure 4-41: **create_clock in the Fanout of Another Clock is not Recommended**

Creating Generated Clocks

A generated clock is a clock derived from another existing clock called the master clock. It usually describes a waveform transformation performed on the master clock by a logic block. Because the generated clock definition depends on the master clock characteristics, the master clock must be defined first. For explicitly defining a generated clock, the `create_generated_clock` command must be used.

Auto-Derived Clocks

Most generated clocks are automatically derived by the Vivado Design Suite timing engine which recognizes the Clock Modifying Blocks (CMB) and the transformation they perform on the master clocks.

In the Xilinx 7 series device family, the CMBs are:

- MMCM*/ PLL*
- BUFR
- PHASER*

In the Xilinx UltraScale device family, the CMBs are:

- MMCM* / PLL*
- BUFG_GT / BUFGCE_DIV
- GT*_COMMON / GT*_CHANNEL / IBUFDS_GTE3
- BITSlice_CONTROL / RX*_BITSlice
- ISERDESE3

For any other combinatorial cell located on the clock tree, the timing clocks propagate through them and do not need to be redefined at their output, unless the waveform is transformed by the cell. In general, you must rely on the auto-derivation mechanism as much as possible as it provides the safest way to define the generated clocks that correspond to the actual hardware behavior.

If the auto-derived clock name chosen by the Vivado Design Suite timing engine does not seem appropriate, you can force your own name by using the `create_generated_clock` command without specifying the waveform transformation. This constraint should be located right after the constraint that defines the master clock in the constraint file. For example, if the default name of a clock generated by a MMCM instance is `net0`, you can add the following constraint to force your own name (`fftClk` in the given example):

```
create_generated_clock -name fftClk [get_pins mmc_m_i/CLKOUT0]
```

To avoid any ambiguity, the constraint must be attached to the source pin of the clock. For more information, see *Vivado Design Suite User Guide: Using Constraints* (UG903) [Ref 18].

User-Defined Generated Clocks

Once all the primary clocks have been defined, you can use the Clock Networks or Check Timing (`no_clock`) reports to identify the clock tree portions that do not have a timing clock and define the generated clocks accordingly.

It is sometimes difficult to understand the transformation performed by a cone of logic on the master clock. In this case, you must adopt the most conservative constraint. For example, the source pin is a sequential cell output. The master clock is at least divided by two, so the proper constraint should be, for example:

```
create_generated_clock -name clkDiv2 -divide_by 2 \
-source [get_pins fd/C] [get_pins fd/Q]
```

Finally, if the design contains latches, the latch gate pins also need to be reached by a timing clock and will be reported by Check Timing (`no_clock`) if the constraint is missing. You can follow the examples above to define these clocks.

Path Between Master and Generated Clocks

Unlike primary clocks, generated clocks must be defined in the transitive fanout of their master clock, so that the timing engine can accurately compute their insertion delay. Failure to follow this rule will result in improper timing analysis and most likely in invalid slack computation. For example, in the following figure `gen_clk_reg/Q` is being used as a clock for the next flop (`q_reg`), and it is also in the fanout cone of the primary clock `c1`. Hence `gen_clk_reg/Q` should have a `create_generated_clock` on it, rather than a `create_clock`.

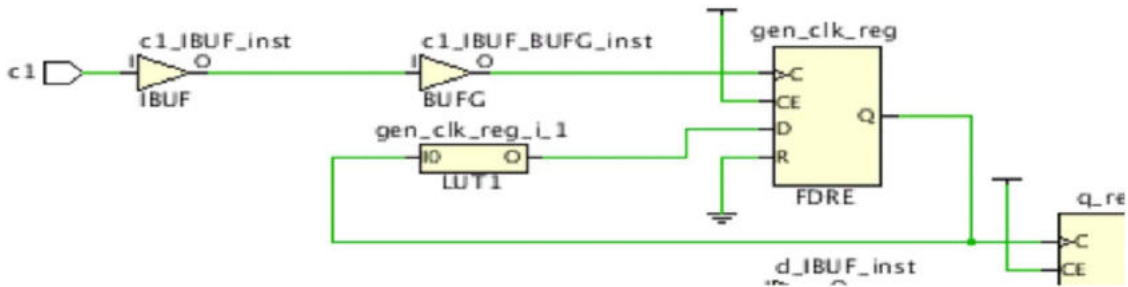


Figure 4-42: Generated Clock in the Fanout Of Master Clock

```
create_generated_clock -name GC1 -source [get_pins gen_clk_reg/C] -divide_by 2
[get_pins gen_clk_reg/Q]
```

Verifying Clocks Definition and Coverage

Once all design clocks are defined and applied in memory, you can verify the waveform of each clock, the relationship between master and generated clocks by using the `report_clocks` command:

```
Clock      Period  Waveform      Attributes  Sources
sysClk    10.00000 {0.00000 5.00000} P           {sysClk}
clkfbout  10.00000 {0.00000 5.00000} P,G        {clkgen/mmcm_adv_inst/CLKFBOUT}
cpuClk    20.00000 {0.00000 10.00000} P,G        {clkgen/mmcm_adv_inst/CLKOUT0}
...
=====
Generated Clocks
=====

Generated Clock   : cpuClk
Master Source    : clkgen/mmcm_adv_inst/CLKIN1
Master Clock     : sysClk
Edges            : {1 2 3}
Edge Shifts     : {0.000 5.000 10.000}
Generated Sources : {clkgen/mmcm_adv_inst/CLKOUT0}
```

You can also verify that all internal timing paths are covered by at least one clock. The Check Timing report provides two checks for that purpose:

- **no_clock**

Reports any active clock pin that is not reached by a defined clock.

- **unconstrained_internal_endpoint**

Reports all the data input pins of sequential cells that have a timing check relative to a clock but the clock has not been defined.

If both checks return zero, the timing analysis coverage will be high.

You can also run the Methodology XDC and Timing DRCs related to clock constraints to verify that all clocks are defined on recommended netlist objects without introducing any constraint conflict or inaccurate timing analysis scenario. To run all Methodology DRCs or only the XDC and Timing DRCs, use the following commands:

```
report_drc -ruledecks methodology_checks
```

or

```
report_drc -checks [get_drc_checks {XDC-* TIMING-*}]
```

For more information, see [Running Methodology DRCs in Chapter 5](#).

Adjusting Clock Characteristics

After defining the clocks and their waveform, the next step is to enter any information related to noise or uncertainty modeling. The XDC language differentiates uncertainty related to jitter and phase error from the one related to skew and delay modeling.

- [Jitter](#)
- [Additional Uncertainty](#)
- [Clock Latency at the Source](#)
- [MMCM or PLL External Feedback Loop Delay](#)

Jitter

For jitter, it is best to use the default values used by the Vivado Design Suite. You can modify the default computation as follows:

- If a primary clock enters the device with a random jitter greater than zero, use the `set_input_jitter` command to specify the jitter value.
- To adjust the global jitter if the device power supply is noisy, use `set_system_jitter`. Xilinx does *not* recommend increasing the default system jitter value.

For generated clocks, the jitter is derived from the master clock and the characteristics of the clock modifying block. The user does not need to adjust these numbers.

Additional Uncertainty

When you need to add extra margin on the timing paths of a clock or between two clocks, you must use the `set_clock_uncertainty` command. This is also the best and safest way to over-constrain a portion of a design without modifying the actual clock edges and the overall clocks relationships. The clock uncertainty defined by the user is additive to the jitter computed by the Vivado tools, and can be specified separately for setup and hold analyses.

For example, the margin on all intra-clock paths of the design clock `clk0` needs to be tightened by 500ps to make the design more robust to noise for both setup and hold:

```
set_clock_uncertainty -from clk0 -to clk0 0.500
```

If you specify additional uncertainty between two clocks, the constraint must be applied in both directions (assuming data flows in both directions). The example below shows how to increase the uncertainty by 250ps between `clk0` and `clk1` for setup only:

```
set_clock_uncertainty -from clk0 -to clk1 0.250 -setup
set_clock_uncertainty -from clk1 -to clk0 0.250 -setup
```

Clock Latency at the Source

It is possible to model the latency of a clock at its source by using the `set_clock_latency` command with the `-source` option. This is useful in two cases:

- To specify the clock delay propagation outside the device independently from the input and output delay constraints.
- To model the internal propagation latency of a clock used by a block during out-of-context compilation. In such a compilation flow, the complete clock tree is not described, so the variation between min and max operating conditions outside the block cannot be automatically computed and must be manually modeled.

This constraint should only be used by advanced users as it is usually difficult to provide valid latency values.

MMCM or PLL External Feedback Loop Delay

When the MMCM or PLL feedback loop is connected for compensating a board delay instead of an internal clock insertion delay, you must specify the delay outside the FPGA device for both best and worst delay cases by using the `set_external_delay` command. Failure to specify this delay will make I/O timing analysis associated with the MMCM or PLL irrelevant and can potentially lead to an impossible timing closure situation. Also, when using external compensation, you must adjust the input and output delay constraint values accordingly instead of just considering the clock trace delay on the board like in normal cases.

Constraining Input and Output Ports

In addition to specifying the location and I/O standard for each port of the design, input and output delay constraints must be specified to describe the timing of external paths to/from the interface of the FPGA device. These delays are defined relatively to a clock that is usually also generated on the board and enters the FPGA device. In some cases, the delays must be defined related to a virtual clock when the I/O path is related to a clock that has a waveform different from the board clock.

System Level Perspective

The I/O paths are modeled like any other `reg-to-reg` paths by the Vivado Design Suite timing engine, except that part of the path is located outside the FPGA device and needs to be described by the user. When analyzing internal paths, minimum and maximum delays are considered for both setup and hold analysis. This is also true for I/O paths. For this reason, it is important to describe both min and max delay conditions. The I/O timing paths are analyzed as single-cycle paths by default, which means:

- For max delay analysis (setup), the data is captured one clock cycle after the launch edge for single data rate interface, and half clock cycle after the launch edge for a double data rate interface.
- For min delay analysis (hold), the data is launched and captured by the same clock edge.

If the relationship between the clock and I/O data must be timed differently, like for example in a source synchronous interface, different I/O delays and additional timing exceptions must be specified. This corresponds to an advanced I/O timing constraints scenario.

Defining Input Delays

The input delay is defined relative to a clock at the interface of the device. Unless `set_clock_latency` has been specified on the source pin of the reference clock, the input delay corresponds to the absolute time from the launch edge, through the clock trace, the external device and the data trace. If clock latency has already been specified separately, you can ignore the clock trace delay.

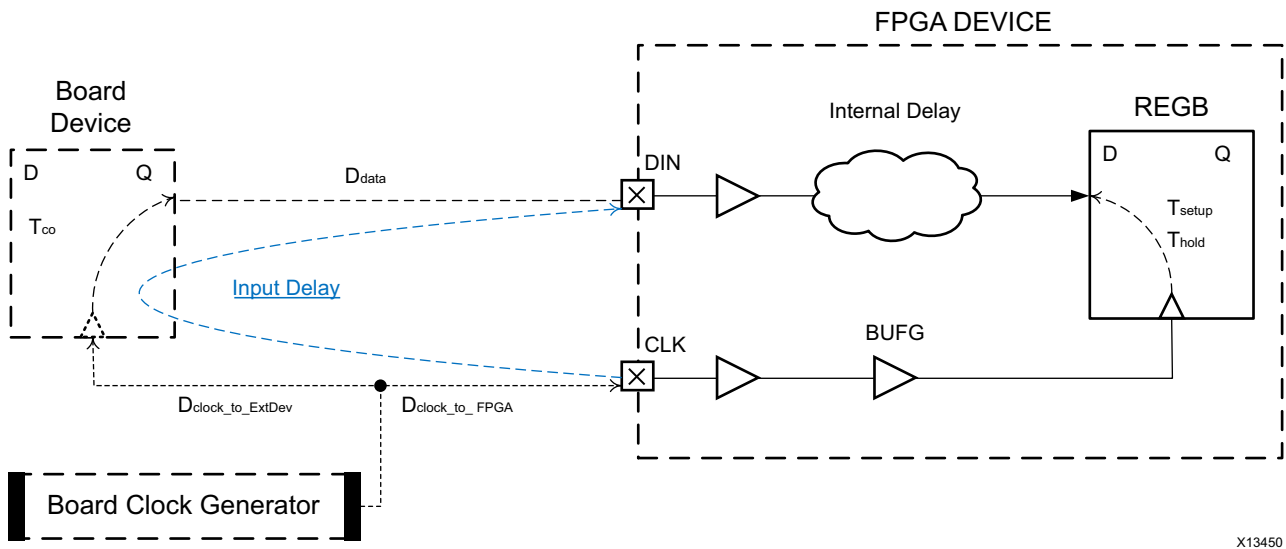


Figure 4-43: Input Delay Computation

The input delay values for the both types of analysis are:

$$\begin{aligned} \text{Input Delay}(\max) &= T_{co}(\max) + D_{data}(\max) + D_{clock_to_ExtDev}(\max) - D_{clock_to_FPGA}(\min) \\ \text{Input Delay}(\min) &= T_{co}(\min) + D_{data}(\min) + D_{clock_to_ExtDev}(\min) - D_{clock_to_FPGA}(\max) \end{aligned}$$

The following figure shows a simple example of input delay constraints for both setup (max) and hold (min) analysis, assuming the `sysClk` clock has already been defined on the CLK port:

```
set_input_delay -max -clock sysClk 5.4 [get_ports DIN]
set_input_delay -min -clock sysClk 2.1 [get_ports DIN]
```




Figure 4-44: Interpreting Min and Max Input Delays

A negative input delay means that the data arrives at the interface of the device before the launch clock edge.

Defining Output Delays

Output delays are similar to input delays, except that they refer to the output path minimum and maximum time outside the FPGA device in order to be functional under all conditions.

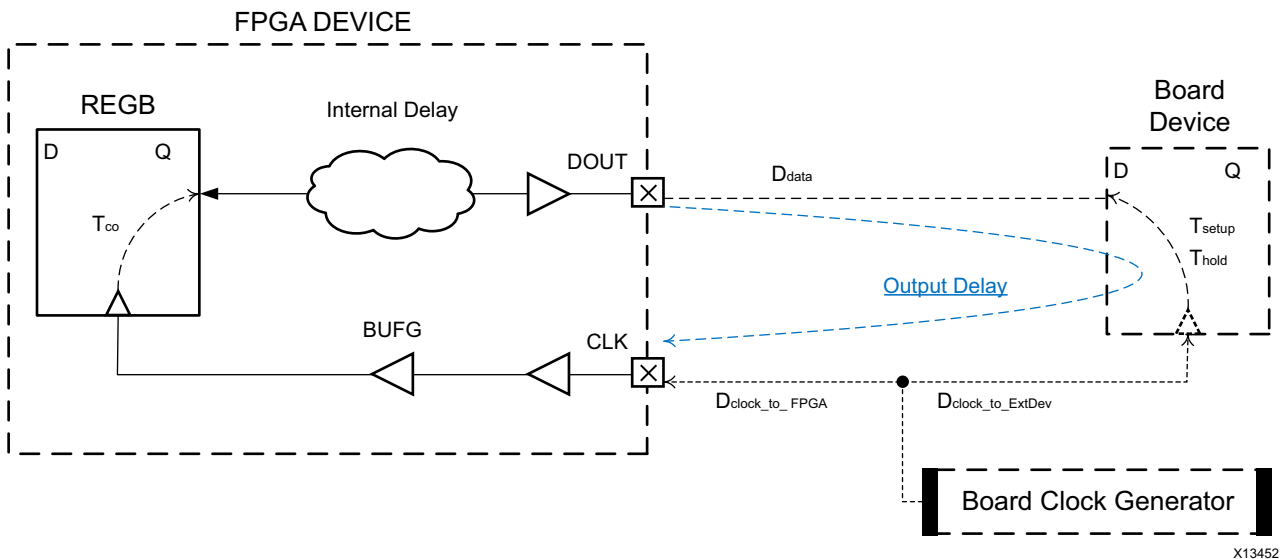


Figure 4-45: Output Delay Computation

The output delay values for the both types of analysis are:

$$\begin{aligned} \text{Output Delay(max)} &= T_{setup} + D_{data(max)} + D_{clock_to_FPGA(max)} - D_{clock_to_ExtDev(min)} \\ \text{Output Delay(min)} &= D_{data(min)} - T_{hold} + D_{clock_to_FPGA(min)} - D_{clock_to_ExtDev(max)} \end{aligned}$$

The following figure shows a simple example of output delay constraints for both setup (max) and hold (min) analyses, assuming the sysClk clock has already been defined on the CLK port:

```
set_output_delay -max -clock sysClk 2.4 [get_ports DOUT]
set_output_delay -min -clock sysClk -1.1 [get_ports DOUT]
```

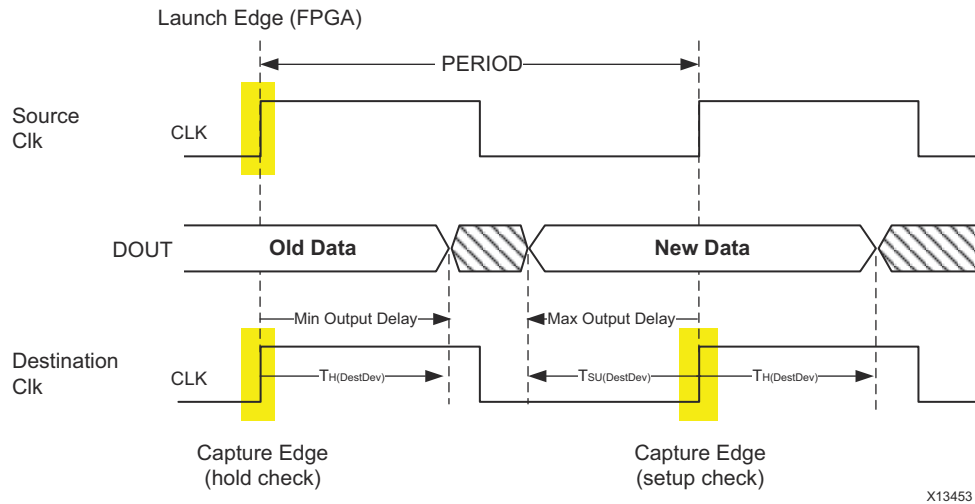


Figure 4-46: Interpreting Min and Max Output Delays

The output delay corresponds to the delay on the board before the capture edge. For a regular system synchronous interface where the clock and data board traces are balanced, the setup time of the destination device defines the output delay value for max analysis. And the destination device hold time defines the output delay for min analysis. The specified min output delay indicates the minimum delay that the signal will incur after coming out of the design, before it will be used for hold analysis at the destination device interface. Thus, the delay inside the block can be that much smaller. A positive value for min output delay means that the signal can have negative delay inside the design. This is why min output delay is often negative. For example:

```
set_output_delay -min -0.5 -clock CLK [get_ports DOUT]
```

means that the delay inside the design until DOUT has to be at least +0.5 ns to meet the hold time requirement.

Choosing the Reference Clock

Depending on the clock tree topology that controls the sequential cells related to input or output ports, you have to choose the most appropriate clock to define the input or output delay constraints. If the clock of the I/O path register is a generated clock, the delay constraint usually needs to be defined relative to the primary clock, which is defined upstream of the generated clocks. There are some exceptions to this rule that are explained in this section.

Identifying the Clocks Related to Each Port

Before defining the I/O delay constraint, you must identify which clocks are related to each port. There are several ways to identify those clocks:

- [Browse the Board Schematics](#)
- [Browse the Design Schematics](#)
- [Report Timing from or to the Port](#)
- [Using Automatically Identified Sampling Clocks](#)

Browse the Board Schematics

For a group of ports connected to a particular device on the board, you can use the same board clock that goes to both the device and the FPGA as the input or output delay reference clock. You need to verify in the device data sheet if the board clock is internally transformed for timing the I/O ports to make sure the FPGA design generates the same clock to control the timing of the related group of ports.

Browse the Design Schematics

For each port, you can expand the path schematics to the first level of sequential cells, and then trace the clock pins of those cells back to the clock source(s). This approach can be impractical for ports that are connected to high fanout nets.

Report Timing from or to the Port

Whether a port is already constrained or not, you can use the `report_timing` command to identify its related clocks in the design. Once all the timing clocks have been defined, you can report the worst path from or to the I/O port, create the I/O delay constraint relative to the clock reported, and re-run the same timing report from/to the other clocks of the design. If it appears that the port is related to more than one clock, create the corresponding constraint and repeat the process.

For example, the `din` input port is related to the clocks `clk0` and `clk1` inside the design:

```
report_timing -from [get_ports din] -sort_by group
```

The report shows that the `din` port is related to `clk0`. The input delay constraint is (for both min and max delay in this example):

```
set_input_delay -clock clk0 5 [get_ports din]
```

Re-run timing analysis with the same command as previously, and observe that `din` is also related to `clk1` thanks to the `-sort_by group` option which reports N paths per endpoint clock. You can add the corresponding delay constraint and re-run the report to validate that the `din` port is not related to another clock.

The same analysis can be done with the Timing Summary report, by looking at the Unconstrained Paths section. With only clock constraints in your design, this section appears as follows:

```

-----
| Unconstrained Path Table
-----
Path Group      From Clock      To Clock
-----
(none)          clk0
(none)                               clk0
(none)                               clk1
  
```

The fields without a clock name (or <NONE> in the Vivado IDE) refer to a group of paths where the startpoints (From Clock) or the endpoints (To Clock) are not associated with a clock. The unconstrained I/O ports fall in this category. You can retrieve their name by browsing the rest of the report. For example in the Vivado IDE, by selecting the Setup paths for the clk0 to NONE category, you can see the ports driven by clk0 in the To column:

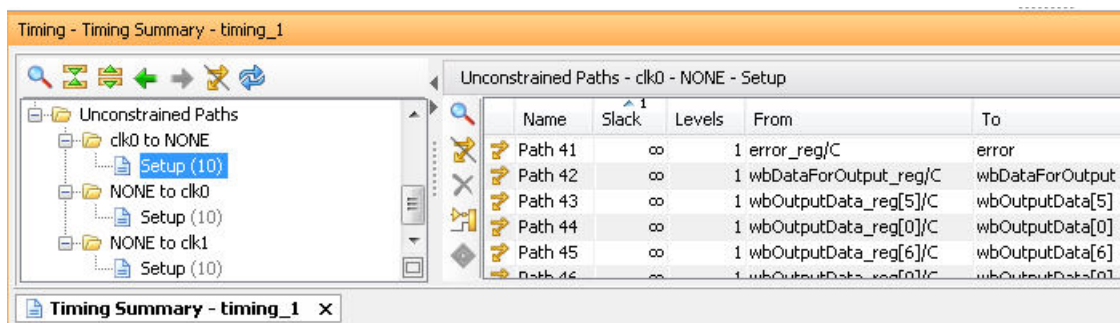


Figure 4-47: Getting a List of Unconstrained Output Ports

After adding the new constraints and applying them in memory, you must re-run the report to determine which ports are still unconstrained. For most designs, you must increase the number of reported paths to make sure all the I/O paths are listed in the report.

Using Automatically Identified Sampling Clocks

You can use the `set_input_delay` and `set_output_delay` constraints without specifying the related clock. The Vivado Design Suite timing engine will analyze the design and associate each port with all the sampling clocks automatically. Then by reporting timing on the I/O paths, you can see how the tool constrained each I/O port. This is convenient for quickly constraining a design, but this type of generic constraints can become a problem if they are too generic and do not model the hardware reality accurately.

Using a Primary Clock

A primary clock (that is, an incoming board clock) should be used when it directly controls the I/O path sequential cells, without traversing any clock modifying block. I/O delay lines are not considered as clock modifying blocks because they only affect the clock insertion

delay and not the waveform. This case is illustrated by the two examples previously provided in Defining Input Delays and Defining Output Delays. Most of the time, the external device also has its interface characteristics defined with respect to the same board clock.

When the primary clock is compensated by a PLL or MMCM inside the FPGA with the zero hold violation (ZHOLD) mode, the I/O paths sequential cells are connected to an internal copy (for example, a generated clock) of the primary clock. Because the waveforms of both clocks are identical, Xilinx recommends using the primary clock as the reference clock for the input/output delay constraints.




Figure 4-48: Input Delay in the Presence of a ZHOLD MMCM in Clock Path

The constraints are identical to the example provided in Defining Input Delays because the ZHOLD MMCM acts like a clock buffer with a negative insertion delay which corresponds to the amount of compensation.

Using a Virtual Clock

When the board clock traverses a clock modifying block which transforms the waveform in addition to compensating the overall insertion delay, it is recommended to use a virtual clock as a reference clock for the input and output delay instead of the board clock. There are three main cases for using a virtual clock:

- The internal clock and the board clock have different period: The virtual clock must be defined with the same period and waveform as the internal clock. This results in a regular single-cycle path requirement on the I/O paths.
- For input paths, the internal clock has a positive shifted waveform compared to the board clock: the virtual clock is defined like the board clock, and a multicycle path constraint of 2 cycles for setup is defined from the virtual clock to the internal clock. These constraints force the setup timing analysis to be performed with a requirement of 1 clock cycle + amount of phase shift.
- For output paths, the internal clock has a negative shifted waveform compared to the board clock: the virtual clock is defined like the board clock and a multicycle path constraint of 2 cycles for setup is defined from the internal clock to the virtual clock. These constraints force the setup timing analysis to be performed with a requirement of 1 clock cycle + amount of phase shift.

To summarize, the use of a virtual clock adjusts the default timing analysis to avoid treating I/O paths as clock domain crossing paths with a very tight and unrealistic requirement.

For example, consider the sysClk board clock which runs at 100MHz and gets multiplied by an MMCM to generate clk266 which runs at 266MHz. An output which is generated by clk266 should use clk266 as the reference clock. If you try to use sysClk as the reference clock (for the set_output_delay specification), it will appear as asynchronous clocks, and the path can no longer be timed as a single cycle.

Using a Generated Clock

For an output source synchronous interface, the design generates a copy of the internal clock and forwards it to the board along with the data. This clock is usually used as the reference clock for the output data delay constraints whenever the intent is to control and report on the phase relationship (skew) between the forwarded clock and the data. The forwarded clock can also be used in input and output delay constraints for a system synchronous interface.

Rising and Falling Reference Clock Edges

The clock edges used in the I/O constraint must reflect the data sheet of the external device connected to the FPGA device. By default, the set_input_delay and set_output_delay commands define a delay constraint relative to the rising reference clock edge. You must use the clock_fall option to specify a delay relative the falling clock edge. You can also specify separate constraints for delays related to both rising and falling clock edges by using the add_delay option with the second constraint on a port.

In most cases, the I/O reference clock edges correspond to the clock edges used to latch or launch the I/O data inside the FPGA. By analyzing the I/O timing paths, you can review which clock edges are used and verify that they correspond to the actual hardware behavior. If by mistake a rising clock edge is used as a reference clock for an I/O path that is only related to the falling clock edge internally, the path requirement is ½-period, which makes timing closure more difficult.

Verifying Delay Constraints

Once the I/O timing constraints have been entered, it is important to review how timing is analyzed on the I/O paths and the amount of slack violation for both setup and hold checks. By using the timing reports from/to all ports for both setup and hold analysis (that is, delay type = min_max), you can verify that:

- The correct clocks and clock edges are used as reference for the delay constraints.
- The expected clocks are launching and capturing the I/O data inside the FPGA device.
- The violations can reasonably be fixed by placement or by setting the proper delay line tap configuration. If this is not the case, you must review the I/O delay values entered

in the constraints and evaluate whether they are realistic, and whether you must modify the design to meet timing.

I/O Path Report Command Lines Example

```
report_timing -from [all_inputs] -nworst 1000 -sort_by group \
-delay_type min_max

report_timing -to [all_outputs] -nworst 1000 -sort_by group \
-delay_type min_max
```

Improper I/O delay constraints can lead to impossible timing closure. The implementation tools are timing driven and work on optimizing the placement and routing to meet timing. If the I/O path requirements cannot be met and I/O paths have the worst violations in the design, the overall design QoR will be impacted.

Input to Output Feed-through Path

There are several equivalent ways to constrain a combinatorial path from an input port to an output port.

Example One

Use a virtual clock with a period greater or equal to the target maximum delay for the feed-through path, and apply max input and output delay constraints as follows:

```
create_clock -name vclk -period 10
set_input_delay -clock vclk <input_delay_val> [get_ports din] -max
set_output_delay -clock vclk <output_delay_val> [get_ports dout] -max
```

where

```
input_delay_val(max) + feedthrough path delay (max) + output_delay_val(max)
<= vclk period.
```

In this example, only the maximum delay is constrained.

Example Two

Use a combination of min and max delay constraints between the feedthrough ports. Example:

```
set_max_delay -from [get_ports din] -to [get_ports dout] 10
set_min_delay -from [get_ports din] -to [get_ports dout] 2
```

This is a simple way to constrain both minimum and maximum delays on the path. Any existing input and output delay constraints on the same ports are also used during the timing analysis. For this reason, this style is not very popular.

The max delay is usually optimized and reported against the Slow timing corner, while the min delay is in the Fast timing corner. It is best to run a few iterations on the feedthrough

path delay constraints to validate that they are reasonable and can be met by the implementation tools, especially if the ports are placed far from one another.

Using XDC Templates - Source Synchronous Interfaces

While most users can properly write timing constraints for system synchronous interfaces, Xilinx recommends using I/O constraint templates for the source synchronous interfaces. The source synchronous constraints can be written in several ways. The templates provided by the Vivado Design Suite are based on the default timing analysis path requirement. The syntax is simpler, but the delay values must be adjusted to account for the fact that the setup analysis is performed with different launch and capture edges (1-cycle or 1/2-cycle) instead of same edge (0-cycle). The timing reports can be more difficult to read as the clock edges do not directly correspond to the active ones in hardware. You can navigate to these templates in Vivado GUI through **Window > Language Templates > XDC > TimingConstraints > Input Delay Constraints > Source Synchronous**.

Defining Clock Groups and CDC Constraints

The Vivado IDE times the paths between all the clocks in your design by default. The `set_clock_groups` command disables timing analysis between groups of clocks that you identify, and not between the clocks within a same group. Unlike `set_clock_groups`, the `set_false_path` constraint ignores timing between the clocks only in the direction specified by the `from` and `to` options. In some specific cases, maximum delay constraints can be set on Clock Domain Crossing (CDC) paths in order to limit the latency of these paths, on one or several signals. If clock groups or false path constraints already exist between the clocks or on the same CDC paths, the maximum delay constraints will be ignored. For this reason, it is important to thoroughly review every path between all clock pairs before choosing one CDC timing constraint over another one in order to avoid constraints collision.



RECOMMENDED: *You should also run the `methodology_check` DRC rule deck to identify when a `set_max_delay -datapath_only` constraint is overridden by a `set_clock_groups` or `set_false_path` constraint. See [Running Methodology DRCs in Chapter 5](#).*

Reviewing Clock Interactions

Clocks that have a logical path between them are timed. The possible clock relationships are:

- Synchronous
- Asynchronous
- Exclusive

Synchronous

Clock relationships are synchronous when two clocks have a fixed phase relationship. This is the case when two clocks share the following:

- Common circuitry (common node)
- Primary clock (same initial phase)

Asynchronous

Clock relationships are asynchronous when they do not have a fixed phase relationship. This is the case when one of the following is true:

- They do not share any common circuitry in the design and do not have a common primary clock.
- They do not have a common period within 1000 cycles (unexpandable) and the timing engine cannot properly time them together.

If two clocks are synchronous but their common period is very small, the setup paths requirement is too tight for timing to be met. Xilinx recommends that you treat the two clocks as asynchronous and implement safe asynchronous CDC circuitry.

Exclusive

Clock relationships are exclusive when they propagate on a same clock tree and reach the same sequential cell clock pins but cannot physically be active at the same time.

Categorizing Clock Pairs

The clock pairs can be categorized by using the following reports:

- [Clock Interaction Report](#)
- [Check Timing Report](#)

Clock Interaction Report

The Clock Interaction report provides a high-level summary of how two clocks are timed together:

- Do the two clocks have a common primary clock? When clocks are properly defined, all clocks that originate from the same source in the design share the same primary clock.
- Do the two clocks have a common period? This shows in the setup or hold path requirement column ("unexpandable"), when the timing engine cannot determine the most pessimistic setup or hold relationship.
- Are the paths between the two clocks partially or completely covered by clock groups or timing exception constraints?

- Is the setup path requirement between the two clocks very tight? This can happen, when two clocks are synchronous, but their period is not specified as an exact multiple (for example, due to rounding off). Over multiple clock cycles, the edges could drift apart, causing the worst case timing requirement to be very tight.

Check Timing Report

The Check Timing report (`multiple_clock`) identifies the clock pins that are reached by more than one clock and a `set_clock_groups` or `set_false_path` constraint has not already been defined between these clocks.

Constraining Exclusive Clock Groups

You can use the regular timing or clock network reports to review the clock paths and identify the situations where two clocks propagate on a same clock tree and are used at the same time in a timing path where the startpoint and endpoint clock pins are connected to the same clock tree. This analysis can be a time consuming task. Instead, you can review the `multiple_clock` section of the Check Timing report. This section returns a list of clock pins and their associated timing clocks.

Based on the clock tree topology, you must apply different constraints:

- [Overlapping Clocks Defined on the Same Clock Source](#)
- [Overlapping Clocks Driven by a Clock Multiplexer](#)

Overlapping Clocks Defined on the Same Clock Source

This occurs when two clocks are defined on the same netlist object with the `create_clock -add` command and represent the multiple modes of an application. In this case, it is safe to apply a clock groups constraint between the clocks. For example:

```
create_clock -name clk_mode0 -period 10 [get_ports clk_in]
create_clock -name clk_mode1 -period 13.334 -add [get_ports clk_in]
set_clock_groups -physically_exclusive -group clk_mode0 -group clk_mode1
```

If the `clk_mode0` and `clk_mode1` clocks generate other clocks, the same constraint needs to be applied to their generated clocks as well, which can be done as follows:

```
set_clock_groups -physically_exclusive \
-group [get_clocks -include_generated_clock clk_mode0] \
-group [get_clocks -include_generated_clock clk_mode1]
```

Overlapping Clocks Driven by a Clock Multiplexer

When two or more clocks drive into a multiplexer (or more generally a combinatorial cell), they all propagate through and become overlapped on the fanout of the cell. Realistically, only one clock can propagate at a time, but timing analysis allows reporting several timing modes at the same time.

For this reason, you must review the CDC paths and add new constraints to ignore some of the clock relationships. The correct constraints are dictated by how and where the clocks interact in the design.

The following figure shows an example of two clocks driving into a multiplexer and the possible interactions between them before and after the multiplexer.

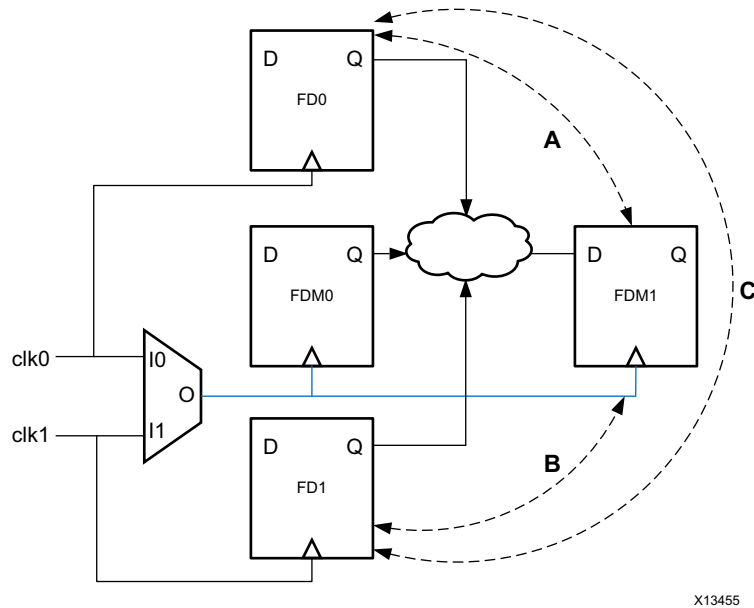


Figure 4-49: Muxed Clocks

- **Case in which the paths A, B, and C do not exist**

clk0 and clk1 only interact in the fanout of the multiplexer (FDM0 and FDM1). It is safe to apply the clock groups constraint to clk0 and clk1 directly.

```
set_clock_groups -logically_exclusive -group clk0 -group clk1
```

- **Case in which only the paths A or B or C exist**

clk0 and/or clk1 directly interact with the multiplexed clock. In order to keep timing paths A, B and C, the constraint cannot be applied to clk0 and clk1 directly. Instead, it must be applied to the portion of the clocks in the fanout of the multiplexer, which requires additional clock definitions.

```
create_generated_clock -name clk0mux -divide_by 1 \
-source [get_pins mux/I0] [get_pins mux/O]
```

```
create_generated_clock -name clk1mux -divide_by 1 \
-add -master_clock clk1 \
-source [get_pins mux/I1] [get_pins mux/O]
```

```
set_clock_groups -physically_exclusive -group clk0mux -group clk1mux
```

Constraining Asynchronous Clock Groups and Clock Domain Crossings

The asynchronous relationship can be quickly identified in the Clock Interaction report: clock pairs with no common primary clock or no common period (unexpanded). Even if clock periods are the same, the clocks will still be asynchronous, if they are being generated from different sources. The asynchronous Clock Domain Crossing (CDC) paths must be reviewed carefully to ensure that they use proper synchronization circuitry that does not rely on timing correctness and that minimizes the chance for metastability to occur. Asynchronous CDC paths usually have high skew and/or unrealistic path requirements. They should not be timed with the default timing analysis, which cannot prove they will be functional in hardware.

Report CDC

The Report CDC (`report_cdc`) command performs a structural analysis of the clock domain crossings in your design. You can use this information to identify potentially unsafe CDCs that might cause metastability or data coherency issues. Report CDC is similar to the Clock Interaction Report, but Report CDC focuses on structures and related timing constraints. Report CDC does not provide timing information because timing slack does not make sense on paths that cross asynchronous clock domains.

Report CDC identifies the most common CDC topologies as follows:

- Single bit synchronizers
- Multi-bit synchronizers for buses
- Asynchronous reset synchronizers
- MUX and CE controlled circuitry
- Combinatorial logic before synchronizer
- Multi-clock fanin to synchronizer
- Fanout to destination clock domain

For more information on the `report_cdc` command, see this [link](#) in the *Vivado Design Suite User Guide: Design Analysis and Closure Techniques* (UG906) [Ref 21] and see [report_cdc](#) in the *Vivado Design Suite Tcl Command Reference Guide* (UG835) [Ref 13].

Specific constraints should be applied to prevent default timing analysis on asynchronous clock domain crossings:

- [Global Constraints Between Clocks in Both Directions](#)
- [Constraints on Individual CDC Paths](#)

Global Constraints Between Clocks in Both Directions

When there is no need to limit the maximum latency, the clock groups can be used. Following is an example to ignore paths between `clkA` and `clkB`:

```
set_clock_groups -asynchronous -group clkA -group clkB
```

When two master clocks and their respective generated clocks form two asynchronous domains between which all the paths are properly synchronized, the clock groups constraint can be applied to several clocks at once:

```
set_clock_groups -asynchronous \  
-group {clkA clkA_gen0 clkA_gen1 ...} \  
-group {clkB clkB_gen0 clkB_gen1 ...}
```

Or simply:

```
set_clock_groups -asynchronous \  
-group [get_clocks -include_generated_clock clkA] \  
-group [get_clocks -include_generated_clock clkB]
```

Constraints on Individual CDC Paths

If a CDC bus uses gray-coding (e.g., FIFO) or if latency needs to be limited between the two asynchronous clocks on one or more signals, you must use the `set_max_delay` constraint with the option `-datapath_only` to ignore clock skew and jitter on these paths, plus override the default path requirement by the latency requirement. It is usually sufficient to use the source clock period for the max delay value, just to ensure that no more than one data is present on the CDC path at any given time.

When the ratio between clock periods is high, choosing the minimum of the source and destination clock periods is also a good option to reduce the transfer latency. A clean asynchronous CDC path should not have any logic between the source and destination sequential cells, so the Max Delay Datapath Only constraint is normally easy to meet for the implementation tools.

For the paths that do not need latency control, you can define a point-to-point false path constraint.

Clock Exceptions Precedence Over `set_max_delay`

When writing the CDC constraints, verify that the precedence is respected. If you use `set_max_delay -datapath_only` on at least one path between two clocks, the `set_clock_groups` constraint cannot be used between the same clocks, and the `set_false_path` constraint can only be used on the other paths between the two clocks.

In the following figure, the clock `clk0` has a period of 5ns and is asynchronous to `clk1`. There are two paths from `clk0` domain to `clk1` domain. The first path is a 1-bit data synchronization. The second path is a multi-bit gray-coded bus transfer.

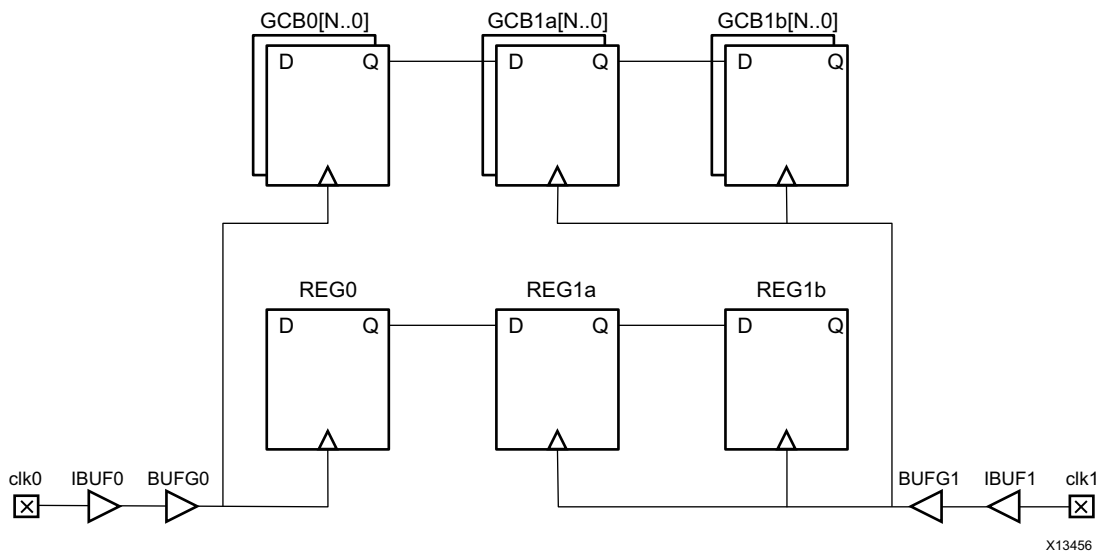


Figure 4-50: Multiple Interactions Between Two Asynchronous Clocks

The designer decides that the gray-coded bus transfer requires a Max Delay Datapath Only to limit the delay variation among the bits, so it becomes impossible to use a Clock Groups or False Path constraint between the clocks directly. Instead, two constraints must be defined:

```
set_max_delay -from [get_cells GCB0[*]] -to [get_cells [GCB1a[*]] \
-datapath_only 5
set_false_path -from [get_cells REG0] -to [get_cells REG1a]
```

There is no need to set a false path from `clk1` to `clk0` because there is no path in this example.

Specifying Timing Exceptions

Timing exceptions are used to modify how timing analysis is done on specific paths. By default, the timing engine assumes that all paths should be timed with a single cycle requirement for setup analysis in order to cover the most pessimistic clocking scenario. For certain paths, this is not true. Following are a few examples:

- Asynchronous Clock Domain Crossing paths cannot be safely timed due to the lack of fixed phase relationship between the clocks. They should be ignored (Clock Groups, False Path), or simply have datapath delay constraint (Max Delay Datapath Only)
- The sequential cells launch and capture edges are not active at every clock cycle, so the path requirement can be relaxed accordingly (Multicycle Path)
- The path delay requirement needs to be tightened in order to increase the design margin in hardware (Max Delay)

- A path through a combinatorial cell is static and does not need to be timed (False Path, Case Analysis)
- The analysis should be done with only a particular clock driven by a multiplexer (Case Analysis).

In any case, timing exceptions must be used carefully and must not be added to hide real timing problems.

Timing Exceptions Guidelines

Use a limited number of timing exceptions and keep them simple whenever possible. Otherwise, you will face the following challenges:

- The runtime of the compilation flow will significantly increase when many exceptions are used, especially when they are attached to a large number of netlist objects.
- Constraints debugging becomes extremely complicated when several exceptions cover the same paths.
- Presence of constraints on a signal can hamper the optimization of that signal. Therefore, including unnecessary exceptions or unnecessary points in exception commands can hamper optimization.

Following is an example of timing exceptions that can negatively impact the runtime:

```
set_false_path -from [get_ports din] -to [all_registers]
```

- If the `din` port does not have an input delay, it is not constrained. So there is no need to add a false path.
- If the `din` port feeds only to sequential elements, there is no need to specify the false path to the sequential cells explicitly. This constraint can be written more efficiently:

```
set_false_path -from [get_ports din]
```

- If the false path is needed, but only a few paths exist from the `din` port to any sequential cell in the design, then it can be more specific (`all_registers` can potentially return thousands of cells, depending upon the number of registers used in the design):

```
set_false_path -from [get_ports din] -to [get_cells blockA/config_reg[*]]
```

Timing Exceptions Precedence and Priority Rules

Timing exceptions are subject to strict precedence and priority rules. The most important rules are:

- The more specific the constraint, the higher the priority. For example:

```
set_max_delay -from [get_clocks clkA] -to [get_pins inst0/D] 12
set_max_delay -from [get_clocks clkA] -to [get_clocks clkB] 10
```

The first `set_max_delay` constraint has a higher priority because the `-to` option uses a pin, which is more specific than a clock.

- The exceptions priority is as follows:
 1. `set_false_path`
 2. `set_max_delay` or `set_min_delay`
 3. `set_multicycle_path`

The `set_clock_groups` command is not considered a timing exception even though it is equivalent to two `set_false_path` commands between two clocks. It has higher precedence than the timing exceptions.

The `set_case_analysis` and `set_disable_timing` commands disable timing analysis on specific portions of the design. They have higher precedence than the timing exceptions.

Adding False Path Constraints

False path exceptions can be added to timing paths to ignore slack computation on these paths. It is usually difficult to prove that a path does not need timing to be functional, even with simulation tools. Xilinx does not usually recommend using a false path unless the risk associated with it has been properly assessed and appear to be acceptable.

Use Cases

The typical cases for using the false path constraint are:

- Ignoring timing on a path that is never active. For example, a path that crosses two multiplexers that can never let the data propagate in a same clock cycle because of the select pins connectivity.




Figure 4-51: Path Cannot be Sensitized

```
set_false_path -through [get_pins MUX0/I0] -through [get_pins MUX1/I1]
```

- Ignoring timing on an asynchronous CDC path. This case is already discussed in the Defining Clock Groups and CDC Constraints section.

- Ignoring static paths in the design. Some registers take a value once during the initialization phase of the application and never toggle again. When they appear to be on the critical path of the design, they can be ignored for timing in order to relax the constraints on the implementation tools and help with timing closure. It is sufficient to define a false path constraint from the static register only, without explicitly specifying the paths endpoints. Example: the paths from a 32-bit configuration register `config_reg[31..0]` to the rest of the design can be ignored by adding the following false path constraint:

```
set_false_path -from [get_cells config_reg[*]]
```

Impact on Synthesis

The false path constraint is supported by synthesis and will only impact max delay (setup/recovery) path optimization. It is usually not needed to use false path exceptions during synthesis except for ignoring CDC paths.

Impact on Implementation

All the implementation steps are sensitive to the false path timing exception.

Adding Min and Max Delay Constraints

The min and max delay exceptions are used to override the default path requirement respectively for hold/removal and setup/recovery checks by replacing the launch and capture edge times with the delay value from the constraint.

Use Cases

Common reasons for using the min or max delay constraints are for:

- Over-constraining a few paths of the design by tightening the setup/recovery path requirement.

This is useful for forcing the logic optimization or placement tools to work harder on some critical path cells, which can provide more flexibility to the router to meet timing later on (after removing the max delay constraint).

- Replacing a multicycle constraint.

This is a valid, but not the recommended way, to relax the setup requirement on a path that has active launch and capture edges every N clock cycles. Although it is the only option to over-constrain a multicycle path by a fraction of a clock period to help with timing closure during the routing step. For example, a path with a multicycle constraint of 3 appears to be the worst violating path after route and fails timing by a few hundred ps.

The original multicycle path constraint can be replaced by the following constraint during optimization and placement:

```
set_max_delay -from [get_pins <startpointCell>/C] \  
-to [get_pins <endpointCell>/D] 14.5
```

where

14.5 corresponds to 3 clock periods (of 5 ns each), minus 500 ps that correspond to amount of extra margin desired.

- Constraining the maximum datapath delay on asynchronous CDC paths.

This technique has already been described in [Defining Clock Groups and CDC Constraints](#).

It is not common or recommended to force extra delay insertion on a path by using the `set_min_delay` constraint. The default min delay requirement for hold or removal is usually sufficient to ensure proper hardware functionality when the slack is positive.

Impact on Synthesis

The `set_max_delay` constraint is supported by synthesis, including the `-datapath_only` option. The `set_min_delay` constraint is ignored.

Impact on Implementation

The `set_max_delay` constraint replaces the setup path requirement and influences the entire implementation flow. The `set_min_delay` constraint replaces the hold path requirement and only affects the router behavior whenever it introduces the need to fix hold.

Avoiding Path Segmentation

Path segmentation is introduced when specifying invalid startpoint or endpoint for the `-from` or `-to` options of the `set_max_delay` and `set_min_delay` commands only. When a `set_max_delay` introduces path segmentation on a path, the default hold analysis no longer takes place. You must constrain the same path with `set_min_delay` if you desire to constrain the hold analysis as well. The same rule applies with the `set_min_delay` command relative to the setup analysis.

Path segmentation must only be used by experts as it alters the fundamentals of timing analysis:

- Path segmentation breaks clock skew computation on the segmented path.
- Path segmentation can break more paths than the one constrained by the segmenting `set_max_delay` or `set_min_delay` command.

Path segmentation is reported by the tools in the log file when the constraints are applied. You must avoid it by using valid startpoints and endpoints:

- **Startpoints**

clock, clock pin, sequential cell (implies valid startpoint pins of the cell), input or inout port

- **Endpoints**

clock, input data pin of sequential cell, sequential cell (implies valid endpoint pins of the cell), output or inout port

Adding Multicycle Path Constraints

Multicycle path exceptions must reflect the design functionality and must be applied on paths that do not have an active clock edge at every cycle, on either the source clock, the destination clock or both clocks. The path multiplier is expressed in term of clock cycles, either based on the source clock when the `-start` option is used, or the destination clock when the `-end` option is used. This is particularly convenient for modifying the setup and hold relationships between the startpoint and endpoint independently of the clock period value.

The hold relationships are always tied to the setup ones. Consequently, in most cases, the hold relationship also needs to be separately adjusted after the setup one has been modified. This is why a second constraint with the `-hold` option is needed. The main exception to this rule is for synchronous CDC paths between phase-shifted clocks: only setup needs to be modified. An example is provided in the Use Cases below.

Multicycle Path Exception Use Cases

There are two main categories of multicycle path exception use cases:

- [Relaxing the Setup Requirement While Keeping Hold Unchanged](#)
- [Adjusting the Setup Edges Analysis on Paths Between Shifted Clocks.](#)

Relaxing the Setup Requirement While Keeping Hold Unchanged

This occurs when the source and destination sequential cells are controlled by a clock enable signal that activates the clock every N cycles. Following are three examples:

- [Example One: Same clock for both startpoint and endpoint, with a clock enable active every 3 cycles](#)
- [Example Two: Path from a slow clock to a fast clock](#)
- [Example Three: Path from a fast clock to a slow clock](#)

Example One: Same clock for both startpoint and endpoint, with a clock enable active every 3 cycles




Figure 4-52: Enabled Flops with Same Clock Signal

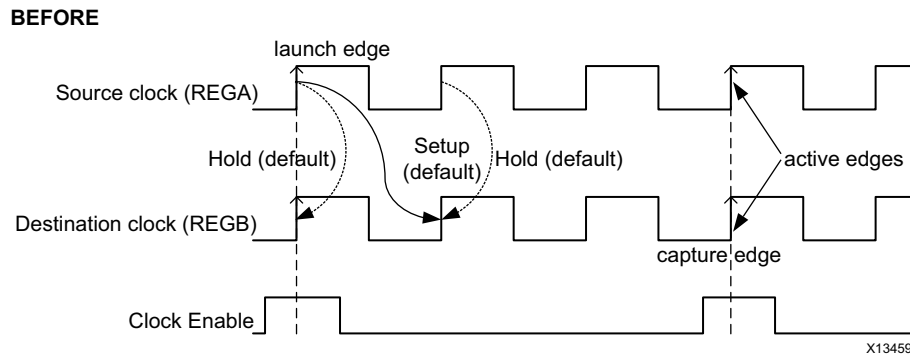


Figure 4-53: Timing Diagram for Setup/Hold Check

Constraints:

```
set_multicycle_path -from [get_pins REGA/C] -to [get_pins REGB/D] -setup 3
set_multicycle_path -from [get_pins REGA/C] -to [get_pins REGB/D] -hold 2
```

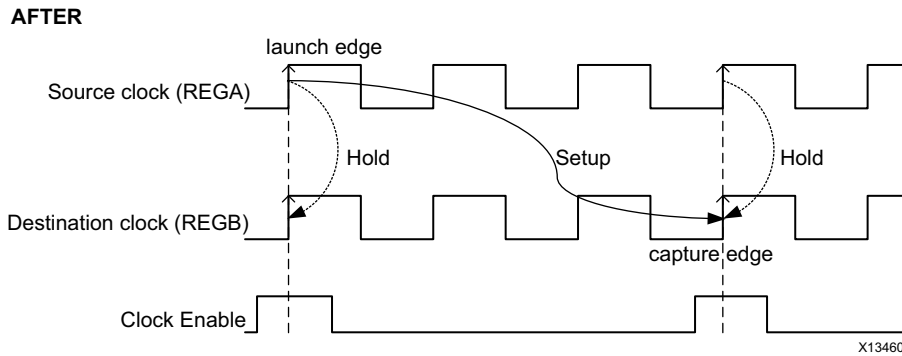


Figure 4-54: Setup/Hold Checks Modified After Multi-cycle Specification

Note: With the first command, as the setup capture edge moved to the third edge (that is, by 2 cycles from its default position), the hold edge also moved by 2 cycles. The second command is for bringing the hold edge back to its original location by moving it again by 2 cycles (in the reverse direction).

Example Two: Path from a slow clock to a fast clock

In this case, assume that only the destination flip-flop is controlled by a clock enable, and the clock enable is always active at the same time as the slow clock rising edge. The path multiplier for setup is 3.




Figure 4-55: Slow to Fast Clock Crossing




Figure 4-56: Timing Diagram for Setup/Hold Check - Slow to Fast Clock

Constraints:

```
set_multicycle_path -from [get_pins REG0/C] -to [get_pins REG1/D] -setup 3 -end
set_multicycle_path -from [get_pins REG0/C] -to [get_pins REG1/D] -hold 2 -end
```

The `-end` option is used only to modify the setup and hold analysis edges with respect to the destination clock (or endpoint clock). The correct source clock edges are already used.




Figure 4-57: Timing Diagram for Setup/Hold Check - Slow To Fast Clock - After Multicycle Constraint

Example Three: Path from a fast clock to a slow clock

This case is similar to the previous case (Path from a slow clock to a fast clock), except that this time only the edges of the source clock must be modified. Example of constraints:

```
set_multicycle_path -from [get_pins REGA/C] -to [get_pins REGB/D] -setup 3 -start
set_multicycle_path -from [get_pins REGA/C] -to [get_pins REGB/D] -hold 2 -start
```

Adjusting the Setup Edges Analysis on Paths Between Shifted Clocks

The main reason for shifting two clocks is to:

- Relax the setup paths from a clock to the late clock at the expense of tightening the paths in the other direction. This is common on I/O interfaces to adjust the timing at the interface of the device.
- Create a 90 degrees phase shift between the forwarded clock and the data of a source synchronous interface.

By default, the timing engine uses the active edges of the source and destination clocks that form the most pessimistic setup relationship. When inserting a positive phase shift in the destination clock definition, the setup relationship corresponds to the phase shift instead of a period plus the phases shift, because this is the tightest positive path requirement. Following is an example:

Source clock waveform: rise @ 0ns, fall @ 5ns, rise @ 10ns

Destination clock waveform: rise @ 2.5ns, fall @ 7.5ns, rise @ 12.5ns

DEFAULT SETUP AND HOLD RELATIONSHIPS:




Figure 4-58: Setup/Hold Edges for Phase Shifted Clocks

If you decide that the capture edge #2 is the valid capture edge for the setup analysis, a multicycle path constraint must be defined. If all the paths between the two phase-shifted clocks must be modified, you can directly specify the constraints on the clocks:

```
set_multicycle_path -from [get_clocks clk] -to [get_clocks clkshift] -setup 2
```




Figure 4-59: Setup/Hold Edges for Phase Shifted Clocks - After Multi-cycle Specification



IMPORTANT: *In this case, it is not necessary to modify the hold relationship with an additional `set_multicycle_path` constraint because it is already properly established relative to the setup relationship and all rising clock edges are active.*

Impact on Synthesis

The `set_multicycle_path` constraint is supported by synthesis and can greatly improve the timing QoR (for setup only) by relaxing long paths that are functionally not active at every clock cycle.

Impact on Implementation

As for synthesis, multicycle path exceptions help the timing-driven algorithms to focus on the real critical paths. The hold requirements are important only during route. If a setup relationship was adjusted with a `set_multicycle_path` constraint but not its corresponding hold relationships, the worst hold requirement may become too hard to meet if it is over 2 or 3 ns. This situation can have a negative impact on setup slack because of the additional delay inserted by the router while fixing hold violations.

Common Mistakes

Following are two typical mistakes that you must absolutely avoid:

- Relaxing setup without adjusting hold back to same launch and capture edges in the case of a multicycle path not functionally active at every clock cycle. The hold requirement can become very high (at least one clock period in most cases) and impossible to meet.
- Setting a multicycle path exception between incorrect points in the design.

This occurs when you assume that there is only one path from a startpoint cell to an endpoint cell. In some cases this is not true. The endpoint cell can have multiple data input pins, including clock enable and reset pins, which are active on at least two consecutive clock edges.

For this reason, Xilinx recommends that you specify the endpoint pin instead of just the cell (or clock). Example: the endpoint cell REGB has three input pins: C, EN and D. Only the REGB/D pin should be constrained by the multicycle path exception, not the EN pin because it can change at every clock cycle. If the constraint is attached to a cell instead of a pin, all the valid endpoint pins are considered for the constraints, including the EN (clock enable) pin.

To be safe, Xilinx recommends that you always use the following syntax:

```
set_multicycle_path -from [get_pins REGA/C] \  
-to [get_pins REGB/D] -setup 3  
set_multicycle_path -from [get_pins REGA/C] \  
-to [get_pins REGB/D] -hold 2
```

Other Advanced Timing Constraints

A few other timing constraints can be set to ignore and modify the default timing analysis:

- [Case Analysis](#)
- [Disable Timing](#)
- [Data Check](#)
- [Max Time Borrow](#)

Case Analysis

The case analysis command is commonly used to describe a functional mode in the design by setting constants in the logic like what configuration registers do. It can be applied to input ports, nets, hierarchical pins, or leaf cell input pins. The constant value propagates through the logic and turns off the analysis on any path that can never be active. The effect is similar to how the false path exception works.

The most common example is to set a multiplexer select pin to 0 or 1 in order to only allow one of the two multiplexer inputs to propagate through. The following example turns off the analysis on the paths through the `mux/S` and `mux/I1` pins:

```
set_case_analysis 0 [get_pins mux/S]
```

Disable Timing

The disable timing command turns off a timing arc in the timing database, which completely prevents any analysis through that arc. The disabled timing arcs can be reported by the `report_disable_timing` command.



CAUTION! Use the disable timing command carefully. It can break more paths than desired!

Data Check

The `set_data_check` command sets the equivalent of a setup or hold timing check between two pins in a design. It is commonly used to constrain and report asynchronous interfaces. This command should be used by expert users.

Max Time Borrow

The `set_max_time_borrow` command sets the maximum amount of time a latch can borrow from the next stage (logic after the latch), and give it the previous stage (logic before the latch). Latches are not recommended in general as they are difficult to test and validate in hardware. This command should be used by expert users.

Creating Block-Level Constraints

When working on a multi-team project, it is convenient to create individual constraint files for each major block of the top-level design. Each of these blocks is usually developed and validated separately before the final integration into one or many top-level designs.

The block-level constraints must be developed independently from the top-level constraints, and must be as generic as possible so that they can be used in various contexts. They must also not affect any logic that is beyond the block boundaries.

Block-Level Constraint Rules

The block-level constraints must comply with the following rules:

1. Do not define clocks in the block-level constraints if they are expected to be created at the top level of the design.

Instead they can be queried inside the block by using the `get_clocks -of_objects` command. This command returns all the clocks that traverse a particular object in the design. Example:

```
set blockClock [get_clocks -of_objects [get_ports clkIn]]
```

If a clock needs to be defined inside the block, it must be on an input/inout port that is driving an instantiated input/inout buffer, or on the output of a cell that creates/transforms a clock (except for MMCM/PLL or special buffers that are automatically handled by the timing tools). Examples:

- Input clock with input buffer
 - Clock Divider
 - GT recovered clock
2. Specify input and output delay only if the port is directly connected to the top-level port and the I/O buffer is instantiated inside the IP. Example:
 - Input data ports with input buffers
 - Output data ports with output buffers
 3. Do not define timing exceptions between two clocks that are not bounded to the IP.
 4. Do not refer to clocks by name as the name may vary based on the top-level clock names or if the block is instantiated multiple times.
 5. Do not add placement constraints if the block can be instantiated multiple times in a same top-level design.

Reading Block-Level Constraints Into a Top-Level Design

The Vivado Design Suite provides a scoping mechanism for reading block-level constraints into a top-level design. This mechanism is based on the `current_instance` command behavior where all name-based queries can only return objects included in the current instance.

When reading in the block-level constraints, the current instance is set to the block instance so that only objects that belong to the block can be constrained. There are a few exceptions to this rule:

- Timing clocks are global and can be queried from anywhere in the design, including from within the block. The `get_clocks` command must be used carefully as it can query clocks outside the block.
- Ports of the block module definition can be queried with the `get_ports` command. Depending on how the block instance is connected in the top level design, the type of objects returned can differ:
 - If a block port is directly connected to a top-level port, then the top-level port is returned by the `get_ports` command.

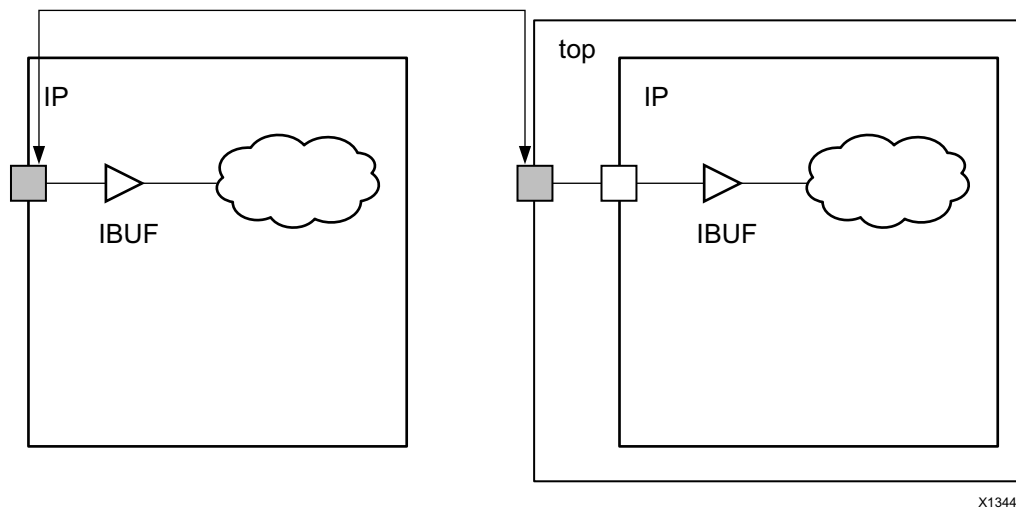


Figure 4-60: `get_ports` for Block Ports Connected Directly to Top-Level Port

- If a block port is not directly connected to a top-level port, then the corresponding hierarchical pin of the block interface is returned by the `get_ports` command.




Figure 4-61: `get_ports` for Block Ports Not Directly Connected to Top- Level Port

This scoping mechanism is used by all Vivado Design Suite IP cores that are delivered with constraints. For more information, see *Vivado Design Suite User Guide: Using Constraints* (UG903) [Ref 18].

Defining Physical Constraints

Physical constraints are used to control floorplan, specific placement, I/O assignments, routers and similar functions. Make sure that each pin has an I/O location and standard specified. Physical Constraints are covered in the following user guides:

- *Vivado Design Suite User Guide: Using Constraints* (UG903) [Ref 18] for locking placement and routing, including relative placement of macros
- *Vivado Design Suite User Guide: Design Analysis and Closure Techniques* (UG906) [Ref 21] for floorplanning
- *Vivado Design Suite User Guide: Programming and Debugging* (UG908) [Ref 24] for configuration

Implementation

Overview of Implementation

Now that you have selected your device, chosen and configured the IP, and written the RTL and the constraints, the next step is implementation. Implementation generates the bitstream that is used to program the device. The implementation process might have some iterative loops, as discussed in [Chapter 1, Introduction](#). This chapter describes the various implementation steps; highlights points for special attention; and gives tips and tricks to identify and eliminate specific bottlenecks.

Synthesis

Synthesis takes in RTL and timing constraints and generates an optimized netlist that is functionally equivalent to the RTL. In general, the synthesis tool can take any legal RTL and create the logic for it. Following are some special points to consider when synthesizing your design. For additional information about synthesis, refer to the following resources:

- *Vivado Design Suite User Guide: Synthesis* (UG901) [[Ref 16](#)]
- [Vivado Design Suite QuickTake Video: Design Flows Overview](#)

Synthesis Attributes During Design Migration

Synthesis attributes such as DONT_TOUCH or MAX_FANOUT can significantly impact your Quality of Results (QoR). Take care in setting these attributes. If your project was originally run with a different synthesis tool, and you are migrating the project to run in Vivado® Design Suite synthesis, you should be aware of the need for any of the attributes already present. Remove all attributes that were added specifically to control the optimization behavior of the previous tool.

Attributes such as KEEP, DONT_TOUCH, and MAX_FANOUT are normally not used when the RTL is initially created. They are typically used to tweak the synthesis tool into giving the best performance. Since all synthesis tools optimize somewhat differently, using these types of attributes can adversely affect your QoR when you migrate to a new tool.



RECOMMENDED: *Start with fresh RTL, and then apply your attributes specifically for your new tools or requirements.*

Accurate Timing Constraints

Since the Vivado Design Suite synthesis tool is timing-driven, be certain that the timing constraints are accurate. See [Baselining the Design](#). If timing exceptions are needed by the design, provide them as well. The tool automatically sends the constraints from the XDC file to synthesis to perform a timing-driven run. If the constraints that are sent to synthesis and to place and route are different, then synthesis and place and route are working on different paths. When this happens, it can be difficult to achieve timing closure.

Check Your HDL Code

If after synthesis, you do not have the desired QoR, check your HDL code and the inferred logic for the following:

- Evaluate signals such as set and reset to determine if they are necessary. And, if really needed, synchronous signals (active-High) should be preferred.
- DSPs and block RAMs have internal registers. Use these registers.

To understand how the HDL code impacts inference and thus QoR, see [Chapter 4, Design Creation](#).

Debugging Your Synthesized Design

If the post-synthesis netlist does not exhibit the same behavior as desired, we would need to debug the source of the problem.

Using the Elaborated Design

The elaborated design is the first step in analyzing or debugging a design. It is a direct representation of the RTL code itself. Using the elaborated design view allows you to debug your design before running synthesis. This allows problems with RTL code to be found earlier in the design flow. For example, consider the elaborated design view shown in the following figure.

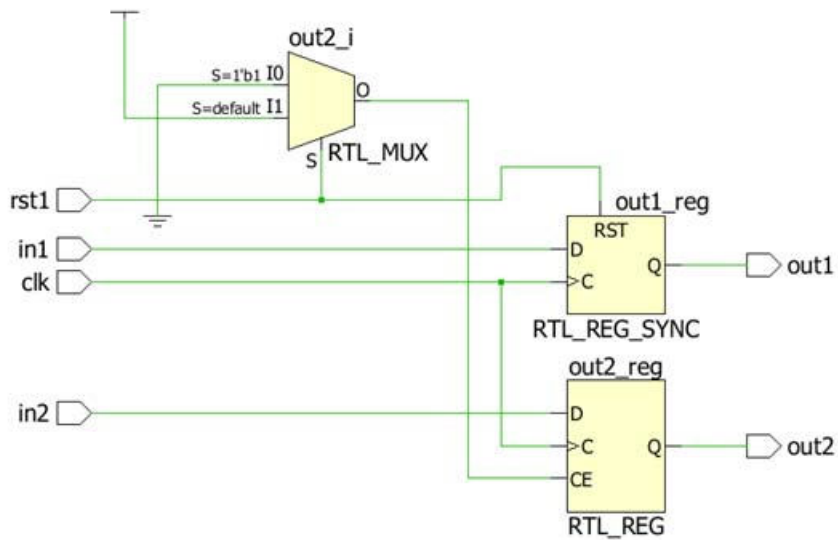



Figure 5-1: Elaborated Design View Example

It is apparent that `out2_reg` is enabled by the `rst1` signal. This is most likely a coding error. The cross probing feature of the view will take you directly to the RTL that created this logic.

The Elaborated View is also used to help designs that are not meeting timing. After running the design through synthesis, and finding the critical path, search for the same path in the elaborated view. This often helps to find RTL changes that can improve the timing of a design. Such constructs as large MUX structures or unpipelined DSP or block RAM structures are easily seen in this view.


To open Elaborated View, click **Open Elaborated Design** in Flow Navigator.



For more information, see this [link](#) in the *Vivado Design Suite User Guide: System-Level Design Entry* (UG895) [Ref 8].

Using the Synthesized Design View

Like the Elaborated View, the Synthesized Design View is also useful in debugging a design that has been synthesized. To open Synthesized Design View after the design has been synthesized, click **Open Synthesized Design** in Flow Navigator.



This view is based on the Xilinx® primitives that were used in creating the netlist. The synthesized design view can be used to view how the RTL was actually translated into the primitives. It will list all the properties of the primitives.

Consider the following example:

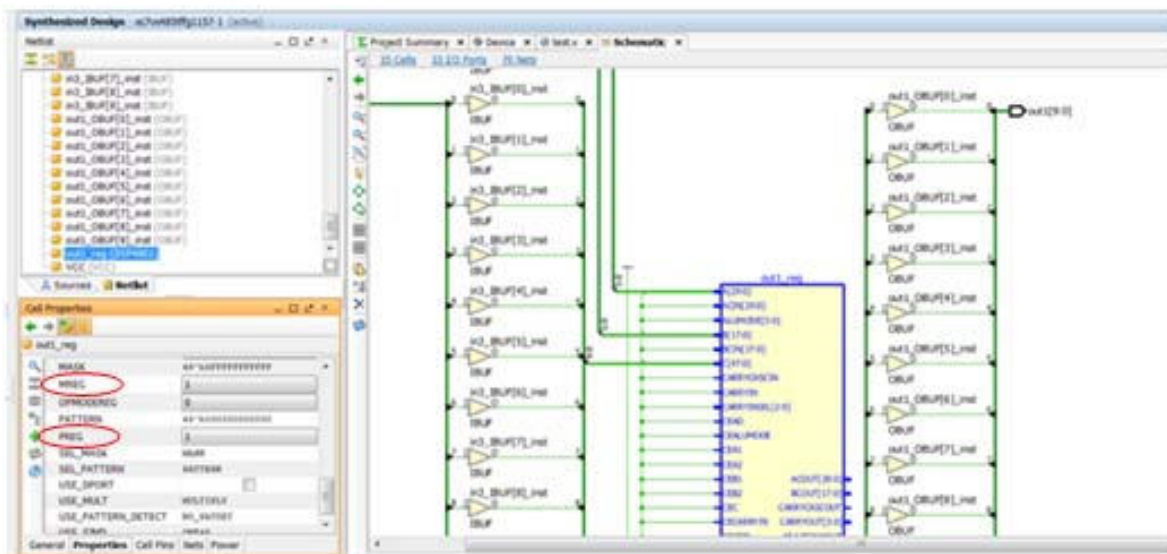


Figure 5-2: Synthesized Design View Example

By viewing the property of `out1_reg` (lower left side of the figure), you can see that DSP48E1 is using the embedded pipeline registers called MREG and PREG. This is useful for determining how pipelined your design is.

Also, the synthesized design is the first place you should look for analyzing the timing of your design. Understanding the post-synthesis timing is very important so that you can have a view of any potential timing bottlenecks before running implementation. Once the design is open, run **Tools > Timing > Report Timing Summary** to view timing information. This report provides much useful information, including, for example:

- Summaries of each clock and inter-clock paths
- Unconstrained paths or I/Os
- Clocks that have not been given timing constraints

From the Tcl Console, you can also check your timing constraints to be sure that they were accepted. Xilinx strongly recommends that you always do so. If the constraints were not accepted, there is no guarantee that the Vivado tools are working on the correct paths.

You can use the `report_timing` command for the path of interest to check if your constraints for the path have been applied. For example, if you intended to apply:

```
set_false_path -from [get_pins inst1/pin1] -through [get_cells inst2]
```

you should run the following after applying the above constraint to confirm that the path slack shows up as infinite.

```
report_timing -from [get_pins inst1/pin1] -through [get_cells inst2]
```

In the Synthesized design, a hierarchy viewer shows a view of the different levels of hierarchy.

This is useful for debugging a design that you suspect has a concern around area. The different blocks are sized relatively to how many primitives they have inside of them. By clicking on each level of hierarchy in the Netlist box, and then looking at the Statistics tab in the Netlist Properties box, you can see exactly the number and type of primitives in the design. These numbers represent the number of primitives in this level and all the levels below. Using this view, you can easily see where the area blow up happens. For more information, see this [link](#) in the *Vivado Design Suite User Guide: Design Analysis and Closure Techniques* (UG906) [Ref 21].

Synthesis Attributes

Synthesis attributes allow you to control the logic inference in a specific way. Although synthesis algorithms are set to give the best results for the largest number of designs, there are often designs with differing requirements. Attributes are then typically used to tweak the tool into making the design a little different for purposes of QoR. For example:

- The `MAX_FANOUT` attribute can enforce a maximum fanout on a specific net.
- The `RAM_STYLE` attribute can force a RAM to be implemented in a specific way.

Xilinx recommends that you allow the tool to operate without using attributes to obtain a first-pass run. Then specific to the design and the results it gives, add synthesis attributes to achieve the desired results.

Take care when putting multiple attributes on one signal, or different attributes on signals that are related to each other. While the tool will try to honor each of those attributes, in some cases it will be unsuccessful, if these attributes are trying to achieve conflicting behavior (for example, `KEEP_HIERARCHY` and `MAX_FANOUT`).

In general, when an attribute is known to the tool (for example, KEEP), it will be used by the tool and the effects will be seen in the netlist. However, the attribute itself will no longer appear in the output netlist. When the attribute is not known to the tool, it is assumed that this attribute will be used later in the flow. In that case, the attribute and value are passed to the output netlist.

For information on the attributes that synthesis supports, see the *Vivado Design Suite User Guide: Synthesis* (UG901) [Ref 16].

A few attributes deserve special mention because they sometimes cause issues that you need to be aware of:

- [KEEP and DONT_TOUCH](#)
- [MAX_FANOUT](#)

KEEP and DONT_TOUCH

KEEP and DONT_TOUCH are valuable attributes for debugging a design. They direct the tool to not optimize the objects on which they are placed.

- KEEP is used by the synthesis tool and is not passed along as a property in the netlist. KEEP may be used to fine-tune the behavior of the synthesis tool, in order to retain a specific signal: namely, to turn off specific optimizations on the specific signal during synthesis.
- DONT_TOUCH is used by the synthesis tool and then passed along to the place and route tools so that it will never be optimized.

There is also a difference between putting DONT_TOUCH on a level of hierarchy or on a signal. If the attribute is placed on a signal, that signal is kept. If the attribute is placed on a level of hierarchy, the tool does not touch the boundaries of that hierarchy and no constant propagation will happen through the hierarchy, but optimizations inside that level are still OK.

Take care when using these attributes. A KEEP attribute on a register that receives RAM output prevents that register from being merged into the RAM, thereby preventing a block RAM from being inferred. Do not use these attributes on a level of hierarchy that is driving a 3-state output or bidirectional signal in the level above. This is very important! If the driving signal and the 3-state condition are in this level of hierarchy, the IOBUF will not be inferred, because in order to do so, the tool must change the hierarchy in order to create the IOBUF.

MAX_FANOUT

MAX_FANOUT forces the synthesis to replicate logic in order to meet a fanout limit. The tool is able to replicate logic, but not inputs or black boxes. Accordingly, if a MAX_FANOUT

attribute is placed on a signal that is driven by a direct input to the design, the tool is unable to handle the constraint.

Take care to analyze the signals on which a MAX_FANOUT is placed. If a MAX_FANOUT is placed on a signal that is driven by a register with a DONT_TOUCH or drives signals that are in a different level of hierarchy when the DONT_TOUCH attribute is on that hierarchy, the MAX_FANOUT attribute will not be honored.



RECOMMENDED: Use MAX_FANOUT sparingly during synthesis. The *phys_opt_design* command in the Vivado tools has a much better idea of the placement of the design, and can do a much better job of replication than synthesis. If a specific fanout is desired, it is often worth the time and effort to manually code the extra registers.

Bottom Up Flow

Often it is desirable to have pre-compiled lower level hierarchies imported into the Vivado tools as a bottom up flow. A bottom up flow can yield quicker run times, since synthesis does not compile and map these blocks every time. On the other hand, QoR can suffer if synthesis is not allowed to do cross boundary optimizations.

Creating the Lower Level Netlist

To do a bottom up flow, you first create the lower level netlist. To do so, set up your project for the lower level of hierarchy specified as the top level of your design and constraints files also specified corresponding to this portion of the hierarchy. Before running synthesis:

1. Open **Synthesis Settings**.
2. In the **More Options** line, enter:

```
-mode out_of_context
```

This tells synthesis not to insert any I/O buffers.

This is necessary because later in the flow when this is inserted into the rest of the design, the tool will error out if there are IBUF or OBUF components that are not touching pads of the design.



CAUTION! Check to see if there are inouts or output 3-states in the design. Since the IOBUF or the OBUFT are the only components in the Xilinx library that can handle 3-states, turning off I/O insertion will cause errors.

If inouts or 3-states are needed in the lower level netlist, instantiate the IOBUF or OBUFT components in the RTL. Even if the *out_of_context* mode is turned on, the synthesis tool does not remove instantiated I/O buffers.



TIP: `out_of_context` is a Vivado Design Suite setting of the `-mode` option. All other synthesis tools support this flow, but in different ways. For information on how to perform this function in other tools, see the third-party synthesis tool documentation.

After running synthesis, create the `.edif` file that will be used as the netlist for this portion of the design:

1. Open the synthesized design.
2. In the Tcl console, enter:

```
write_edif <name>.edif
```

Running the Top Level When Using a Lower Level Netlist

On the other end, when running the top level and instantiating a lower level netlist (often referred to as a black box), you must do the following. First, the netlist must be instantiated. The ways to do this differ in VHDL and Verilog. In both cases, the lower level ports must be described to the synthesis tool. For VHDL, a component statement is used to describe the black box.

```
component <name>
port (in1, in2 : in std_logic;
out1 : out std_logic);
```

Since Verilog does not have an equivalent of a component that VHDL does, a wrapper file is used to communicate the ports to the tool. This wrapper file looks like normal Verilog, but only contains the list of ports.

```
module <name> (in1, in2, out1);
input in1, in2;
output out1;
endmodule
```

In both cases, make sure that the port definitions are correct. If they do not match, errors will result while trying to insert the lower level netlists after synthesis.

Assembling the Design

Now that the lower level netlists have been created and the top level is instantiating the netlists correctly, add the lower level netlists to the Vivado Design Suite project as you would any other source file. The tool inserts them into the flow after synthesis is run on the top level. Place and route then happens normally.

If there are lower level IOBUFs in the netlist, the synthesis tool must be told not to insert any IOBUFs on the pins of the black box in question. You need to specify `BUFFER_TYPE` attribute in order to prevent IOBUF insertion on specific ports during synthesis.

Other synthesis tools also support bottom-up flows. For information on how to do this with other third-party synthesis tools, see the documentation for that tool.

Moving Past Synthesis

Be sure that the netlist you obtained during synthesis is of good quality so that it does not create problems downstream. Important items to check before proceeding with the rest of the implementation flow include:

- [Review and Clean DRCs](#)
- [Review Synthesis Log](#)
- [Review Timing Constraints](#)
- [Meet Post-Synthesis Timing](#)

Review and Clean DRCs

The `report_drc` command runs Design Rule Checks (DRCs) to look for common design issues and errors. There are multiple rule decks. The default rule deck for the command:

- Checks for DRCs related to the post synthesis netlist.
- Checks for I/O, BUFG, and other placement specific requirements.
- Performs basic checks on the attributes and wiring on MGTs, IODELAYs, MMCMs, PLLs and other primitives.

In addition to running the default rule deck, also run the `methodology_checks` and `timing_checks` rule decks.



RECOMMENDED: Review and correct DRC violations as early as possible in the design process to avoid timing or logic related issues later in the implementation flow. Make sure to run the methodology checks rule deck (see [Running Methodology DRCs](#)).

Review Synthesis Log

You must review the synthesis log files and confirm that all messages given by the tool match your expectations in terms of the design intent. Pay special attention to Critical Warnings and Warnings. In most cases, Critical Warnings need to be cleaned up for a reliable synthesis result.



CAUTION! If a message appears more than 100 times, the tool writes only the first 100 occurrences to the synthesis log file. You can change the limit of 100 through the Tcl command `set_param messaging.defaultLimit`.

Review Timing Constraints

You must provide clean timing constraints, along with timing exceptions, where applicable. Bad constraints result in long runtime, performance issues, and hardware failures.



RECOMMENDED: Review all Critical Warnings and Warnings related to timing constraints which indicate that constraints have not been loaded or properly applied.

For more information, see [Organizing the Design Constraints in Chapter 4](#).

Meet Post-Synthesis Timing

The following sections discuss how to meet post-synthesis timing:

- [Guidelines Regarding Remaining Violations](#)
- [Dealing with High Levels of Logic](#)
- [Reviewing Utilization](#)
- [Reviewing Clock Trees](#)

Guidelines Regarding Remaining Violations



IMPORTANT: Analyze timing post-synthesis to identify the major design issues that must be resolved before you move forward in the flow.

HDL changes tend to have the biggest impact on QoR. You are therefore better off solving problems before implementation to achieve faster timing convergence. When analyzing timing paths, pay special attention to the following:

- Most frequent offenders (that is, the cells or nets that show up the most in the top worst failing timing paths)
- Paths sourced by unregistered block RAMs
- Paths sourced by SRL
- Paths containing unregistered, cascaded DSP blocks
- Paths with large number of logic levels
- Paths with large fanout

For more information see [Timing Closure](#).

Dealing with High Levels of Logic

Identifying long logic paths is useful to diagnose difficult QOR challenges. Estimated net delays post-synthesis are close to the best possible placement. To evaluate if a path with high logic-level delay is meeting timing, you can generate timing reports with no net delay. Timing closure cannot be achieved on paths that are still violating timing with no net delays.

For more information, see [Timing Closure](#).

Reviewing Utilization

It is important to review utilization for LUT, FF, RAMB, and DSP components independently. A design with low LUT/FF utilization might still experience placement difficulties if RAMB utilization is high. The `report_utilization` command generates a comprehensive utilization report with separate sections for all design objects.

Reviewing Clock Trees

This section discusses reviewing clock trees and includes:

- [Clock Buffer Utilization](#)
- [Clock Tree Topology](#)

Clock Buffer Utilization

The `report_clock_utilization` command provides details on clock primitive utilization. Observe the architecture clocking rules to avoid downstream placement issues. For example, a BUFH can only fanout to loads in its clock region. Invalid placement constraints or very high fanout for regional clock buffers might cause issues in the placer. For designs with very high clock buffer utilization, it might be necessary to lock the clock generators and some regional clock buffers to aid placement.

For some interfaces needing very tight timing relationship, it is sometimes better to lock specific resources for these signals which need very tight timing relationship - for example, source synchronous interfaces. In general, as a starting point for your design, lock only the I/Os - unless, there are specific reasons as cited above.

For more information on recommended placement constraints, see [Timing Closure](#).

Clock Tree Topology

- Run the `report_clock_networks` command to show the clock network in detail tree view.
- Utilize clock trees in a way to minimize skew.
- For the outputs of PLLs and MMCMs, use the same clock buffer type to minimize skew.
- Look for unintended cascaded BUFG elements that can introduce additional delay, skew, or both.

Implementing the Design

Vivado Design Suite implementation includes all steps necessary to place and route the netlist onto the FPGA device resources, while meeting the design's logical, physical, and timing constraints. For additional information about implementation, refer to the following resources:

- *Vivado Design Suite User Guide: Implementation* (UG904) [Ref 19]
- [Vivado Design Suite QuickTake Video: Design Flows Overview](#)

Project Mode vs. Non-Project Mode Options

Implementation can be achieved in Project Mode or Non-Project Mode. Project Mode provides the project infrastructure such as runs management, file sets management, reports generation, and cross probing. Non-Project Mode provides easy integration and is driven by a Tcl script which must explicitly call all the desired reports along the flow. For additional information about these modes, see this [link](#) in the *Vivado Design Suite User Guide: Design Flows Overview* (UG892) [Ref 5].

Project Mode

Project Mode is based on runs. You can create and launch new implementation runs that use different synthesis results, design constraints, or both, to increase the implementation solution space and find the best results. In Project Mode, the Vivado IDE allows you to run multiple strategies on a single design; customize implementation strategies to tune the algorithms to your design; and save customized implementation strategies to use in other projects. Once you have found the best strategy for your design, you can use it for future designs with similar characteristics.

Non-Project Mode

In Non-Project Mode, implementation is run using a Tcl script that defines the design flow.

Recommended Flow

Below is a minimal list of commands that you must run after reading in the design to generate a valid bitstream:

- `link_design`
- `opt_design`
- `place_design`
- `route_design`
- `report_drc`
- `report_timing_summary`
- `write_bitstream`

The timing constraints should be complete and correct. These constraints should be met with a positive slack to ensure a working design in hardware.

Iterative Flows

In Non-Project Mode, you can iterate between various optimization commands with different options. For example, you can run `place_design -post_place_opt` after `route_design` to run post placement optimization on a routed design that is not meeting timing on a few critical paths. The placer uses the actual timing delays to do post-placement optimization. You need to follow this step by running `route_design` again.

Running `phys_opt_design` iteratively can provide timing improvement. The `phys_opt_design` command attempts to optimize the top timing problem paths. By running `phys_opt_design` iteratively, lower level timing problems may benefit from the optimization. Invocation of `phys_opt_design` at post-route stage will reroute any nets that might have been unrouted. So, `phys_opt_design` at post-route need not be followed by another explicit run of `route_design`.

Running Methodology DRCs

Due to the importance of methodology, the Vivado tools provide a set of Design Rule Checks (DRCs) that specifically check for compliance with methodology. There are different types of DRCs depending on the stage of the design process. RTL lint-style checks are run on the elaborated RTL design; netlist-based logic and constraint checks are run on the synthesized design; and implementation and timing checks are run on the implemented design.

To run these checks at the Tcl prompt, open the design to be validated and enter following Tcl command:

```
report_drc -ruledeck methodology_checks
```

To run these checks from the IDE, open the design to be validated and run the Report DRC command. Once the dialog appears, select the methodology checks Rule deck, as shown in the following figure.

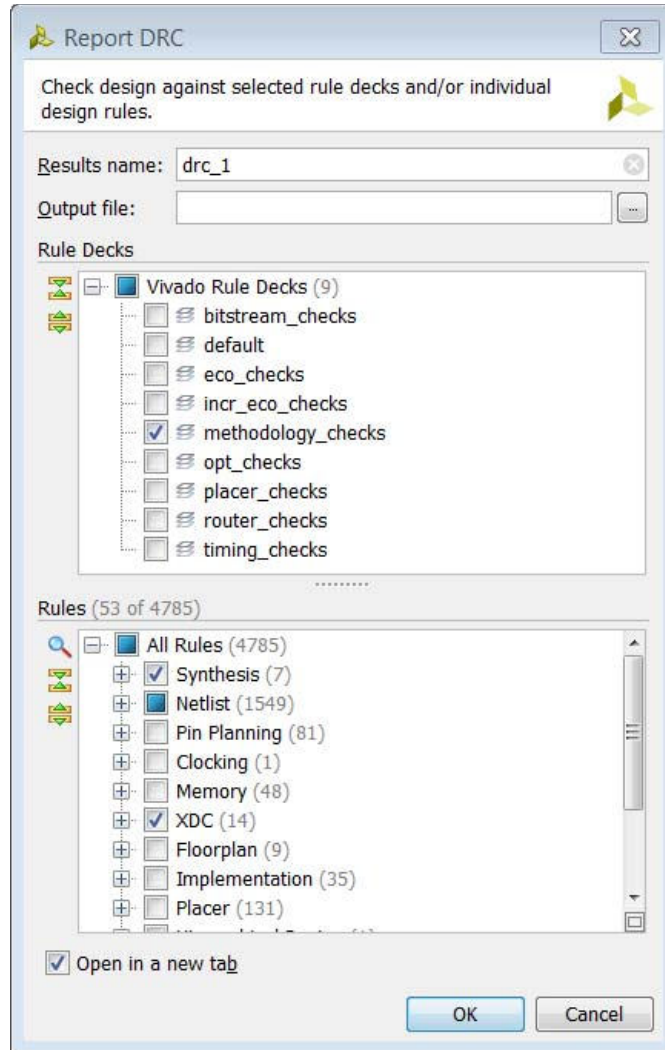


Figure 5-3: Report DRC Dialog Box

Violations (if there are any) are listed in the DRC window, as shown in the following figure.

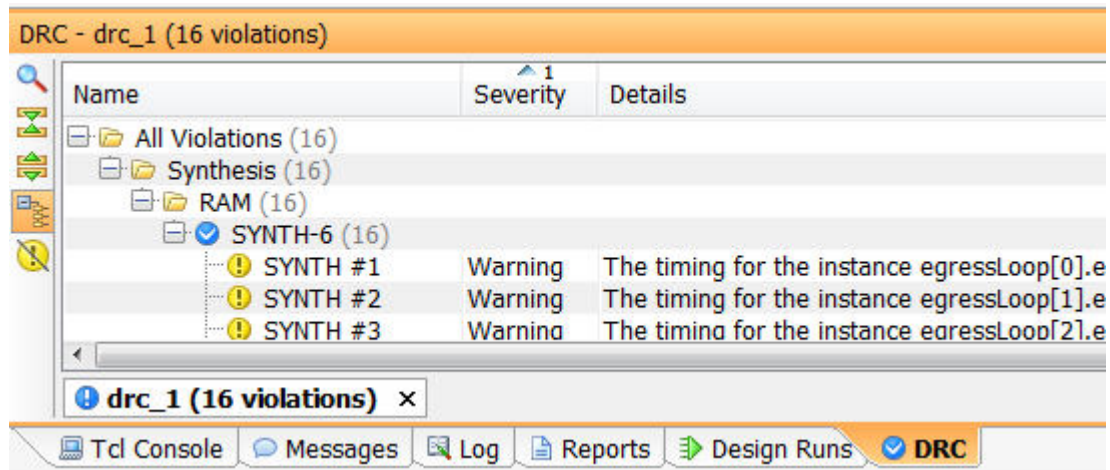


Figure 5-4: DRC violations

For more information on running design methodology DRCs, refer to *Vivado Design Suite User Guide: System-Level Design Entry* (UG895) [Ref 8] and *Vivado Design Suite User Guide: Design Analysis and Closure Techniques* (UG906) [Ref 21].

If a specific DRC violation does not need to be cleaned for your design, make sure that you understand the violation and its implication clearly and why the violation does not negatively impact your specific design. For Xilinx supplied IP cores, the DRC violations have been already reviewed and checked.

Strategies

Strategies are a defined set of Vivado Design Suite implementation options that control the behavior of runs in Project Mode. The strategies behavior in turn is controlled by the directives that are applied to the individual implementation commands. For more information, see [Directives](#).



RECOMMENDED: Try the default strategy Vivado Design Suite implementation defaults first. It provides a good trade-off between runtime and design performance.

Strategies are tool and version specific. Each major release of the Vivado Design Suite includes version-specific strategies.

The strategies are broken into categories according to their purposes, with the category name as a prefix. See the following table.

The Performance strategies aim to improve design performance at the expense of runtime. For example, the Performance_Explore strategy can help improve results for a large variety of designs at the expense of a large runtime increase.

Table 5-1: Strategy Categories

Category	Purpose
Performance	Improve design performance.
Area	Reduce LUT count.
Power	Add full power optimization.
Flow	Modify flow steps.
Congestion	Reduce congestion and related problems.



IMPORTANT: Strategies containing the terms SLL or SLR provide additional control for SSI technology devices.

Directives

Directives provide different modes of behavior for the following implementation commands:

- `opt_design`
- `place_design`
- `phys_opt_design`
- `route_design`

Use the default directive initially. Use other directives when the design nears completion to explore the solution space for a design. Only one directive may be specified at a time. The directive option is incompatible with other options.

For more information on strategies and directives, see the *Vivado Design Suite User Guide: Implementation* (UG904) [Ref 19].

Intermediate Steps and Checkpoints

The Vivado Design Suite uses a physical design database to store placement and routing information. Design checkpoint files (.dcp) allow you to save (`write_checkpoint` command) and restore (`read_checkpoint` command) this physical database at key points in the design flow. Checkpoints are a snapshot of the design at a specific point in the flow.

This design checkpoint file includes the following:

- Current netlist, including any optimizations made during implementation
- Design constraints
- Implementation results

Checkpoint designs can be run through the remainder of the design flow using Tcl commands. They cannot be modified with new design sources.

A few common examples for the use of checkpoints are:

- Saving results so you can go back and do further analysis on that part of the flow.
- Trying `place_design` using multiple directives and saving the checkpoints for each. This would allow you to select the `place_design` checkpoint with the best timing results for the subsequent implementation steps.

Incremental Flows

Incremental place and route in the Vivado Design Suite reuses existing placement and routing data to reduce implementation runtime and produce more predictable results when similarities exist between a previously implemented design and the current design. Incremental place and route can achieve an average of a twofold improvement over normal place and route runtimes when designs have at least 95 percent similar cells, nets, and ports.

The average runtime improvement decreases as similarity between the reference design and the current design decreases.

Below 80 percent, there may be little or no benefit to using the incremental place and route feature.

Good use-models for incremental flows include:

- Fixing implemented designs that are close to meeting timing and that require small localized fixes.
- Adding debug cores to an implemented design.
- Reworking critical localized paths that are impacting a limited amount of logic.
- Creating a new revision of design.

Note: This tends to have a lower level of similarities.

Besides the runtime savings, incremental compile also causes minimal disruptions to the portions of the design that have not changed, thereby, reducing timing variation.

Effective reuse of the placement and routing from the reference design depends on the design differences between the two variants. Sometimes, small differences in source can have a large impact in the final result, making reuse difficult or less effective.

For more information, see this [link](#) in the *Vivado Design Suite User Guide: Implementation* (UG904) [Ref 19].

Impact of Small RTL Changes

Although synthesis tries to minimize netlist name changes, small RTL changes such as the following can sometimes lead to large design changes:

- Increasing the size of an inferred memory
- Widening an internal bus
- Changing data types from unsigned to signed

Impact of Changing Constraints and Synthesis Options

Similarly, changing constraints and synthesis options such as the following can also have a large impact on incremental placement and routing:

- Changing timing constraints and resynthesizing
- Preserving or dissolving logical hierarchy
- Enabling register re-timing

For more information, see this [link](#) in the *Vivado Design Suite User Guide: Implementation* (UG904) [Ref 19].

Validating the Netlist Quality

To ensure the best possible implementation results, it is important to check the quality of that starting netlist. See [Moving Past Synthesis](#), for instructions on inspecting netlist quality, if these checks were not already made during synthesis stage itself, or if you are not sure of the netlist quality.

Depending on the nature of source files containing the design description, and the state of the design, the following Tcl commands can be used to read the synthesized design into memory:

- `synth_design/launch_runs synth_1`
- `open_checkpoint`
- `open_run`
- `link_design`

Table 5-2: Modes in Which Tcl Commands Can Be Used

Command	Project Mode	Non-Project Mode
synth_design	X (launch_runs synth_1)	X (synth_design)
open_checkpoint		X
open_run	X	
link_design		X

For more information, see *Vivado Design Suite User Guide: Implementation* (UG904) [Ref 19].

Logic Optimization (opt_design)

Vivado Design Suite logic optimization optimizes the current in-memory netlist. Since this is the first view of the assembled design (RTL and IP blocks), the design can usually be further optimized. By default the `opt_design` command performs logic trimming, removing of cells with no loads, propagating constant inputs, and block RAM power optimization. It also optionally performs other optimizations such as remap, which combines LUTs in series into fewer LUTs to reduce path depth.

Constraints and Attributes Affecting Logic Optimization

The Vivado Design Suite respects the `DONT_TOUCH` and `MARK_DEBUG` properties during logic optimization, and does not optimize away nets with these properties. For more information, see the *Vivado Design Suite User Guide: Synthesis* (UG901) [Ref 16].

- `MARK_DEBUG` is placed on nets that are candidates for probing with the Vivado Logic Analyzer tool. A net with `MARK_DEBUG` is connected to a slice boundary to ensure it can be probed.
- The `DONT_TOUCH` property is typically placed on leaf cells to prevent them from being optimized. `DONT_TOUCH` on a hierarchical cell preserves the cell boundary, but optimization may still occur within the cell.
- The `DONT_TOUCH` property might be applied to your design-portions and IP cores that have scoped constraint to make sure that the objects to which the constraints are applied are not optimized out.

Logic Optimization Directives

Directives exist that change the behavior of the `opt_design` command to run multiple passes with and without emphasis on area reduction and to add LUT remapping to the default flow.

For more information on logic optimization directives, see the *Vivado Design Suite User Guide: Synthesis* (UG901) [Ref 16].

Optimization Analysis

The `opt_design` command generates messages detailing the results for each optimization phase. After optimization you can run `report_utilization` to analyze utilization improvements. To better analyze optimization results, use the `-verbose` option to see additional details of the logic affected by `opt_design` optimization.

Power Optimization

For optimizing your design for power, see [Power Optimization](#).

Placement (place_design)

The Vivado Design Suite placer engine positions cells from the netlist onto specific sites in the target Xilinx device. Like the other implementation commands, the Vivado Design Suite placer works from, and updates the in-memory design.

Constraints Affecting Placement

The following constraints affect placement of design objects in the Vivado Design Suite placer:

- I/O Constraints (Example: IOB, IOSTANDARD)
- Location Constraints (Example: LOC, PBLOCK, PROHIBIT)
- Timing Constraints (Example: create_clock)
- Netlist Constraints (Example: LOCK_PINS, CLOCK_DEDICATED_ROUTE)
- RPM and XDC Macros (Example: create_macro)

For more information on placement constraints, see the *Vivado Design Suite User Guide: Using Constraints* (UG903) [\[Ref 18\]](#).

Placement Analysis

Use the timing summary report after placement to check the critical paths.

- Paths with very large negative setup time slack may require that you check the constraints for completeness and correctness, or logic restructuring to achieve timing closure.
- Paths with very large negative hold time slack are most likely due to incorrect constraints or bad clocking topologies and should be fixed before moving on to route design.
- Paths with small negative hold time slack are likely to be fixed by the router. You can also run `report_clock_utilization` after `place_design` to view a report that breaks down clock resource and load counts by clock region.

For more information, see [Timing Closure](#).

If the Vivado Design Suite placer fails to find a solution for the clock and I/O placement, the placer reports the placement rules that were violated; and briefly describes the affected cells.

Placement can fail for several reasons, including:

- Clock tree issues caused by conflicting constraints
- Clock tree issues that are too complex for the placer to resolve
- RAM and DSP block placement conflicts with other constraints such as Pblocks
- Overutilization of resources
- I/O bank requirements and rules

To fix any placement issues, carefully analyze the generated error messages. In many cases, messages refer to intermediate placement that could not lead to a valid solution. Try removing placement constraints that could cause placement issues, or for complex clock tree issues constraining the clock buffers might lead to a successful placement.

SSI Placement

Placement Strategies

Using the built-in placement algorithms, the tools attempt to:

1. Place the design in a way that does not exceed SLL resources.
2. Limit the number of timing critical paths that must cross SLR components.
3. Balance the resources in a way that does not overly fill an SLR with a given resource.
4. Limit the number of SLL crossings to a minimum.

By following these strategies, the tools try to strike a balance placement while meeting performance requirements.

Other Factors That Influence SLR Selection

Other design and implementation factors can also influence SLR selection. These factors include:

1. Pin placement
2. Clock selection
3. Resource type
4. Physical constraints such as floorplanning (PBlocks) and LOC constraints
5. Timing constraints
6. I/O Standards and other constraints

Xilinx recommends that you allow the tools to assign SLR components while making intelligent pin placement, clock selection, and other design choices.

For additional information, see the following sections in this chapter:

- [Placer Directives](#) mentions about directive types that are specifically useful for SSI technology-based designs.
- [Strategies](#) mentions strategies that are specifically useful for SSI technology-based designs.

Manual SLR Assignment

Manual SLR assignment might be necessary when the tools do not find a solution that meets design requirements, or when run-to-run repeatability is important.

Performing Manual SLR Assignment

To perform manual SLR assignment:

1. Create large PBlocks (area groups).
2. Assign portions of the design to those area groups.

To assign large sections of the design to a single SLR:

1. Create a PBlock that encompasses a single SLR.
2. Assign the associated hierarchy of the logic to that PBlock.

While you can assign logic to multiple adjacent SLR components, you must ensure that the PBlock encompasses the entire SLR.

Do not create PBlocks that cross SLR boundaries without constraining the entire SLR. Doing so can make it difficult for the automatic SLR placement algorithms to legalize placement.

Manual SLR Assignment Guidelines

When you manually assign logic to the SLR components, Xilinx recommends that you:

1. Place the design in a way that does not exceed SLL resources.
2. Limit the number of timing critical paths that must cross SLR components.
3. Balance the resources in a way that does not overly fill an SLR with a given resource.
4. Limit the number of SLL crossings to a minimum.

Placer Directives

Because placement typically has a major impact on overall design performance, several Placer directives exist that let you explore the solution space for different scenarios.



TIP: Use the default directive initially. Use other directives when the design nears completion to explore the solution space for a design.

The following table shows which directives may benefit which types of designs.

Table 5-3: Common Scenarios

Directive Type	Designs Benefitted
Block Placement	Designs with many block RAM, DSP blocks, or both
NetDelay	Designs that anticipate many long-distance net connections and nets that fan out to many different modules
SpreadLogic	Designs with very high connectivity that tend to create congestion
ExtraPostPlacement Opt	All design types
SSI	SSI designs that may benefit from different styles of partitioning to relieve congestion or improve timing.

For more information on placer directives, see the *Vivado Design Suite User Guide: Implementation* (UG904) [Ref 19].

Physical Optimization (phys_opt_design)

Physical optimization is an optional step of the flow. It performs timing-driven optimization on the negative-slack paths of a design. Optimizations involve replication, retiming, hold fixing, and placement improvement. Because physical optimization automatically performs all necessary netlist and placement changes, `place_design` is not required after `phys_opt_design`.

Need for Physical Synthesis

To determine if a design would benefit from physical synthesis, evaluate timing after placement. Analyze failing paths for fanout. High fanout critical paths can benefit from fanout optimization. Additionally, high-fanout data, address and control nets of large RAM blocks involving multiple block RAMs that fail timing after `route_design` might benefit from Forced Net Replication. For more information on physical synthesis, see the *Vivado Design Suite User Guide: Implementation* (UG904) [Ref 19].

Constraints Affecting Physical Optimization

Timing constraints impact physical optimization. Most physical optimizations are performed on timing paths that have a negative slack within a percentage of the WNS. The netlist is modified and the changes are incrementally placed. Changes are committed only after slack, area, and power are evaluated.

The Vivado Design Suite respects the `DONT_TOUCH` and `MARK_DEBUG` properties during physical optimization for the same reason it respects them during logic optimization.

Physical Optimization Directives

Several Physical Optimization directives let you explore the solution space for different scenarios.



TIP: Use the default directive initially. Use other directives when the design nears completion to explore the solution space for a design.

For more information on physical optimization directives, see the *Vivado Design Suite User Guide: Implementation* (UG904) [Ref 19].

Routing (`route_design`)

The Vivado Design Suite router performs routing on the placed design, and performs optimization on the routed design to resolve hold time violations. It is timing-driven by default, although this can be disabled.

Constraints Affecting Routing

The following constraints affect routing in the Vivado Design Suite router:

- Fixed Routing (Example: `FIXED_ROUTE`)
- Pin Locking Constrains (Example: `LOCK_PINS`)
- Timing constraints (Example: `create_clock`)

Conflicting constraints will cause errors in the router.

For more information on routing constraints, see the *Vivado Design Suite User Guide: Using Constraints* (UG903) [Ref 18].

Route Analysis

Nets that are routed sub-optimally are often the result of incorrect timing constraints. Before you experiment with router settings, make sure that you have validated the constraints and the timing picture seen by the router. Validate timing and constraints by reviewing timing reports from the placed design before routing.

Common examples of poor timing constraints include cross-clock paths and incorrect multi-cycle paths causing route delay insertion for hold fixing. Congested areas can be addressed by targeted fanout optimization in RTL synthesis or through physical optimization. You can preserve all or some of the design hierarchy to prevent cross-boundary optimization and reduce the netlist density. Or you can use floorplan constraints to ease congestion.

For more information, see [Timing Closure](#).

Intermediate Route Results

When routing fails, the Vivado Design Suite router continues and tries to provide a design that is as complete as possible to aid in debug. If the routing is not complete, you might have to provide manual intervention. Use the following tips to help identify the next steps:

- Run `report_route_status` and check the section "Nets with Routing Errors." Find the net and create a schematic, then look for areas like high fanout nets or clock rule violations. Running the DRC checker can sometimes identify clock rule violations.
- If a physical placement constraint (Pblock) is causing the problem, generate a build with all Pblock constraints removed.
- Review the `vivado.log` file routing section for "Phase 3.2 Budgeting." The amount of congestion is outlined in "levels" where "7" is the highest. Level 7 congestion indicates that a region spanning 2^7 (128) tiles has routing utilization close to 100%. The route direction (north, east, south and west) is reported. The "INT_xxx" numbers are the coordinates of the interconnecting routing tiles that are visible in the device routing resource view.

- Open the routed design check point file (.dcp) and turn on the Device Metric Vertical and Horizontal congestion overlay. This view provides an estimation of routing utilization per logic tile (for example: a CLB). When the estimate is over 100%, the router has to take detours to successfully route all nets passing in the congested area. Look for hot spot areas in the tile area reported in the log file congestion report. From within the device view, select all the cells in the hot spot and generate a schematic. Look for nets with large fanouts that can be buffered by a global clock buffer to free local routing resources, or identify the major blocks to create floorplan constraints or to synthesize with different options or attributes.
- Rerun the design using placement directives that focus on congestion.

For more information on rerouting only specific nets, see [Using Re-Entrant Route Mode](#).

Router Directives

Several Router directives let you explore the solution space for different scenarios.



TIP: Use the default directive initially. Use other directives when the design nears completion to explore the solution space for a design.

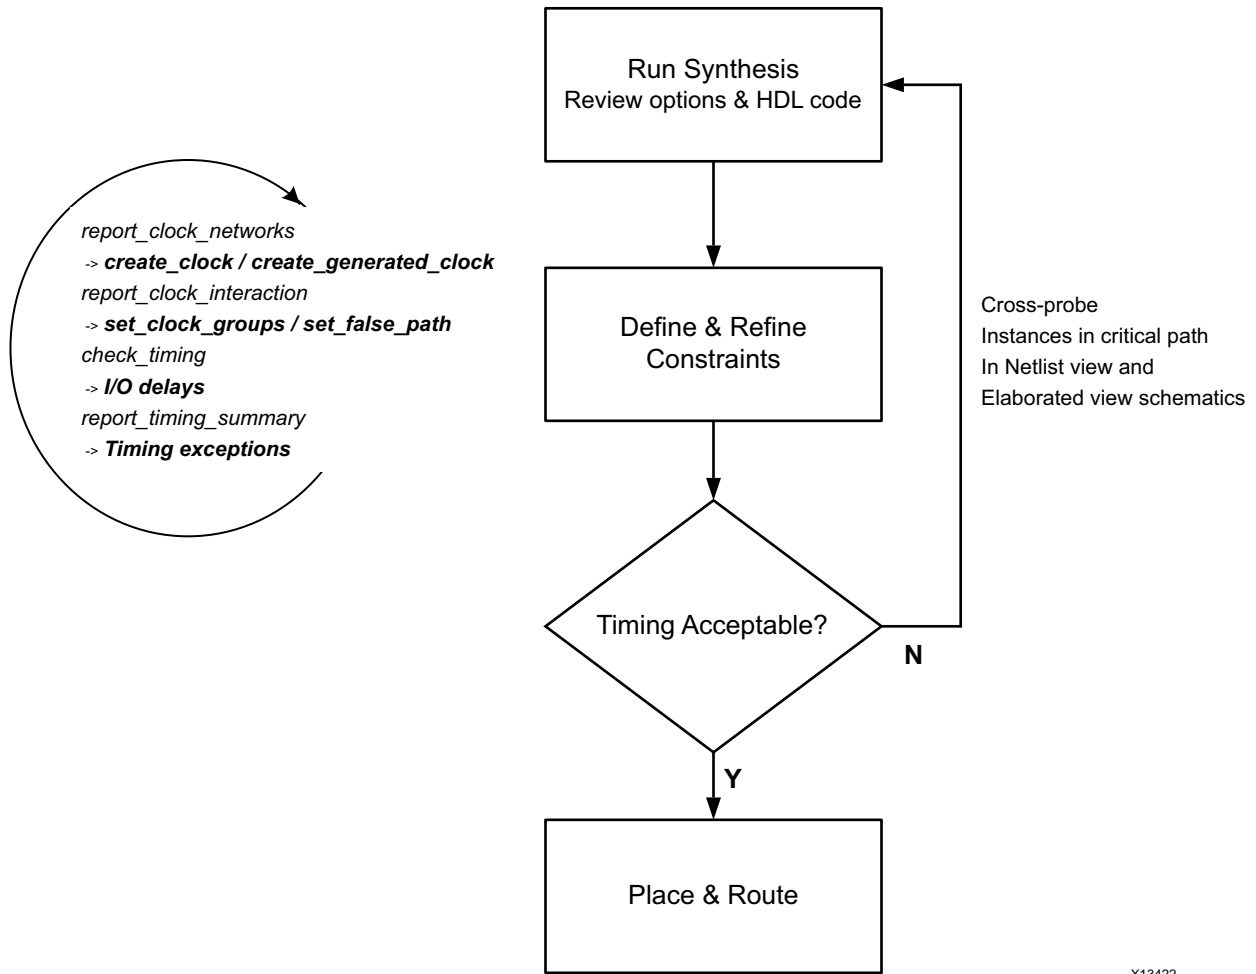
For more information on router directives, see the *Vivado Design Suite User Guide: Implementation* (UG904) [\[Ref 19\]](#).

Using Re-Entrant Route Mode

The `route_design` command is re-entrant in nature. For a partially routed design, the Vivado Design Suite router uses the existing routes as the starting point, instead of starting from scratch. Re-entrant mode is usually run interactively to address specific routing issues such as pre-routing critical nets and locking down resources before a full route; and manually unrouting non-critical nets to free up routing resources for more critical nets. For more information on re-entrant mode, see the *Vivado Design Suite User Guide: Implementation* (UG904) [\[Ref 19\]](#).

Timing Closure

Timing Closure refers to the design being able to meet all its timing requirements. This section explains how to achieve timing closure on your design. Often, users attempt to close their timing through implementation stages only. However, as explained in [Chapter 1, Introduction](#), timing closure would be easier if we have the right HDL and constraints while going into synthesis itself. [Figure 1-3](#) is repeated here to recapitulate the importance of iterating over synthesis stages with improved HDL, constraints, and synthesis options. For more information, see this [link](#) in the *Vivado Design Suite User Guide: Design Analysis and Closure Techniques* (UG906) [\[Ref 21\]](#).



X13422

Figure 5-5: Design Methodology for Rapid Convergence

Follow these guidelines:

- When not meeting timing at first, evaluate timing throughout the flow.
- Focus on WNS (of each clock) as the main way to improve TNS.
- Review large WHS violations (<-1 ns) to identify missing or inappropriate constraints.
- Revisit the trade-offs between design choices, constraints and target architecture.
- Know how to use the tool options and XDC.
- Be aware that the tools do not try to further improve timing (additional margin) once timing is met.

Baselining the Design

Creating baseline constraints means generating the simplest set of timing constraints. Once clocks (including generated clocks) are completely constrained, all paths where the start and the end-points are within the design (all register-to-register paths) are automatically constrained. This allows for an easy mechanism to identify internal device timing challenges, even while the design is evolving. Since the design might also have clock domain crossings, baseline constraints should also include the relationship among the specified clocks (and the generated clocks).

The primary concept behind baselining is to create a minimalistic set of constraints that are correct and that cover most of the timing paths, rather than waiting until all the constraints can be completely specified. Thus, I/O timing definition and closure is deferred until later, when the design has evolved significantly and I/O timings are better known. Accordingly, the baseline constraints do not include I/O timing constraints.

Xilinx recommends that you create the baseline constraints very early in the design process. Any major change to the design HDL should be timed against these baseline constraints. Timing the design updates regularly ensures that any timing bottleneck is caught almost as soon as it is introduced.



TIP: See [Defining Baseline Constraints](#) to create baseline constraints. See [Understanding Timing Reports](#) to understand and interpret the timing reports. See [Debugging and Fixing Timing Issues](#) if you face timing issues with baseline constraints.

The I/O constraints may be added when I/O timing requirements are known and determined.

All Xilinx IPs or Partner IPs are delivered along with specific XDC constraints that comply with Xilinx constraints methodology. The IP constraints are automatically included during synthesis and implementation. You must keep them when creating the design baselining constraints.

As you go through the design flow and refine your constraints, fill in the questionnaire provided in [Appendix A, Baselining and Timing Constraints Validation Procedure](#). This procedure helps track your progress towards timing closure and helps identify potential bottlenecks.

Understanding Timing Reports

The Timing Summary report provide high-level information on the timing characteristics of the design compared to the constraints provided. Review the timing summary numbers during signoff:

- TNS (Total Negative Slack) is the sum of the setup/recovery violations for each endpoint in the entire design or for a particular clock domain. The worst setup/recovery slack is the WNS (Worst Negative Slack).
- THS (Total Hold Slack) is the sum of the hold/removal violations for each endpoint in the entire design or for a particular clock domain. The worst hold/removal slack is the WHS (Worst Hold Slack).
- TPWS (Total Pulse Width Slack) is the sum of the violations for each clock pin in the entire design or a particular clock domain for the following checks:
 - minimum low pulse width
 - minimum high pulse width
 - minimum period
 - maximum period
 - maximum skew (between two clock pins of a same leaf cell)
- WPWS (Worst Pulse Width Slack) is the worst slack for all pulse width, period, or skew checks on any given clock pin.

The Total Slack (TNS, THS or TPWS) only reflects the violations in the design. When all timing checks are met, the Total Slack is null.

The timing path report provides detailed information on how the slack is computed on any logical path for any timing check. In a fully constrained design, each path has one or several requirements that must all be met in order for the associated logic to function reliably.

The main checks covered by WNS, TNS, WHS, and THS are derived from the sequential cell functional requirements:

- The *setup time* is the time before which the new stable data must be available before the next active clock edge in order to be safely captured.
- The *hold requirement* is the amount of time the data must remain stable after an active clock edge to avoid capturing an undesired value.
- The *recovery time* is the minimum time required between the time the asynchronous reset signal has toggled to its inactive state and the next active clock edge.
- The *removal time* is the minimum time after an active clock edge before the asynchronous reset signal can be safely toggled to its inactive state.

A simple example is a path between two flip-flops that are connected to the same clock net.

After a timing clock is defined on the clock net, the timing analysis performs both setup and hold checks at the data pin of the destination flip-flop under the most pessimistic, but reasonable, operating conditions. The data transfer from the source flip-flop to the destination flip-flop occurs safely when both setup and hold slacks are positive.

For more information on timing analysis, see this [link](#) in the *Vivado Design Suite User Guide: Design Analysis and Closure Techniques* (UG906) [Ref 21].

Timing Closure Criteria

Timing closure starts with writing valid constraints that represent how the design will operate in hardware. The following criteria should be met:

- [Clean Constraints](#)
- [No Timing Violation](#)

Clean Constraints

- All active clock pins are reached by a clock definition.
- All active path endpoints have requirement with respect to a defined clock (setup/hold/recovery/removal).
- All active input ports have an input delay constraint.
- All active output ports have an output delay constraint.
- Timing exceptions are correctly specified.



CAUTION! *Excessive use of wildcards in constraints can cause the actual constraints to be different from what you intended.*

Except for constraints delivered with Xilinx IP, input and output delay constraints as well as some path specific timing exceptions are not applicable for baselining constraints.

No Timing Violation

- Setup/Recovery (max analysis): $WNS > 0ns$ and $TNS = 0ns$
- Hold/Removal (min analysis): $WHS > 0ns$ and $THS = 0ns$
- Pulse Width: $WPWS > 0ns$ and $TPWS = 0ns$

Checking That Your Design is Properly Constrained

Before looking at the timing results to see if there are any violations, be sure that every synchronous endpoint in your design is properly constrained.

Run `check_timing` to identify unconstrained paths. This command can be run as a stand-alone command, but it is also part of the `report_timing_summary`.

The `check_timing` command reports the following kinds of situations. These situations indicate something missing or wrong in the timing definition, or implication on correctly meeting timing:

- `no_clock` or `constant_clock`
- `unconstrained_internal_endpoints`
- `no_input_delay`
- `no_output_delay`
- `multiple_clock`
- `generated_clocks`
- `loops`
- `partial_input_delay`
- `partial_output_delay`
- `latch_loops`

no_clock* or *constant_clock

Number of clock pins not reached by a defined timing clock. Constant clock pins are also reported. Check if some clock constraints are missing. Or, if some clock pins have been inadvertently connected to constant.

unconstrained_internal_endpoints

Number of path endpoints (excluding output ports) without a timing requirement. This number is directly related to missing clock definitions, reported by the `no_clock` check.

no_input_delay

Number of non-clock input ports without at least one input delay constraint. Check if some `set_input_delay` is missing.

no_output_delay

Number of non-clock output ports without at least one output delay constraint.

multiple_clock

Number of clock pins reached by more than one timing clock. This can happen if there is a clock multiplexer in one of the clock trees. The clocks that share the same clock tree are

timed together by default, which does not represent a realistic timing situation. Only one clock can be present on a clock tree at any given time.

If you do not believe that the clock tree is supposed to have a MUX, review the clock tree to understand how and why multiple clocks are reaching the specific clock pins.

generated_clocks

Number of generated clocks that refer to a master clock source which is not in the fanin cone of the same clock tree.

loops

Number of combinational loops found in the design. These loops are automatically broken by the Vivado Design Suite timing engine in order to report timing.

partial_input_delay

Number of input ports with only a min input delay or max input delay constraint, but not both. These ports are not analyzed for both setup and hold analysis.

partial_output_delay

Number of output ports with only a min output delay or max output delay constraint, but not both. These ports are not analyzed for both setup and hold analysis.

latch_loops

Checks for and warns of loops passing through latches in the design. These loops will not be reported as part of combinational loops, and will affect latch time borrowing computation on the same paths.

Fixing Issues Flagged by `check_timing`

Not all checks are equally important. The following checks are sorted by importance (most important to least important) when reviewing and fixing the issues flagged by `check_timing`.

No Clock and Unconstrained Internal Endpoints

These are the most important checks. These checks allow you to determine whether the internal paths in the design are completely constrained. You must ensure that the unconstrained internal endpoints are at zero as part of the Static Timing Analysis signoff quality review.

Zero unconstrained internal endpoints should not give a false sense of security. This indicates only that all internal paths are constrained for timing analysis. However, the correct value of the constraints is not yet guaranteed.

Generated Clocks

`Generated_clocks` are a normal part of a design. However, if a generated clock is derived from a master clock which is not part of the same clock tree, this can be a serious problem. The timing engine cannot properly calculate the generated clock tree delay. This results in erroneous slack computation. In the worst case situation, the design meets timing according to the reports, but does not work in hardware.

Loops and Latch Loops

A good design does not have any combinational loops. The timing loop will be broken by the timing engine. The broken paths are not reported during timing analysis, or evaluated during implementation. This can lead to incorrect behavior in hardware, even if the overall timing requirements are met.

No Input/Output Delays and Partial Input/Output Delays

All I/O ports must be properly constrained.



RECOMMENDED: *Start with baselining constraints. Once those have been validated, complete the constraints with the I/O timing.*

Multiple Clocks

Multiple clocks are usually acceptable. Xilinx recommends that you ensure that these clocks are expected to propagate on the same clock tree. You must also verify that the paths requirement between these clocks does not introduce tighter requirements than needed in order for the design to be functional in hardware.

If this is the case, you must use `set_clock_groups` or `set_false_path` between these clocks on these paths. Any time that you use timing exceptions, you must ensure that they affect only the intended paths.



IMPORTANT: *Since the XDC is a Tcl program, the order of constraints matters.*

Debugging and Fixing Timing Issues

The following table provides a quick guidance in terms of systematic approach to debug and fix timing failures (if any) reported in the timing report.

Table 5-4: Steps to Debug and Fix Timing Issues

Step	Section Containing More Details
Check that all clocks and their relationships are defined correctly	Defining Baseline Constraints
Check that clock skew and uncertainty are not too high	Clock Skew and Uncertainty
Check that number of logic levels in the path is not too high	Datapath Delay and Logic Levels
Check that the path uses optimal resources (cells/pins)	MMCM Frequency Synthesis
Check that you do not have too many control sets unnecessarily	Control Sets

After the above aspects have been determined to be good, examine the remaining violations to determine the next course of action involving different options with backend, including manually creating a floorplan, when all else fails.

For information on additional design analysis techniques and how to use the `report_design_analysis` Tcl command, see this [link](#) and this [link](#) in the *Vivado Design Suite User Guide: Design Analysis and Closure Techniques* (UG906) [Ref 21].

Defining Baseline Constraints

If you are unsure of your clock constraints, the Vivado IDE can be used to create a complete set of clock constraints on the post-synthesized netlist. The graphical interface of the IDE and the reporting capabilities of the Vivado Design Suite show precisely what must be constrained.

- [Step 1: Identify Which Clocks Must be Created](#)
- [Step 2: Verify That No Clocks Are Missing](#)
- [Step 3: Identify Asynchronous Clock Domains](#)

Step 1: Identify Which Clocks Must be Created

Begin by loading the post synthesized netlist or checkpoint into the Vivado IDE. In the Tcl console, reset the timing to ensure that all timing constraints are removed. In this way, you can be certain that the slate is clean.

A report of clock networks can be generated in order to create a list of all the primary clocks that must be defined in the design. The resulting list of clock networks shows which clock constraints should be created. Use the clock creation wizard to specify the appropriate parameters for each clock.

Step 2: Verify That No Clocks Are Missing

Once the clock network report shows that all clock networks have been constrained, verification of the accuracy of the generated clocks can begin. Since the Vivado tools automatically propagate clock constraints through clock-modifying blocks such as MMCMs, PLLs and BUFGCTRLs, it is important to review the constraints that were generated. Use `report_clocks` to show which clocks were created with a `create_clock` constraint, and which clocks were generated.

The `report_timing` results show that all clocks are propagated. The difference between the primary clocks (created with `create_clock`) and the generated clocks (generated by clock-modifying-block) is displayed in the attributes field.

- Clocks that are propagated only (P) are primary clocks.
- Clocks that were generated are shown to be both propagated (P) and Generated (G).

You can also create generated clocks using the `create_generated_clock` constraint. For more information, see the *Vivado Design Suite User Guide: Using Constraints* (UG903) [Ref 18].

```

Attributes
P: Propagated
G: Generated
V: Virtual
I: Inverted

Clock      Period      Waveform      Attributes      Sources
sysClk     10.00000    {0.00000 5.00000}  P                {sysClk}
clkfbout   10.00000    {0.00000 5.00000}  P,G              {clkgen/mmcm_adv_inst/CLKFBOUT}
cpuClk     20.00000    {0.00000 10.00000} P,G              {clkgen/mmcm_adv_inst/CLKOUT0}

```

Figure 5-6: Report_Clocks Shows which Clocks were Generated from Primary Clocks

Step 3: Identify Asynchronous Clock Domains

Upon verification of the clocking constraints, asynchronous clock domain crossing paths must be identified.

Note: This section does not explain how to properly cross clock region boundaries, but explains how to identify which crossings exist and how to constrain them.

The clock domain interactions are best viewed using `report_clock_interaction`. The report shows a matrix of source clocks and destination clocks. The color in each cell indicates the nature of interaction among clocks represented by the corresponding row and the column. The following figure shows a sample clock interaction report.



Figure 5-7: Sample Clock Interaction Report

The following table explains the meaning of each color in this report.

Table 5-5: report_clock_interaction Colors

Color	Meaning	What Next
Black	No interaction among these clock domains.	Primarily for information unless you expected these clock domains to be interacting.
Green	There is interaction among these clock domains, and the paths are getting timed.	Primarily for information unless you do not expect any interaction among the clock domains.
Cyan	Some of the paths for the interacting domains are not being timed due to user exceptions.	Ensure that the timing exceptions are really desired.
Red	There is interaction among these clock domains, and the paths are being timed. However, the clocks appear to be independent (and hence, asynchronous)	Check whether these clocks should have been declared as asynchronous, or whether they should be sharing a common primary source.
Orange	There is interaction among these clock domains. The clocks appear to be independent (and hence, asynchronous). However, only some of the paths are not timed due to exceptions.	Check why are only a few paths getting user exception? Should all the paths be getting the exception?
Blue	There is interaction among these clock domains, and the paths are not being timed.	Confirm that these clocks are supposed to be asynchronous. Also, check that the corresponding HDL code has been written correctly to ensure proper synchronization and reliable data transfer across clock domains.
Light blue	There is interaction among these clock domains, and the paths are getting timed through: <code>set_max_delay -datapath only.</code>	Confirm that the clocks are asynchronous and that the specified delay is correct.

Before the creation of any false paths or clock group constraints, the only colors that appear in the matrix are black, red, and green. Because all clocks are timed by default, the process of decoupling asynchronous clocks takes on a high degree of significance. Failure to decouple asynchronous clocks often results in a vastly over-constrained design.

Identify Clock Pairs That Do Not Share Common Primary Clocks

The clock interaction report indicates whether or not each pair of interacting clocks has a common primary clock source. Clock pairs that do not share a common primary clock are frequently asynchronous to each other. As such, it is helpful to identify these pairs by sorting the columns in the report using the Common Primary Clock field. The report does not determine whether clock-domain crossing paths are or are not designed properly. For information about the proper design of clock-domain crossing paths, [Chapter 4, Design Creation](#).

Identify Tight Timing Requirements

For each clock pair, the clock interaction report also shows the path requirement for all paths that cross from source clock to destination clock. Sort the columns by path requirement (WNS) to view a list of the tightest requirements in the design. [Figure 5-7](#) shows the timing report sorted by WNS column. Review these requirements to ensure that no invalid tight requirements exist.

The Vivado tools identify the path requirements by expanding each clock out to one thousand cycles, then determining where the closest, non-coincident edge alignment occurs:

Consider a timing path that crosses from a 250 MHz clock to a 200 MHz clock.

- The positive edges of the 200 MHz clock are {0, 5, 10, 15, 20 ...}.
- The positive edges of the 250 MHz clock are {0, 4, 8, 12, 16, 20 ...}.

The tightest requirement for this pair of clocks occurs when:

- The 250 MHz clock has a rising edge at 4 ns, and
- The next rising edge of the 200 MHz clock is at 5 ns.

This results in all paths timed from the 250 MHz clock domain into the 200 MHz clock domain being timed at 1 ns.

Note: The simultaneous edge at 20 ns is NOT the tightest requirement in this example, because the capture edge cannot be the same as the launch edge.

Because this is a rather tight timing requirement, additional steps must be taken.

Depending on the design, one of the following constraints might be the correct way to handle these crossings:

- `set_clock_groups`
- `false_path`
- `max_delay_path`
- `multicycle path`

If nothing is done, the design may exhibit timing violations that cross these two domains. Furthermore, all of the best optimization, placement and routing may end up being dedicated to these paths instead of given to the critical paths in the design. It is critical that these types of paths be identified before any timing-driven implementation step.




Figure 5-8: Clock Domain Crossing from 250MHz to 200 MHz

Using Report_Clock_Networks to Decouple Primary Clocks and Generated Clocks

Before any timing exceptions are created, it is helpful to go back to report_clock_networks to identify which primary clocks exist in the design. If all primary clocks are asynchronous to each other, as is often the case, a single constraint can be used to decouple the primary clocks from each other, and to decouple their generated clocks from each other. Using the primary clocks in report_clock_networks as a guide, each clock group and associated clocks can be decoupled as shown in the following figure.




Figure 5-9: Report Clock Networks

```

### Decouple asynchronous clocks
set_clock_groups -asynchronous \
-group [get_clocks sysClk -include_generated_clocks] \
-group [get_clocks gt0_txusrclk_i -include_generated_clocks] \
-group [get_clocks gt2_txusrclk_i -include_generated_clocks] \
-group [get_clocks gt4_txusrclk_i -include_generated_clocks] \
-group [get_clocks gt6_txusrclk_i -include_generated_clocks]
    
```

Limiting I/O Constraints and Timing Exceptions

Most timing violations are on internal paths. I/O constraints are not needed during the first baselining iterations, especially for I/O timing paths in which the launching or capturing register is located inside the I/O bank. The I/O timing constraints can be added back once the design and other constraints are stable and the timing is nearly closed.

Timing exceptions must be limited, based on recommendations of the RTL designer, and must not be used to hide real timing problems. The false path or clock groups between clocks must have already been reviewed and finalized at this point.

IP constraints must be entirely kept. When IP timing constraints are missing, known false paths can end up being reported as timing violations.

Evaluating Design WNS Before and After Each Step

You must evaluate the design WNS after each implementation step. Tcl users of Tcl command line flow can easily incorporate `report_timing_summary` after each implementation step in their build script. IDE users can make use of simple `tcl.post` scripts to run `report_timing_summary` after each step. In both cases, when a significant degradation in WNS is noted, you must analyze the checkpoint immediately preceding that step.

In addition to evaluating the timing for the entire design before and after each implementation step, a more targeted approach can be taken for individual paths in order to evaluate the impact of each step in the flow on the timing. For example, the estimated net delay for a timing path after the optimization step may differ significantly from the estimated net delay for the same path after placement. Comparing the timing of critical paths after each step is an effective method for highlighting where the timing of a critical path diverges from closure.

Post Synthesis and Post Logic Optimization

Estimated net delays are close to the best possible placement for all paths. To fix violating paths any of the following:

- Change the RTL
- Use different synthesis options
- Add timing exceptions such as multicycle paths (if appropriate and safe for the functionality in hardware)

Pre- and Post-Placement

After placement, the estimated net delays are close to the best possible route, except for long and medium-to-high fanout nets, which use more pessimistic delays. In addition, congestion or hold fixing impact are not accounted for in the net delays at this point, which can make the timing results optimistic.

Clock skew is accurately estimated and can be used to review imbalanced clock trees impact on slack.

Hold fixing can be estimated by running min delay analysis. High violations require clock tree modification. Small violations are acceptable and will likely be fixed by the router.

Pre- and Post-Physical Optimization

Evaluate the need for running physical optimization in order to fix timing problems related to:

- Nets with high fanout (`report_high_fanout_nets` shows highest fanout non-clock nets)
- Nets with targets located far apart
- DSP and RAMB with sub-optimal pipeline register usage

Pre and Post-Route

Slack is reported with actual routed net delays except for the nets that are not completely routed. Slack reflects the impact of hold fixing on setup; and the impact of congestion.

No hold violation should remain after route, regardless of the worst setup slack (WNS) value. If the design fails hold, further analysis is needed. This is typically due to very high congestion, in which case the router gives up on optimizing timing. This can also occur for high hold violations (over 4 ns) which the router does not fix by default. High hold violations are usually due to improper clock constraints, high clock skew or, improper I/O constraints which can already be identified after placement or even after synthesis.

If hold is met ($WHS > 0$) but setup fails ($WNS < 0$), follow the analysis steps described below.

Identifying Timing Violations Root Cause

The timing-driven algorithms focus on the worst violations. Understanding and fixing problems related to the worst violation will likely resolve most, if not all, smaller violations on re-running the implementation flow.

For setup, you must first analyze the worst violation of each clock group.

- Clock group = all intra, inter and asynchronous paths captured by a given clock

For hold, all violations must be reviewed, starting with the worst one.

Several factors can impact the setup and hold slacks. You can easily identify each factor by reviewing the setup and hold slack equations when written in the following simplified form:

Slack (setup/recovery) = setup path requirement
 - datapath delay(max)
 + clock skew
 - clock uncertainty
 - setup/recovery time.

Slack (hold/removal) = hold path requirement
 + datapath delay(min)
 - clock skew
 - clock uncertainty
 - hold/removal time.

For timing analysis, clock skew is always calculated as follows:

Clock Skew = destination clock delay - source clock delay (after the common node if any)

During the analysis of the violating timing paths, you must review the relative impact of each variable to determine which variable contributes the most to the violation. Then you can start analyzing the main contributor to understand what characteristic of the path influences its value the most and try to identify a design or constraint change to reduce its impact. If a design or constraint change is not practical, you must do the same analysis with all other contributors starting with the worst one. The following list shows the typical contributor order from worst to least.

For setup/recovery:

- Datapath delay: Subtract the timing path requirement from the datapath delay. If the difference is comparable to the (negative) slack value, then either the path requirement is too tight or the datapath delay is too large.
- Datapath delay + setup/recovery time: Subtract the timing path requirement from the datapath delay plus the setup/recovery time. If the difference is comparable to the (negative) slack value, then either the path requirement is too tight or the setup/recovery time is larger than usual and noticeably contributes to the violation.
- Clock skew: If the clock skew and the slack have similar negative values and the skew absolute value is over a few 100 ps, then the skew is a major contributor and you must review the clock topology.
- Clock uncertainty: If the clock uncertainty is over a few 100 ps, then you must review the clock topology and jitter numbers to understand why the uncertainty is so high.

For hold/removal:

- Clock Skew: If the clock skew is over 500 ps, you must review the clock topology.
- Clock uncertainty: If the clock uncertainty is over a few 100 ps, then you must review the clock topology and jitter numbers to understand why the uncertainty is so high.
- Hold/removal time: If the hold/removal time is over a few 100 ps, you can review the primitive data sheet to validate that this is expected.
- Hold path requirement: The requirement is usually zero. If not, you must verify that your timing constraints are correct.

Assuming all timing constraints are accurate and reasonable, the most common contributors to timing violations are usually the datapath delay for setup/recovery timing paths, and skew for hold/removal timing paths. At the early stage of a design cycle, you will fix most timing problems by analyzing these two contributors. But after the design and constraints have been improved and refined, the remaining violations, if any, usually come from a combination of factors in which case you must review all contributors in parallel and identify which one you can still improve.

There are multiple ways to perform design analysis. For information, see this [link](#) in the *Vivado Design Suite User Guide: Design Analysis and Closure Techniques* (UG906) [Ref 21].

Datapath Delay and Logic Levels

In general, the number of LUTs and other primitives in the path is most important factor in contributing to the delay.

If the path delay is dominated by:

- **Cell delay is 50% to 100%**

Can the path be modified to be shorter or to use faster logic cells? See [Reviewing Technology Choices](#).

For information on using the elaborated view to optimize the RTL, see this [link](#) in the *Vivado Design Suite User Guide: Design Analysis and Closure Techniques* (UG906) [Ref 21].

- **Route delay is 50%-100%**

Was this path impacted by hold fixing? Use the corresponding analysis technique.

- Yes - Is the impacted net part of a CDC path?
 - Yes - Is the CDC path missing a constraint?
 - No - Do the startpoint and endpoint of that hold-fixed path use a balanced clock tree? Look at the skew value.
- No - See the following information on congestion.

Was this path impacted by congestion? Look at each individual net delay, the fanout and observe the routing in the Device view with routing details enabled (post-route analysis only). You can also turn on the congestion metrics to see if the path is located in or near a congested area.

- Yes - For the nets with the highest delay value, is the fanout low (<10)?
 - Yes - If the routing seems optimal (straight line) but driver and load are far apart, the sub-optimal placement is related to congestion. Try to move manually the driver, or the load, and re-run timing analysis on the same path to see if that slack improves without degrading other paths. After doing the same exercise for several nets, create floorplanning constraints which will ensure that a similar placement solution will be used the next time the implementation tools are run.
 - No - Try to use physical logic optimization to duplicate the driver of the net. Once duplicated, each driver can automatically be placed closer to its loads, which will reduce the overall datapath delay.
- No - The design is spread out too much. Start working on floorplanning to identify portions of the design that must be kept in particular region based on their connection to I/O components (if any) or any other particular anchor point. See the floorplanning section for more information.

For more information on congestion and timing path characteristics, see this [link](#) in the *Vivado Design Suite User Guide: Design Analysis and Closure Techniques* (UG906) [Ref 21].

Clock Skew and Uncertainty

Xilinx FPGA devices use various types of routing resources to support most common clocking schemes and requirements such as high fanout clocks, short propagation delays, and extremely low skew. Clock skew affects any register-to-register path with either a combinational logic or interconnect between them.



RECOMMENDED: Run a design analysis report (`report_design_analysis`) to generate a timing report, which includes information on clock skew data. Verify that the clock nets do not contain excessive clock skew.

Clock skew in high performance clock domains (+300 MHz) can impact performance. The clock skew should be no more than 15% of the period. In the example of 300 MHz, the maximum should be 500 ps in a single clock domain. In cross domain clock paths the skew can be higher, because the clocks use different resources and the common node is located further up the clock trees. SDC-based tools time all clocks together unless constraints specify that they should not be, for example:

```
set_clock_groups/set_false_path/set_max_delay -datapath_only
```

If you suspect high clock skew, conduct a timing analysis on that path in the Vivado IDE and create a schematic to investigate the clocking topology.

Debugging Timing Reports with High Clock Skew

You must first understand the source clock and destination clock and their relationship:

- [When the Source and Destination Clocks Are Synchronous](#)
- [When the Source and Destination Clocks Are Asynchronous](#)

When the Source and Destination Clocks Are Synchronous

When the source and destination clocks are the same or derived from the same primary clock (synchronous clocks), the tools use the common node on the clock path to determine the clock skew. Most synchronous paths have a common node. When analyzing the clock path in the timing report, the delays before and after the common node are not provided separately because the common node only exists in the physical database of the design and not in the logical view. For this reason, you can see the common node in the device view of the Vivado IDE when routing details are turned on, but not in the schematic view. The timing report only provides a summary of skew calculation with source clock delay, destination clock delay, and credit from clock pessimism removal (CPR) up to the common node.

For some synchronous clocks, the common node is located prior to a clock modifying block, such as an MMCM or a PLL, so the skew can also be very high. Xilinx recommends avoiding such synchronous timing paths because they can cause difficulty in meeting timing. Instead you must consider treating them as asynchronous paths by adding timing exceptions and implementing asynchronous clock domain crossing circuitry.

When the Source and Destination Clocks Are Asynchronous

When the source and destination clocks are not the same and originate from different primary clocks, there is no common node. The skew corresponds to the difference between the entire tree delays from the primary clock sources (usually input ports) to the sequential cells of the path. Depending on the clock topology and the placement, the skew can be very large and make it impossible to meet timing. You must review all timing paths between asynchronous clocks and add timing exceptions to either completely ignore the timing analysis by adding `set_clock_groups` or `set_false_path` constraints or just ignore the clock skew and uncertainty by adding `set_max_delay -datapath_only` constraints.



TIP: *The best way to analyze the clock paths is to use the schematic viewer in the Vivado IDE and cross probe with the timing report.*

Causes of High Clock Skew

High clock skew can be caused by:

- [Clock Signal Driven From a Gated Logic Source](#)
- [Serially Connected BUFG Components Driving Synchronous Elements](#)
- [BUFG Drives Synchronous Elements](#)
- [IBUFG Drives Multiple MMCMs \(Related Clocks\)](#)
- [BUFG Drives Register Elements and MMCMs \(Related Clocks\)](#)
- [BUFR/BUFIO/BUFH Drives Register Elements in Several Clock Regions](#)
- [Using the CLOCK_DEDICATED_ROUTE=FALSE Constraint](#)

Clock Signal Driven From a Gated Logic Source

This method is not recommended and can lead to excess clock skew. Since the gated logic driver buffer does not have direct access to the global clock lines, it uses local fabric routing resources. In some cases, although the gated logic can be connected to a BUFG, this also leads to excessive route delay. Check the `report_clock_utilization` results for excessive clock skew.

When verifying the clock report, there might be additional high skew clocks automatically inserted by the place and route algorithms. This is a common practice to prevent long term

silicon metastability in unused clock generators. The clocks operate in the low frequency (Hz) range with minimal resources and do not impact design performance.

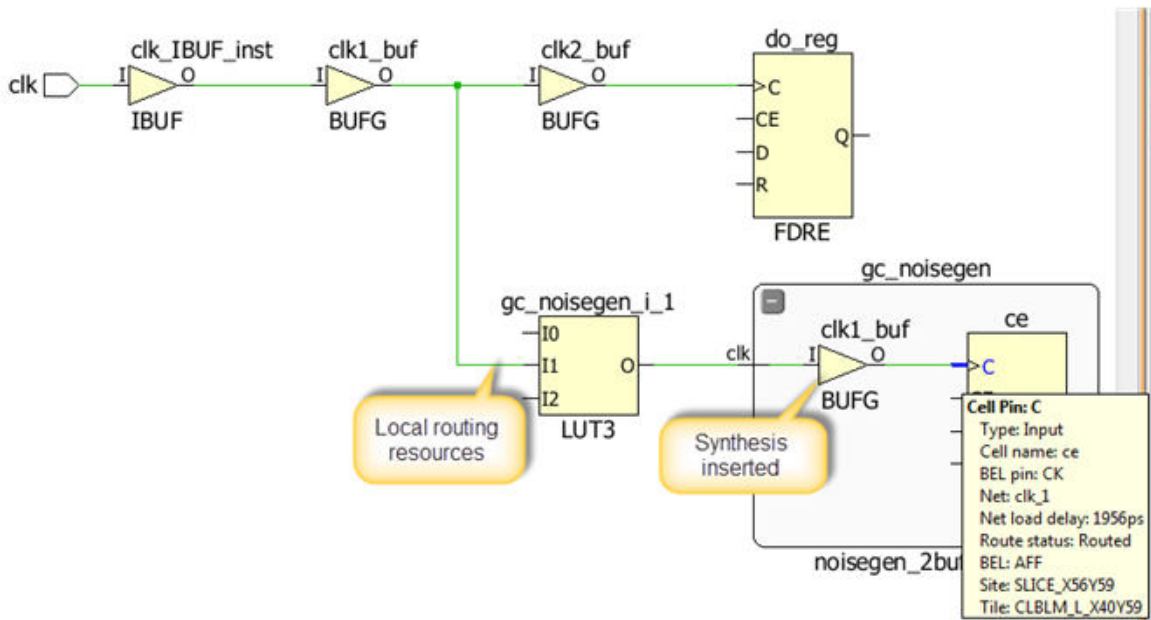


Figure 5-10: Skew Due to Local Routing on Clock Network

In Figure 5-10, the first BUFG (`clk1_buf`) is used in LUT3 to create a gated clock condition. This practice is not recommended. In order to comply with the requirement, the connection uses slower local routing connections. The second BUFG in block `gc_noisegen` was inserted automatically by the synthesis algorithm.

Serially Connected BUFG Components Driving Synchronous Elements

When adding synthesized IP netlists, verify that the number of global clock buffers inserted by the synthesis tool is correct. A common mistake when importing black box IP is that the synthesizer automatically inserts a BUFG when it detects a signal connected to the CLK port of a cell. If the downstream black box IP contains a global clock buffer, the two BUFG components will increase the amount of clock skew.

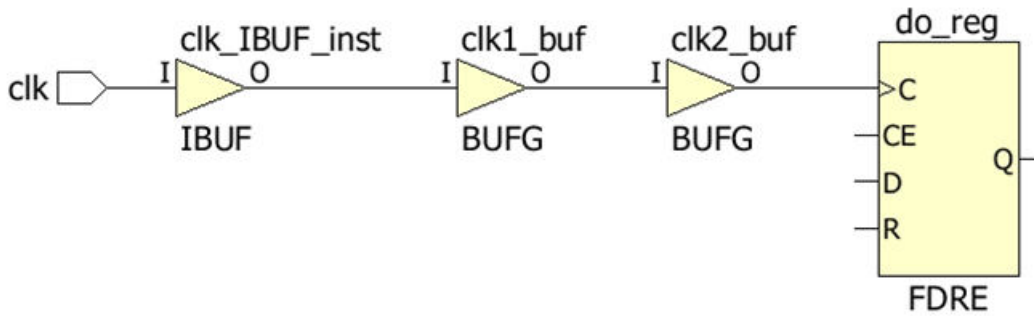


Figure 5-11: Skew Due to Cascaded BUFG

In the above example, the clock net delay is 2.362 ns at the clock pin of the register. If the BUFG is not driven by an MMCM, no PVT and fabric skew is compensated for.



TIP: If extra MMCM is available, use it to reduce clock skew.

BUFG Drives Synchronous Elements

Each clock region contains identical clock routing. The relative location of the source and destination clock pins on the clock tree determines the difference in clock skew. If the source and destination clock delay from the common node are the same, clock skew will be minimal.

Higher than normal clock skew is common if the source and destination are in different clock regions, or in different SLRs. Keeping the source and destination within one clock region helps to minimize clock skew. AREA GROUPS or PBLOCKS may be used to force the source and the destination elements to lie in the same clock region.

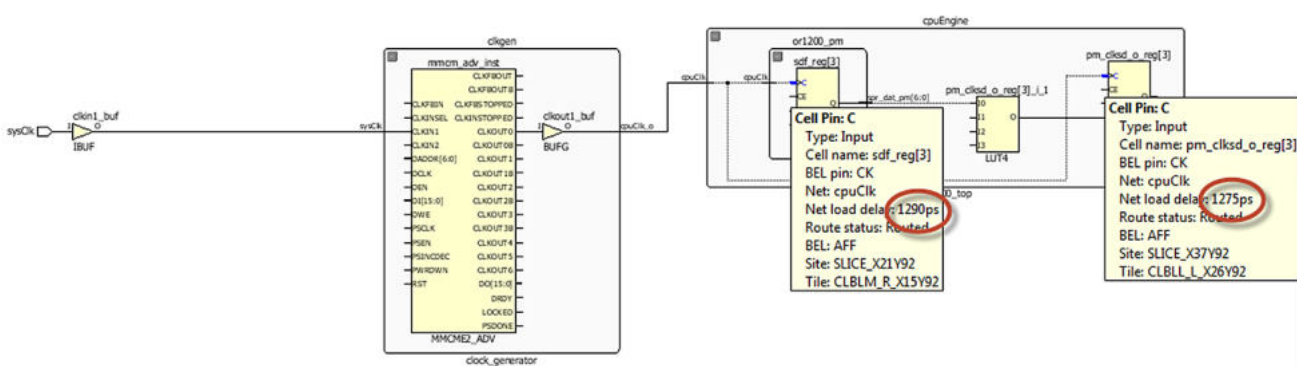


Figure 5-12: Low Skew Due to Source and Destination in the Same Clock Region

In Figure 5-12, the clock net delay is 1.290 ns at the clock pin of the source register and 1.275 ns at the destination register. This results in a clock skew of only 15 ps over PVT. This clock net has a max skew of 0.287 ns across all destinations.

```

Num Loads
Index BUFG cell Net Name BELs Sites Locked MaxDelay (ns) Skew (ns)
-----
10 clkgen/clkout1_buf clkgen/cpuClk_o 3297 1298 no 1.37 0.287

```

IBUFG Drives a Single MMCM with Multiple Outputs (Related Clocks)

The period constraint is defined on the driver pin or port of its tree root. If the clock signal drives an MMCM to generate multiple common output frequencies, the skew from each related clock is the same to the output BUFG. The source and destination may be located in different clock regions. Xilinx recommends that you use AREA GROUPS or PBLOCKS if it impacts your timing performance. In the above example, the clock net skew between the two clock domains is 36 ps.



TIP: *The Clocking Wizard provides performance guidelines (jitter and phase error) based on your clocking requirements.*

IBUFG Drives Multiple MMCMs (Related Clocks)

Xilinx recommends using a simplified clocking topology if possible. Consolidating the number of clock domains helps with performance, resources, and timing closure.



TIP: *Be careful of using a MMCM output BUFG that drives other MMCMs and other registers. The extra BUFG can lead to additional skew.*

BUFG Drives Register Elements and MMCMs (Related Clocks)

Make sure the MMCM CLKIN BUFG is used to drive only the MMCM if possible. You can use the CLK0 output of the MMCM to drive your registered elements. The MMCM provides clock stability over PVT where the BUFG does not. Given the complexity of today's designs, it is possible that the timing engine will detect a cross domain clock path somewhere downstream that might not be detected by the designer.

In the following figure, the lock signal from the MMCM is monitored with the input clock. This is a common practice if the clock source is interrupted. The number of resources on signal `clockRef` is minimal.

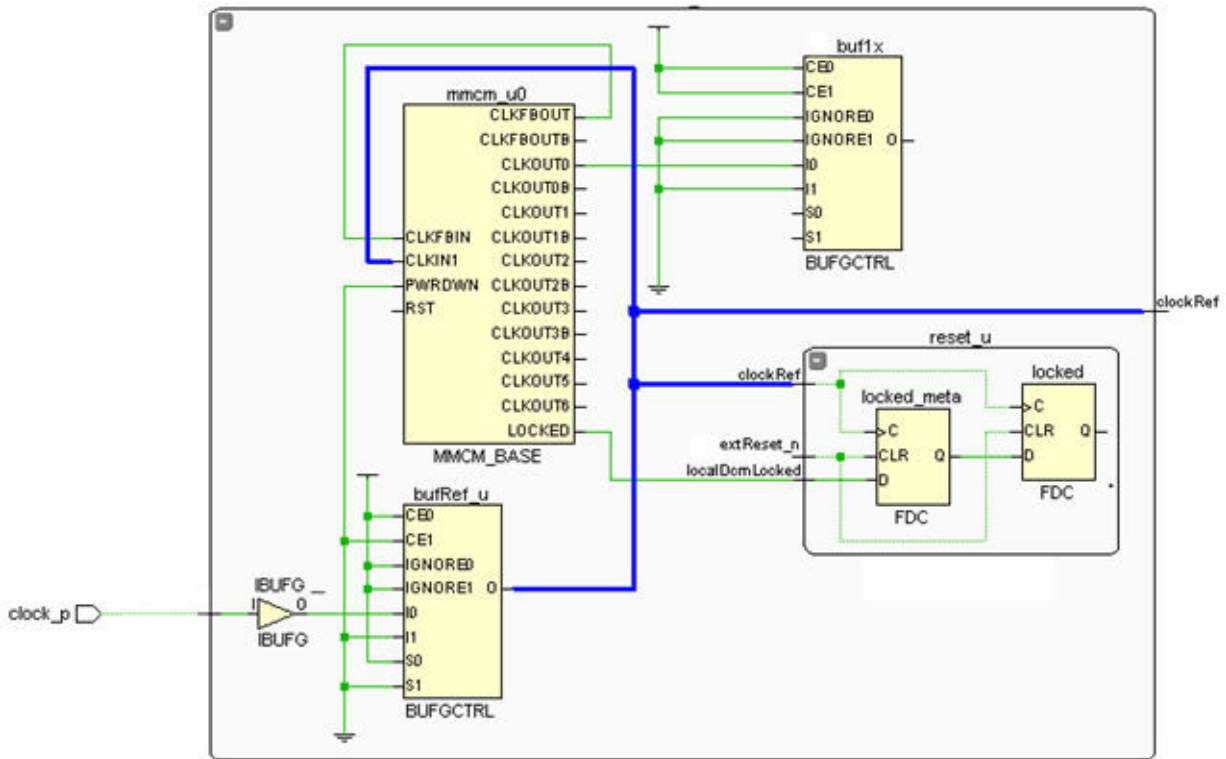


Figure 5-13: Clocks Driven by MMCM

BUFR/BUFIO/BUFH Drives Register Elements in Several Clock Regions

The `clock_report_utilization` reports all types of regional clock buffers. Verify that the clock skew for each regional clock is reasonable ($\ll 1$ ns.). In the following example, BUFR clock skew is very high and shows that the destination elements violated the clocking rules. (BUFR can only drive resources in the region it is located.)

Details of Regional Clocks

```

-----
Num Loads
Index BUFR cell Net Name BELs Sites Locked MaxDelay (ns) Skew (ns)
-----
1 u0_pcie/txoutclk_i u0_pcie/refclk 1 2 no 0.594 0.055
2 u0_pcie/usrclk1_i1 u0_pcie/pipe_userclk1_in 11 25 no 5.93 5.36
3 u0_pcie/usrclk2_i1 u0_pcie/pipe_userclk2_in 463 160 no 0.728 0.202
4 u0_pcie/pclk_i1 u0_pcie/pipe_bclk_in 557 248 no 0.952 0.396
-----

```

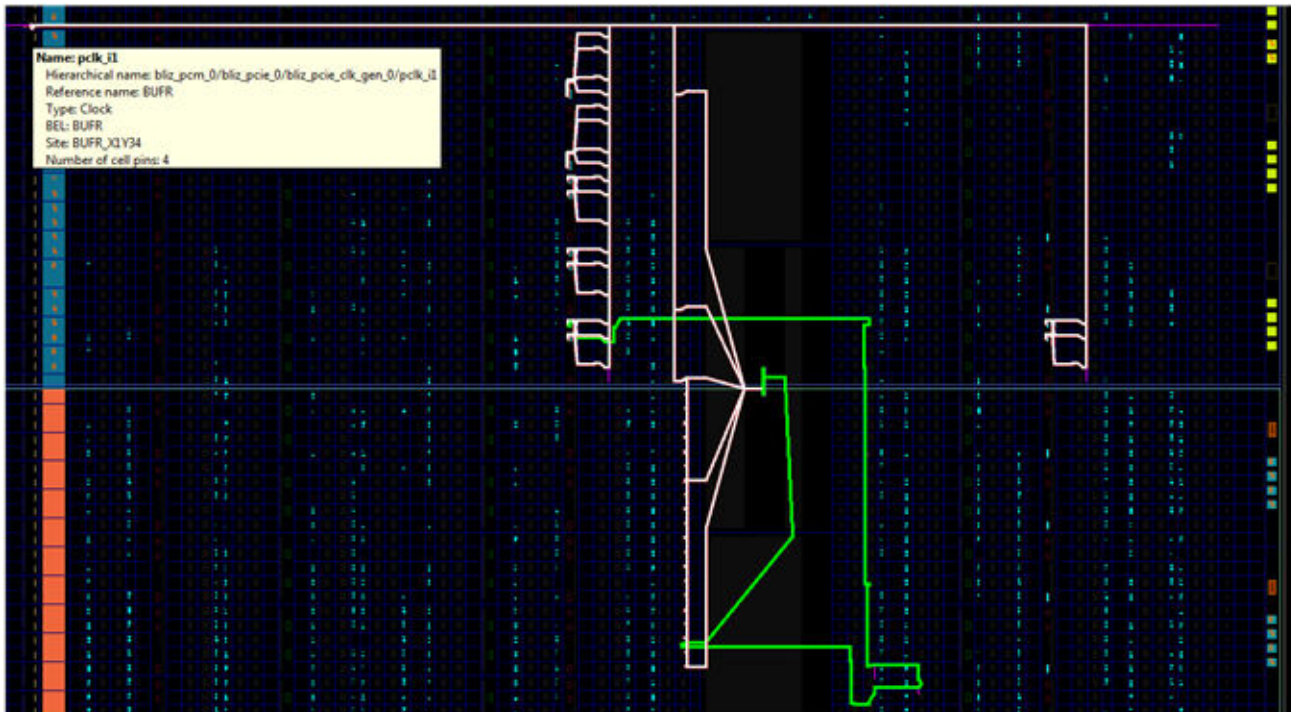


Figure 5-14: BUFR Driving Flops in Two Regions

Using the CLOCK_DEDICATED_ROUTE=FALSE Constraint

Do not use the CLOCK_DEDICATED_ROUTE=FALSE constraint in a production design.

Use CLOCK_DEDICATED_ROUTE=FALSE only as a temporary workaround to a clock failure ONLY to obtain the design through the place and route in order to view the clocking topology in the device and schematic viewer for debugging. These types of paths can have high clock skew leading to poor performance or non-functional designs. In the following figure, the right side has a dedicated clock route, while on the left side, the dedicated route is disabled for clock.




Figure 5-15: Use of Clock Dedicated Routing

Causes of High Uncertainty

Uncertainty is the total amount of uncertainty (relative to an ideal clock) that results from user-specified external clock uncertainty, jitter, or duty cycle distortion. Clock modifying blocks such as MMCM and PLL generate clock uncertainty.

The Clocking Wizard provides accurate uncertainty data for the specified device. The Clocking Wizard can also generate various MMCM clocking configurations for comparing different topologies.

It is common to see clocking topologies created for older FPGA architectures from legacy code that has not been migrated to newer device technologies. Xilinx recommends recreating the clocking section using the target device so that system performance parameters and DRC rules are calculated and verified.

MMCM

MMCM filters input clock uncertainty as they regenerate the required clocks. The MMCM generates some clock uncertainty composing of System Jitter, Discrete Jitter and Phase Error if multiple related clocks are used.

Designs in which phase alignment does not affect system performance, such as logic clocked by the same output from an MMCM, as shown in the following figure, will not be affected.




Figure 5-16: No Impact of Phase Error Through MMCM

MMCM and I/O Timing

Designs in which phase alignment between the input and output of the MMCM is important should be checked to ensure all timing constraints are still met (that is, `set_input_delay`, and `set_output_delay`). See the following figure.

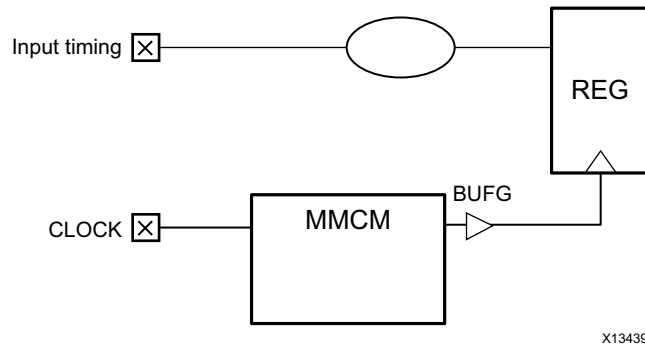


Figure 5-17: Phase Alignment Through MMCM Might Impact Timing

MMCM Frequency Synthesis

When configuring the MMCM for frequency synthesis, the target frequency may have several M (multiplier) and D (divider) values. To minimize clock uncertainty, use values that generate a higher VCO frequency remembering not to exceed the maximum MMCM VCO frequency switching characteristics of that device. If you are migrating your design from an older technology, make sure that you are modifying the M and D values to provide the highest VCO frequency for the current technology.

The MMCM frequency synthesis example below uses an input clock of 62.5 MHz to generate an output clock of around 40 MHz. There are two solutions, but only one (MMCM_2) generates less jitter and lower clock uncertainty.

Table 5-6: MMCM Frequency Synthesis Example

	MMCM_1	MMCM_2
Input clock	62.5 MHz	62.5 MHz
Output clock	40.0 MHz	39.991 MHz
CLKFBOUT_MULT_F	16	22.875
CLKOUT0_DIVIDE_F	25	35.750
VCO Frequency	1000.000 MHz	1429.688
Jitter (ps)	167.542	128.632
Phase Error (ps)	384.432	123.641

When using Clocking Wizard from the IP Catalog, make sure that Jitter Optimization Setting is set to Minimum Output Jitter, which will provide the higher VCO frequency.

Reviewing Technology Choices

It is important to be aware of how design and synthesis choices impact the overall timing, utilization, and power of a design. There are often many different resource types to implement the same logic function and the choice of resources can have a significant impact. For example, a RAM implemented using distributed RAM performs differently than when using block RAM. When designing with attention to technological details, good trade-offs can be made to help improve the quality of results.

The logic fabric is constructed of configurable blocks with each block sharing the same control signals. The smallest block entity is a slice or CLB (Configurable Logic Block), depending on the architecture. The following discussions refer to fabric logic as CLBs, except when considering specific technologies such as Xilinx series FPGA devices, which use slices.

In addition to clocks, sequential primitives require control signals such as resets, sets, and clock enables. The fact that many resources share the same control signals limits their use. Inefficient use of control signals can lead to inefficient use of device resources and packing of logic. This can later lead to other problems such as routing congestion and overutilization of slices and CLBs.

This section gives an overview of the impact of technical choices of combinational and sequential logic resources and control signal implementation. Although the examples and figures are based on general characteristics of Xilinx 7 series FPGA device technology, similar analysis can be applied to UltraScale™ devices and future technologies. For full details on timing parameters, see *AC Switching Characteristics* in the device data sheet.

Combinational Logic: LUT Pin Delays

Not all paths through a LUT have the same delay. In a timing report, each input pin delay appears the same, but there is some extra wire delay consolidated into the net delay

leading to each input. This is due to the physical implementation of the related interconnect. Although the tools try to implement logic such that critical signals use the fastest inputs, awareness of this aspect can help improve complex timing failures, and analyze suboptimal LOCK_PINS constraints.

LUTs can be described as logical pins or physical pins. Logical pins are named I0, I1, I2, I3, I4, and I5, and reflect the netlist connections. These are also the names used in the timing reports. Physical pins are named A1, A2, A3, A4, A5, and A6, and may use different letters B, C, or D depending on the BEL used. The physical pins represent the actual device pins, and are typically only seen when analyzing the physical implementation at the device level.

Logical pins are mapped to physical pins with mapping chosen by the placer, router, or by LOCK_PINS properties. In general, the physical A6 LUT input is the fastest path, followed by A5, A4, and so on, down to A1 which is the slowest. The A6 path is typically a few hundred picoseconds faster than the slowest path through A1.

For each LUT, the logical pin to physical pin mapping can be seen after placement. For example, a LUT6 has the following pin mappings by default as can be seen in the Cell Pins tab of the cell properties, or by using the `get_site_pins` command.

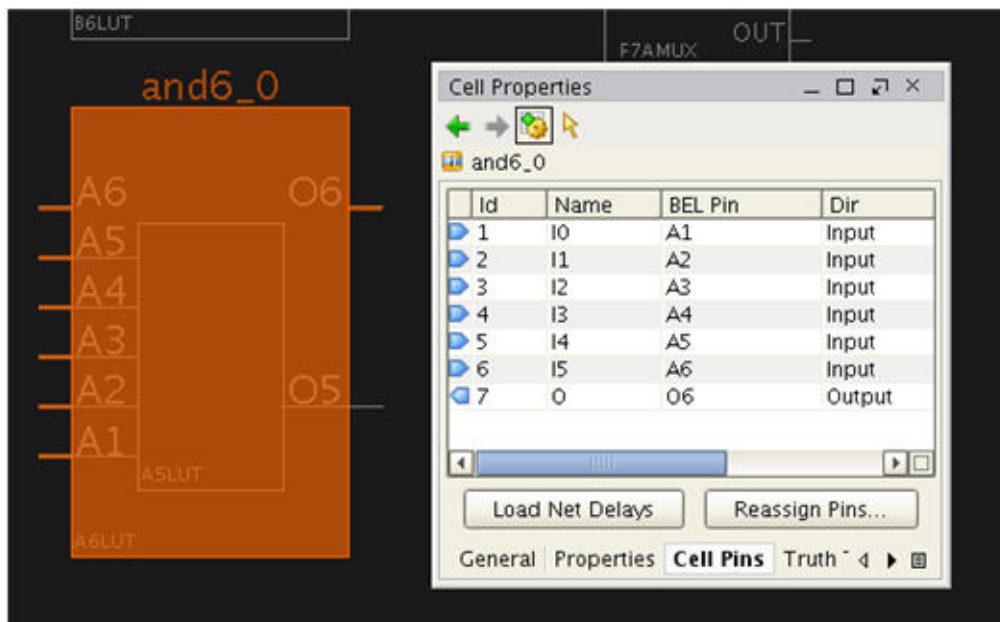


Figure 5-18: LUT Logical and Physical Pins

The BEL Pin column shows the physical pin mapped from the logical pin. Accordingly:

```
I0 maps to A1
I1 maps to A2
.
.
.
I5 maps to A6
```

During implementation, placement, physical optimization, and routing may swap LUT input pins to improve critical path timing. A timing-critical logical pin is moved to a faster physical pin such as A6 while the slower logical pin moves to a slower physical pin. For a critical path that traverses several LUTs, the difference between using the fastest physical pins versus the slowest can be quite pronounced. Pin-swapping can be overridden by setting a LOCK_PINS property on the cell to define its explicit mapping.

Combinational Logic: Combining LUTs

The logic LUT of Xilinx 7 series FPGA devices is designed to be flexible to accommodate more than one single 6-input function. It has two outputs (O6 and O5) that allow two logical LUT functions to be combined to fit into a single resource. The internal logical representation consists of two 5-input LUTs each sharing common inputs. One LUT is used to generate the O5 output, while the O6 combines the LUT5 function with the sixth input A6.




Figure 5-19: Multiple Outputs from the Same LUT6

Following are example LUT combinations:

- A LUT2 and LUT3 that are completely unrelated
- Two LUT3 with at least one common input
- Two LUT4 with at least two common inputs
- Two LUT5 with all common inputs
- A LUT5 and LUT6 in which the LUT6 is a combination of the LUT5, which generates the O5 output, and the A6 input

Because the fastest A6 input is dedicated to O6 muxing, it is important to realize that its use is limited when LUTs are combined. The A6 pin either remains unused; or, if the combined LUT includes a 6-input function, the A6 pin must be used for the uncommon input.

Sequential Logic: Registers

A register can be mapped to one of several types of resources in the device:

- CLB register
- CLB LUTRAM as an SRL
- ILOGIC
- OLOGIC
- DSP and block RAMs (if the register is adjacent to arithmetic or memory functionality)

Not all sequential logic will have the flexibility to be mapped into any of the above resources. However specific sequential logic can sometimes be mapped into more than one of these resources. In case such a choice exists, you may want to exploit the fastest resource available to implement that specific sequential logic.

The setup requirement for a CLB register driven by a LUT is generally very small, usually a negligible contribution to the path delay. The same goes for a register placed in an ILOGIC. The setup increases slightly when the data enters the CLB through an X input that bypasses the LUT. If the data passes through an FMUX or carry logic, the setup requirement suddenly jumps to a few hundred picoseconds. The same goes for a register mapped to an OLOGIC or LUTRAM. In these cases the setup requirement can be a significant contribution to a critical path.

Table 5-7: Relative Setup Requirement of Registers Placed in Different Resources

FF location	
ILOGIC	Faster
CLB FF driven from LUT	
CLB FF driven from X input	
LUTRAM (as SRL)	
CLB FF driven from MUXFX or CARRY	
OLOGIC	Slower

The CLB register typically has the fastest clock to output delay, nominally around 250 ps for the fastest speed grade of Xilinx 7 series FPGA devices. There are two BEL positions for a CLB register. The lower Q position is slightly faster than the MUX position which requires traversing a MUX before exiting.

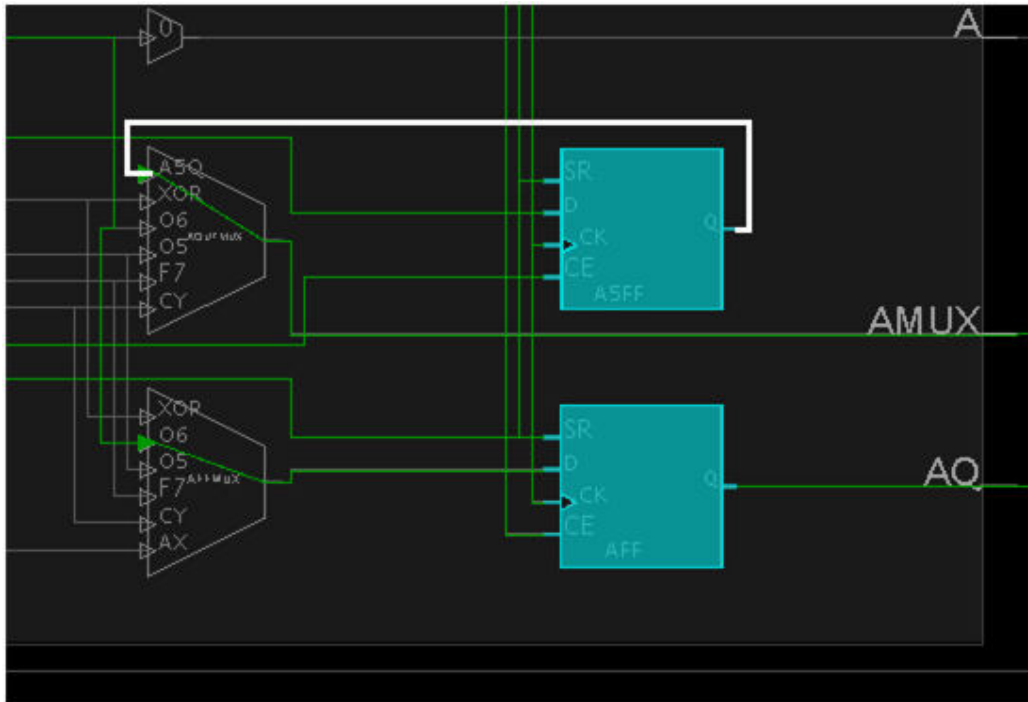


Figure 5-20: Lower Register Has a Lesser Delay

The clock to output delay of a register mapped to a LUTRAM as an SRL is significantly slower, on the order of a nanosecond. If the SRL output drives a critical path, it may be necessary to move the final register stage from the LUTRAM to its CLB pair register, to reduce the clock to output delay contribution.

The clock to output delays of an ILOGIC and OLOGIC are somewhat slower than the CLB register clock to output.

Table 5-8: Relative Clock to Out Delays in Different Resources

FF location	
CLB FF in Q BEL	Faster
CLB FF in MUX BEL	
ILOGIC	
OLOGIC	
LUTRAM (as SRL)	
DSP	
BRAM	Slower

Memory

Memory in Xilinx devices is implemented in either block RAM or distributed RAM. Either type of RAM can be inferred by synthesis; generated using the IP Catalog; or instantiated as

UNISIM. Most single-port and dual-port RAM and ROM functions can be implemented using either style of memory. Often functional requirements dictate the type of RAM needed. For example, an asynchronous read path requires a distributed RAM. Sometimes requirements steer toward one or the other. Very deep RAMs often need multiple block RAM. Narrow data widths are often more efficient in distributed RAM. Regardless of the choice, you must be aware of the design impact of each type.

Block RAM

Block RAM is a dedicated hardware resources used for RAM, ROM, and FIFO. These are organized into columns that span the height of the device with the columns fairly evenly distributed between CLBs.

Because they are dedicated blocks, block RAMs are more suited to higher capacity. block RAMs also tend to have smaller power consumption compared to distributed RAM of similar capacity. However there is generally higher delay getting to and from the block RAM columns. The following figure shows two such routing paths.



Figure 5-21: Example: Routing to and from a Block RAM

The read access time for a block RAM is relatively slow: about 1.5 to 2 ns for the clock-to-output delay followed by another 400-500 ps routing delay to obtain to CLB logic.

Block RAMs have an optional data output register which can reduce the clock-to-output by more than half. Setup and hold times also significantly impact high-speed paths. Each range from 500 to 700 ps, but each is reduced by more than half when using READ_FIRST mode.



TIP: *When manually placing block RAMs, place RAMs sharing the same address lines in columns. The columns have access to fast, dedicated routing of address lines for cascading block RAMs to create deeper RAMs.*

Distributed RAM

Distributed RAM is implemented using CLB logic (LUTRAM, registers, LUT, MUX). Because distributed RAM is implemented using CLB logic, it is more suited to smaller capacity. Compared to block RAMs, larger distributed RAM sizes consume more CLB resources and power. However, smaller sizes can give very good performance, because there is much smaller routing delay getting to and from the RAM.

LUTRAMs are the LUTs used for distributed RAM storage. The LUTRAMs have similar setup requirements to block RAMs, but hold requirements are about half. Read access times are also about half, but exiting the CLB requires additional delay, a nominal amount 200-400 ps for routing, and possibly up to another nanosecond for propagating through muxing logic. The following figure shows a typical minimum path for reading from a distributed RAM.

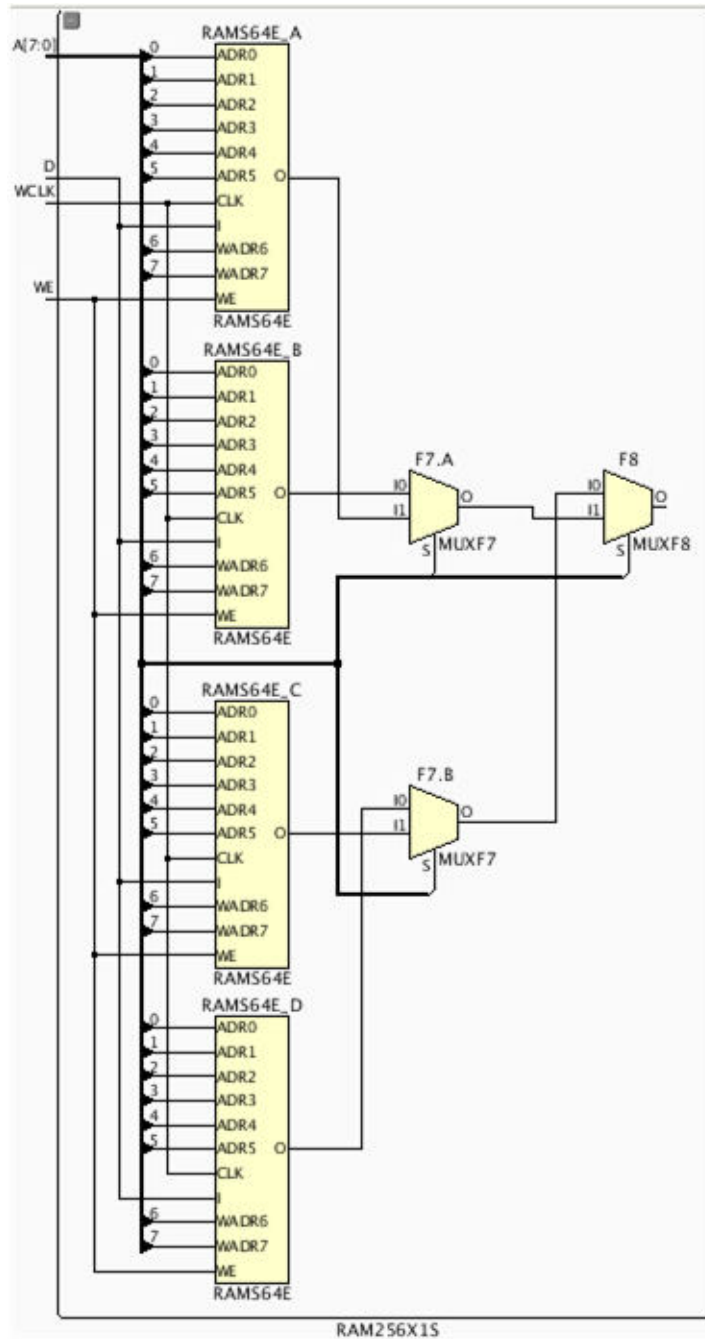


Figure 5-22: Read Delay for Distributed RAM

Comparison Between Block RAM and Distributed RAM

The following examples highlight the differences between block RAM and distributed RAM. Two different sizes of single port RAMs are created using IP Catalog with the IP defaults for a Virtex®-7 -2 speed grade. The post-route results reflect ideal conditions. Actual performance may vary depending on the surrounding logic of the containing design.

Table 5-9: Comparison of 8kx32 RAM Implementations

	Block RAM	Distributed RAM
Fmax	Over 500 MHz	250 MHz
Area	8 RAMB36, 18 CLBs	2043 CLBs
Power	370 mW	440 mW

The results indicate that for this depth and width, a block RAM implementation gives better overall results. The distributed RAM critical path is shown in the following figure.

Note: Fmax can be increased by adding pipeline stages to balance delays throughout the decoding logic, but will be limited by the delay through the RAM primitive.

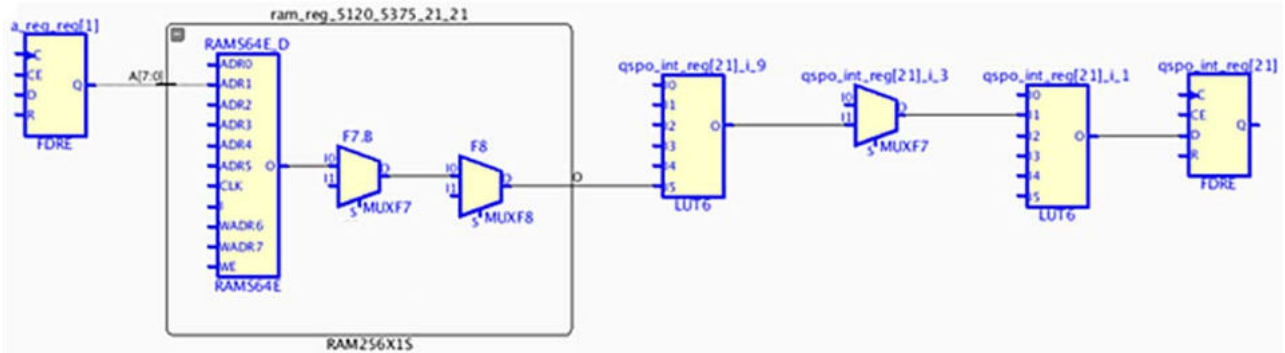


Figure 5-23: Critical Path through Distributed RAM

Table 5-10: Comparison Between 128x4 RAM Implementations

	Block RAM	Distributed RAM
Fmax	About 400 MHz	Over 500 MHz
Area	1 RAMB18, 3 slices	4 slices
Power	260 mW	260 mW

For this relatively small size, a distributed RAM implementation is fast, compact, and uses less routing resources.

DSP48E1 Blocks

In most cases, synthesis and IP Catalog determine the best implementation for arithmetic functions. Most advanced functions (especially those that depend on wide, high-speed multiplication) are best implemented in DSP48E1 blocks which have dedicated hardware

multipliers and ALUs to offload CLBs. The DSP48E1 is not only highly optimized internally, but also has dedicated high-speed routing along the DSP columns where the blocks are arranged. This enables multiple DSP48E1s to implement much wider multipliers and cascaded circuits such as pipelined FIR filters, all while running in excess of 500 MHz.

CLB carry logic is usually more appropriate for certain circuits such as multiplication by a constant and small-width multipliers. When resources of a certain type are overutilized or highly utilized, functions can be moved from one type to another. DSP48E1-based functions can be moved to CLB logic when running low on DSP blocks.

Similarly, many CLB-logic based functions can be moved to DSP48E1 when CLBs are overutilized. The latter is useful for addressing areas of congestion. The DSP48E1 not only implements multipliers and multiply-accumulate functions, but can implement adder-subtractors, counters, and even wide parallel logic gates.

The DSP48E1 block is pipelined with input and output registers and an intermediate register between the multiplier and ALU as shown in the following figure.

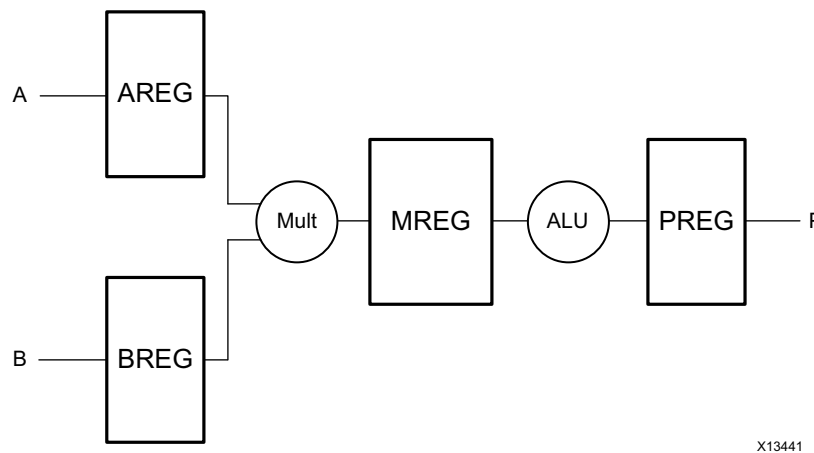


Figure 5-24: Pipelining Registers Available Within DSP48

All register stages must be used to achieve the highest performance which corresponds to a latency of three cycles. Following is an example of how the number of stages used affects timing. A 16x16 signed multiplier with 32-bit output is implemented on a Virtex-7 device, using the middle -2 speed grade.

Table 5-11: Impact of DSP48 Registers on FMax

Latency	AREG/BREG	MREG	PREG	setup path	clock to output path	Fmax
0	No	No	No	n/a	n/a	250 MHz
1	No	No	Yes	2.65 setup + 400 ps routing	350 ps clk->out + 770 ps routing	300 MHz
2	Yes	No	Yes	260 setup + 760 ps routing	350 ps clk->out + 700 ps routing	360 MHz
3	Yes	Yes	Yes	260 setup + 760 ps routing	350 ps clk->out + 700 ps routing	over 500 MHz

There is some routing delay getting to and from the DSP block.

- With *one* stage (the PREG stage), the delay getting to the register input is substantial, over 3ns when including the routing getting from CLB logic to the DSP column.
- With *two* stages (the input registers and output registers), the Fmax is limited by the internal register to register path.
- With *three* stages, the DSP block internally can operate over 500 MHz. If logic connecting the DSP block is well placed, the containing system can achieve an Fmax of 500 MHz.

If the three-stage multiplier must be moved to CLB logic, the equivalent implementation achieves around 440 MHz and requires about 143 slices, including 15 carry chains to add the partial products. The carry chains are five to six slices tall, and must be placed in vertically adjacent CLBs. The placer must be able to integrate these tall macros into existing CLB logic. See the following figure.



Figure 5-25: DSP48 Implemented in CLBs

An additional pipeline stage is required to be able to achieve similar performance.

Control Signals and Control Sets

Often not much consideration is given to control signals such as resets or clock enables. Many designers start HDL coding with if reset statements without deciding whether the reset is needed or not. While all registers support resets and clock enables, their use can significantly affect the end implementation in terms of performance, utilization, and power. The following sections define control signals and control sets.

Control Signals

The following table lists the control signals of library primitives that are placed on CLB resources.

Table 5-12: Control Signals

Clocks	Enables	Resets
<ul style="list-style-type: none"> • Clock and gate (for latches) 	<ul style="list-style-type: none"> • Clock enable • Write enable • Gate enable (for latches) 	<ul style="list-style-type: none"> • Logic 0 <ul style="list-style-type: none"> ◦ reset (synchronous) ◦ clear (asynchronous)
		<ul style="list-style-type: none"> • Logic 1 <ul style="list-style-type: none"> ◦ set (synchronous) ◦ preset (asynchronous)

Control Sets

A control set is the group of clock, enable, and set/reset signals used by a sequential cell. This includes lack of enable and lack of set/reset. For example, two cells clocked by the same clock actually have different control sets if only one cell has a reset, or only one cell has a clock enable.

The number of unique control sets affects how many registers can be put together in a Slice, because all eight registers share the same clock, reset, and clock enable signal. This means that if two registers have a different clock, reset, or enable signal (including not having one), then they cannot share the same Slice. This can negatively impact not only utilization, but placement as well. This in turn can negatively impact performance and power.

In a Xilinx 7 series FPGA Slice (which is effectively half a CLB, eight registers per Slice), all share the same control signals, and therefore the same control set. If the number of registers in the control set does not divide cleanly by eight, some registers must go unused. This is mainly of concern in designs that have several very low fanout control signals, such as a clock enable that feeds a single register. A design with a large number of control sets potentially can often show lower utilization of registers due to the registers that must go unused.

For more information about control signals, see [Control Signals and Control Sets in Chapter 4](#).



TIP: *If you are facing issues such as utilization or congestion issues, use the `report_control_sets Tcl` command to see if your design has too many control signals.*

If the number of control sets is high, use the following tips to reduce the number of controls sets, with special emphasis on those control sets which have a relatively low fanout.

- Avoid using sequential elements with both asynchronous and synchronous resets.
- Avoid asynchronous assignments to non-constant values. This causes many issues:
 - It results in a larger circuit than most realize, two registers, a latch and a LUT.
 - The sequential each have different control sets, occupying a minimum of 3 CLBs.
 - Several asynchronous timing paths result that if not analyzed properly, may cause additional timing hazards affecting overall design stability.
- Use active high control signals when possible.
- Avoid asynchronous sets/resets. Each asynchronous reset is a control signal that cannot be moved to the data path. A resulting increase in control sets cannot be alleviated by dissolving the asynchronous reset or set into the data path logic, unlike when using synchronous resets or sets. This allows for greater flexibility in packing and placement.
- Only use set / reset when necessary:
 - Often data paths contain many registers that automatically flush uninitialized values.
 - Registers of I/O, State machines and critical control signals that must be reset to known values should use sets and resets.
 - With a synchronous reset, timing is assessed for both assertion and desertion of the signal. In general, synchronous signals are more predictable and suggested for use unless an asynchronous reset is absolutely needed.
- Occasionally clock enables are actually redundant logic that may not be reduced by synthesis or logic optimization.
- Use caution with fanout controls for synthesis and logic and physical optimization. Low fanout limits may unnecessarily introduce too many control sets.

Tuning the Compilation Flow

The default compilation flow provides a quick way to obtain a baseline of the design and start analyzing the design if timing is not met. After initial implementation, tuning the compilation flow might be required to achieve timing closure.

- [Strategies and Directives](#)
- [Optimization Iterations](#)
- [Incremental Compilation](#)
- [Overconstraining the Design](#)

Strategies and Directives

Strategies and directives can be used to increase the implementation solution space and find the most optimal solution for your design. The strategies are applied globally to a project implementation run, while the directives can be set individually on each step of the implementation flow in both project and non-project modes. The pre-defined strategies should be tried first before trying to customize the flow with directives. Xilinx does not recommend running the SSI technology strategies for a non-SSI technology device.

If timing cannot be met with a default strategy, you can manually explore a custom combination of directives. Because placement typically has a large impact on overall design performance, it can be beneficial to try various placer directives with only the I/O location constraints and with no other placement constraints. By reviewing both WNS and TNS of each placer run (these values can be found in the placer log), you can select two or three directives that provide the best timing results as a basis for the downstream implementation flow.

For each of these checkpoints, several directives for `phys_opt_design` and `route_design` can be tried and again only the runs with the best estimated or final WNS/TNS should be kept. In Non-Project Mode, you must explicitly describe the flow with a Tcl script and save the best checkpoints. In Project Mode, you can create individual implementation runs for each placer directive, and launch the runs up to the placement step. You would continue implementation for the runs that have the best results after the placer step (as determined by the Tcl-post script).

Physical constraints (Pblocks and DSP and RAM macro constraints) can prevent the placer from finding the most optimal solution. Xilinx therefore recommends that you run the placer directives without any Pblock constraints. The following Tcl command can be used to delete any Pblocks before placement with directives commences:

```
delete_pblock [get_pblocks *]
```

Running `place_design -directive <directive>` and analyzing placement of the best results can also provide a template for floorplanning the design to stabilize the flow from run to run.

Optimization Iterations

Sometimes it is advantageous to iterate through a command multiple times to obtain the best results. For example, it might be helpful to first run `phys_opt_design` with the `force_replication_on_nets` option in order to optimize some critical nets that appear to have an impact on TNS during route:

```
phys_opt_design -force_replication_on_nets
```

Next run `phys_opt_design` with any of the directives to improve the overall WNS of the design.

In Non-Project Mode, use the following commands:

```
phys_opt_design -force_replication_on_nets [get_nets -hier *phy_reset*]  
phys_opt_design -directive <directive name>
```

In Project Mode, the same results can be achieved by running the first `phys_opt_design` command as part of a Tcl-pre script for a `phys_opt_design` run step which will run using the `-directive` option.

Incremental Compilation

Incremental compile yields the best results in preserving QOR for a design when the critical path of the reference design is not affected by the changes in the current design. For more details on using incremental compilation, see [Incremental Flows](#).

Overconstraining the Design

When the design fails timing by a small amount after route, it is usually due to a small timing margin after placement. It is possible to increase the timing budget for the router by tightening the timing requirements during placement and physical optimization. The recommended way to do this is to use the `set_clock_uncertainty` constraint for the following reasons:

- It does not modify the clock relationships (clock waveforms remain unchanged).
- It is additive to the tool-computed clock uncertainty (jitter, phase error).
- It is specific to the clock domain or clock crossing specified by the `-from` and `-to` options.
- It can easily be reset by applying a null value to override the previous clock uncertainty constraint.

In any case, Xilinx recommends that you:

- Overconstrain only the clocks or clock crossing that cannot meet setup timing.
- Reset the extra uncertainty before running the router step.

See the following example:

A design misses timing by -0.2 ns on paths with the `clk1` clock domain and on paths from `clk2` to `clk3` by -0.3 ns before and after route.

1. Load netlist design and apply the normal constraints.
2. Apply the additional clock uncertainty to overconstrain certain clocks.
 - a. The value should be at least the amount of violation.
 - b. The constraint should be applied only to setup paths.

```
set_clock_uncertainty -from clk0 -to clk1 0.3 -setup
set_clock_uncertainty -from clk2 -to clk3 0.4 -setup
```

3. Run the flow up to the router step. It is best if the pre-route timing is met.
4. Remove the extra uncertainty.

```
set_clock_uncertainty -from clk1 -to clk1 0 -setup
set_clock_uncertainty -from clk2 -to clk3 0 -setup
```

5. Run the router.

After the router, you can review the timing results to evaluate the benefits of overconstraining. If timing was met after placement but still fails by some amount after route, you can increase the amount of uncertainty and try again.



RECOMMENDED: *Do not overconstrain beyond 0.5 ns.*

Considering Floorplan

Floorplanning allows you to guide the tools, either through high-level hierarchy layout, or through detail placement. This can provide improved QOR and more predictable results. You will achieve the greatest improvements by fixing the worst problems or the most common problems. For example if there are outlier paths that have significantly worse slack, or high levels of logic, fix those paths first by grouping them in a same region of the device through a Pblock. Limit floorplanning only to portions of design that need additional human help through floorplanning, rather than floorplanning the entire design.

Floorplanning logic that is connected to the I/O to the vicinity of the I/O can sometimes yield good results in terms of predictability from one compilation to the next. In general, it is best to keep the size of the Pblocks to a clock region. This provides the most flexibility for the placer. Avoid overlapping Pblocks, as these shared areas could potentially become more congested. Minimize the number of nets that cross Pblocks.

There are extra considerations for Stacked Silicon Interconnect (SSI) technology devices. The SSI technology devices are made of multiple Super Logic Regions (SLRs), joined by an interposer. The interposer connections are called Super Long Lines (SLLs). There is some delay penalty when crossing from one SLR to another. To minimize the impact of the SLL

delay on your design, floorplan the design so that SLR crossings are not part of the critical path. Minimizing SLR crossings through floorplanning by keeping a Pblock within one SLR only can also improve timing and routability of the design targeting SSI technology devices. For more information, see this [link](#) in *Vivado Design Suite User Guide: Design Analysis and Closure Techniques* (UG906) [Ref 21].



VIDEO: For information on using floorplanning techniques to address design performance issues, see the [Vivado Design Suite QuickTake Video: Design Analysis and Floorplanning](#).

Preserving Placement and Routing

Once you have results that meet your timing constraints, you might want to lock down the placement and possibly the routing of the critical portions of your design. This can help preserve the performance of your circuitry, and also provide more predictable results for future runs.

It is fairly easy to reuse the placement of:

- I/Os
- Global Clock Resources
- Block RAM macros
- DSP macros

Reusing this placement helps reduce the variability in results from one netlist revision to the next. These primitives generally have stable names and the placement is usually easy to maintain.

It is sometimes desired to preserve the routing of a critical net or a portion of a critical net to guarantee the same timing from run to run. Routing for critical nets can be preserved by setting the property `is_route_fixed` of the net to 1. This can be done by using the Vivado IDE or through a Tcl command. To fix the routing on a net in the Vivado Design Suite, select the net in the device view, right-click and select **Fix Routing** from the context menu.

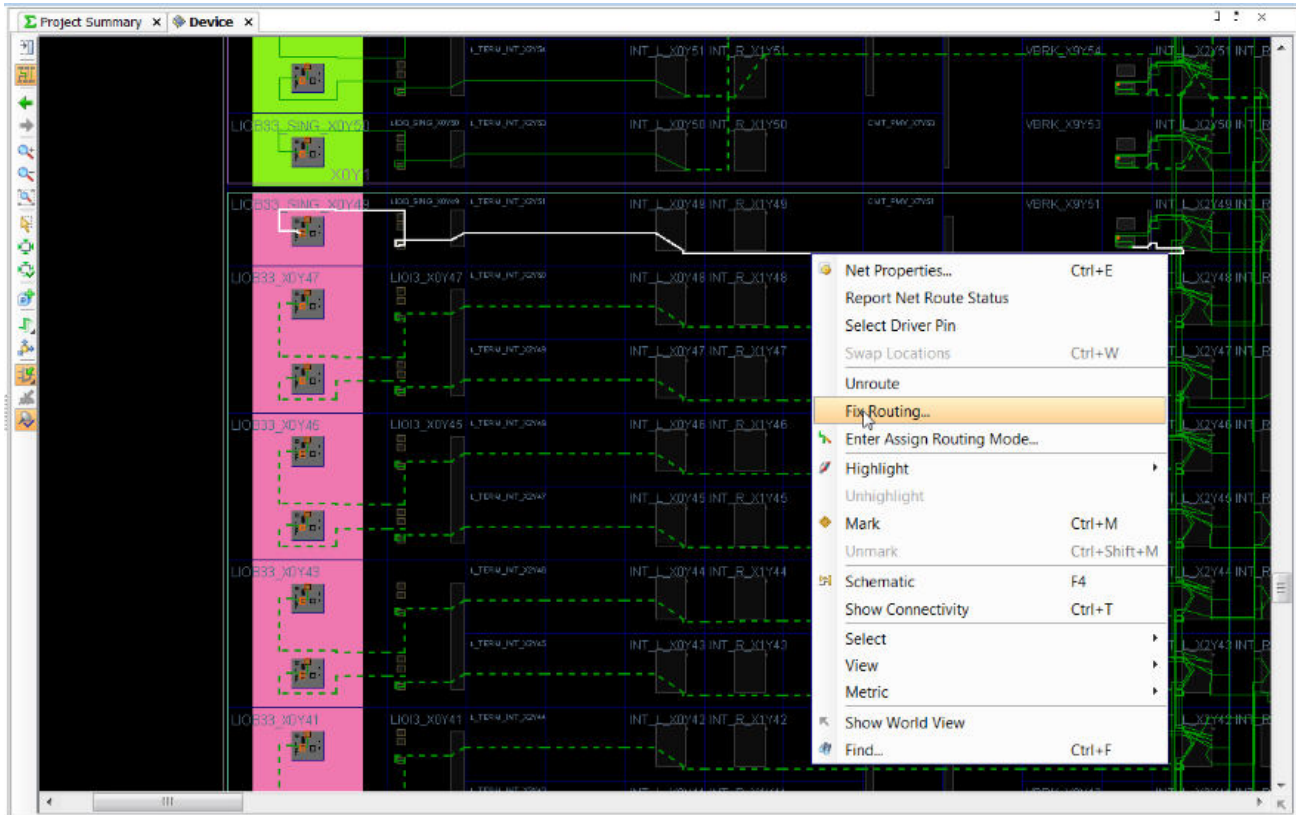


Figure 5-26: Preserving a Specific Route

The Tcl command example below fixes the routing for all nets in \$net:

```
set_property IS_ROUTE_FIXED 1 $net
```

This marks the route as fixed and adds a constraint to the in-memory design. Nets with fixed routing are shown as dashed lines in the Device View.

Sometimes, when a design has very high utilization or very tight timing requirements, even a small change in the design or constraints might cause a significant change in the placement, routing, or both. Preserving the relevant portions of the design might be especially helpful in such cases.

Power

Given the importance of power, the Vivado tools support methods for obtaining an accurate estimate for power, as well as it providing some power optimization capabilities. For additional information refer to *Vivado Design Suite User Guide: Power Analysis and Optimization* (UG907) [Ref 22].

Estimating Power Through All Stages in the Vivado Design Suite Flow

As your design flow progresses through synthesis and implementation, you must regularly monitor and verify the power consumption to be sure that thermal dissipation remains within budget. You can then take prompt remedial actions if power approaches your budget too closely.

The accuracy of the power estimates varies depending on the design stage when the power is estimated. To estimate power post-synthesis through implementation, run the `report_power` command, or open the Power Report in the Vivado IDE.

- **Post Synthesis**

The netlist is mapped to the actual resources available in the target device.

- **Post Placement**

The netlist components are placed into the actual device resources. With this packing information, the final logic resource count and configuration becomes available.

This accurate data can be exported to the Xilinx Power Estimator spreadsheet. This allows you to:

- Perform what-if analysis in XPE.
- Provide the basis for accurately filling in the spreadsheet for future designs with similar characteristics.

For more information, see [Chapter 3, Board and Device Planning](#).

- **Post Routing**

After routing is complete all the details about routing resources used and exact timing information for each path in the design are defined.

In addition to verifying the implemented circuit functionality under best and worst case gate and routing delays, the simulator can also report the exact activity of internal nodes and include glitching. Power analysis at this level provides the most accurate power estimation before you actually measure power on your prototype board.

Types of Power Estimation Supported by `report_power`

The `report_power` command supports two modes of power analysis. Depending upon the accuracy desired, you can use the appropriate style.

Vector-Based Estimation

In parallel with all stages of the design development, you will generally perform simulations to verify that the design behaves as expected. Different verification techniques are available depending on the design development state; the design complexity; and your company's policy.

The following paragraphs highlight the valuable data you can capture and common pitfalls related to using this data to perform power analysis. An important factor for getting accurate power estimation is that the design activity must be realistic. It should represent the typical or worst case scenario for data coming into the simulated block. This type of information is not necessarily provided while performing verification or validating functions.

Invalid data can be input to verify that the system can handle the data, and that it remains stable even when invalid data or commands are entered. Using such test cases to perform power analysis may result in inaccurate power estimation, because the design logic is not stimulated as it would be under typical system operation.

System Transaction Level

Early in the design cycle, you may have created a description of the transactions that occur between devices on a PCB, or between the different functions of your FPGA application. You can extract from this the expected activity per functional block for certain I/O ports and most of the clock domains. This information can help you fill in the Xilinx Power Estimator spreadsheet.

FPGA Description Level

While defining the RTL for your application, you can verify the functionality by performing behavioral simulations. This helps you verify the data flow and the validity of calculations to the clock cycle. The following are not yet available:

- Exact FPGA resources used
- Count
- Configuration

You can manually extrapolate resource utilization and extract activity for I/O ports or internal control signals (set, reset, clock enable). This information can be applied to refine the Xilinx Power Estimator spreadsheet information.

Your simulator should be able to extract node activity, and export it in the form of a SAIF file. You can save this file for more accurate power analysis in the Vivado Design Suite design flow (for example after place and route) if you do not plan to run post-implementation simulations.

FPGA Implementation Level

Simulation may be performed at different stages in the implementation process with different outcomes in terms of the power-related information that can be extracted. This additional information may also be used to refine the Xilinx Power Estimator spreadsheet and the Vivado Design Suite power analysis. It may also save I/O ports and specific module activity, which can be reused in the Vivado Design Suite power analysis feature.

Simulation File

The Vivado Design Suite Report Power matches nets in the design database with names in the simulation results netlist. The simulation data for power is stored in a Switching Activity Interchange Format (SAIF) file.

Because `report_power` does net name matching, it is best to obtain simulation results on the same design view (such as post-synthesis and post-routing) on which the power analysis is being carried out.



IMPORTANT: *In the Vivado IDE, specify a SAIF file name in the Input Files tab of the Report Power dialog box to read a SAIF simulation output file. Alternatively, use the `read_saif` Tcl command to read the SAIF simulation output file. To generate a SAIF file from the Mentor Graphics ModelSim simulator for power analysis within the Vivado Design Suite, see Answer Record 53544.*

Vectorless Estimation

When design node activity is not provided (either from you or from simulation results), the vectorless power estimation algorithms can predict this activity.

The vectorless engine assigns initial seeds (default signal rates and static probability) to all undefined nodes. Starting from the design primary inputs, it then propagates activity to the internal nodes through various circuit elements; and repeats this operation until the primary outputs are reached.

The algorithm understands the design connectivity and resource functionality and configuration. Its heuristics can even approximate the glitching rate for any nodes in the netlist. Glitching occurs when design elements change states multiple times in between active clock edges before settling to a final value.

While the vectorless propagation engine is not as accurate as a post-route simulation with a reasonably long duration and realistic stimulus, it is an excellent compromise between accuracy and compute efficiency.

Best Practices for Accurate Power Analysis

Use the following for accurate power analysis:

- [Accurate Clock Constraints](#)
- [Accurate I/O Constraints](#)
- [Accurate Signal Rate and %High on Top Level Control Signals](#)

Accurate Clock Constraints

Because power is significantly dependent on the frequency of operation, clock frequency must be specified accurately.

In any design, you will typically know the activity of some specific nodes, since they are imposed by the system specification or the interfaces with which the FPGA device communicates. Providing this information to the tools helps guide the power estimation algorithms.

This information is especially helpful for nodes that drive multiple cells in the FPGA device:

- Set
- Reset
- Clock enable
- Clock signals

You will typically know the exact frequency of all FPGA clock domains, whether they are externally provided (input ports); or internally generated; or externally supplied to the printed circuit board (output ports).

Accurate I/O Constraints

With your knowledge of the exact protocols and format of the data flowing in and out of the FPGA device, you can usually specify signal transition rate and/or signal percentage high rate in the tools for at least some of the I/O components.

For example, some protocols have a DC balanced requirement (signal percentage high rate =50%), or you may know how often data is written or read from your memory interface. This allows you to set the data rate of strobe and data signals.

The board and other external capacitance driven by the output ports are typically known.

Enter the following in the Tcl prompt to set the load on all the output ports:

```
set_load <value in pF> [all_outputs]
```

Accurate Signal Rate and %High on Top Level Control Signals

With your knowledge of the system and the expected functionality, you may be able to predict the activity on control signals such as Set, Reset, and Clock Enable. Because these signals typically can turn on or turn off large pieces of the design logic, providing this activity information significantly increases the power estimation accuracy.

If you know the data patterns of your I/O interfaces, specify this activity. Unless you are calculating the total power per supply in a separate tool (such as a spreadsheet), specify the termination technique for your outputs to allow Report Power to include the amount of power the FPGA device supplies to these external components.

Accurate signal characteristics for all of the above type of signals can be provided through `set_switching_activity` in the Tcl prompt; or Signal Rate and Static Probability (% High) in the Power Properties window of the Vivado IDE.

Project Device Settings

Review the different user-editable selections in the Environment and Power Supply tabs of the Report Power dialog box. Make sure the process, voltage, and environment data closely match your expected environment. These settings have a significant influence on the total estimated power.

The user-editable selections in these tabs are:

- [Device Settings](#)
- [Design Thermal Settings](#)
- [Voltage Supply Settings](#)

Device Settings

- **Temp Grade**

Select the appropriate grade for the device (typically Commercial or Industrial). Some devices may have different device static power specifications depending on this setting. Setting this properly allows for the proper display of junction temperature limits for the chosen device.

- **Process**

The recommended process setting is Maximum for a worst-case analysis. Although the default setting of Typical gives a more accurate picture of the statistical measurements, changing the setting to Maximum modifies the power specification to worst-case values.

In the Tcl prompt, use the following to set the temp grade and process selections:

```
set_operating_conditions -process maximum
set_operating_conditions -grade industrial
```

Design Thermal Settings

Review the different user-editable selections in the Environment tab of the Report Power dialog box.

- **Ambient Temperature (°C)**

Specify the maximum possible temperature expected inside the enclosure that will house the FPGA design. This, along with airflow and other thermal dissipation paths (for example, the heatsink), allows an accurate calculation of Junction Temperature. This in turn allows a more accurate calculation of device static power.

In the Tcl prompt, enter the following to set the ambient temperature selection:

```
set_operating_conditions -ambient_temp 75
```

- **Airflow (LFM)**

The airflow across the chip is measured in Linear Feet per Minute (LFM). LFM can be calculated from the fan output in Cubic Feet per Minute (CFM) divided by the cross sectional area through which the air passes.

Specific placement of the FPGA device or fan may impact the effective air movement across the device, and thus the thermal dissipation. The default for this parameter is 250 LFM. If you plan to operate the FPGA device without active air flow (still air operation), change the 250 LFM default to 0 LFM.

In the Tcl prompt, enter the following to set the airflow selection:

```
set_operating_conditions - airflow 250
```

- **Heat Sink (if available)**

If a heatsink is used and more detailed thermal dissipation information is not available, choose an appropriate profile for the type of heatsink. This (along with other entered parameters) is used to help calculate an effective Theta_{JB} (printed circuit board thermal resistance), resulting in a more accurate junction temperature and quiescent power calculation. Some types of sockets may act as heatsinks, depending on the design and construction of the socket.

In the Tcl prompt, enter the following to set the heatsink selection:

```
set_operating_conditions - heatsink low
```

- **Board Selection and # of Board Layers (if available)**

Selecting an approximate size and stack of the board will help calculate the effective ThetaJB by taking into account the thermal conductivity of the board itself.

- **ThetaJB (printed circuit board thermal resistance)**

If more accurate thermal modeling of the board and system is available, use to specify the amount of heat dissipation expected from the FPGA device.

The more accurately custom ThetaJB can be specified, the more accurate the estimated junction temperature will be, thus affecting device static power calculations.



IMPORTANT: *In order to specify a custom ThetaJB, the Board Selection must be set to Custom. If you specify a custom ThetaJB, you must also specify a Board Temperature for an accurate power calculation.*

In the Tcl prompt, enter the following to set the board selection:

```
set_operating_conditions - board jedec
```

In the Tcl prompt, enter the following to set the ThetaJB selection:

```
set_operating_conditions - thetajib 3
```

Voltage Supply Settings

Review the different user-editable selections in the Power Supply tab of the Report Power dialog box.

- **Power Supply**

If this information is known, in the Power Supply tab make sure all voltage levels are set correctly for the different supply sources. Voltage is a large factor contributing to both static and dynamic power.

In the Tcl prompt, enter the following to set the voltage on the VccAux rail:

```
set_operating_conditions -voltage {Vccaux <value>}
```

Reviewing the Design Power Distribution After Running Vivado Design Suite Power Analysis

Open the Summary view to review the Total On-Chip Power and thermal properties. The On-Chip Power graph shows the power dissipated in each of the device resource types. With this high-level view, you can determine which parts of your design contribute most to the total power. The Power Supply tab shows the current drawn for each supply source, and breaks down this total between static and dynamic power.

From the Utilization Details tab, to see more details of the power at the resource level, click the different resource types in the graph. The different resources views are organized as a tree table. Drag a column header to reorder the column arrangement. Click on a column header to change the sorting order.

Further Refining Control Signal Activity After Running Vivado Design Suite Power Analysis

When SAIF-based annotation has not been used for accurate power analysis, you can fine-tune the power analysis after doing the first level analysis.

Report Power extracts and lists all the different control signals in the Signal view. You may know from the expected behavior of your application that some Set/Reset signals are not active in normal design operation. In that case, you may want to adjust the activity for these signals. Similarly, some signals in your application may disable entire blocks of the design when the blocks are not in use. Adjust their activity according to the expected functionality.

Because synthesis tool and place and route algorithms can infer or remap control signals to optimize your RTL description, many of the signals listed in these views may be unfamiliar. If you are unsure of what these signals are, let the tool determine the activity.

Writing Out a Power Report Text File

Open the Report Power dialog box from the Flow Navigator window in the Vivado IDE. Use this dialog box to review power settings and adjust activity for known elements in your design.

For project documentation, you may want to save the power estimation results in an output text file.

In other circumstances, you may be experimenting with different mapping, placement, and routing options to close on performance or area constraints. Saving power results for each experiment can help you select the most power-effective solution when several experiments meet your requirements.

In the Tcl prompt, enter the following:

```
report_power -file report.pwr
```

The power engine writes out a file `report.pwr` in the current working directory. This file contains the power estimation results.

Exporting the Power Estimate from the Vivado Tools to XPE

Open the Report Power dialog box from the Vivado IDE. In this dialog box, you can review power settings and adjust activity for known elements in your design.

The Xilinx Power Estimator (XPE) output file saves all environment information, device usage, and design activity in a file (.xpe) which you can later import into the XPE spreadsheet. This is useful when your power budget is exceeded, and you do not think that software optimization features alone will be able to meet your budgets.

In this case, you can:

- Import the current implementation results into XPE.
- Explore different mapping, gating, folding, and other strategies.
- Estimate their impact on power before modifying the RTL code or rerunning the implementation.

Compare your assumptions in the XPE spreadsheet with these final results. This helps provide more accurate inputs for XPE for future designs.

In the Tcl prompt, enter the following:

```
report_power -xpe report.xpe
```

The power engine writes out a file report.xpe in the current working directory. This file can now be imported into the XPE spreadsheet.

Power Optimization

If the power estimates are outside the budget, you must take steps to reduce power.

- [Analyzing Your Power Estimation and Optimization Results](#)
- [Running Power Optimization](#)
- [Using the Power Optimization Report](#)
- [Using the Timing Report to Determine the Impact of Power Optimization](#)

Analyzing Your Power Estimation and Optimization Results

Once you have generated the power estimation, Xilinx recommends the following:

- Examine the total power in the Summary section. Does the total power and junction temperature fit into your thermal and power budget?
- If the results are substantially over budget, review the power summary distribution by block type and by the power rails. This provides an idea of the highest power consuming blocks.
- Review the Hierarchy section. The breakdown by hierarchy provides a good idea of the highest power consuming module. You can drill down into a specific module to determine the functionality of the block. You can also cross-probe in the GUI to

determine how specific sections of the module have been coded, and whether there are power efficient ways to recode it.

Running Power Optimization



TIP: To maximize the impact of power optimizations, see [Coding Styles to Improve Power in Chapter 4](#).

Power optimization works on the entire design or on portions of the design (when `set_power_opt` is used) to minimize power consumption.

Power optimization can be run either pre-place or post-place in the design flow, but not both. The pre-place power optimization step focusses on maximizing power saving. This can result (in rare cases) in timing degradation. If preserving timing is the primary goal, Xilinx recommends the post-place power optimization step. This step performs only those power optimizations that preserve timing.

In cases where portions of the design should be preserved due to legacy (IP) or timing considerations, use the `set_power_opt` command to exclude those portions (such as specific hierarchies, clock domains, or cell types) and rerun power optimization.

Using the Power Optimization Report

To determine the impact of power optimizations, run the following command in the Tcl console to generate a power optimization report:

```
report_power_opt -file myopt.rep
```

Using the Timing Report to Determine the Impact of Power Optimization

Power optimization works to minimize the impact on timing while maximizing power savings. However, in certain cases, if timing degrades after power optimization, you can employ a few techniques to offset this impact.

Where possible, identify and apply power optimizations only on non-timing critical clock domains or modules using the `set_power_opt` XDC command. If the most critical clock domain happens to cover a large portion of the design or consumes the most power, review critical paths to see if any cells in the critical path were optimized by power optimization.

Objects optimized by power optimization have an `IS_CLOCK_GATED` property on them. Exclude these cells from power optimization.

To locate clock gated cells, run the following Tcl command:

```
get_cells -hier -filter {IS_CLOCK_GATED==1}
```

Configuration and Debug

Overview of Configuration and Debug

Configuration is the process of loading application-specific data (a bitstream) into the internal memory of the FPGA device. Debug is required if the design does not meet expectations on the hardware. After successfully completing the design implementation, the next step is to load the design into the FPGA and run it on hardware.

Refer to the following two user guides for details on configuration and debug software flows and commands:

- *Vivado Design Suite User Guide: Programming and Debugging* (UG908) [\[Ref 24\]](#)
 - *Vivado Design Suite Tcl Command Reference Guide* (UG835) [\[Ref 13\]](#)
-

Configuration

This section includes tips to successfully implement the targeted configuration solution after you have selected the configuration mode, and you are ready to load the design into the FPGA device. For the common configuration modes and for recommendations during the initial planning stage, see the following resources:

- *Vivado Design Suite User Guide: Programming and Debugging* (UG908) [\[Ref 24\]](#)
- *UltraScale FPGA BPI Configuration and Flash Programming* (XAPP1220) [\[Ref 52\]](#)
- *BPI Fast Configuration and iMPACT Flash Programming with 7 Series* (XAPP587) [\[Ref 53\]](#)
- *Using SPI Flash with 7 Series FPGAs* (XAPP586) [\[Ref 54\]](#)
- *SPI Configuration and Flash Programming in UltraScale FPGAs* (XAPP1233) [\[Ref 55\]](#)
- *Using Encryption to Secure a 7 Series FPGA Bitstream* (XAPP1239) [\[Ref 56\]](#)
- Programming and Debug video tutorials available from the [Vivado Design Suite Video Tutorials](#) page on the Xilinx website

Your design must be successfully synthesized and implemented before a bitstream (.bit) image can be created. Once the bitstream is created, it can be loaded onto the FPGA device through one of the two methods:

- **Direct Programming**

The bitstream is loaded directly to the FPGA device by means of a cable, processor, or custom solution.

- **Indirect Programming**

The bitstream can be loaded into an external flash memory. The flash memory then loads the bitstream into the FPGA device.

Xilinx provides software tools that:

- Create the FPGA bitstream (.bit or .rbit)
- Format the bitstream into flash programming files (.mcs)
- Directly program the FPGA device
- Indirectly program the attached configuration flash device

Bitstream Generation

A bitstream (.bit) is a binary file that represents a user design. The bitstream contains configuration data that can be loaded into the FPGA device. There are several bitstream file format options and a variety of bitstream input options that enable features by initializing FPGA internal configuration registers. Before generating the bitstream file, it is important to review the `set_property` settings in the XDC file or the Tcl shell for bitstream generation to ensure they are set correctly for the target configuration mode selected.

For example, in master configuration modes (Master SPI or Master BPI, etc.), you can use an internal configuration clock (CCLK) or an external master configuration clock (EMCCLK). Use the appropriate `set_property` (for example, `BITSTREAM_CONFIG_CONFIGRATE` or `BITSTREAM.CONFIG.EXTMASTERCCLK_EN`) to enable the clock source targeted.

For details on bitstream generation properties, see the following resources:

- *Vivado Design Suite User Guide: Programming and Debugging* (UG908) [Ref 24]
- *Vivado Design Suite Tcl Command Reference Guide* (UG835) [Ref 13]
- For more information on how to use the bitstream generation `write_bitstream` command, see the [Vivado Design Suite QuickTake Video: How To Use the "write_bitstream" Command in the Vivado Design Suite](#), or refer to help for the command in the Vivado® Design Suite Tcl shell.

Use `list_property_value` to see the possible values that the property can take or refer to this [link](#) in *Vivado Design Suite User Guide: Programming and Debugging* (UG908) [Ref 24].

Flash File Generation

The Vivado Design Suite Tcl command `write_cfgmem` supports converting an FPGA bitstream (`.bit`) into a standard format flash programming file (`.mcs`).

The proper `write_cfgmem` interface option must be selected when generating `.mcs` programming files to ensure the data order for the bus-width is used and that the FPGA configuration from the flash is successful. For available options, see the *Vivado Design Suite User Guide: Programming and Debugging* (UG908) [Ref 24].

Vivado Design Suite Device Programmer - In-System JTAG Programming

The Vivado Design Suite Device Programmer has several functions. The most common uses of the programmer are: program the FPGA device by means of JTAG with a Xilinx supported download cable, indirectly program external SPI flash and parallel NOR flash, or program FPGA device eFUSE AES key or user code.

For more information on support and commands see:

- *Vivado Design Suite Tcl Command Reference Guide* (UG835) [Ref 13]
- *Vivado Design Suite User Guide: Programming and Debugging* (UG908) [Ref 24]

The JTAG cable connector is required for JTAG-based programming mode. When programming the FPGA device using the Vivado Design Suite Device Programmer, keep in mind that the JTAG maximum frequency is limited by the slowest device in the JTAG chain.

Indirectly Programming SPI NOR or Parallel NOR Flash

For the basic configuration solution using external flash, the FPGA device automatically retrieves the bitstream from a flash memory at power-on. Since the FPGA device is connected to the flash memory for configuration this enables the FPGA device to program the flash through the connected interface.

Indirectly programming flash in-system is a popular option. Vivado Design Suite (version 2014.1 and later) supports indirectly programming select SPI NOR and parallel NOR flash. The flash is indirectly programmed by loading an indirect programming bitstream. See *Vivado Design Suite User Guide: Programming and Debugging* (UG908) [Ref 24] for a list of supported flash memory devices.

The Vivado Design Suite bitstream image controls or accesses only the FPGA configuration flash interface pins. All other FPGA SelectIO™ technology pins are unused. You can program

the unused SelectIO pins to have all internal pull-ups, all internal pull-downs, or all pull-none. For unexpected board signal activity during indirect flash programming, check the Vivado Design Suite selection for these unused SelectIO technology pins.

Table 6-1: Recommendations for Flash Indirect Programming

Operation	Recommendation
Erase	Flash devices are non-volatile devices and must be erased before programming. Unless a full chip erase is specified, only the address range covered by the assigned MCS is erased.
BlankCheck	Verify the erase operation
Readback	Read back the contents of the flash into a file for comparison against the original flash programming file. This operation will read back the entire flash contents, not just the address range covered by the assigned MCS

For additional information on indirectly programming SPI flash or parallel NOR flash and related commands, see the following resources:

- *UltraScale FPGA BPI Configuration and Flash Programming (XAPP1220)* [Ref 52]
- *BPI Fast Configuration and iMPACT Flash Programming with 7 Series (XAPP587)* [Ref 53]
- *Using SPI Flash with 7 Series FPGAs (XAPP586)* [Ref 54]

Basic Configuration Debug

The best practices discussed in this section will help enable debug and resolution if you encounter an issue when implementing a configuration solution.

Before you embark on a full debug of configuration solution, create and test a simple design using the bitstream defaults (for example, a counter or LED output pattern). This simple design test will help eliminate any potential issues with advanced bitstream settings or board interfaces.

File Generation Review

If configuration does not complete successfully, review that the bitstream properties and flash programming file options were selected correctly. To verify the bitstream generation options used by an image, run the following Tcl command:

```
report_property [current_design]
```

This command displays all those properties that were changed from their default settings while being applied to a design. Where there are no values displayed, the default is applied.

Also, review the `write_cfgmem` flash programming file generation options.

Status Register

If configuration has not completed properly, the status register provides important information about what errors may have caused the failure. For more information, see:

- FPGA family configuration user guide
- Xilinx Configuration Solution Center [Answer Record 34909](#)

The FPGA Status register data on Xilinx® FPGA devices can be read by the Vivado Design Suite Device Programmer by means of JTAG. In case of a configuration failure, this register captures the specific error conditions that can help identify the cause of a failure. In addition, the Status register allows you to verify the Mode pin settings M[2:0] and the bus width detect. For details on the Status register, see the *Configuration User Guide* [Ref 37] for your device.

Verification and Readback

If configuration is not successful, a JTAG readback/verify operation on the FPGA device can determine whether the intended configuration data was loaded correctly into the device. In the event of a discrepancy, the case can be investigated. In order to perform a JTAG verify operation with the Vivado Design Suite Device Programmer, a mask (.msk) file is required. The file is created during bitstream generation.



IMPORTANT: For UltraScale™ devices, if you attempt to download an encrypted bitstream (which uses the BBR as the key source) before the key is programmed into the BBR register, the FPGA device will lock up and you will not be able to load the BBR key. You can still download unencrypted bitstreams, but you will not be able to download encrypted bitstreams because the FPGA device will prevent you from downloading a key into BBR. You must power-cycle the board to unlock the UltraScale device and then reload the BBR key.

Configuration Sequence

During configuration, some basic checks can help isolate issues. Xilinx FPGA bitstreams have a unique header. The header includes a special Sync word, and can include an auto detect, configuration clock type and rate setting. The Sync word is 32 bits (bit [31:0] = AA995566) for most Xilinx FPGA devices and is a valuable debug parameter.

The FPGA configuration state machine does not begin until the Sync word is recognized at the pins of the FPGA. In the event of a configuration that does not start, you can observe the configuration data pins to ensure the Sync word is being received correctly. Additionally, if the bitstream header is seen properly, any increase in the configuration clock due to a configuration or EMCCLK option setting should be seen or the header was not recognized.

Configuration Startup

A common configuration sequence is followed for the FPGA device power up. This is described in detail in the FPGA device configuration user guides. The following FPGA features can extend the configuration startup sequence:

- Wait for PLL
- MMCM lock
- DCI match

If any of these options are used, be sure that the images in the configuration source are properly spaced for MultiBoot images. For more information on MultiBoot image handling, see the *Configuration User Guide* [Ref 37] for your device. In addition, when using Slave modes, be sure that enough clock cycles are supplied to complete the startup sequence.

When the startup clock (JTAGCLK, CCLK, EMCCLK) is not clocked to the end of the startup sequence, the following symptoms may indicate an incorrect or incomplete startup:

- I/O remains 3-stated.
- Dual mode pins operate in LVCMOS rather than the specified I/O standard.
- ICAP interface cannot be accessed from the FPGA device fabric because the configuration logic is locked.

Successful completion of startup is indicated by the EOS signal being driven High. This can be observed in the STATUS register, or detected in the FPGA device fabric using the STARTUP primitive.

For designs accessing the ICAP, Xilinx recommends that you instantiate the STARTUP primitive. This primitive has an EOS pin, which will indicate that: (1) the configuration process has completed; and (2) the ICAP is available for read and write access.

Remote Update

Xilinx FPGA devices support MultiBoot and fallback features that make updating systems in the field more robust. Bitstream images can be upgraded dynamically in the field. The MultiBoot and fallback feature can be used with all master configuration modes.

The MultiBoot feature enables switching between images by the user application. If an error is detected during the MultiBoot configuration process, the device can trigger a fallback mechanism to retrieve a known bitstream from a different flash address.

Implementation of a robust in-system update solution involves a set of decisions around the initial configuration method, the update method, and any fallback mechanisms.

A MultiBoot solution requires a flash large enough to hold all required bitstreams. Because compression results can vary, Xilinx recommends considering the maximum uncompressed bitstream when planning the flash memory map.

There are some special cases when using advanced options with fallback.

- Master SPI configuration mode falls back to x1 mode in 7 series FPGAs.
- BPI configuration mode synchronous read falls back to asynchronous read mode in 7 series FPGAs. This means that the higher clock speeds intended for synchronous read may fail if fallback is utilized. The clock frequency used must be able to accommodate both modes.

Debugging

In-system debugging allows you to debug your design in real time on your target device. This step is needed if you encounter situations that are extremely difficult to replicate in a simulator.

For debug, you provide your design with special debugging hardware that allows you to observe and control the design. After debugging, you can remove the instrumentation or special hardware to increase performance and logic reduction.

Debugging an FPGA design is a multistep, iterative process. Like most complex problems, it is best to break the FPGA design debugging process down into smaller parts by focusing on getting smaller sections of the design working one at a time rather than trying to get the whole design to work at once.

Though the actual debugging step comes after you have successfully implemented your design, Xilinx recommends planning how and where to debug early in the design cycle. You can run all necessary commands to perform programming of the FPGA devices and in-system debugging of the design from the Program and Debug section of the Flow Navigator window in the Vivado IDE.

The steps involved in debug are:

1. **Probing:** Identify the signals in your design that you want to probe and how you want to probe them.
2. **Implementing:** Implement the design that includes the additional debug IP attached to the probed nets.
3. **Analyzing:** Interact with the debug IP contained in the design to debug and verify functional issues.
4. **Fixing phase:** Fix any bugs and repeat as necessary.

For more information, see *Vivado Design Suite User Guide: Programming and Debugging* (UG908) [Ref 24].

Probing the Design

The Vivado tools provide several methods to add debug probes in your design. The table below explains the various methods, including the pros and cons of each method.

Table 6-2: Debugging Flows

Debugging Flow Name	Flow Steps	Pros/Cons
HDL instantiation probing flow	Explicitly attach signals in the HDL source to an ILA debug core instance.	<ul style="list-style-type: none"> You have to add/remove debug nets and IP from your design manually, which means that you will have to modify your HDL source This method provides the option to probe at the HDL design level. It is easy to make mistakes when generating, instantiating, and connecting debug cores.
Netlist insertion probing flow (recommended) This method works only with ILA 2.1 or later.	<p>Use one of the following two methods to identify the signal for debug:</p> <ul style="list-style-type: none"> Use the MARK_DEBUG attribute to mark signals for debug in the source RTL code. Use the MARK_DEBUG right-click menu option to select nets for debugging in the synthesized design netlist. <p>Once the signal is marked for debug, use the Set up Debug wizard to guide you through the Netlist Insertion probing flow.</p>	<ul style="list-style-type: none"> This method is the most flexible with good predictability. This method allows probing at different design levels (HDL, synthesized design, system design). This method does not require HDL source modification.
Tcl-based netlist insertion probing flow	Use the <code>set_property</code> Tcl command to set the MARK_DEBUG property on debug nets then use Netlist insertion probing Tcl commands to create debug cores and connect them to debug nets. See Modifying the Netlist for post-synthesis insertion of ILA core.	<ul style="list-style-type: none"> This method provides fully automatic netlist insertion You can turn debugging on or off by modifying the Tcl commands. This method does not require HDL source modification.

Debug Insertion is a Two-Step Process

Using MARK_DEBUG on nets of interest and adding the debug cores into the design is a two-step process:

1. Identify the nets for debug by right-clicking on nets and setting the MARK_DEBUG attribute or BY using the properties window and/or Tcl commands to set this property.
2. After synthesis, use the Set up Debug wizard. You can access the wizard from the Flow Navigator Synthesis tab or by selecting the **Tools > Set up Debug** menu item if you are in the design.

Choosing Debug Nets

Xilinx makes the following recommendations for choosing debug nets:

- Probe nets at the boundaries (inputs or outputs) of a specific hierarchy. This method helps isolate problem areas quickly. Subsequently, you can probe further in the hierarchy if needed.
- Do not probe nets in between combinatorial logic paths. If you add MARK_DEBUG on nets in the middle of a combinatorial logic path, none of the optimizations applicable at the implementation stage of the flow are applied, resulting in sub-par QOR results.
- Probe nets that are synchronous in order to get cycle accurate data capture.

Retaining Names of Debug Probe Nets Using MARK_DEBUG

You can mark a signal for debug either at the RTL stage or post-synthesis. The presence of the MARK_DEBUG attribute on the nets ensures that the nets are not replicated, retimed, removed, or otherwise optimized. You can apply the MARK_DEBUG attribute on top level ports, nets, hierarchical module ports and nets internal to hierarchical modules. This method is most likely to preserve HDL signal names post synthesis. Nets marked for debugging are shown in the Unassigned Debug Nets folder in the Debug window post synthesis.

Add the mark_debug attribute to HDL files as follows:

VHDL:

```
attribute mark_debug : string;
attribute keep : string;
attribute mark_debug of sine : signal is "true";
```

Verilog:

```
(* mark_debug = "true" *) wire sine;
```

You can also add nets for debugging in the post-synthesis netlist. These methods do not require HDL source modification. However, there may be situations where synthesis might not have preserved the original RTL signals due to netlist optimization involving absorption or merging of design structures. Post-synthesis, you can add nets for debugging in any of the following ways:

- Select a net in any of the design views (such as the Netlist or Schematic windows), then right-click and select **Mark Debug**.
- Select a net in any of the design views, then drag and drop the net into the Unassigned Debug Nets folder.
- Use the net selector in the Set Up Debug Wizard.
- Set the MARK_DEBUG property using the properties window or the Tcl Console.

```
set_property mark_debug true [get_nets -hier [list {sine[*]}]]
```

This applies the mark_debug property on the current, open netlist. This method is flexible, because you can turn MARK_DEBUG on and off through the Tcl command.

Using ILA Cores

The Integrated Logic Analyzer (ILA) core allows you to perform in-system debugging of post-implementation designs on an FPGA device. Use this core when you need to monitor signals in the design. You can also use this feature to trigger on hardware events and capture data at system speeds.

Xilinx recommends inserting ILA cores after synthesis so that you do not have to modify HDL source files and to avoid the need to reverify the design.

To add nets to the debug cores, open the synthesized design and select **Set up Debug** from the Flow Navigator window or select the **Tools > Set up Debug** menu item.

For more information, see *Vivado Design Suite User Guide: Programming and Debugging* (UG908) [Ref 24].

ILA Core and Timing Considerations

The configuration of the ILA core has an impact in meeting the overall design timing goals. Follow the recommendations below to minimize the impact on timing:

- Choose probe width judiciously. The bigger the probe width the greater the impact on both resource utilization and timing.
- Choose ILA core data depth judiciously. The bigger the data depth the greater the impact on both block RAM resource utilization and timing.
- Ensure that the clocks chosen for the ILA cores are free-running clocks. Failure to do so could result in an inability to communicate with the debug core when the design is loaded onto the device.
- Ensure that the clock going to the `dbg_hub` is a free running clock. Failure to do so could result in an inability to communicate with the debug core when the design is loaded onto the device. You can use the `connect_debug_port` Tcl command to connect the `clk` pin of the debug hub to a free-running clock.
- Close timing on the design prior to adding the debug cores. Xilinx does not recommend using the debug cores to debug timing related issues.
- If you still notice that timing has degraded due to adding the ILA debug core and the critical path is in the `dbg_hub`, perform the following steps:
 - a. Open the synthesized design.
 - b. Find the `dbg_hub` cell in the netlist.
 - c. Go to the Properties of the `dbg_hub`.
 - d. Find property `C_CLK_INPUT_FREQ_HZ`.
 - e. Set it to frequency (in Hz) of the clock that is connected to the `dbg_hub`.
 - f. Find property `C_ENABLE_CLK_DIVIDER` and enable it.
 - g. Re-implement design.
- Make sure the clock input to the ILA core is synchronous to the signals being probed. Failure to do so results in timing issues and communication failures with the debug core when the design is programmed into the device.
- Make sure that the design meets timing before running it on hardware. Failure to do so results in unreliable results.

The following table shows the impact of using specific ILA features on design timing and resources.

Note: This table is based on a design with one ILA and does not represent all designs.

Table 6-3: Impact of ILA Features on Design Timing and Resources

ILA Feature	When to Use	Timing	Area
Capture Control/ Storage Qualification	<ul style="list-style-type: none"> To capture relevant data To make efficient use of data capture storage (BRAM) 	Medium to High Impact	<ul style="list-style-type: none"> No additional BRAMs Slight increase in LUT/FF count
Advanced Trigger	<ul style="list-style-type: none"> When BASIC trigger conditions are insufficient To use complex triggering to focus in on problem area 	High Impact	<ul style="list-style-type: none"> No additional BRAMs Moderate increase in LUT/FF count
Number of Comparators per Probe Port <i>Note:</i> Maximum is 4.	To use probe in multiple conditionals: <ul style="list-style-type: none"> 1-2 for Basic 1-4 for Advanced +1 for Capture Control 	Medium to High Impact	<ul style="list-style-type: none"> No additional BRAMs Slight to moderate increase in LUT/FF count
Data Depth	To capture more data samples	High Impact	<ul style="list-style-type: none"> Additional BRAMs per ILA core Slight increase in LUT/FF count
ILA Probe Port Width	To debug a large bus versus a scalar	Medium Impact	<ul style="list-style-type: none"> Additional BRAMs per ILA core Slight increase in LUT/FF count
Number of Probes Ports	To probe many nets	Low Impact	<ul style="list-style-type: none"> Additional BRAMs per ILA core Slight increase in LUT/FF count

For designs with high-speed clocks, consider the following:

- Limit the number and width of signals being debugged.
- Pipeline the input probes to the ILA (C_INPUT_PIPE_STAGES), which enables extra levels of pipe stages.

For designs with limited MMCM/BUFG availability, consider clocking the debug hub with the lowest clock frequency in the design instead of using the clock divider inside the debug hub.

Enabling ILA Cross-Trigger Mode

ILA Cross Triggering feature enables cross triggering between ILAs and between ILAs and a processor (for example, Zynq®-7000 AP SoC or MicroBlaze™ devices).

For using Cross-Trigger mode, at core generation time, you should configure the ILA core to have dedicated trigger input ports (TRIG_IN and TRIG_IN_ACK) and dedicated trigger output ports (TRIG_OUT and TRIG_OUT_ACK). If you want to use the ILA trigger input or output signals, consider using the HDL instantiation method of adding ILA cores to your design. Note that you need to use the netlist change commands to connect these ports to nets in your design if you insert the ILA into your design post-synthesis. See [Modifying the Netlist](#) for post-synthesis insertion of ILA core.

For a detailed tutorial that covers using the cross-trigger feature between the FPGA fabric and the Zynq-7000 AP SoC processor, see the *Vivado Design Suite Tutorial: Embedded Processor Hardware Design* (UG940) [[Ref 33](#)].

Debugging in Hardware

Once you have the debug cores in your design, you can use the runtime logic analyzer features to debug the design in hardware. To use the Vivado Design Suite logic analyzer to interact with the ILA debug cores instantiated in your design, select **Flow Navigator > Program and Debug > Open Hardware Session**.

In this step, you do the following:

- Connect with your target hardware
- Program the bitstream into the device
- Set up the ILA debug core trigger and probe conditions
- Arm the ILA debug core trigger
- Analyze the data captured from the ILA debug core in the Waveform window.

ILA Trigger Mode Settings

Trigger modes control the detection of real-time hardware events that are represented by trigger markers in the capture window.

- **BASIC_ONLY**: Use the ILA basic trigger mode to trigger the ILA core when a basic AND/OR functionality of debug probe comparison result is satisfied.
- **ADVANCED_ONLY**: Use the ILA advanced trigger mode to trigger the ILA core as specified by a user-defined state machine.
- **TRIG_IN_ONLY**: Use the ILA TRIG_IN trigger mode to trigger the ILA core when the TRIG_IN pin of the ILA core transitions from a low to high.

- **BASIC_OR_TRIG_IN:** Use the ILA BASIC_OR_TRIG_IN trigger mode to trigger the ILA core when you want a logical ORing of the TRIG_IN pin of the ILA core and BASIC_ONLY trigger mode.
- **ADVANCED_OR_TRIG_IN:** Use the ILA ADVANCED_OR_TRIG_IN trigger mode to trigger the ILA core when you want a logical ORing of the TRIG_IN pin of the ILA core and ADVANCED_ONLY trigger mode.

You can set the trigger position to a specific position in the captured data buffer. For example, in the case of a captured data buffer that is 1024 samples deep:

- Sample number 0 corresponds to the first (left-most) sample in the captured data buffer.
- Sample number 1023 corresponds to the last (right-most) sample in the captured data buffer.
- Samples numbers 511 and 512 correspond to the two center samples in the captured data buffer.

ILA Core Capture Modes

Capture modes control how the data is captured by the ILA core.

- **ALWAYS:** Captures probe data on every clock cycle.
- **BASIC:** Captures probe data when a basic AND/OR functionality of debug probe comparison result is satisfied.

Sharing Waveform for Data Captured by ILA Core

You can use the `write_hw_ila_data` Tcl command to save the data that was captured by the ILA into a file archive. To restore captured data from the file archive and display it in a waveform viewer use the Tcl command `display_hw_ila_data`.

This two-command sequence allows you to view in waveform the data that was captured by the ILA core. You need to run an `open_hw` command prior to running the above commands.

The waveform configuration settings (dividers, markers, colors, probe radices, etc.) for the ILA data waveform window are also saved in the ILA captured data archive file. Restoring and displaying any previously saved ILA data uses these stored waveform configuration settings.

Multiple Capture Windows

You can capture signals on multiple occurrences of trigger through the multiple capture feature. To use this feature, select the number of capture windows at run-time, as shown in the following figure, and arm the trigger as usual.

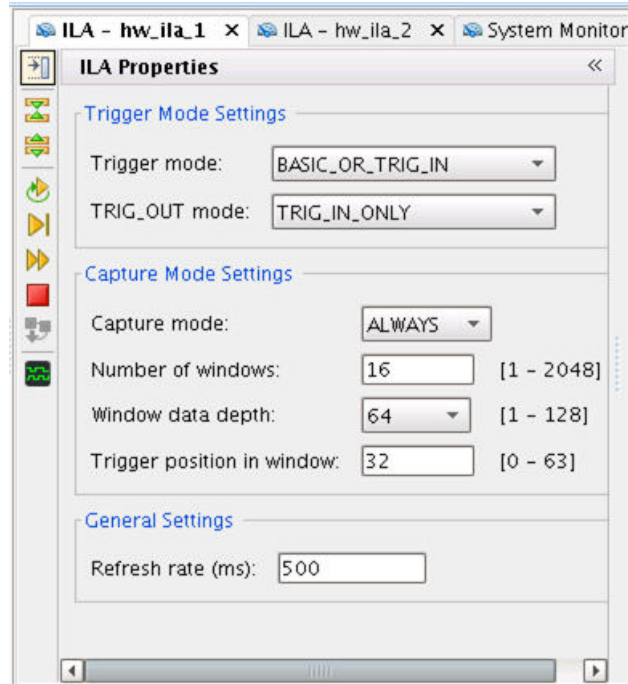


Figure 6-1: Setting ILA for Multiple Capture Windows

The following figure shows an example multiple capture window, triggered at each occurrence of falling edge of signal `fast_cnt_reset_1`. Observe the multiple windows with trigger marks and alternating “checkerboard” window background.

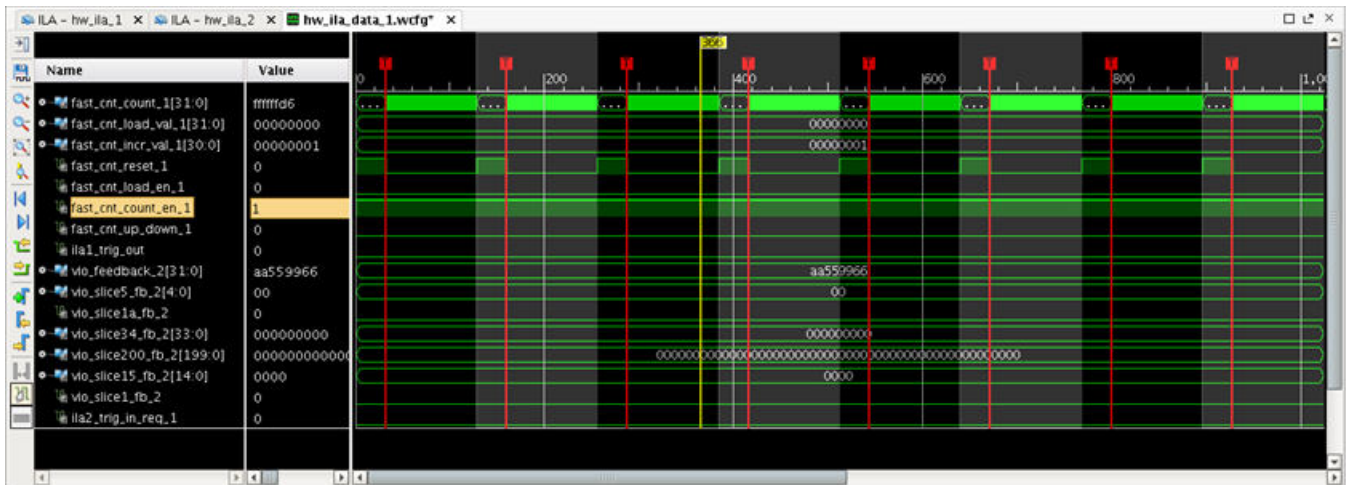


Figure 6-2: Sample Multiple Capture Windows

Virtual I/O (VIO)

The Virtual Input/Output (VIO) core is a customizable core that can both monitor and drive internal FPGA signals in real time. In the absence of physical access to the target hardware,

you can use this IP to drive and monitor signals that are present on the real hardware. Two different kinds of inputs and two different kinds of outputs are available, and both are customizable in size to interface with the design. If you need a VIO core, choose it from the IP Catalog and instantiate it in the design.

The VIO core output probes are used to write values to a design that is running on an FPGA device. The VIO output probes are typically used as low-bandwidth control signals for a design under test. Similarly, VIO input pins are used to read values from a design running on an FPGA device.

The VIO core does the following:

- Provides virtual LEDs and other status indicators through its input ports.
- Includes optional activity detectors on its input ports to detect rising and falling transitions between samples.
- Provides virtual buttons and other controls through its output ports.
- Includes custom output initialization that allows you to specify the value of the VIO core outputs immediately following device configuration and start-up.
- Runs time reset of the VIO core to its initial values.

Debugging Designs in Vivado IP Integrator

The Vivado IP Integrator provides different ways to set up your design for debugging. You can use one of the following flows to add debug cores to your IP Integrator design. The flow you choose depends on your preference and the types of nets and signals that you want to debug.

- HDL instantiation flow

Use this flow to:

- Perform hardware-software co-verification using the cross-trigger feature of a MicroBlaze device or Zynq-7000 AP SoC.
- Verify the interface-level connectivity.

- Netlist insertion flow

Use this flow to analyze I/O ports and internal nets.

Note: You can also use a combination of both flows to debug your design.

For more information on using ILA in your IP Integrator design, see the *Vivado Design Suite User Guide: Designing IP Subsystems Using IP Integrator* (UG994) [Ref 26].

Running Debug-Related DRCs

The Vivado Design Suite provides debug-related DRCs, which are selected as part of the default rule deck when `report_drc` is run. The DRCs check for the following:

- Block RAM resources for the device are exceeded because of the current requirements of the debug core.
- Non-clock net is connected to the clock port on the debug core.
- Port on the debug core is unconnected.

Generating AXI Transactions

Use the JTAG-to-AXI debug core to generate AXI transactions that interact with various AXI full and AXI lite slave cores in a system that is running on hardware. Instantiate this core in your design from the IP Catalog to generate AXI transactions and debug/drive AXI signals internal to your FPGA at run time. You also can use this core in designs without processors.

Using the Probes File

The probes file is an `.ltx` file that corresponds to the `.bit` file associated with the device. The tool automatically generates this probes file during the implementation process. The Vivado tools also automatically associate the debug probes file with the hardware device if the tools are in Project Mode and if a probes file (`debug_nets.ltx`) is located in the same directory as the bitstream programming (`.bit`) file that is associated with the device.

In case of a suspected mismatch between the bit file programmed into the device and the probes file associated with the `.bit` file, ensure that the `.bit` file and probes file are up to date.

To write out the debug probes information to a file, use the `write_debug_probes` Tcl command on a synthesized design. Use the following procedure to specify the location of the probes file:

1. Select the FPGA device in the Hardware window.
2. Set the Probes file location in the Hardware Device Properties window.
3. Click **Apply**.

You also can set the location using the `set_property` Tcl command as well:

```
set_property PROBES.FILE {location} [lindex [get_hw_devices] 0]
```

Slowing Down JTAG clock

The JTAG chain is as fast as the slowest device in the chain. Therefore, to lower the JTAG clock frequency, connect to a device target whose JTAG clock frequency is less than the default JTAG clock frequency.

You should attempt to open with a default JTAG clock frequency that is 15 MHz for the Digilent cable connection and 6 MHz for the USB cable connection. If it is not possible to connect at these speeds, Xilinx recommends that you lower the default JTAG clock frequency even further as described below.

To change the JTAG clock frequency, use the Open New Hardware Target wizard, from the Vivado IDE, as shown in the following figure.

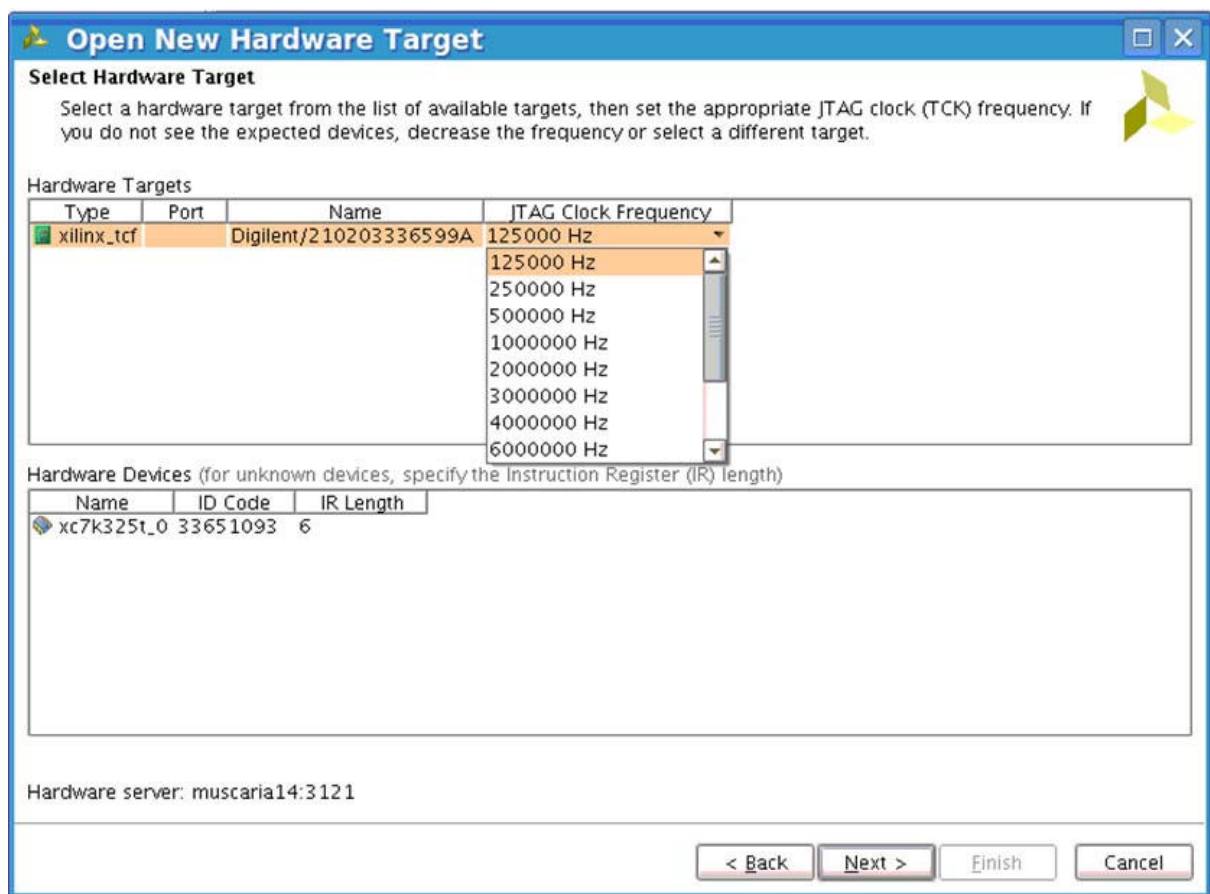


Figure 6-3: Setting JTAG frequency

Alternately, you can use the following sequence of Tcl commands:

```
open_hw
connect_hw_server -host localhost -port 60001 -url machinename:3121
current_hw_target [get_hw_targets */xilinx_tcf/Digilent/210203327962A]
set_property PARAM.FREQUENCY 250000 [get_hw_targets
*/xilinx_tcf/Digilent/210203327962A]
open_hw_target
```


Vivado Debug Layouts

The Vivado IDE provides specific layouts associated with the debug process when the design is loaded onto the device. Set the Vivado IDE to the correct layout to ensure that the right windows are displayed in the debug context. When debugging the design using the ILA and VIO cores, use the Logic Analyzer layout. When debugging the design using the IBERT core, use the Serial I/O Analyzers layout.

Tcl Objects and Commands

Table 6-4: Tcl Objects and Commands

Tcl Object	Represents
hw_server	Hardware server
hw_target	Hardware target (cable) connection. A hw_target is a live and physically link to a server.
hw_device	A real device in the hardware. The device may be automatically detected from a live server connected to hardware or manually created from an hw_part query.
hw_ila	An ILA core.
hw_ila_data	A captured ILA data, together with probe information.
hw_probe	A probe to debug. A probe is associated with a list of nets in your design.
hw_probeset	A collection of hw_probe objects.
hw_propset	A property set in the hardware core.

Many of the actions described in this chapter can be accomplished by the use of Tcl commands. For a description of the Tcl commands for debug, see the *Vivado Design Suite Tcl Command Reference Guide* (UG835) [Ref 13].

Recommended Design Practices

1. Plan for debug well ahead during the design phase by following these guidelines:
 - a. Reserve logic slices and block RAMs for in-system debug.
 - b. Make sure that the JTAG interface to the FPGA device is accessible and used for debug.
 - c. Identify the appropriate debugging flow that you will use to debug your design. This flow may be:
 - HDL instantiation
 - Netlist insertion (recommended)
 - A combination of both flows

- d. Consider probing outputs of synchronous cells such as flip-flops and block RAM as opposed to combinational logic. This minimizes the impact on design optimization, and improves your chances of meeting timing.
 - e. Add custom debug logic to your design.
2. When performing debug, follow these guidelines:
 - a. Do not test the entire design at once.
 - b. Make incremental changes to the design, and test features one at a time.
 - c. Consider design implementation using incremental compilation.
 - d. Route control and high-speed data signals to pins for analysis on a logic analyzer or scope.
 - e. In order to capture data based on events, add debug cores to your design.
 3. Consider recreating the problem in RTL simulation and validate that the fix works in simulation as well.
 4. After you have successfully debugged your design, consider removing the debug cores before entering production in order to:
 - a. Reduce the possibility of unauthorized access to the design using JTAG.
 - b. Reduce power consumption.
 - c. Take into account any effects that adding debug cores to your design might have on your design timing constraints.

For more information, see:

- *Vivado Design Suite User Guide: Programming and Debugging* (UG908) [\[Ref 24\]](#)
- *Vivado Design Suite Tcl Command Reference Guide* (UG835) [\[Ref 13\]](#)
- *Configuration User Guide* for your device [\[Ref 37\]](#)

Modifying the Netlist

Netlists sometimes require changes to fix functional logic bugs, meet timing closure, or insert debug logic. This is commonly referred to as Engineering Change Order (ECO). You can modify an existing netlist using Tcl commands post-synthesis, post-place, and post-route. For information on netlist modifying commands, see this [link](#) in the *Vivado Design Suite User Guide: Implementation* (UG904) [\[Ref 19\]](#).

Baselining and Timing Constraints Validation Procedure

Introduction

As the design progresses through the implementation stages and you keep refining your constraints, fill in the following questionnaire. This questionnaire helps track your progress/deviation from timing closure, any potential bottlenecks, and any constraints which need fixing.

Procedure

1. Open the synthesized design.
2. Run `report_timing_summary -delay_type min_max` and record the information shown in the following table.

	WNS	TNS	Num Failing Endpoints	WHS	THS	Num Failing Endpoints
Synth						

3. Open the post-synthesis `report_timing_summary` text report and record the `no_clock` section of `check_timing`.

Number of missing clock requirements in the design: _____

4. Run `report_clock_networks` to identify primary clock source pins/ports in the design. (Ignore QPLLOUTCLK, QPLLOUTREFCLK because they are pulse-width only checks.)

Number of unconstrained clocks in the design: _____

5. Run `report_clock_interaction -delay_type min_max` and sort the results by WNS path requirement.

Smallest WNS path requirement in the design: _____

- Sort the results of `report_clock_interaction` by WHS to see if there are large hold violations (>500 ps) after synthesis.

Largest negative WHS in the design: _____

- Sort results of `report_clock_interaction` by Inter-Clock Constraints and list *all* the clock pairs that show up as unsafe:

- Upon opening the synthesized design, how many CRITICAL_WARNINGS exist?

Number of synthesized design CRITICAL WARNINGS: _____

- What types of CRITICAL WARNINGS exist?

Record examples of each type.

- Run `report_high_fanout_nets -timing -load_types -max_nets 25`.

Number of high fanout nets NOT driven by FF: _____

Number of loads on highest fanout net NOT driven by FF: _____

Do any high fanout nets have negative slack?- If yes, WNS = _____

- Implement the design. After each step, run `report_timing_summary` and record the information shown in the following table.

	WNS	TNS	Num Failing Endpoints	WHS	THS	Num Failing Endpoints
Opt						
Place						
Physopt						
Route						

- Run `report_exceptions -ignored` to identify if there are constraints that overlap in the design. Record the results.

Additional Resources and Legal Notices

Xilinx Resources

For support resources such as Answers, Documentation, Downloads, and Forums, see [Xilinx Support](#).

Solution Centers

See the [Xilinx Solution Centers](#) for support on devices, tools, and intellectual property at all stages of the design cycle. Topics include design assistance, advisories, and troubleshooting tips.

References

These documents provide supplemental material useful with this guide.

1. [Vivado® Design Suite Documentation](#)
2. [UltraFast™ Design Methodology Checklist](#)

Vivado Design Suite User and Reference Guides

3. *Vivado Design Suite User Guide: Release Notes, Installation, and Licensing* ([UG973](#))
4. *Vivado Design Suite User Guide: I/O and Clock Planning* ([UG899](#))
5. *Vivado Design Suite User Guide: Design Flows Overview* ([UG892](#))
6. *Vivado Design Suite User Guide: Using the Vivado IDE* ([UG893](#))
7. *Vivado Design Suite User Guide: Using Tcl Scripting* ([UG894](#))
8. *Vivado Design Suite User Guide: System-Level Design Entry* ([UG895](#))
9. *Vivado Design Suite User Guide: Designing with IP* ([UG896](#))

10. *Vivado Design Suite User Guide: Embedded Processor Hardware Design* ([UG898](#))
11. *Vivado Design Suite User Guide: Logic Simulation* ([UG900](#))
12. *Vivado Design Suite User Guide: Getting Started* ([UG910](#))
13. *Vivado Design Suite Tcl Command Reference Guide* ([UG835](#))
14. *Vivado Design Suite Properties Reference Guide* ([UG912](#))
15. *AXI BFM Cores LogiCORE IP Product Guide* ([PG129](#))
16. *Vivado Design Suite User Guide: Synthesis* ([UG901](#))
17. *Vivado Design Suite User Guide: High-Level Synthesis* ([UG902](#))
18. *Vivado Design Suite User Guide: Using Constraints* ([UG903](#))
19. *Vivado Design Suite User Guide: Implementation* ([UG904](#))
20. *Vivado Design Suite User Guide: Hierarchical Design* ([UG905](#))
21. *Vivado Design Suite User Guide: Design Analysis and Closure Techniques* ([UG906](#))
22. *Vivado Design Suite User Guide: Power Analysis and Optimization* ([UG907](#))
23. *Xilinx Power Estimator User Guide* ([UG440](#))
24. *Vivado Design Suite User Guide: Programming and Debugging* ([UG908](#))
25. *Vivado Design Suite User Guide: Partial Reconfiguration* ([UG909](#))
26. *Vivado Design Suite User Guide: Vivado Design Suite User Guide: Designing IP Subsystems Using IP Integrator* ([UG994](#))
27. *Vivado Design Suite User Guide: Creating and Packaging Custom IP* ([UG1118](#))

Vivado Design Suite Tutorials

28. *Vivado Design Suite Tutorial: High-Level Synthesis* ([UG871](#))
29. *Vivado Design Suite Tutorial: Design Flows Overview* ([UG888](#))
30. *Vivado Design Suite Tutorial: Revision Control Systems* ([UG1198](#))
31. *Vivado Design Suite Tutorial: I/O and Clock Planning* ([UG935](#))
32. *Vivado Design Suite Tutorial: Logic Simulation* ([UG937](#))
33. *Vivado Design Suite Tutorial: Embedded Processor Hardware Design* ([UG940](#))
34. *Vivado Design Suite Tutorial: Partial Reconfiguration* ([UG947](#))

Other Xilinx Documentation

35. *7 Series FPGAs PCB Design Guide* ([UG483](#))

- UltraScale Architecture PCB Design User Guide* ([UG583](#))
- Zynq-7000 All Programmable SoC PCB Design Guide* ([UG933](#))
36. *UltraFast Embedded Design Methodology Guide* ([UG1046](#))
37. *7 Series FPGAs Configuration User Guide* ([UG470](#))
- UltraScale Architecture Configuration User Guide* ([UG570](#))
38. *7 Series FPGAs SelectIO Resources User Guide* ([UG471](#))
- UltraScale Architecture SelectIO Resources User Guide* ([UG571](#))
39. *7 Series Clocking Resources Guide* ([UG472](#))
- UltraScale Architecture Clocking Resources User Guide* ([UG572](#))
40. *UltraScale Architecture GTH Transceivers User Guide* ([UG576](#))
- UltraScale Architecture GTY Transceivers Advance Specification User Guide* ([UG578](#))
41. *UltraScale Architecture Gen3 Integrated Block for PCI Express LogiCORE IP Product Guide* ([PG156](#))
42. *7 Series FPGAs Memory Resources User Guide* ([UG473](#))
43. *7 Series FPGAs DSP48E1 Slice User Guide* ([UG479](#))
44. *UltraScale Architecture DSP Slice User Guide* ([UG579](#))
45. *7 Series FPGAs and Zynq-7000 All Programmable SoC XADC Dual 12-Bit 1 MSPS Analog-to-Digital Converter User Guide* ([UG480](#))
46. *Reference System: Kintex-7 MicroBlaze System Simulation Using IP Integrator* ([XAPP1180](#))
47. *Zynq-7000 SoC and 7 Series FPGAs Memory Interface Solutions User Guide* ([UG586](#))
48. *UltraScale Architecture FPGAs Memory IP LogiCORE IP Product Guide* ([PG150](#))
49. *Xilinx White Paper: Simulating FPGA Power Integrity Using S-Parameter Models* ([WP411](#))
50. *7 Series Schematic Review Recommendations* ([XMP277](#))
51. *UltraScale Architecture Schematic Review Checklist* ([XTP344](#))
52. *UltraScale FPGA BPI Configuration and Flash Programming* ([XAPP1220](#))
53. *BPI Fast Configuration and iMPACT Flash Programming with 7 Series FPGAs* ([XAPP587](#))
54. *Using SPI Flash with 7 Series FPGAs* ([XAPP586](#))
55. *SPI Configuration and Flash Programming in UltraScale FPGAs* ([XAPP1233](#))
56. *Using Encryption to Secure a 7 Series FPGA Bitstream* ([XAPP1239](#))



TIP: A complete set of Xilinx documents can be accessed from Documentation Navigator. For more information, see [Using the Documentation Navigator](#).

Training Resources

Xilinx provides a variety of training courses and QuickTake videos to help you learn more about the concepts presented in this document. Use these links to explore related training resources:

1. [Essentials of FPGA Design Training Course](#)
2. [Xilinx Video Training: UltraFast Vivado Design Methodology](#)
3. [Vivado Design Suite QuickTake Video: Vivado Design Flows Overview](#)
4. [Vivado Design Suite QuickTake Video: Designing with Vivado IP Integrator](#)
5. [Vivado Design Suite QuickTake Video: Targeting Zynq Using Vivado IP Integrator](#)
6. [Vivado Design Suite QuickTake Video: Partial Reconfiguration in Vivado Design Suite](#)
7. [Vivado Design Suite QuickTake Video: Creating Different Types of Projects](#)
8. [Vivado Design Suite QuickTake Video: Managing Sources With Projects](#)
9. [Vivado Design Suite QuickTake Video: Using Vivado Design Suite with Revision Control](#)
10. [Vivado Design Suite QuickTake Video: Managing Vivado IP Version Upgrades](#)
11. [Vivado Design Suite QuickTake Video: I/O Planning Overview](#)
12. [Vivado Design Suite QuickTake Video: Configuring and Managing Reusable IP in Vivado](#)
13. [Vivado Design Suite QuickTake Video: How To Use the "write_bitstream" Command in the Vivado Design Suite](#)
14. [Vivado Design Suite QuickTake Video: Customizing and Instantiating IP](#)
15. [Vivado Design Suite QuickTake Video: Design Analysis and Floorplanning](#)
16. [Vivado Design Suite QuickTake Video: Introducing the UltraFast Design Methodology Checklist](#)
17. [Vivado Design Suite Video Tutorials](#)

Please Read: Important Legal Notices

The information disclosed to you hereunder (the "Materials") is provided solely for the selection and use of Xilinx products. To the maximum extent permitted by applicable law: (1) Materials are made available "AS IS" and with all faults, Xilinx hereby DISCLAIMS ALL WARRANTIES AND CONDITIONS, EXPRESS, IMPLIED, OR STATUTORY, INCLUDING BUT NOT LIMITED TO WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, OR FITNESS FOR ANY PARTICULAR PURPOSE; and (2) Xilinx shall not be liable (whether

in contract or tort, including negligence, or under any other theory of liability) for any loss or damage of any kind or nature related to, arising under, or in connection with, the Materials (including your use of the Materials), including for any direct, indirect, special, incidental, or consequential loss or damage (including loss of data, profits, goodwill, or any type of loss or damage suffered as a result of any action brought by a third party) even if such damage or loss was reasonably foreseeable or Xilinx had been advised of the possibility of the same. Xilinx assumes no obligation to correct any errors contained in the Materials or to notify you of updates to the Materials or to product specifications. You may not reproduce, modify, distribute, or publicly display the Materials without prior written consent. Certain products are subject to the terms and conditions of Xilinx's limited warranty, please refer to Xilinx's Terms of Sale which can be viewed at <http://www.xilinx.com/legal.htm#tos>; IP cores may be subject to warranty and support terms contained in a license issued to you by Xilinx. Xilinx products are not designed or intended to be fail-safe or for use in any application requiring fail-safe performance; you assume sole risk and liability for use of Xilinx products in such critical applications, please refer to Xilinx's Terms of Sale which can be viewed at <http://www.xilinx.com/legal.htm#tos>.

© Copyright 2013–2015 Xilinx, Inc. Xilinx, the Xilinx logo, Artix, ISE, Kintex, Spartan, Virtex, Vivado, Zynq, and other designated brands included herein are trademarks of Xilinx in the United States and other countries. AMBA, AMBA Designer, ARM, ARM1176JZ-S, CoreSight, Cortex, and PrimeCell are trademarks of ARM in the EU and other countries. MATLAB and Simulink are registered trademarks of The MathWorks, Inc. OpenCL and the OpenCL logo are trademarks of Apple Inc. used by permission by Khronos. PCI, PCIe and PCI Express are trademarks of PCI-SIG and used under license. All other trademarks are the property of their respective owners.