Application Report
SPRABI2C—August 2013

Hardware Design Guide for KeyStone I Devices

High-Performance and Multicore Processors

Abstract

This document describes hardware system design considerations for the KeyStone I family of processors. This design guide is intended to be used as an aid during the development of application hardware. Other aids including, but not limited to, device data manuals and explicit collateral should also be used.

Contents

List of Tables

List of Figures

Purpose

This document is intended to aid in the hardware design and implementation of a KeyStone I-based system. This document should be used along with the respective data manual and other relevant user guides and application reports.

Trademarks

SmartReflex is a trademark of Texas Instruments.
All other trademarks belong to their respective owners.

Terms and Abbreviations

AIF	Antenna Interface
AMI	IBIS Algorithmic Modeling Interface
BGA	Ball Grid Array
CML	Current Mode Logic, I/O type
Data Manual	Also referred to as the Data Sheet
DDR3	Double Data Rate 3 (SDRAM Memory)
DSP	Digital Signal Processor
EMIF	External Memory Interface
EVM	Evaluation Module
FC-BGA	Flip-Chip BGA
GPIO	General-Purpose I/O
I2C	Inter-IC Control Bus
IBIS	Input Output Buffer Information Specification, or ANSI/EIA-656-A
IO	Input / Output
JEDEC	Joint Electronics Device Engineering Council
LJCB	Low Jitter Clock Buffer: Differential clock input buffer type, compatible with LVDS & LVPECL
LVDS	Low Voltage Differential Swing, I/O type
McBSP	Multi-Channel Buffered Serial Port
MDIO	Management Data Input/Output
NSMD	Non-Solder Mask Defined BGA Land
OBSAI	Open Base Station Architecture Initiative
PCIe	Peripheral Component Interconnect Express
PHY	Physical Layer of the Interface
PTV	Process / Temperature / Voltage
RIO	Rapid IO, also referred to as SRIO
Seating Plane	The maximum compression depth of the BGA for a given package design
SerDes	Serializer/De-Serializer
SGMII	Serial Gigabit Media Independent Interface
SPI	Synchronous Serial Input/Output (port)
SRIO	Serial RapidIO
TBD	To Be Determined. Implies something is currently under investigation and will be clarified in a later version of the specification.
UART	Universal Asynchronous Receiver / Transmitter
UI	Unit Interval
XAUI	10 Gigabit (X) Attachment Unit Interface standard

1 Introduction

The *Hardware Design Guide for KeyStone I Devices Application Report* provides a starting point for the engineer designing with one of the KeyStone I devices. It shows a road map for the design effort and highlights areas of significant importance that must be addressed. This document does not contain all the information that is needed to complete the design. In many cases, it refers to the device-specific data manual or to various user guides as sources for specific information.

This guide is generic to the entire family of KeyStone I devices, and as such, may include information for subsystems that are not present on all devices. Designers should begin by reviewing the device-specific data manual for their intended device to determine which sections of this guide are relevant.

The guide is organized in a sequential manner. It moves from decisions that must be made in the initial planning stages of the design, through the selection of support components, to the mechanical, electrical, and thermal requirements. For the greatest success, each of the issues discussed in a section should be resolved before moving to the next section.

1.1 Before Getting Started

The KeyStone I family of processors provides a wide variety of capabilities, not all of which will be used in every design. Consequently, the requirements for different designs using the same device can vary widely depending on how that part is used. Many designers simply copy the evaluation module (EVM) for the device without determining their requirements. While the EVM is a good example, it is not a reference design and is not be the optimum design for every customer. You will need to understand your requirements before determining the details of the design. In addition, your design may require additional circuitry to operate correctly in the target environment. Take some time to review the data manual for your KeyStone I device and determine the following:

- Which peripherals will be used to move data in and out of the processor
- What is the speed and organization of the DDR3 memory interface that will be used
- How much processing will each of the cores in your KeyStone I device be performing
- How you will boot the KeyStone I device
- What are the expected environmental conditions for your KeyStone I device

1.2 Design Documentation

Throughout this guide, we will periodically recommend that a design document be generated based on your requirements. Generating and storing this information will provide you with the foundation for your documentation package and this design document will be needed if you are seeking support from TI. Examples of many of these can be found in the schematic package provided with the EVM boards for your KeyStone I device.

2 Power Supplies

The first requirement for a successful design is to determine the power needs for your KeyStone I device. All KeyStone I devices operate with four main voltage levels requiring four power supply circuits. Some devices will require an additional voltage level, requiring a fifth power supply. Check the device-specific data manual to determine all the voltages needed.

2.1 Determining Your Power Requirements

The maximum and minimum current requirements for each of these voltage rails are not found in the data manual for the part. These requirements are highly application-dependent and must be calculated for your specific product. There is a power consumption model in the product folder for each KeyStone I device. **You must use the model for the KeyStone I device that you have selected to get accurate results.**

There is a link in the model spreadsheet to an application note that explains how to populate the necessary parameters. Review this application note carefully and determine the values needed for your application. Once you have entered these values you will have the maximum power requirement for each of the four voltage rails. These values are divided into *activity* and *baseline* components. The *baseline* power portion is associated with leakage, clock tree, and phase-locked loop (PLL) power. This value will not change based on the processor utilization. The *activity* power reflects the additional power required due to the processing load defined in the spreadsheet. The *baseline* plus the *activity* defines your maximum power requirement. Power supplies should be designed to transition from the *baseline* level to the *baseline* + *activity* level within one CPU clock.

Note that power consumption can vary greatly based on process and temperature. The values provided by the spreadsheet are the maximum average power consumption for all production devices. Measurements of power consumption on a single device may be considerably lower than the values provided by the model, but may not be representative of a wide sample of devices. To ensure that your power supply design is adequate, use the values provided by the model.

The power consumption model populated with the values for your application should be saved with your design documentation.

2.2 Power Rails

The power consumption model will provide the requirements for four separate voltage rails as shown in Figure 1.

Figure 1 KeyStone I Power Supply Planes (Rails)

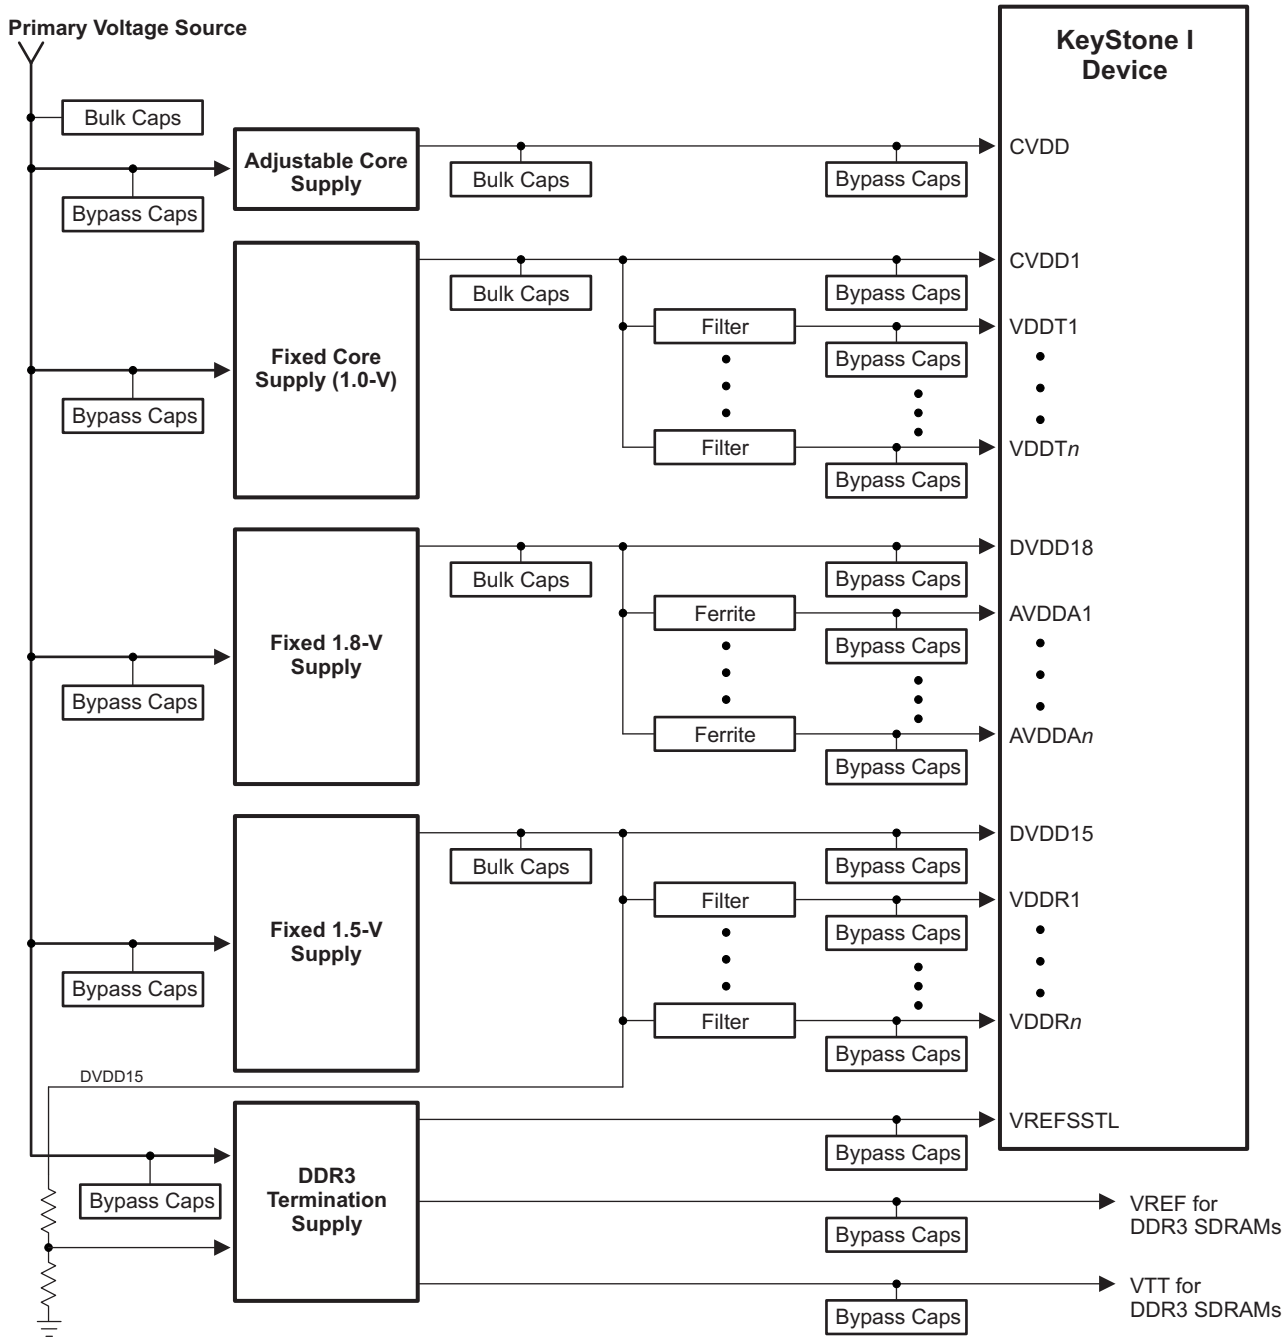

These four rails (CVDD, CVDD1, DVDD18, and DVDD15) provide the basic power requirement for all KeyStone I devices. In addition, the KeyStone I device requires a number of filtered connections to the main rails to provide power for noise-sensitive portions of the device. These include the AVDDA*n*, VDDT*n*, and VDDR*n* pins. The power requirements for each of the filtered voltages are included in the values

calculated by the power consumption models. Figure 1 also includes a DDR3 termination power supply. DDR3 address and command signals are terminated using a source/sink-tracking LDO designed to provide VTT and a low-noise reference. Figure 1 includes this power supply as a source for the VREFSSTL reference voltage. Each of these voltage rails are described in the following sections.

2.2.1 CVDD

CVDD is the adjustable supply used by the core logic for the KeyStone I device. KeyStone I devices use adaptive voltage scaling (AVS) to compensate for variations in performance from die to die, and from wafer to wafer. The actual voltage for CVDD can vary across the range specified in the device data manual and will be different for each component. Each KeyStone I device needs an independent power supply to generate the CVDD voltage needed by that device. That CVDD power supply must use a SmartReflex-compliant circuit. SmartReflex is described in the next section. This power supply circuit must initialize to a level of 1.1 V and then adjust to the voltage requested by the SmartReflex circuit in the KeyStone I device.

2.2.1.1 SmartReflex

To reduce device power consumption, SmartReflex allows the core voltage to be optimized (scaled) based on the process corners of each device. KeyStone I devices use Class 0 SmartReflex. This class of operation uses a code for a single ideal voltage level determined during manufacturing tests. At the end of these tests, the code for the lowest acceptable voltage (while still meeting all performance requirements) is established and permanently programmed into each die. This 6-bit code, referred to as the VID value, represents the optimal fixed voltage level for that device.

Table 1 shows the voltage level associated with each of the possible 6-bit VID values.

Table 1 SmartReflex VID Value Mapping

VID #	VID (5:0)	CVDD	VID #	VID (5:0)	CVDD	VID #	VID (5:0)	CVDD	VID #	VID (5:0)	CVDD
0	00 0000b	0.7	16	01 0000b	0.802	32	10 0000b	0.905	48	11 0000b	1.007
1	00 0001b	0.706	17	01 0001b	0.809	33	10 0001b	0.911	49	11 0001b	1.014
2	00 0010b	0.713	18	01 0010b	0.815	34	10 0010b	0.918	50	11 0010b	1.02
3	00 0011b	0.719	19	01 0011b	0.822	35	10 0011b	0.924	51	11 0011b	1.026
4	00 0100b	0.726	20	01 0100b	0.828	36	10 0100b	0.93	52	11 0100b	1.033
5	00 0101b	0.732	21	01 0101b	0.834	37	10 0101b	0.937	53	11 0101b	1.039
6	00 0110b	0.738	22	01 0110b	0.841	38	10 0110b	0.943	54	11 0110b	1.046
7	00 0111b	0.745	23	01 0111b	0.847	39	10 0111b	0.95	55	11 0111b	1.052
8	00 1000b	0.751	24	01 1000b	0.854	40	10 1000b	0.956	56	11 1000b	1.058
9	00 1001b	0.758	25	01 1001b	0.86	41	10 1001b	0.962	57	11 1001b	1.065
10	00 1010b	0.764	26	01 1010b	0.866	42	10 1010b	0.969	58	11 1010b	1.071
11	00 1011b	0.77	27	01 1011b	0.873	43	10 1011b	0.975	59	11 1011b	1.078
12	00 1100b	0.777	28	01 1100b	0.879	44	10 1100b	0.982	60	11 1100b	1.084
13	00 1101b	0.783	29	01 1101b	0.886	45	10 1101b	0.988	61	11 1101b	1.09
14	00 1110b	0.79	30	01 1110b	0.892	46	10 1110b	0.994	62	11 1110b	1.097
15	00 1111b	0.796	31	01 1111b	0.898	47	10 1111b	1.001	63	11 1111b	1.103

End of Table 1

Note—Not all ranges or voltage levels are supported by every KeyStone I device. The intended range of operation is defined in the device-specific data manual. Operation outside that range may impact device reliability or performance.

2.2.1.2 SmartReflex VCNTL Interface

The VCNTL[3:0] interface is used to transmit the 6-bit VID value to the SmartReflex power supply circuit. The VCNTL[3:0] pins are open drain IOs, and require 4.7-kΩ pull-up resistors to the DVDD18 rail of the KeyStone I device. VCNTL[3] acts as a command signal while VCNTL[2:0] act as a 3-bit data bus. VCNTL[3] will transition low and then high to indicate the presence of the lower three bits of the VID value, followed by the upper three bits. The timing of this interface can be found in the data manual for your KeyStone I device. It is important to note that the VID values are not latched on the falling and rising edge of the VCNTL[3] signal. The VID values will be available only after a specified amount of time following the transition, as defined in the data manual for the KeyStone I device. As stated previously, this transition will occur only a single time after a power-up reset has occurred. The connection of the VCNTL interface to the power supply circuit is discussed in later sections.

2.2.1.3 Supported Power Supply Solutions

Texas Instruments provides a number of power supply solutions for CVDD, which fall into one of two classes. The first uses the LM10011 to interface the VCNTL pins with a number of analog power supply controllers. The second is a solution based on the UCD92xx family of digital power supply controllers.

Texas Instruments provides reference designs for both possible power supply solutions. These can be accessed through the power management page on TI.com:

http://www.ti.com/processorpower

Choose C6000 from the processor type pull down menu and then choose the KeyStone I part number from the processor family pull down menu. This will bring you to a page containing the documentation for the possible power supply solutions that are supported for that part. The page will include power supply requirements. These are generally the requirements calculated for the EVM platform and do not reflect the maximum requirements for the part. Your requirements should be calculated using the Power Consumption Model for your KeyStone I device.

The power management solutions are managed by a different group in TI. Questions on the power supply solutions should be directed to the appropriate power group. Question about the LM10011-based solution can be posted to the non-isolated DC/DC forum and questions about the UCD92xx bases solutions can be posted to the Digital Power Forum. If you need additional information, please contact your local Texas Instruments FAE.

2.2.1.4 LM10011 and Analog Controller Solutions

The LM10011 captures the VID value presented on the VCNTL interface and uses it to control the voltage level of a DC/DC converter. The VID value is used to define the output of a current DAC that is connected to the feedback pin of a DC/DC switching regulator. This will adjust the output voltage of the regulator to the level required by the KeyStone I device. The LM10011 can be used to control any DC/DC regulator with a compatible feedback circuit. The current requirements calculated using the power consumption model should be used to determine the best device for your application.

The LM10011 VID interface should be connected as shown in Figure 2.

Figure 2 LM10011 CVDD Supply Block Diagram

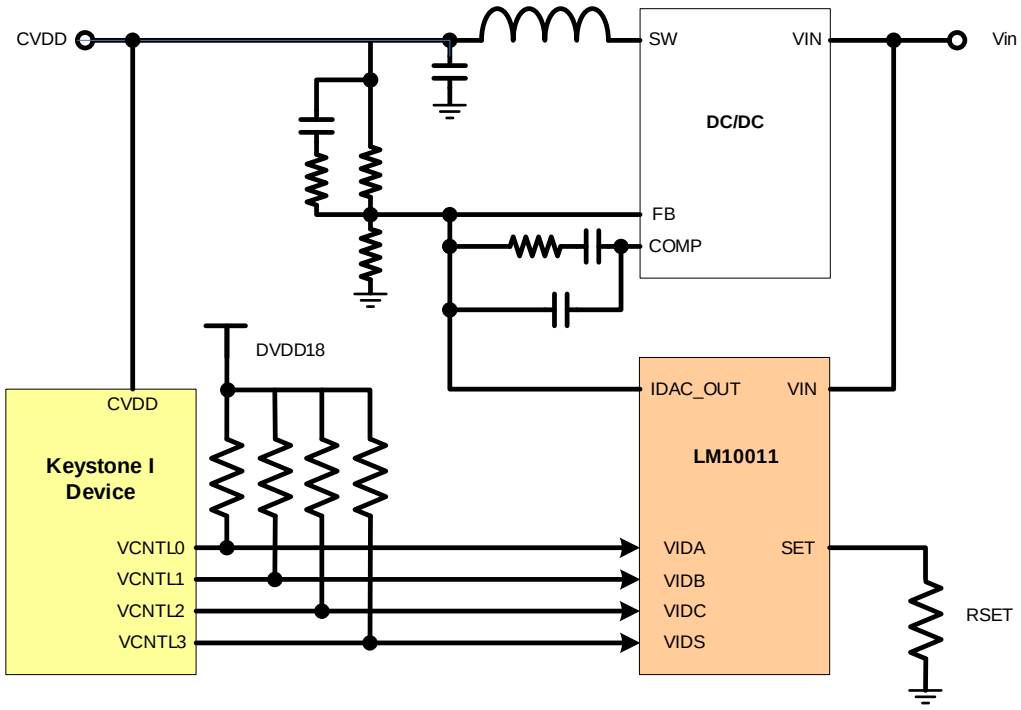

VIDA, VIDB, VIDC, and VIDS are connected to VCNTL0, VCNTL1, VCNTL2, and VCNTL3, respectively. The LM10011 uses a supply voltage in a range from 3 V to 5.5 V, however, the VID inputs are compatible with the 1.8-V logic levels used by the KeyStone I device and no voltage translator is needed. Remember that the VCNTL outputs are open-drain and require a 4.7-kΩ pull-up resistor to the KeyStone I device DVDD18 rail.

2.2.1.5 UCD92xx Digital Controller Solutions

The UCD9222 and the UCD9244 are synchronous buck digital PWM controllers designed for DC/DC power applications. The UCD9222 provides control for two separate voltage rails and the UCD9244 can control four rails. The controller can be programmed to support either an AVS voltage rail controlled by the KeyStone I device or a fixed voltage rail. Because they include multiple VID interfaces that are compatible with the timing for the KeyStone I VCNTL interface, the UCD92xx can be used to control the AVS voltage for multiple KeyStone I devices if more then one is used.

The UCD92xx controller is designed to be paired with UCD7xxx family synchronous buck power drivers. One UCD7xxx is needed for each voltage rail controlled by the UCD92xx device, therefore the UCD9222 requires two UCD7xxx devices and the UCD9244 requires four UCD7xxx devices. Each device in the UCD7xxx family supports different maximum current levels. The power supply solutions link will help you select the best component for your application.

See the respective data sheets, application notes, user guides, and this document for additional details. TI continues to evolve its power supply support for the devices, so check the product folder, updates to this document, and your local FAE for additional details and updates.

The UCD92xx components are highly programmable, allowing you to optimize the operation of your supply using the Fusion Digital Power Designer software. The software may be downloaded from:

http://focus.ti.com/docs/toolsw/folders/print/fusion_digital_power_designer.html

This Windows-based graphical user interface allows the designer to configure the system operating parameters and store the configuration into non-volatile memory. Examples of the XML configuration files for the KeyStone I device EVMs can be found in the technical documents section for the EVM. Note that these files are specific to the EVM design and will require modifications depending upon the components selected. Communication between the Fusion software and the UCD92xx is achieved using a three-wire PMBUS interface. A connector for this interface should be included in your design to allow programming and monitoring of the operation of the UCD92xx. The Fusion software uses a USB-to-GPIO evaluation pod to facilitate this communication. Information for this pod can be found at:

http://www.ti.com/tool/usb-to-gpio

2.2.1.5.1 KeyStone I/UCD92xx Communications

The KeyStone I device communicates the VID value to the UCD92xx device using the VCNTL interface. The VCNTL operates at 1.8 V and the UCD92xx VID interface operates at 3.3 V, requiring a voltage translator. The voltage translator is shown in Figure 3 with VIDA, VIDB, VIDC, and VIDS on the 3.3-V side connected to VCNTL0, VCNTL1, VCNTL2, and VCNTL3 on the 1.8-V side, respectively. The open-drain VCNTL interface requires 4.7-kΩ pull-up resistors be connected to DVDD18.

Figure 3 VCNTL Level Translation

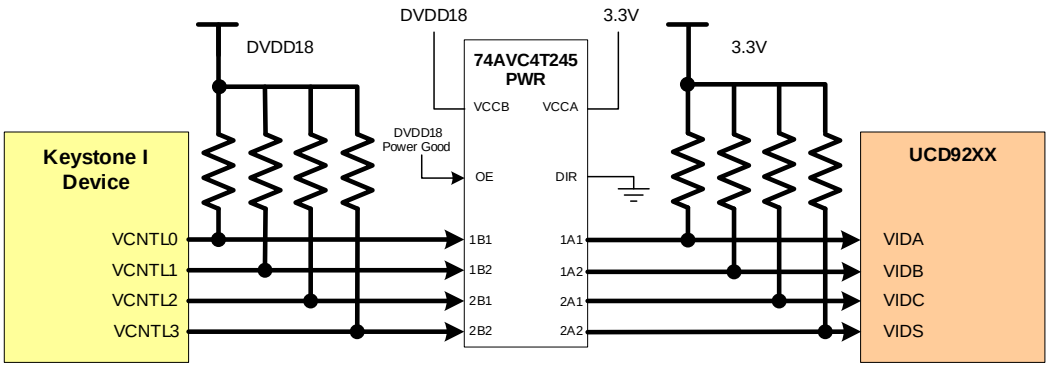

To avoid any glitch on the VID interface, 4.7-kΩ pull-up resistors connected to 3.3 V should be placed on the 3.3-V side of the converter and the converter output enable should be held high until DVDD18 has reached a valid level. If the output enable is not held off until DVDD18 has stabilized, the converter will drive the VID interface low, followed by a transition to high as DVDD18 ramps to 1.8 V. This will cause an invalid voltage code to be latched into the UCD92xx and may prevent the KeyStone I device from booting correctly.

The UCD92xx components can be used to control multiple rails. A rail used for the CVDD supply should be configured to monitor the VID interface. This is done using the Fusion Digital Power Designer software. In the configuration section, there is a VID Config tab. In this tab, select the VID mode labeled 6-Bit VID Code via VID Pins. Set the VID Vout Low to 0.7 V and the VID Vout High to 1.103 V. Note that this defines the full range of VID values shown in Table 1 in this document and not the CVDD voltage range of the KeyStone device as defined in the data manual. These values must be used to ensure that the correct voltage is presented for the VID code received. Set the VID Code Init value to 63, which will provide an initial voltage of 1.103 V. This voltage will be generated until the KeyStone I device presents the required VID value on the VCNTL interface.

2.2.1.6 CVDD for Designs with Multiple KeyStone I Devices

As stated previously, each KeyStone I device must have a separate power supply voltage for CVDD even if there are multiple KeyStone I devices on the same board. If a single device is used, then the LM10011-based solution or the UCD9222 can be used. The UCD9222 is a dual controller but the other digital controller in the device can be used to generate one of the other voltages needed by the KeyStone I device. On the KeyStone I EVM platforms, the second controller is used to generate the CVDD1 1-V fixed voltage. If your design includes multiple KeyStone I devices, you have the option of using any of the recommended power supply designs. A single UCD9222 has two independent VID interfaces and can be used to generate the CVDD voltage for two separate KeyStone I devices. The UCD9244 includes four VID interfaces and can be used for up to four KeyStone I devices.

2.2.2 CVDD1

CVDD1 is the fixed 1-V supply for the internal memory arrays and as a termination voltage for the SerDes interfaces of the KeyStone I device. This supply must meet the requirements for stability specified in the data manual. This supply cannot be connected to the AVS supply used for CVDD and must be generated by a separate power circuit. If multiple KeyStone I devices are used in your design, a single 1-V supply may be used for all of the devices as long as it is scaled to provide the current needed.

2.2.2.1 VDDT*n*

The VDDT pins provide the termination voltage for the SerDes interfaces in the KeyStone I devices. The VDDT pins are grouped under separate names VDDT1 to VDDTx. The number of groups will vary from device to device depending on the number of SerDes interfaces supported. Each group of pins provides the termination voltage for one or more SerDes interfaces. Each group of pins should be connected to CVDD1 through a filter circuit. A separate filter should be used for each group of pins. For details on the filter, see the Power Supply Filters section later in this section. Note that the power estimate generated by the power consumption model includes the power for the VDDT pins in the total for CVDD1.

2.2.3 DVDD18

DVDD18 is the 1.8-V supply for the LVCMOS buffers and for the PLLs. This supply must meet the requirements for stability specified in the data manual. If multiple KeyStone I devices are used in your design, a single 1.8-V supply may be used for all of the devices as long as it is scaled to provide the current needed.

2.2.3.1 AVDDA*n*

The AVDDA pins provide the supply voltage for the PLL modules in the KeyStone I device. The number of AVDDA pins will vary from device to device depending on the number of PLLs present. Each AVDDA pin should be connected to DVDD18 through a filter circuit. A separate filter should be used for each pin. For details on the filter, see 2.3 "Power Supply Filters" on page 16. Note that the power estimate generated by the power consumption model includes the power for the AVDDA pins in the total for DVDD18. 50 mA can be used as a current limit for each AVDDA pin when selecting filter components. Recommendations for the filter can be found later in this document.

2.2.4 DVDD15

DVDD15 is the 1.5-V supply for the DDR3 IO buffers and for the SerDes regulation in the KeyStone I device. This supply must meet the requirements for stability specified in the data manual. If multiple KeyStone I devices are used in your design, a single 1.5-V supply may be used for all of the devices as long as it is scaled to provide the current needed.

2.2.4.1 VDDR*n*

The VDDR pins provide the supply voltage for the SerDes Regulator in the KeyStone I device. The number of VDDR pins will vary from device to device depending on the number of SerDes interfaces present. Each VDDR pin should be connected to DVDD15 through a filter circuit. A separate filter should be used for each pin. For details on the filter, see 2.3 "Power Supply Filters" on page 16. Note that the power estimate generated by the power consumption model include the power for the VDDR pins in the total for DVDD15. 50mA can be used as a current limit for each VDDR pin when selecting filter components. Recommendations for the filter can be found later in this document.

2.2.5 Other Supplies

The four supplies listed above provide the majority of the power required for KeyStone I devices, however, your design may require additional supplies. Information for these supplies is provided here to advise the designer of their necessity.

2.2.5.1 VTT Termination Supply

The DDR3 interface requires a VTT termination at the end of the flyby chain for the DDR3 address and control signals. The 0.75-V termination voltage is generated using a special push/pull termination regulator specifically designed to meet the VTT requirements for DDR3. The Texas Instruments TPS51200 is one device that meets these requirements. Although the VTT supply is not connected to the KeyStone I devices it is necessary for the DDR3 interface connected to the device.

2.2.5.2 VREF Supply

The KeyStone I device requires a 0.75-V reference voltage connected to the VREFSSTL pin for the DDR3 SDRAM interface. Many VTT termination supplies will provide a separate supply output pin specifically for a reference voltage. This is true for the TPS51200, which provides the reference voltage on the REFOUT pin. If the reference

voltage is provided by the VTT termination supply, it should be connected to the VREFSSTL pin. If it is not supplied, a simple voltage divider as shown in Figure 4 can be used.

Figure 4 DDR3 VREFSSTL Voltage Divider

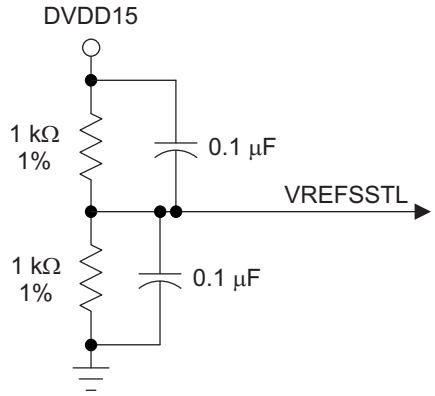

All resistive components must be 1% or better tolerance. The VREFSSTL voltage created by the voltage divider should be routed to the KeyStone I device and the memory components using a trace width of 20 mil or greater. Additional bypass capacitors should be placed next to the connections at the KeyStone I device and the memories.

If a resistor divider is used, the VREFSSTL source voltage MUST be generated from the DVDD15. Never connect the VREFSSTL to the VTT voltage rail.

2.2.5.3 VPP

Some KeyStone I devices include a VPP voltage rail. For unsecure devices, the VPP pin should be left unconnected. See the *Security Addendum for KeyStone I Devices* for information on connecting VPP in designs using secure devices.

2.3 Power Supply Filters

The following section defines the recommended power plane filters for the KeyStone I devices as well as the implementation details and requirements.

Filters are recommended for AVDDA*n*, VDDR*n*, and VDDT*n* KeyStone I device voltage rails. An overview of the recommended power supply generation architecture is shown in Figure 5 (only AVDDAn supply rails are shown).

2.3.1 Filters Versus Ferrite Beads

In previous devices, EMI filters (commonly referred to as *T-filters*) have been used to condition power rails that may be susceptible to induced noise. These filters acted as low-pass, band-pass, or high-pass filters (depending on selection), limiting parasitic coupled noise on each respective power plane/supply. These EMI filters solve many problems typically overlooked by designers, which may result in device-related problems on production boards.

Ferrite beads also were evaluated based on performance, physical form factor, and cost. Ferrite beads did offer a similar benefit at a lower cost point and therefore are required by TI for the PLL supplies. See Table 2 for recommended ferrite beads. Key to selecting the EMI filter and ferrite beads are cutoff frequency and current rating.

Table 2 Power Rail Filter Recommendations

Power Rail	Voltage	Filter (p/n)	Mfg	Current	Insertion Loss (dB)/ Impedance	Capacitance/DCR
AVDDA*n*	1.8	BLM18HE601SN1D	Murata	200 mA	470 Ω or greater @ 100 MHz**	250 mΩ or less
VDDR*n**	1.5	BLM18HE601SN1D	Murata	200 mA	470 Ω or greater @ 100 MHz**	250 mΩ or less
VDDT*n**	1.0	NFM18CC223R1C3	Murata	1000 mA	65 dB @ ~140 MHz	22,000 pF
End of Table 2						

Note *: Filter not required if peripheral not used.
Note **: Requires the addition of a 100-nF and 10-nF ceramic capacitor between the ferrite and respective device pins. Routing must be kept short and clean.

Table 2 shows the characteristics necessary for selecting the proper T-filter. The following specifications must be adhered to for proper functionality. See the power supply bulk and decoupling capacitor section for additional details necessary to design an acceptable power supply system.

2.3.2 AVDDA*n* Filters

The AVDDA*n* pins provide the power for the PLL modules in the KeyStone I device. The number of AVDDA pins on the device will equal the number of PLL modules and each is a power source for an individual PLL. If a PLL is used, the associated AVDDA pin should be connected to the fixed 1.8-V supply through a ferrite bead with a 10-nF and a 100-nF decoupling capacitor placed between the ferrite and the KeyStone I AVDDA pin as shown in Figure 5. If the PLL is not used by the KeyStone I device, the associated AVDDA pin can be connected directly to the fixed 1.8-V supply without the ferrite or the associated decoupling capacitors.

Figure 5 AVDDAx Power Supply Filter

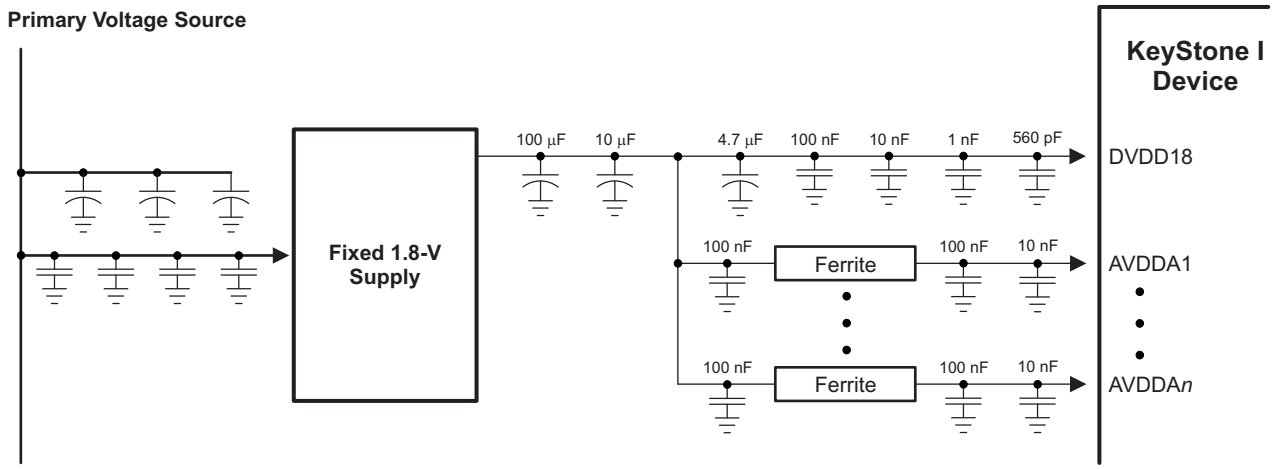

2.3.3 VDDR*n* Filters

The VDDR*n* pins on the device provide the power for the PLLs within the SerDes controllers in the KeyStone I device. The number of VDDR*n* pins on the device equal the number of SerDes interfaces and each is a power source for an individual SerDes PLL. If a SerDes interface is used, the associated VDDR*n* pin should be connected to the fixed 1.5-V supply through a ferrite bead with a 10-nF and a 100-nF decoupling capacitor placed between the ferrite bead and the Keystone I VDDR*n* pin as shown in Figure 6. If the PLL is not used by the KeyStone I device, the associated VDDR*n* pin can be connected directly to the fixed 1.5-V supply without the ferrite bead or the associated decoupling capacitors.

Figure 6 VDDR*n* Power Supply Filter

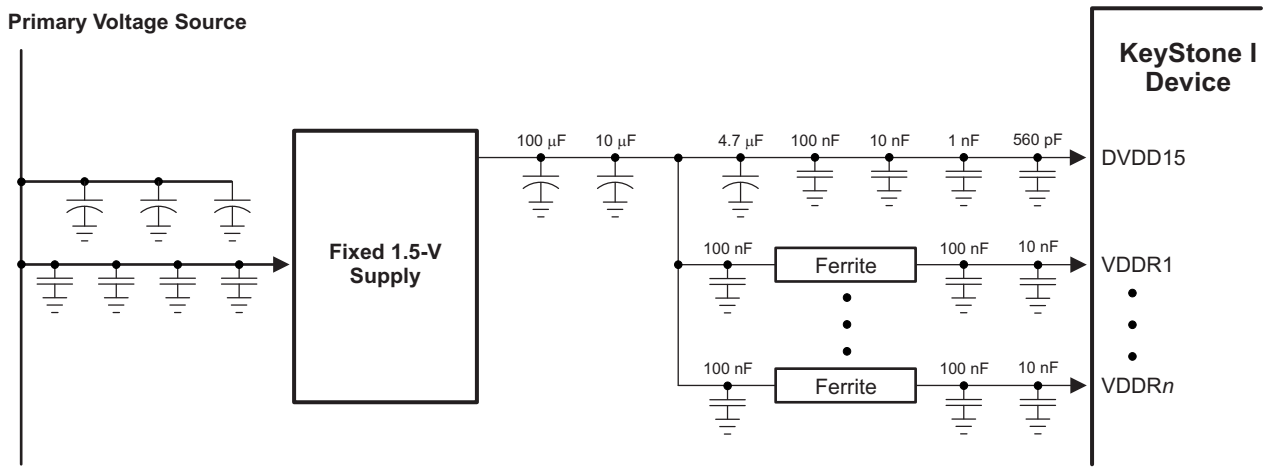

2.3.4 VDDT*n* Filters

The VDDT*n* pins on the device provide the analog termination voltage for the SerDes interfaces in the KeyStone I device. The VDDT*n* pins are grouped into two or more VDDT*n* groups. The number of groups varies from device to device. Each group of pins provides the termination voltage for one or more SerDes interfaces. If any of the SerDes interfaces in a group are implemented then the VDDT*n* group should be connected to the fixed 1.0-V supply through a capacitive filter. The number of bypass capacitors between the filter and the VDDT*n* pins will vary depending on the number of pins on the KeyStone I part. If none of the SerDes in a group are used then the VDDT*n* pins can be connected directly to the DVDD15 supply.

Figure 7 VDDT*n* Power Supply Filter

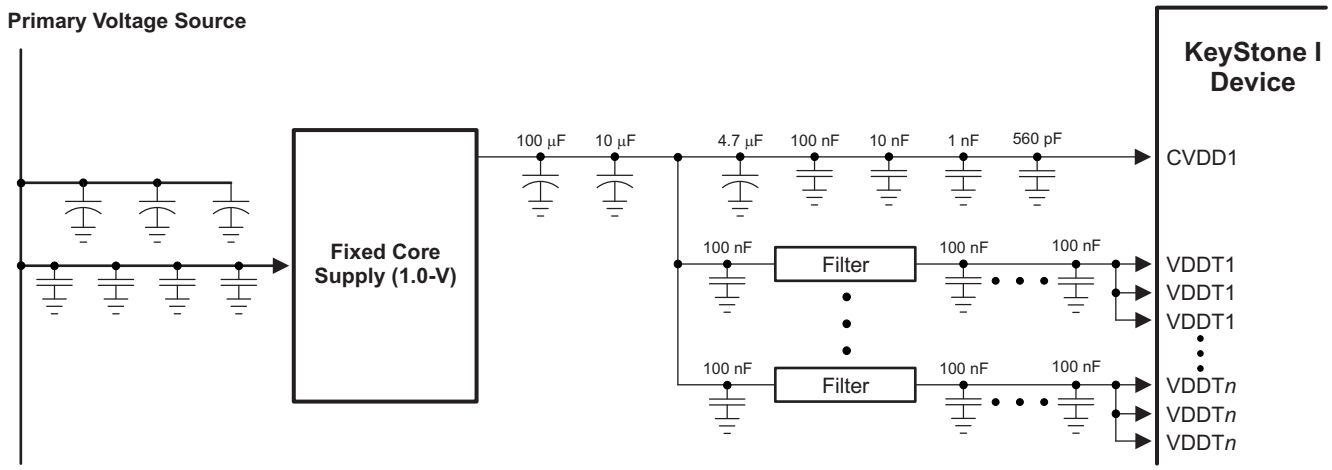

2.4 Power-Supply Decoupling and Bulk Capacitors

To properly decouple the supply planes from system noise, decoupling and bulk capacitors are required. Component parasitics play an important role on how well noise is decoupled. For this reason, bulk capacitors should be of the low ESR/ESL type and all decoupling capacitors (unless otherwise specified) should be ceramic and 0402 size (largest recommended). As always, power rating and component and PCB parasitics must be taken into account when developing a successful power supply system. Proper board design and layout allow for correct placement of all capacitors (see the reference table in Section 2.4.4 for a minimum set of capacitor recommendations and values).

Bulk capacitors are used to minimize the effects of low-frequency current transients (see Section 2.4.1) and decoupling or bypass capacitors are used to minimize higher frequency noise (see Section 2.4.3). Proper printed circuit board design is required to assure functionality and performance.

One key element to consider during the circuit board (target) design is added lead inductance, or the pad-to-plane length. Where possible, attachment for decoupling and bypass capacitors to the respective power plane should be made using multiple vias in each pad that connects the pad to the respective plane. The inductance of the via connect can eliminate the effectiveness of the capacitor so proper via connections are important. Trace length from the pad to the via should be no more than 10 mils (0.01") or 0.254 mm, and the width of the trace should be the same width as the pad.

As with selection of any component, verification of capacitor availability over the product's production lifetime should be considered. In addition, the effects of the intended operating environment (temperature, humidity, etc.) should also be considered when selecting the appropriate decoupling and bulk capacitors.

Note—All values and recommendations are based on a single KeyStone I device, TI high performance SWIFT power supplies, and the New UCD9222/44 digital controller. The use of alternate, non-specified on-board power supply modules, alternate power supplies, and alternate decoupling/bulk capacitor values and configurations require additional evaluation.

2.4.1 Selecting Bulk Capacitance

This section defines the key considerations necessary to select the appropriate bulk capacitors for each rail:

- Effective impedance for the power plane to stay within the voltage tolerance
- Amount of capacitance needed to provide power during the entire period when the voltage regulator cannot respond (sometimes referred to as the transient period)

The effective impedance of the core power plane is determined by:

(Allowable Voltage Deviation due to Current Transients) / (Max Current)

Calculating for the variable core supply (*as an example only*), the allowable deviation (transient tolerance) is 25 mV (based on 2.5% AC [ripple] and 2.5% DC tolerance [voltage]) of the nominal 1-V rail. Using an allowable 25-mV tolerance and a maximum transient current of 10 amps, the maximum allowable impedance can be calculated as follows:

25 mV/10 Amps = 2.5 mΩ

The effective ESR for all bulk capacitors (in the above example) should not exceed this impedance value. The combination of good-quality and multiple bulk capacitors in parallel help to achieve the overall ESR required. Therefore, to limit the maximum transient voltage to a peak deviation of 25 mV, the power supply output impedance, which is a function of the power supply bandwidth and the low impedance output capacitance, should not exceed 2.5 mΩ.

The following three plots show the recommended maximum impedance to current transient for three different tolerances (15 mV, 25 mV, and 35 mV). The first plot covers an allowable current between 500 mA and 3500 mA. The second plot covers 3700 mA to 6900 mA, and the final plot covers 7000 mA to 13000 mA.

The expected maximum current for the KeyStone I device (core) will depend heavily on the use case and design of the application hardware. The recommended components are designed to support up to 17 A (at 86% efficiency) and up to 19 A when using the recommended components and implementing them correctly at 86% efficiency. It does not include current transients that occur during power on (yet these must be considered). Given the recommended topology, only a reduced amount of bulk capacitors is possible, provided they are chosen for there ESR/ESL and voltage rating properly. The majority of all core bulk capacitors must be located in close proximity to the UCD7242 dual FETs.

Capacitance values should not be less than those specified in the applicable table in Section 2.4.4. Final capacitor selection is determined using the provided capacitor selection tool (see spreadsheet and application note).

Some intermediate size ceramic bulk capacitors (i.e., 22 µF and 47 µF as listed under Section 2.4.2) are recommended to cover the response time between the bypass capacitors and the larger bulk capacitors.

Figure 8 ESR Plot for Delta Current to Tolerance #1

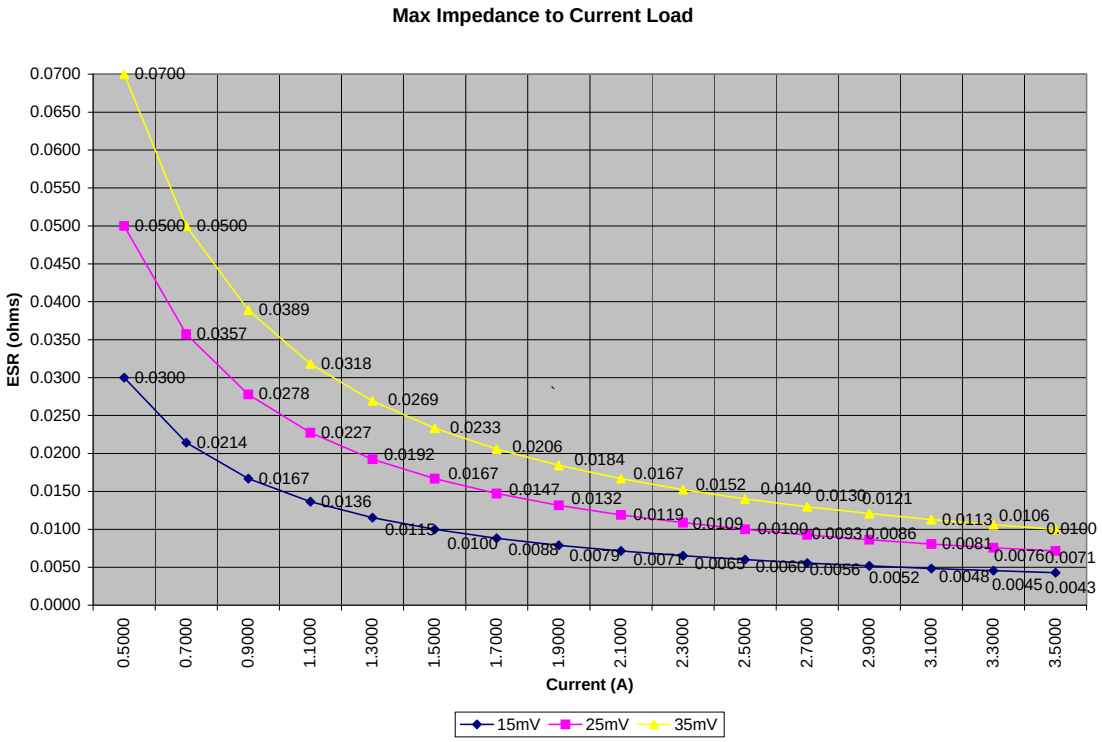

Figure 9 ESR Plot for Delta Current to Tolerance #2

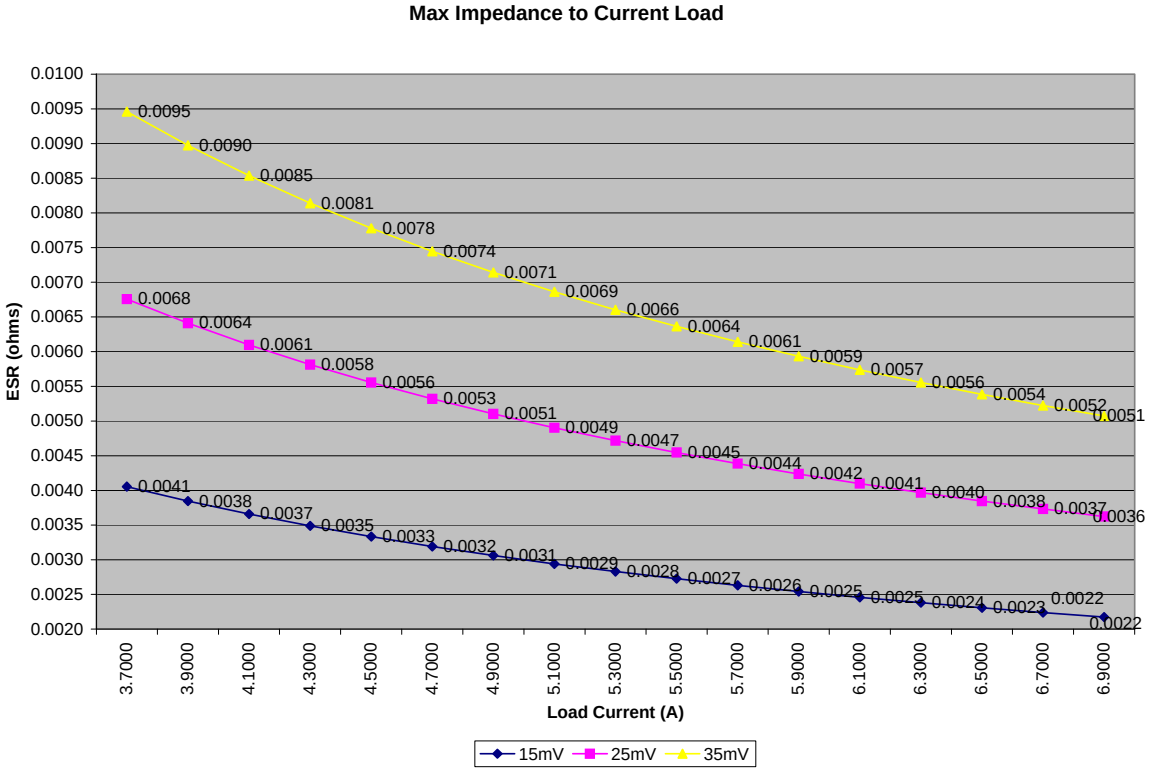

Figure 10 ESR Plot for Delta Current to Tolerance #3

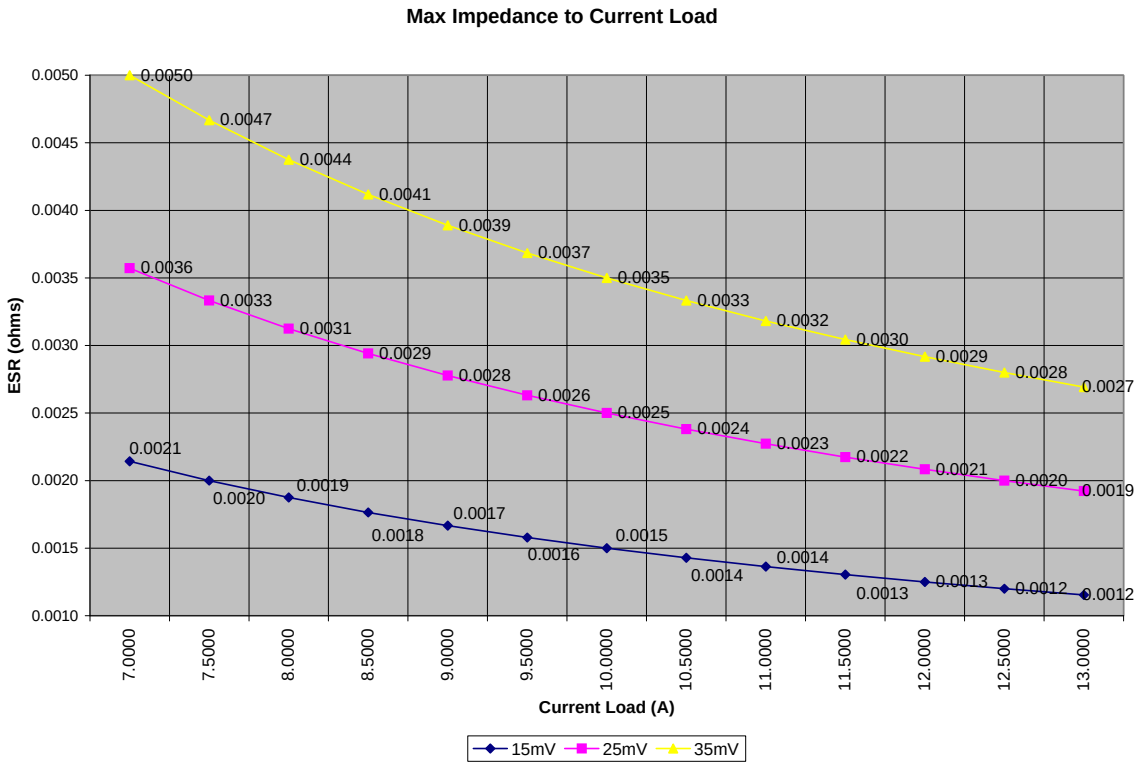

2.4.1.1 Bulk Capacitor Details and Placement

Placement of the bulk capacitors is highly dependent on the power supply solution selected. It is important to review the application documents associated with the power supply selected to determine the best placement of the bulk capacitors for your design. Generally, the bulk capacitors are placed in close proximity to the power supply. For the purpose of this document and related devices, a bulk capacitor is defined as any capacitor ≥ 22 µF, unless otherwise noted.

2.4.2 Selecting Intermediate Value Capacitors

All intermediate capacitors must be positioned as close to the device as possible, decoupling capacitors taking precedence where placement between the two form factors and values come into question. Intermediate capacitor values are considered to be those between 22 µF and 47 µF, and they should be low ESR and ceramic, where possible.

2.4.3 Selecting Decoupling Capacitors

All decoupling or bypass capacitors must be positioned close to the device. In practice, each decoupling capacitor should be placed a maximum distance of 10 mm from the respective pin(s). Where possible, each decoupling capacitor should be tied directly between the respective device pin and ground without the use of traces. Any parasitic inductance limits the effectiveness of the decoupling capacitors; therefore, the physically smaller the capacitors are (0402 or 0201 are recommended), the better the power supply will function.

Proper capacitance values are also important. All small decoupling or bypass capacitors (560 pF, 0.001 µF, 0.01 µF, and 0.1 µF) must be placed closest to the respective power pins on the target device. Medium bypass capacitors (100 nF or as large as can be obtained in a small package such as an 0402) should be the next closest. TI recommends placing decoupling capacitors immediately next to the BGA vias, using the *interior* BGA space and at least the corners of the *exterior*.

The inductance of the via connect can eliminate the effectiveness of the capacitor so proper via connections are important. Trace length from the pad to the via should be no more than 10 mils and the width of the trace should be the same width as the pad. If necessary, placing decoupling capacitors on the back side of the board is acceptable provided the placement and attachment is designed correctly.

2.4.3.1 Decoupling Capacitor Details and Placement

All decoupling capacitors should be placed in close proximity to the device power pins. For the purpose of this document and related devices, a decoupling or bypass capacitor shall be defined as any capacitor < 1 µF, unless otherwise noted.

Table 3 shows minimum recommended decoupling capacitor values. Each decoupling or bypass capacitor should be directly coupled to the device via. Where direct coupling is impractical, use a trace 10 mil (0.010"/0.254 mm) or shorter, and having the same width as the capacitor pad, is strongly recommended. The table values are minimums; many variables will have an effect on the performance of the power supply system. Re-evaluate the individual power supply design for all variables before implementing the final design.

2.4.4 Example Capacitance

Table 3 establishes the bulk and decoupling capacitance requirements for the various device voltage rails. In many cases these rails originate from a common source, which accounts for a reduction in bulk capacitance for many rails. This table does not cover the additional capacitance (decoupling or bulk) required by any other components or the filters outlined in previous sections. All capacitance identified assumes the use of the recommended power supplies (especially for the variable core), and switching frequencies as indicated (see the note section). The recommended values are in addition to those required for proper operation of the recommended power supplies.

The values identified in the following table *are per device* and assume a typical loading per device IO. In the event multiple devices are used and a voltage rail is common or shared (where applicable), the capacitor requirements must be re-evaluated (on a case-by-case basis) and take into account key concepts including current (load), minimum and maximum current requirements, voltage step, power supply switching frequency, and component ESR. Other constraints also apply.

Table 3 Bulk, Intermediate, and Bypass Capacitor Recommendations (Absolute Minimums) (Part 1 of 2)

Voltage Supply	Capacitors	Qty	ESR (mΩ)*	Total Capacitance[5]	Description	Notes
CVDD***	560 pF Ceramic	20		1662.9112 µF	Variable Core Supply	1, 5, 6, 12
	100 nF Ceramic	8				5, 6, 12
	10 nF Ceramic	10				5, 6, 12
	22 µF Low ESR	1	9 mΩ*			14
	47 µF Low ESR	4	3 mΩ*			5, 6, 12
	47 µF Low ESR	2	3 mΩ*			14
	220 µF Low ESR	2	7 mΩ*			5, 6, 12
	470 µF Low ESR	2	12 mΩ*			5, 6, 12
DVDD18	560 pF Ceramic	8		125.03748 µF	1.8-V IO Supply	2, 5, 6, 12
	100 nF Ceramic	3				2, 5, 6, 12
	10 nF Ceramic	3				2, 5, 6, 12
	1 nF Ceramic	3				2, 5, 6, 12
	1 µF Ceramic	0				
	4.7 µF Ceramic	1	10 mΩ*			2, 5, 6, 12
	10 µF Ceramic	2	9 mΩ*			2, 5, 6, 12
	47 µF Low ESR	0				
	100 µF Low ESR	1	9 mΩ*			2, 5, 6, 12
	220 µF Low ESR	0				
	10 µF Ceramic	0				9
CVDD1	560 pF Ceramic	10		277.4556 µF	1.0-V Fixed Supply	3, 5, 6, 12
	100 nF Ceramic	4				3, 5, 6, 12
	10 nF Ceramic	5				3, 5, 6, 12
	1 µF Ceramic	0				
	10 µF Ceramic	3	9 mΩs*			14
	47 µF Low ESR	1	3 mΩs*			3, 5, 6, 12
	100 µF Low ESR	2	9 mΩs*			3, 5, 6, 12
DVDD15	560 pF Ceramic	8		5.16448 µF	1.5-V DDR Supply	4, 5, 6, 10, 12
	100 nF Ceramic	4				4, 5, 6, 10, 12
	10 nF Ceramic	6				4, 5, 6, 10, 12
	4.7 µF Ceramic	1	10 mΩs*			4, 5, 6, 10, 12
	10 µF Low ESR	0				
	22 µF Low ESR	0				
	47 µF Low ESR	0				
	100 µF Low ESR	0				
	220 µF Low ESR	0				
VREFSSTL	100 nF Ceramic	2		0.2000 µF	1.5-V DDR V_{REF} Supply	

Table 3 Bulk, Intermediate, and Bypass Capacitor Recommendations (Absolute Minimums) (Part 2 of 2)

Voltage Supply	Capacitors	Qty	ESR (mΩ)*	Total Capacitance[5]	Description	Notes
TOTAL				2281.40528 µF		
End of Table 3						

Notes:
1. Estimates provided assume 10 W @ 1 V for the variable (scalable) core
2. Estimates provided assume 0.75 W @ 1.8 V for the 1.8-V IOs
3. Estimates provided assume 5 W @ 1 V for the fixed core
4. Estimates provided assume 1.0 W @ 1.5 V for the DDR3 IOs (SDRAM not included in estimations)
5. Always denotes a minimum capacitance
6. Fsw is 1e6 Hz.
7. Fsw is 550e3 Hz
8. Not required if common input to rail to supply (bulk is applied to AVDDA1)and distance is short between supply and device pin
9. Not required if adequate bulk capacitance is supplied for CVDD1 common input
10. Additional capacitance required when adding SDRAM (must be common supply per device)
11. To be placed on the input side of the ferrite or filter
12. Total value is calculated on a 1.5-pF device pin capacitance, variations in pin capacitance require re-calculation
13. Not required if adequate bulk capacitance is supplied for DVDD15 common input
14. Required for all devices within the same family having higher current requirements beyond that identified (future devices)
15. Not required if supplied for VDDT1 and is a common input source

Note: ***: capacitor selection for CVDD variable core (as well as all other supplies) is highly dependant upon design and must be evaluated on a case by case scenario. The bulk capacitors shown are with respect to a 470-nH inductor for 10 A with a maximum delta current of 85%, other configurations may apply.

Note: *: refers to each individual capacitor ESR Value (and is a maximum value per capacitor)

Note—The capacitor values provided are a *MINIMUM* recommended amount.

The final capacitor value should not be less than the value specified in Section 2.4.4. Final selection for the core rails are determined using the provided capacitor selection tool and power estimation application guide, which take into account the board impedance, variation in the CVDD supply voltage, use case, and the ESR of the bulk and decoupling capacitors selected.

2.4.4.1 Calculating the Value of Bypass Capacitors

The following is a brief example of how the decoupling (bypass) capacitance was obtained. Final capacitor values should not be less than the values specified in Table 3. Final selection must take into account key variables including PCB board impedance, supply voltage, PCB topologies, component ESR, layout and design, as well as the end-use application, which should include crosstalk and overall amount of AC ripple allowed.

1. Assume all gates are switching at the same time (simultaneously) in the system, identify the maximum expected step change in power supply current ΔI.
2. Estimate the maximum amount of noise that the system can tolerate.
3. Divide the noise voltage by the current change results in the maximum common path impedance or:

$$Z_{\max} = V_n / \Delta_I$$

4. Compute the inductance (L_{PSW}) of the power supply wiring. Use this and the maximum common path impedance ($Z_{\max}$) to determine the frequency,

$f_{PSW} = f_{corner} = f_{3db} = Z_{max} / (2\Pi L_{PSW})$ below which the power supply wiring is fine (power supply wiring noise <V_n.

5. Calculate the capacitance (C_{bypass}) of the bypass capacitor:

$$C_{bypass} = 1/2\Pi * f_{PSW} * Z_{max}$$

Note—If the operating frequency $f_{oper} < f_{PSW}$, then no bypass capacitor(s) are necessary. If $f_{oper} \geq f_{PSW}$ then bypass or decoupling capacitors are needed. The following example is provided as a sample calculation.

Example: (Decoupling Capacitance)

Table 4 provides an example of the capacitance needed for each of the power supply inputs of the TMS320C6670. The decoupling needed for your application must be calculated based on the KeyStone I device that you have selected and the application that you are running. Examples for nominal usage for each KeyStone I device may be found in the schematic for the EVM.

Table 4 Total Decoupling Capacitance

Rail	Voltage	Pins	nF	Notes
CVDD	1	81	773.490	SmartReflex rail
AVDDA1 to AVDDA*n*	1.8	N	110.00 each	
DVDD18	1.8	18	14.3239	
VDDT1 to VDDT*n*	1	X		
CVDD1	1	43	76.3944	Fixed 1-V rail
VDDR1 to VDDR*n*	1.5	N	110.00 each	
DVDD15	1.5	31	309.397	
VrefSSTL	0.75	1	17.9049	
End of Table 4				

Assumptions (Core):

- 81 core power pins
- 2.5% ripple allowed
- Supply voltage is 1.0 V
- Each pin has a capacitive load of 1.5 pF
- Power supply inductance is 100 nH (approx. from PCB stack up)
- Power supply switching frequency is 1000 kHz

1. Calculate the current: $\Delta I = nC * Vcc/\Delta t \rightarrow \Delta I = (81)(1.5^{-12})(1.0)/1^{-9} \rightarrow$

$\Delta I = (1.066^{-10})/1^{-9} \rightarrow \Delta I = 0.1215$

2. Maximum impedance; $Z_{max} = 25mV/0.1215 \rightarrow Z_{max} = 0.205761316872\Omega$

3. Insert the power supply switching frequency: $f_{PSW} = 1000^3$

4. Calculate the bypass value(s) needed:

$$C_{bypass(\min)} = 1/(2 * \Pi * f_{PSW} * Z_{\max}) \rightarrow$$

$$C_{bypass\ (\min)} = 1/(2 * pi * 1e6 * 0.2057613168\ \ 7242) \rightarrow$$

$$C_{bypass(\min)} = 1/1292836.483 \rightarrow C_{bypass(\min)} = 7.7349302^{-7} \text{ or } 0.77349302\ \mu F$$

In theory, decoupling capacitors should be placed on all pins, or at a minimum, on pins that share a common via to the power rail.

2.5 Power Supply Sequencing

A detailed representation of the power supply sequencing for KeyStone I devices can be found in the device-specific data manual. All power supplies associated with the KeyStone I devices should allow for controlled sequencing using on-board logic or a power supply sequence controller such as the UCD9090. Two sequencing orders are supported with the KeyStone I devices.

The first is a core voltage before IO voltage sequence with the power supplies enabled in the following order:

CVDD -> CVDD1 -> DVDD18 -> DVDD15

The second sequencing is an 1.8-V IO before core voltage sequence with the power supplies enabled in the following order:

DVDD18 -> CVDD -> CVDD1 -> DVDD15

A power supply should reach a valid voltage level and declare the supply output in a power good state before enabling the next supply in the sequence.

2.6 Power Supply Block Diagram

Now that the power supply components have been selected, create a power supply block diagram similar to the one found in the EVM schematic for the KeyStone I devices. This block diagram should include all of the power supplies in the design, the output voltage associated with each, the current limit, and the signal used to enable the supply. A second diagram showing the sequencing of the supplies should also be generated. These documents will be requested by Texas Instruments if support is needed. They should be added to the design folder and it is good practice to include them in the schematics for reference.

3 Clocking

The next requirement the KeyStone I design is proper clocking for the needs of your design. KeyStone I devices include a number of internal PLLs that are used to generate the clocking within the part. These include PLLs for generating system clocks and PLLs for the high speed SerDes interfaces. All clocks must meet the described stability and jitter requirements for successful operation of the design. In addition, there are a number of clocks associated with other logic within the part. This section describes the clocks found in KeyStone I devices and the requirements for these clocks. Detailed clock input requirements are found in the device-specific data manual. Not all of the clocks described in this section will appear on every KeyStone I device. See the device-specific data manual to determine which clocks are present in your KeyStone I device.

3.1 System PLL Clock Inputs

The system PLL clock inputs for KeyStone I devices are shown in Table 6.

Table 5 KeyStone I System PLL Clock Inputs

Clock	Frequency Range (MHz)	Drivers Supported	Description
CORECLKp CORECLKn	40-312.5	LVDS, LVPECL (AC Coupled)	Used to clock the Core PLL
ALTCORECLKp ALTCORECLKn	40-312.5	LVDS, LVPECL (AC Coupled)	Used to clock the Core PLL if CORECLKSEL = 1
DDRCLKp DDRCLKn	40-312.5	LVDS, LVPECL (AC Coupled)	Used to clock the DDR PLL
PASSCLKp PASSCLKn	40-312.5	LVDS, LVPECL (AC Coupled)	Used to clock the packet accelerator subsystem PLL if PASSCLKSEL = 1
End of Table 5			

All clock inputs are differential and must be driven by one of the specified differential driver types. All differential clock inputs are implemented with Texas Instruments low jitter clock buffers (LJCBs). These input buffers include a 100-Ω parallel termination (P to N) and common mode biasing. Because the common mode biasing is included, the clock source must be AC coupled. Low voltage differential swing LVDS and LVPECL clock sources are compatible with the LJCBs. Clock drivers may only source one clock input so a separate clock driver must be provided for each clock input. In addition the proper termination for the clock driver selected must be included. For additional information on AC termination schemes, see the *AC-Coupling Between Differential LVPECL, LVDS, and CML Application Report* (SCAA059). Note that the LJCB clock input is assumed to be a CML input in this document.

These clock inputs are used by system PLLs to provide the internal clocks needed to operate the device. These PLLs are compatible with a wide range of clock input frequencies, which can be multiplied to the valid operating frequencies of the device. Detailed information on programming the PLLs can be found in the *Phase-Locked Loop (PLL) for KeyStone Devices Users Guide*.

The Main PLL is driven in one of two manners depending on the device. It is driven by CORECLKP/N if an exclusive clock input is present on the device. Some KeyStone I devices allow one of two clock inputs to be used as a source for the Main PLL. For these components the clock is provided by either the ALTCORECLKP/N or the SYSCLKP/N depending on the state of the CORECLKSEL configuration input. SYSCLKP/N is used as a clock source for the AIF interface and is considered a SerDes PLL clock source.

Table 6 KeyStone I System PLL Clock Inputs

Clock	Frequency Range (MHz)	Drivers Supported	Description
CORECLKp CORECLKn	40-312.5	LVDS, LVPECL (AC Coupled)	Used to clock the Core PLL
ALTCORECLKp ALTCORECLKn	40-312.5	LVDS, LVPECL (AC Coupled)	Used to clock the Core PLL if CORECLKSEL = 1
DDRCLKp DDRCLKn	40-312.5	LVDS, LVPECL (AC Coupled)	Used to clock the DDR PLL
PASSCLKp PASSCLKn	40-312.5	LVDS, LVPECL (AC Coupled)	Used to clock the packet accelerator subsystem PLL if PASSCLKSEL = 1
End of Table 6			

3.1.1 System Reference Clock Jitter Requirements

Table 7 shows the specific input clocking requirements for the KeyStone I device, including maximum input jitter, rise and fall times, and duty cycle. When discrepancies occur between the data manual and this application note, always use the *latest* device-specific data manual.

Table 7 KeyStone I System PLL Clocking Requirements - Jitter, Duty Cycle, t_r / t_f

Clock	Input Jitter[4]	t_{rise} / t_{fall}[2]	Duty Cycle %	Stability
CORECLKp CORECLKn	2.0% of CORECLK input period (pk-pk)[3]	50 – 350 ps[1]	45 / 55	± 100 PPM
ALTCORECLKp ALTCORECLKn	2.0% of ALTCORECLK input period (pk-pk)[3]	50 – 350 ps[1]	45 / 55	± 100 PPM
DDRCLKp DDRCLKn	2.0% of DDRCLK input period (pk-pk)[3]	50 – 350 ps[1]	45 / 55	± 100 PPM
PASSCLKp PASSCLKn	2.0% of PASSCLK input period (pk-pk)[3]	50 – 350 ps[1]	45 / 55	± 100 PPM
End of Table 7				

Notes for Table 7 Note 1: Swing is rated for a 250 mV (pk-pk) at zero crossing where 10% – 90% of t_{rise} and t_{fall} must occur within the prescribed 50-ps to 350-ps time frame. The minimum slew specified is relative to the minimum input signal value. An input signal must transition through:

Minimum

$$Vtransition = 250mV * (90\% - 10\%)$$
$$Vtransition = 250mV * 80\%$$
$$Vtransition = 200mVin \xrightarrow{in} 350\,pS$$

and

Maximum

$$Vtransition = 250mV * (90\% - 10\%)$$
$$Vtransition = 250mV * 20\%$$
$$Vtransition = 200mVin \xrightarrow{in} 50\,pS$$

Note 2: t_{RISE}/t_{FALL} values are given for 10% to 90% of the voltage swing.

Note 3: Example: 312.5MHz (~3.2 ns) * 2.0% = 64 ps pk-pk input jitter.

3.2 SerDes PLL Reference Clock Inputs

The Serializer/Deserializer (SerDes) PLL reference clock inputs for KeyStone I devices are shown in Table 8.

Table 8 KeyStone I SerDes PLL Reference Clock Inputs

Clock	Valid Input Frequencies (MHZ)	Drivers Supported	Description
SYSCLKp SYSCLKn	122.88, 153.60, 307.20	LVDS, LVPECL (AC Coupled)	Used to clock the AIF PLL. Used to clock the Main PLL if CORECLKSEL= 0.
SRIOSGMIICLKp SRIOSGMIICLKn	156.25, 250, 312.5	LVDS, LVPECL (AC Coupled)	Used to clock the SGMII PLL and the SRIO PLL.
PCIECLKp PCIECLKn	100, 156.25, 250, 312.5	LVDS, HCSL LVPECL (AC Coupled)	Used to clock the PCIexpress PLL.
MCMCLKp MCMCLKn	156.25, 250, 312.5	LVDS, LVPECL (AC Coupled)	Used to clock the HyperLink PLL.
End of Table 8			

Not all clock inputs are present on all KeyStone devices. Check the device-specific data manual for a complete list of clock inputs. All clock inputs are differential and must be driven by one of the specified differential driver types. All differential clock inputs are implemented with Texas Instruments low jitter clock buffers (LJCBs). These input buffers include a 100-Ω parallel termination (P to N) and common-mode biasing. Because the common mode biasing is included, the clock source must be AC coupled. Low voltage differential swing LVDS and LVPECL clock sources are compatible with the LJCBs. Clock drivers may source only one clock input, so a separate clock driver must be provided for each clock input. In addition, the proper termination for the clock driver selected must be included. For additional information on AC termination schemes, see the *AC-Coupling Between Differential LVPECL, LVDS, and CML Application Report* (SCAA059). Note that the LJCB clock input is assumed to be a CML input in this document.

In addition, the LJCB will support a PCI express-compliant HCSL clock input. This clock is terminated in the same manner as an LVDS clock driver. Spread spectrum clocks are commonly used with PCI express and a may be present on a PCI express bus connector. The PCIECLKP/N input is compatible with PCI express-compliant spread-spectrum clocking. Note that a common clock for a PCI express root complex and endpoint are not required unless a spread-spectrum clock is used. If a spread-spectrum clock is used by one device at the far end of a PCI express link, the same clock must be used as a reference by the KeyStone I device.

These clock inputs are used by PLLs in the reference SerDes interface. These interfaces require high-quality clocks with low phase noise and are specified only to operate with specific reference clock input frequencies. Detailed information on programming the PLL associated with a reference clock input can be found in the user guide for that SerDes interface.

3.2.1 SerDes Reference Clock Jitter Requirements

KeyStone I devices contain several high-performance serial interfaces commonly known as SerDes interfaces. The SerDes architecture allows for reliable data transmission without a need for common synchronized clocks. However, the architecture does require high-quality clock sources that have very little phase noise. Phase noise is also commonly referred to as clock jitter.

Note that PCIe can be operated with a common clock and this allows the use of spread-spectrum clock sources that have higher levels of bounded phase noise. HyperLink is also defined as an interface requiring a common clock because it is expected to be a very short-reach interface between DSPs located on the same board.

3.2.1.1 Random Jitter

Phase noise amplitude is not the only concern. The frequency content of this phase noise is also significant. Therefore, masks are provided as a means of indicating an acceptable phase noise tolerance. Each SerDes interface has a slightly different mask. SerDes data rate, input clock rate, and required bit error rate affect the phase noise tolerance.

Figure 11 shows the basic phase jitter tolerance masks for the KeyStone I SerDes interfaces. The mask shown is for the highest performance operating mode defined for each SerDes interface. Board designers must provide a reference clock that has clock jitter below the appropriate mask across the entire frequency range.

Figure 11 SerDes Reference Clock Jitter Masks

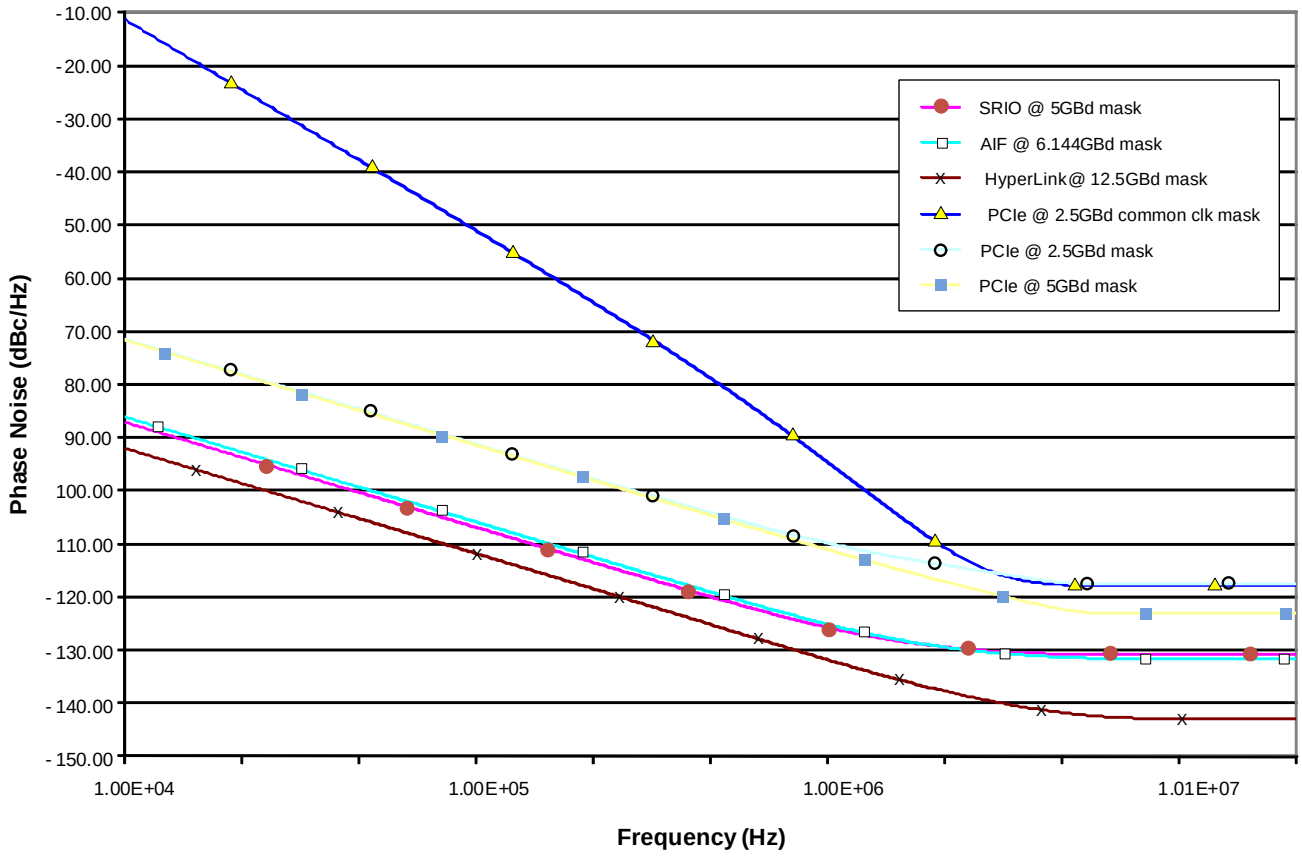

This information is also shown as a straight-line approximation with the end points and the knee defined for ease of understanding. This template is shown in Figure 12 below and the endpoints in Table 9. The mask level at specific frequencies can now be estimated based on the straight-line approximation. This is useful when comparing clock sources from various vendors who specify jitter masks in different ways.

Figure 12 Phase Jitter Template

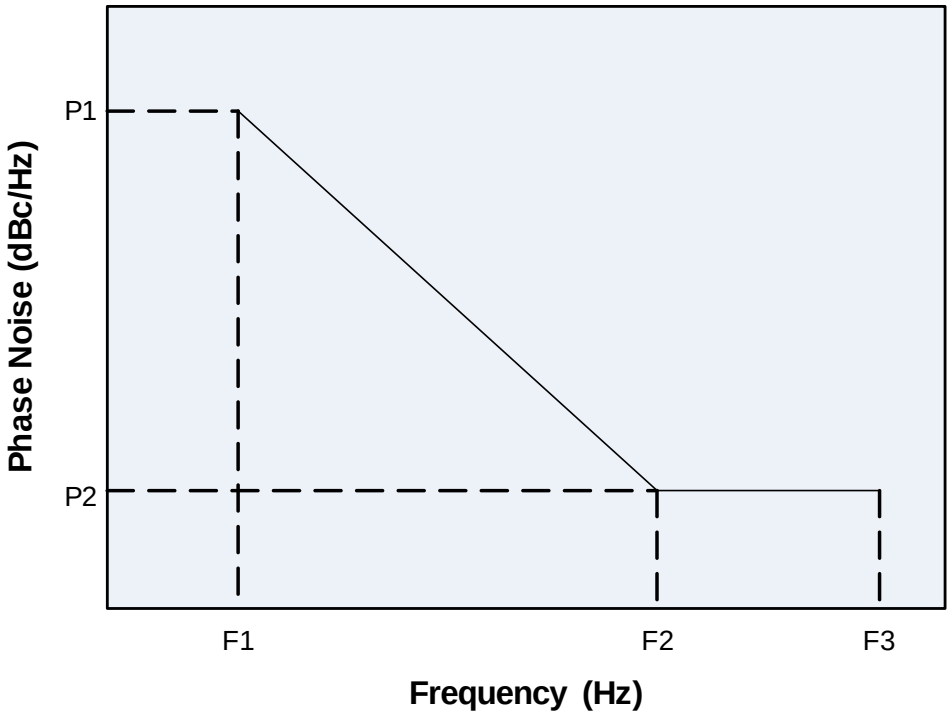

Table 9 Phase Jitter Template Endpoints

Interface & Rate (GBd)	Reference Clock Frequency (MHz)	BER	F1 (kHz)	P1 (dBc/Hz)	F2 (MHz)	P2 (dBc/Hz)	F33 (MHz)
SRIO @ 5	250	1^{-15}	10	-87.1	1.56	-130.8	20
AIF @ 6.144	256	1^{-12}	10	-86.1	1.91	-131.6	20
HyperLink @ 12.5	250	1^{-17}	10	-92.1	3.43	-142.9	20
PCIe @ 2.5	250 (common)	1^{-12}	10	-11.4	2.96	-117.7	20
PCIe @ 2.5	250	1^{-12}	10	-71.5	2.18	-117.7	20
PCIe @ 5	250	1^{-12}	10	-71.7	3.76	-123.1	20
End of Table 9							

A single value is also commonly provided as the limit for phase noise. This is an integration of the phase noise across a frequency range. Table 10 shows the integrated value for each of the phase noise masks across the 10 kHz to 20 MHz frequency range and across the 1 MHz to 20 MHz frequency range. These limits are also useful for evaluating high-performance clocks sources from some vendors.

Table 10 Phase Noise Mask Integrated Values

Interface & Rate (GBd)	Reference Clock Frequency (MHz)	Bit Error Rate (BER)	Rj (ps) 10kHz to 20MHz	Rj (ps) 1MHz to 20MHz
SRIO @ 5	250	1^{-15}	4.12	1.17
AIF @ 6.144	256	1^{-12}	4.56	1.09
HyperLink @ 12.5	250	1^{-17}	2.26	0.33
PCIe @ 2.5	250 (common)	1^{-12}	14030	8.99
PCIe @ 2.5	250	1^{-12}	24.59	5.62
PCIe @ 5	250	1^{-12}	23.65	3.37
End of Table 10				

Please note that not all clock sources with jitter values below the single values shown in Table 10 will meet the templates shown above. These figures were generated assuming that the reference clock spectrum is shaped in the same manner as the template. Further verification may be needed.

All of the masks shown in Table 10 assume the input clock rate is about 250 MHz. The input clock rate affects the phase noise tolerance. This is because the reference clock phase noise is multiplied within the SerDes PLL. Higher levels of phase noise can be tolerated if the input reference clock rate is higher. The entire mask shifts proportional to the frequency difference. This mask offset is approximated by the equation 20*log(F2/F1). The integrated phase noise value also shifts by the same ratio. It is approximated by the equation F2/F1. Because all of the masks shown above are for 250 MHz reference clocks, the mask and phase noise adjustments from 250 MHz are shown in Table 11.

Table 11 Mask and Phase Noise Adjustments

Reference Clock Frequency (MHz)	Phase Noise Mask Offset (dB)	Rj Multiplier
100[1]	-8.0	0.40
122.88	-6.2	0.49
153.60	-4.2	0.61
156.25	-4.1	0.63
250	0.0	1.00
312.5	1.9	1.25
End of Table 11		

1. Only supported for PCIe

The phase noise masks provided are at the maximum supported data rates for SRIO, AIF, and HyperLink. Additional phase noise margin is gained when the data rate is reduced. This is because the margin for jitter increases as the data eye gets longer when the data rate decreases. Once again, the mask offset is approximated by the equation 20*log(F2/F1) and the integrated phase noise value is approximated by the equation F2/F1.

Table 12 shows the mask and phase noise adjustments for the available SRIO operating rates. Similar numbers can be computed for AIF and HyperLink. Note that this does not apply for PCIe. The phase noise masks are defined within the PCIe specification and the masks converge at low frequency regardless of data rate as shown in Figure 12.

Table 12 Mask and Phase Noise Adjustments

SerDes Data Rate (GBd)	Phase Noise Mask Offset (dB)	Rj Multiplier
5	0.00	1.00
3.125	4.08	1.60
2.5	6.02	2.00
1.25	12.04	4.00

3.2.1.2 Deterministic Jitter

Reference clock sources may also contain deterministic jitter in the form of spurs that must also be accounted for. The effects of spurs vary depending upon whether they are correlated or uncorrelated. An exact determination of the effects of the spurs would be very complicated. (There are published papers concerning the complexities of combining the noise power of spurs.) A simplified formula is shown below:

$$Jitter_{spurs} = \frac{2}{2 \cdot pi \cdot F_{reference}} \sqrt{\sum_{1}^{N} 10^{\frac{Spurs_N\,(dBc)}{10}}} \leq 0.1ps$$

The limit of 0.1 ps is a level chosen to provide a reasonable margin without complex analysis. This equation has sufficient margin to be used for all interface types and baud rates.

3.3 Unused Clock Inputs

Any unused LJCB or LVDS differential clock inputs should be connected to the appropriate rails to establish a valid logic level.

The recommended connections are shown in Figure 13. The added 1-kΩ resistor is designed to reduce power consumption. The positive terminal should be connected to the respective power rail.

Figure 13 Unused Clock Input Connection

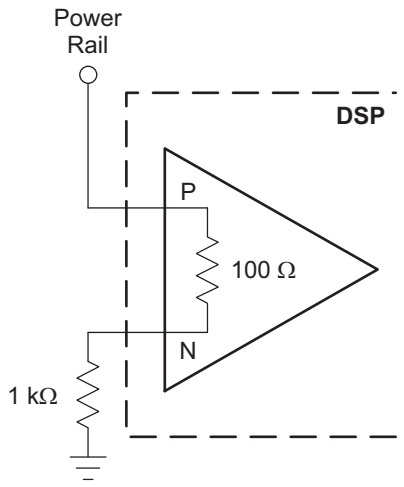

Table 13 shows the specific rails to which each unused clock input should be connected (in accordance with Figure 13).

Table 13 Managing Unused Clock Inputs

Clock Input	Power Supply Rail[1]	LJCB or LVDS
SYSCLKP SYSCLKN	CVDD GND	LJCB
PASSCLKP PASSCLKN	CVDD GND	LJCB
ALTCORECLKP ALTCORECLKN	CVDD GND	LJCB
SRIOSGMIICLKP SRIOSGMIICLKN	CVDD GND	LJCB
DDRCLKP DDRCLKN	CVDD GND	LJCB
PCIECLKP PCIECLKN	CVDD GND	LJCB
MCMCLKP MCMCLKN	CVDD GND	LJCB
RP1CLKP RP1CLKN	DVDD1V8 GND	LVDS
RP1FBP RP1FBN	DVDD1V8 GND	LVDS
End of Table 13		

1. NOTE: The power rails must be identical to those directly supporting the intended device

3.4 Input Clock Requirements

Table 14 shows the specific input clocking requirements for the KeyStone I device including maximum input jitter, rise and fall times and duty cycle. Where discrepancies occur (between data manual and this application note), always defer to the *latest* device-specific data manual. Note that all LJCB and LVDS differential input clock buffers contain a 100-Ω internal termination resistor (no external resistor is required).

Table 14 Clocking Requirements - Jitter, Duty Cycle, t_r / t_f

Clock Input	Logic	Input Jitter[8, 9, 10]	T_{rise} / T_{fall}[3]	Duty Cycle %	Stability
SRIOSGMIICLKp SRIOSGMIICLKn (if SRIO is used)	LJCB or LVDS[1]	4 ps RMS 56 ps pk-pk @ 1x10E-12 BER[7]	50 – 350 ps[2]	45 / 55	± 100 PPM
SRIOSGMIICLKp SRIOSGMIICLKn (SRIO not used)	LJCB or LVDS[1]	8 ps RMS 112 ps pk-pk @ 1x10E-12 BER[7]	50 – 350 ps[2]	45 / 55	± 100 PPM
PCIECLKp PCIECLKn	LJCB or LVDS[1]	4 ps RMS 56 ps pk-pk	50 – 350 ps[2]	45 / 55	± 100 PPM
RPICLKp RP1CLKn	LVDS	42 ps RMS 600 ps pk-pk	350	45 / 55	± 100 PPM
MCMCLKp MCMCLKn	LJCB or LVDS[1]	4 ps RMS 56 ps pk-pk @ 1x10E-15 BER[7]	50 – 350 ps[2]	45 / 55	± 100 PPM
ALTCORECLKp ALTCORECLKn	LJCB or LVDS[1]	100 ps pk-pk	50 – 350 ps[2]	45 / 55	± 100 PPM
PASSCLKp PASSCLKn	LJCB or LVDS[1]	100 ps pk-pk	50 – 350 ps[2]	45 / 55	± 100 PPM
SYSCLKp SYSCLKn	LJCB or LVDS[1]	4 ps RMS (56 ps pk–pk), period[5]	50 – 350 ps[2]	45 / 55	± 100 PPM
DDRCLKp DDRCLKn	LJCB or LVDS[1]	2.5% of DDRCLK output period (pk-pk)[4]	50 – 350 ps[2]	45 / 55	± 100 PPM
CORECLKp CORECLKn	LJCB or LVDS[1]	100 ps pk-pk	50 – 350 ps[2]	45 / 55	± 100 PPM
End of Table 14					

Notes for Table 14

Note 1: Can be DPECL or LVPECL if properly terminated and biased.

Note 2: Swing is rated for a 250 mV (pk-pk) at zero crossing where 10% – 90% of T_{RISE} and T_{FALL} must occur within the prescribed 50-350 pS time frame. The minimum slew specified is relative to the minimum input signal value. An input signal must transition through:

Minimum

$$Vtransition = 250mV * (90\% - 10\%)$$
$$Vtransition = 250mV * 80\%$$
$$Vtransition = 200mVin \xrightarrow{in} 350\,pS$$

and

Maximum

$$Vtransition = 250mV * (90\% - 10\%)$$
$$Vtransition = 250mV * 20\%$$
$$Vtransition = 200mVin \xrightarrow{in} 50\,pS$$

Note 3: Trise/Tfall values are given for 10% to 90% of the voltage swing.

Note 4: Example 1.600 GHz (~1.25e-9 nS) * 2.5% = 32.25 pS pk-pk input jitter.

Note 5: 2 ps RMS to cover future devices in the same family, (at the release of this document the requirements are only 4 pS RMS. Specified values are when AIF2 is used, when AIF2 is not used the allowable jitter is 7.1 ps RMS or 100ps pk-pk.

Note 6: 10 kHz to 20 MHz Jitter integration bandwidth over a minimum of 10K samples.

Note 7: BER rate value specified is related to the expected output error rate based on maximum input jitter requirements listed.

Note 8: Represents total input jitter (random + deterministic) for all clock inputs.

Note 9: Final jitter numbers are based on input jitter mask values at a given frequency. The numbers provided are first approximation. Always use the jitter mask to assure input clocking requirements are met.

The two major concerns pertaining to differential reference clocks are low jitter and having the proper termination. LVDS, CML, or LVPECL clock sources can be used but they require different termination strategies. The input buffer sets its own common mode voltage, so AC coupling is necessary. It also includes a 100-Ω differential termination resistor, eliminating the need for an external 100-Ω termination when using an LVDS driver. For additional information on AC termination schemes, see the *AC-Coupling Between Differential LVPECL, LVDS, and CML Application Report* (SCAA059). For information on DC coupling, see the *DC-Coupling Between Differential LVPECL, LVDS, HSTL, and CM Application Report* (SCAA062).

3.4.1 BER and Jitter

Excessive jitter can cause date errors, for this reason (and usually in high-performance data transmission protocols) a BER (bit error rate) is usually specified to minimize the risk of an error occurring. Jitter is the sum of both random and deterministic sources. In an ideal design, all jitter is identified as undesirable.

Random jitter can be defined as the sum of all noise sources including components used to construct and implement the data path between source and target. Random jitter is considered unbounded and will continue to increase over time. Deterministic jitter (unlike Random Jitter) is the total jitter induced by the characteristics of the data path between source and target.

Deterministic jitter, also called bounded jitter, obtains its minimum or maximum phase deviation within the respective time interval. Sources of deterministic jitter include (most prominently):

- [EMI] Electromagnetic Interference radiation (from power supplies, AC power lines, and RF-signal sources).
- Crosstalk (occurs when incremental inductance from one signal line (conductor) converts an induced magnetic field from an adjacent signal line into induced current resulting in either an increase or decrease in voltage.
- Inter-Symbol Interference (ISI) increases in deterministic jitter of this nature can be caused by either *simultaneous switching* (inducing current spikes on power or ground planes possibly causing a voltage threshold level shift) or a *reflection* in the signal or transmission line. Reflections result in energy flowing back through the conductor that sum with the original signal causing an amplitude variation on each conductor within a differential pair resulting in a specific time variation of the crossover (crossing) points.

Random jitter, also called unbounded jitter, can theoretically reach infinity. Random or unbounded jitter sources include:

- Shot noise [electron and hole noise in a semiconductor - increasing due to bias current and measurement bandwidth].
- Thermal noise [associated with electron flow in conductors, variations due to temperature, bandwidth, and noise resistance].
- Flicker or *pink* noise [spectral noise related to $1/f$].

Given a specific BER and an equivalent RMS jitter value, a peak-to-peak (pk-pk) jitter value can be determined (assuming a truly Gaussian or Random distribution). The following calculation can be used to approximate total jitter:

$$Jitter_{Total} = Jitter_{Random^{PK-PK}} + Jitter_{Deterministic}$$

The Random pk-pk jitter is a function of the BER and RMS jitter as denoted in the following equation:

$$Jitter_{Random^{PK-PK}} = \propto * Jitter_{RandomRMS} \qquad Or \qquad Jitter_{RMS} = Jitter_{pk-pk} / \propto$$

Table 15 Scaling Factors (Alpha) for Different BER Tolerances

BER Rate	Number $\propto$	BER Rate	Number $\propto$
10^{-4}	7.438	10^{-11}	13.412
10^{-5}	8.53	10^{-12}	14.069
10^{-6}	9.507	10^{-13}	14.698
10^{-7}	10.399	10^{-14}	15.301
10^{-8}	11.224	10^{-15}	15.883
10^{-9}	11.996	10^{-16}	16.444
10^{-10}	12.723		

Note— $\propto$ is a calculation based on the formula: $1/2\,erfc\left(\sqrt{2*\alpha}\right) = BER$

$$Jitter_{Random^{PK-PK}} = 14.069 * 4\,pS$$ (*in this example 4pS refers to total RMS input Jitter to device*)

$$Jitter_{Random^{PK-PK}} = 56.267\,pS$$

$$Jitter_{Total} = Jitter_{Random^{PK-PK}} + Jitter_{Deterministic}$$

$$Jitter_{Total} = 56.267\,pS + 4\,pS$$ (*in this example 4pS represents the total deterministic jitter expected*)

$$Jitter_{Total} = 60.267\,pS$$

Example: SYSCLK input requirement of 4 pS (RMS jitter) and a BER rate of 10^{-15}

$$Jitter_{Random^{PK-PK}} = 15.883 * 4\,pS$$ *(10e-15 BER * 4ps RMS Jitter)*

$$Jitter_{Random^{PK-PK}} = 63.532\,pS$$

$$Jitter_{Total} = Jitter_{Random^{PK-PK}} + Jitter_{Deterministic}$$

$$Jitter_{Total} = 63.532\,pS + 4\,pS$$

$$Jitter_{Total} = 67.532\,pS$$

Note 1: *It should be noted that a bit error rate can infer a specific limit on errors that are possible however where the error occurs and whether two errors can occur back-to-back is still possible.*

Table 16 shows the minimum and maximum input clock frequencies, termination, input clock type, and termination location for the device. Where discrepancies occur between the device-specific data manual and this application report, always defer to the *latest* data manual. See the termination section of this document for additional termination details. Items in **bold** represent the recommended or suggested clock input frequencies. The DDRCLKP/N input frequency will be based upon expected operating frequency (e.g., 66.667 MHz for 1333 MHz DDR3).

Table 16 Clocking Requirements - Input Frequency, Termination

Clock Input	Logic	Frequency Input (MHz) Min / Typical / Max	Termination	Term. Location
SRIOSGMIICLKp SRIOSGMIICLKn (if SRIO is used)	LJCB or LVDS[1]	156.25 / 250.00 / **312.50**	AC Couple 0.1µF	Destination Side
SRIOSGMIICLKp SRIOSGMIICLKn (SRIO not used)	LJCB or LVDS[1]	156.25 / 250.00 / **312.50**	AC Couple 0.1µF	Destination Side
PCIeCLKp PCIeCLKn	LJCB or LVDS[1]	100.00 / **156.25** / 250 / 312.50	AC Couple 0.1µF	Destination Side
RP1CLKp RP1CLKn	LVDS	**30.72**	AC Couple 0.1µF	Destination Side
MCMCLKp MCMCLKn	LJCB or LVDS[1]	156.25 / 250.00 / **312.50**	AC Couple 0.1µF	Destination Side
ALTCORECLKp ALTCORECLKn	LJCB or LVDS[1]	40.00 / **122.88** / 312.50	AC Couple 0.1µF	Destination Side
PASSCLKp PASSCLKn	LJCB or LVDS[1]	40.00 / **122.88** / 312.50	AC Couple 0.1µF	Destination Side
SYSCLKp SYSCLKn	LJCB or LVDS[1]	122.88 / **153.60** / 307.20	AC Couple 0.1µF	Destination Side
DDRCLKp DDRCLKn	LJCB or LVDS[1]	40.00 / **66.667** / 312.50	AC Couple 0.1µF	Destination Side
CORECLKp CORECLKn	LJCB or LVDS[1]	40.00 / **122.88** / 312.50	AC Couple 0.1µF	Destination Side
End of Table 16				

1. Can be DPECL or LVPECL if properly terminated and biased.

Additional SerDes clock requirements that must be met and followed in order to assure a functional system are identified in Section 3.6.6. Texas Instruments has designed and evaluated specific clock sources (see Section 3.6) that meet or exceed the performance requirements necessary for a functional system.

Calculation for proper AC termination is based on the following formula:

$$C \geq 1/2 * pi * R * Freq_{MHz}$$

Where:

- R = impedance of the net (50 Ω single ended and 100 Ω differential) – each net is 50 Ω
- C = AC capacitor value derived from formula (results are in nF)
- Freq = input frequency in MHz
- Pi = rounded to 3.141592654

Note—The pole associated with the high-pass response should be placed at a minimum of two decades below the input clock frequency.

Two examples are provided, the first for an input clock of 50 MHz and the second for an input clock frequency of 312.50 MHz.

Although the results denoted in the following examples provide a minimum value, Texas Instruments recommends a minimum of 10 nF or a 0.1 µF (100 nF) ceramic capacitor be used as the AC termination value.

Example 1: **50 MHz**

$$C \geq 1/2 * pi * R_{ohms} * Freq_{MHz/100}$$

$$C \geq 1/(2 * 3.141592654 * (50) * (50e6/100))$$

$$C \geq 1/157079632.7$$

$$C \geq 6.36619nF$$

Example 2: **312.50 MHz**

$$C \geq 1/(2 * pi * R_{ohms} * Freq_{MHz/100})$$

$$C \geq 1/(2 * 3.141592654 * 50 * (312.50e6/100))$$

$$C \geq 1/981747704.375$$

$$C \geq 1.01859e-9F$$

or 1.01859 nF

Figure 14 shows the typical LVDS-based solution, including the appropriate AC termination.

Figure 14 LJCB LVDS Clock Source

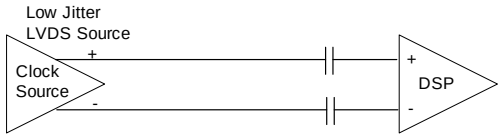

AC termination values will vary depending on many conditions, but should be in the range of 0.01 µF to 1.0 µF to help ensure minimal amplitude and phase degradation of the incoming clock. Capacitor placement is critical and highly dependent upon design topology. Determination on whether the AC coupling (DC blocking) capacitors are to be staggered or in parallel and whether they are to be placed near the receiver end or elsewhere in the nets should be determined through simulation and modeling or a signal-analysis tool typically incorporated in most common layout packages.

Ultra high speed interfaces such as SerDes require special consideration. The parasitic capacitance of the capacitor bodies (and mounting pads) to one another and to the reference plane beneath may require a change in traditional layout techniques. Placement of the AC capacitors should also be evaluated for impact on signal integrity due to reflections.

For additional clocking solutions, see the *Clocking Design Guide for KeyStone Devices Application Report* (SPRABI4).

3.5 Single-ended Clock Sources

Unlike some previous Texas Instruments DSPs, KeyStone architecture DSPs do not allow single-ended clock input sources (see note). Differential clock inputs or clock sources provide better noise immunity and signal integrity, and therefore, are the clock input choice for the KeyStone architecture DSP.

Note—It is possible to use a single-ended clock source, however, the complexity and difficulty in properly biasing the circuit as well as the potential effects on signal quality, slew rates, overshoots and undershoots make the single-ended solution undesirable.

3.6 Clocking and Clock Trees (Fan out Clock Sources)

For systems with multiple high-performance processor devices, it is not recommended that a single clock source be used (single output only). Excessive loading, reflections, and noise will adversely impact performance. Instead, it is recommended that differential clocking be used, which includes a differential clock tree.

Texas Instruments has developed a specific line of clock sources to meet the challenging requirements of today's high performance devices. In most applications the use of these specific clock sources eliminates the need for external buffers, level translators, external jitter cleaners, or multiple oscillators.

A few of the recommended parts include:

- CDCM6208 - Single PLL, 8 Differential Output Clock Tree with 4 Fractional Dividers
- CDCE62005 - Single PLL, 5 Differential Output Clock Tree with Jitter Cleaner
- CDCE62002 - Single PLL, 2 Differential Output Clock Tree with Jitter Cleaner
- CDCE62005 - Five Output Low Jitter Clock Tree with Jitter Cleaner

Note—These parts are specifically designed for the KeyStone I devices. To minimize cost and maximize performance, both parts accept a differential, single-ended, or crystal input clock source.

See the device-specific data manuals for clock tree input requirements and functionality.

3.6.1 Clock Tree Fanout Solution

The following section provides examples for routing between a CDCM6208/12 to one and multiple KeyStone I devices. A similar topology should be used regardless of the end-use application hardware. Figure 15 shows an example of a CDCM6208 used to provide all the clocks needed for a KeyStone I device.

Figure 15 Clock Fan Out - for Single Device

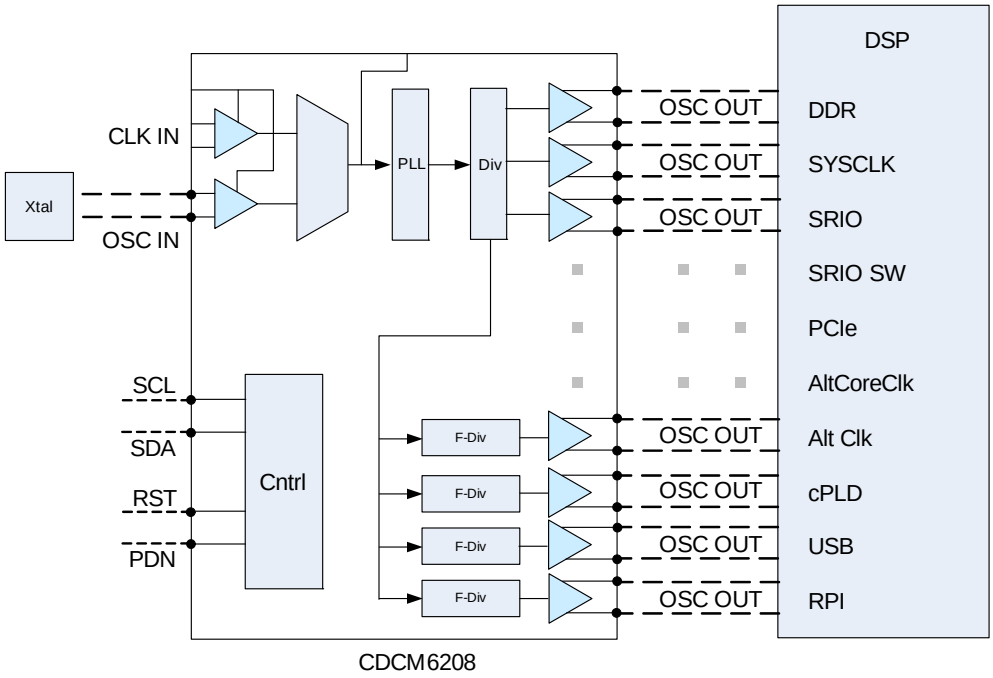

The CDCM6208 incorporates fractional dividers. When assigning fractional divider outputs from the CDCM6208, it will be necessary to verify that the input jitter and performance meet or exceed the requirements for the respective device input (select the correct clock outputs for the proper clock inputs).

Figure 16 shows the recommended clock source as applied to multiple DSPs using alternate TI clock sources (CDCL6010 and CDCE62005). Terminations are not shown, but must be included where required. Additional application hardware topologies may dictate different configurations based on trace lengths and routing rules. It is always recommended to model the clocking topologies to confirm that the design has been optimized.

Figure 16 Clock Fan Out - Multiple DSPs

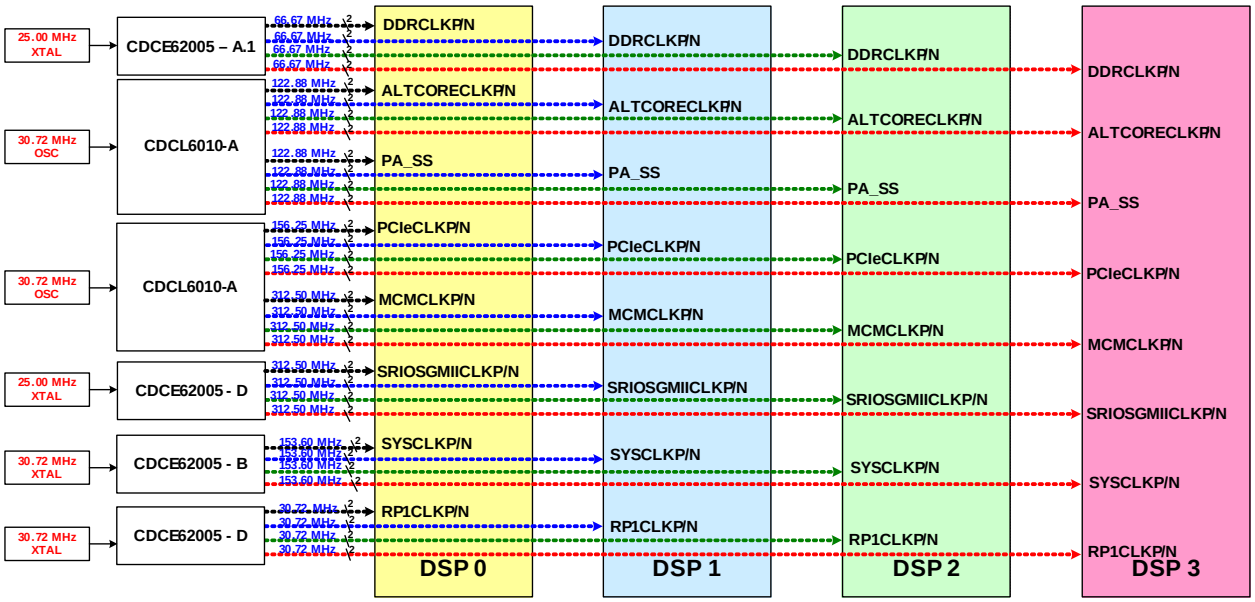

3.6.1.1 Clock Tree Layout Requirements

Key considerations when routing between clock sources and the device include the problems associated with reflections, crosstalk, noise, signal integrity, signal levels, biasing, and ground references. To establish the optimal clocking source solution, the following requirements (at a minimum) should be followed.

As with all high-performance applications, it is strongly recommended that you model (simulate) the clocking interface to verify the design constraints in the end-use application. Application board stack up, component selection, and the like all have an impact on the characteristics identified below.

See the *Clocking Design Guide for KeyStone Devices Application Report* (SPRABI4) for additional details pertaining to routing and interconnection.

3.6.2 Clock Tree Trace Width

Clocking trace widths are largely dependent on frequency and parasitic coupling requirements for adjacent nets on the application board. As a general rule of thumb, differential clock signals must be routed in parallel and the spacing between the differential pairs should be a minimum 2 times (2×) the trace width. Single-ended nets should have a spacing of 1.5× the distance of the widest parallel trace width.

3.6.3 Clock Tree Spacing

Given the variation in designs, the minimum spacing between any adjacent signals should be a minimum of 1× the width. A spacing of 1× is typically suitable for control signals (or static signals) and not data or clock nets. Single-ended or differential nets should be spaced from other nets a minimum of 1.5× and 2× respectively. A detailed cross-talk and coupling analysis (signal integrity) should be performed (modeled) before any design is released to production. An additional rule of thumb would be to maintain trace spacing ≥ three times the dielectric height (reduces ground bounce and crosstalk).

3.6.4 Clock Transmission Lines - Microstrip versus Stripline

See Section 10.6 for additional detailed information regarding microstrip and stripline features and selection criteria.

3.6.5 Clock Component Selection & Placement

A proper clock source is critical to ensure the high-performance device maintains its intended functionality. TI strongly recommends the use of the clock sources developed to compliment the device (see Section 3.6). These clock sources have been optimized for performance, jitter, and form-factor. The recommended clock source (CDCM6208) has been designed to accept either a low frequency crystal, LVPECL, LVDS, HCSL, or LVCMOS input and are capable of sourcing either a LVPECL, LVDS, HCSL, or LVCMOS clock to the device and peripheral logic.

All clock sources should be placed as close to the targeted device(s) as possible. Excessive routing induces many problems that are costly to correct in the end product and usually require a respin of the application hardware.

Alternative clock sources have been identified and include the CDCE62002 and CDCE62005. These alternative clock sources will provide functionally identical resources when implemented correctly. Additional information regarding the CDCE62002/5 is available in "Appendix A" of this document.

3.6.6 Clock Termination Type and Location

All input clocks are presumed to be of the differential type and require proper terminations or coupling. Each differential pair must be AC-Coupled using a 0.1-µF ceramic capacitor (0402 size or smaller recommended). In some cases, depending on topology, a 1-µF capacitor can be used. All AC-coupling capacitors should be located at the destination end as close to the device as possible.

Figure 17 shows the intended connection between the clock source and DSP.

Figure 17 Clock Termination Location

3.7 Clock Sequencing

Table 17 shows the requirements for input clocks to the KeyStone I device. Time frames indicated are minimum; maximum timeframes are defined in the data manual and shall not be violated. (See the device-specific data manual for specific timing requirements.)

Note—All clock drivers must be in a high-impedance state until CVDD (at a minimum) is at a valid level. After CVDD is at a valid level, all clock inputs must either be active or in a static state.

Table 17 Clock Sequencing

Clock	Condition	Sequencing
DDRCLK	None	Must be present 16 µs before $\overline{\text{POR}}$ transitions high.
SYSCLK	CORECLKSEL = 0	SYSCLK used to clock the core PLL. It must be present 16 µs before $\overline{\text{POR}}$ transitions high.
	CORECLKSEL = 1	SYSCLK only used for AIF. Clock most be present before the reset to the AIF is removed.
ALTCORECLK	CORECLKSEL = 0	ALTCORECLK is not used and should be tied to a static state.
	CORECLKSEL = 1	ALTCORECLK is used to clock the core PLL. It must be present 16 µs before $\overline{\text{POR}}$ transitions high.
PASSCLK	PASSCLKSEL = 0	PASSCLK is not used and should be tied to a static state.
	PASSCLKSEL = 1	PASSCLK is used as a source for the PA_SS PLL. It must be present before the PA_SS PLL is removed from reset and programmed.
SRIOSGMIICLK	An SGMII port will be used	SRIOSGMIICLK must be present 16 µs before $\overline{\text{POR}}$ transitions high
	SGMII will not be used. SRIO will be used as a boot device	SRIOSGMIICLK must be present 16 µs before $\overline{\text{POR}}$ transitions high.
	SGMII will not be used. SRIO will be used after boot	SRIOSGMIICLK is used as a source to the SRIO SerDes PLL. It must be present before the SRIO is removed from reset and programmed.
	SGMII will not be used. SRIO will not be used.	SRIOSGMIICLK is not used and should be tied to a static state.
PCIECLK	PCIE will be used as a boot device.	PCIECLK must be present 16 µs before $\overline{\text{POR}}$ transitions high.
	PCIE will be used after boot.	PCIECLK is used as a source to the PCIE SerDes PLL. It must be present before the PCIE is removed from reset and programmed.
	PCIE will not be used.	PCIECLK is not used and should be tied to a static state.
MCMCLK	MCM will be used as a boot device.	MCMCLK must be present 16 µs before $\overline{\text{POR}}$ transitions high.
	MCM will be used after boot	MCMCLK is used as a source to the MCM SerDes PLL. It must be present before the MCM is removed from reset and programmed.
	MCM will not be used.	MCMCLK is not used and should be tied to a static state.
End of Table 17		

4 JTAG

The JTAG interface on KeyStone I devices is used to communicate with test and emulation systems. Although JTAG is not required for operation, it is strongly recommended that a JTAG connection be included in all designs.

4.1 JTAG / Emulation

Relevant documentation for the JTAG/Emulation:

- *Emulation and Trace Headers Technical Reference Manual* (SPRU655) (but note differences defined below)
- Boundary Scan Test Specification (IEEE-1149.1)
- AC Coupled Net Test Specification (IEEE-1149.6)
- *Clocking Design Guide for KeyStone Devices Application Report* (SPRABI4)

4.1.1 Configuration of JTAG/Emulation

The IEEE Standard 1149.1-1990, IEEE Standard Test Access Port and Boundary-Scan Architecture (JTAG) interface can be used for boundary scan and emulation. The boundary scan implementation is compliant with both IEEE-1149.1 and 1149.6 (for SerDes ports). Boundary scan can be used regardless of the device configuration.

As an emulation interface, the JTAG port can be used in various modes:

- **Standard emulation:** requires only five standard JTAG signals
- **HS-RTDX emulation:** requires five standard JTAG signals plus EMU0 and/or EMU1. EMU0 and/or EMU1 are bidirectional in this mode.
- **Trace port:** The trace port allows real-time dumping of certain internal data. The trace port uses the EMU[18:00] pins to output the trace data, however, the number of pins used is configurable.

Emulation can be used regardless of the device configuration.

For supported JTAG clocking rates (TCLK), see the device-specific data manual. The EMU[18:00] signals can operate up to 166 Mbps, depending on the quality of the board-level design and implementation.

Any unused emulation port signals can be left floating.

Note—Not all KeyStone I devices contain the same number of emulation or trace pins. See the device-specific data manual for the number of emulation pins provided in that particular device.

4.1.2 System Implementation of JTAG / Emulation

For most system-level implementation details, see the *Emulation and Trace Headers Technical Reference Manual* (SPRU655). However, there are a few differences for KeyStone I device implementation compared to the information in this document:

Although the document implies 3.3-V signaling, 1.8-V signaling is supported as long as the TVD source is 1.8 V.

For a single device connection in which the trace feature is used, the standard non-buffered connections can be used along with the standard 14 pin connector. If the trace feature is used (which requires the 60-pin emulator connector), the five standard JTAG signals should be buffered and TCLK and RTCLK should be buffered separately. It is recommended to have the option for an AC parallel termination on TCLK because it is critical that the TCLK have a clean transition. EMU0 and EMU1 should not be buffered because these are used as bidirectional signals when used for HS-RTDX.

For a system with multiple DSPs that do not use the trace analysis features, the JTAG signals should be buffered as described above but the standard 14 pin connector can be used.

There are two recommended solutions if trace analysis is desired in a system with multiple DSPs:

- Emulator with trace, solution #1: trace header for each DSP (see Figure 18).
  - Pros
    - Most simple
    - Most *clean* solution electrically
  - Cons
    - Expensive (multiple headers)
    - Takes up board real estate
    - No global breakpoints, synchronous run/halt

Figure 18 Emulator with Trace Solution #1

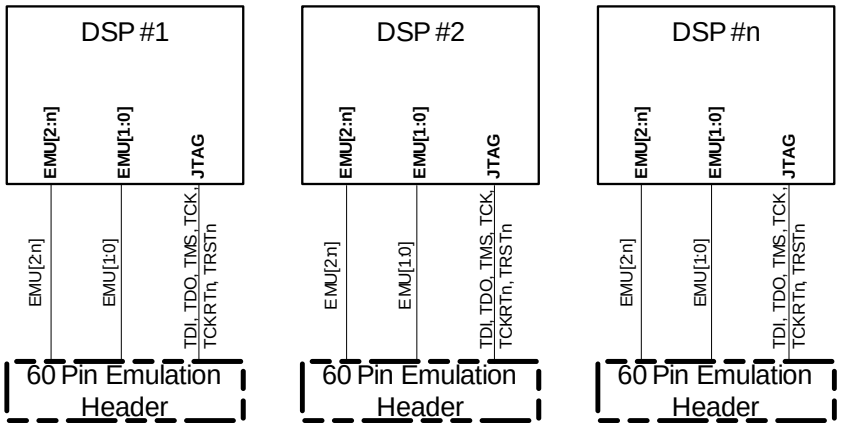

- Emulator with trace, solution #2: single trace header (see Figure 19).
  - Pros
    - Fairly *clean* solution electrically
    - Supports global breakpoints, synchronous run/halt
  - Cons
    - Supports trace on only one device
    - Less bandwidth for trace (EMU0 used for global breakpoints)
    - Loss of AET action points on EMU1 (only significant if EMU1 has been used as a trigger input/output between devices. EMU0 can be used instead if needed).

Figure 19 Emulator with Trace Solution #2

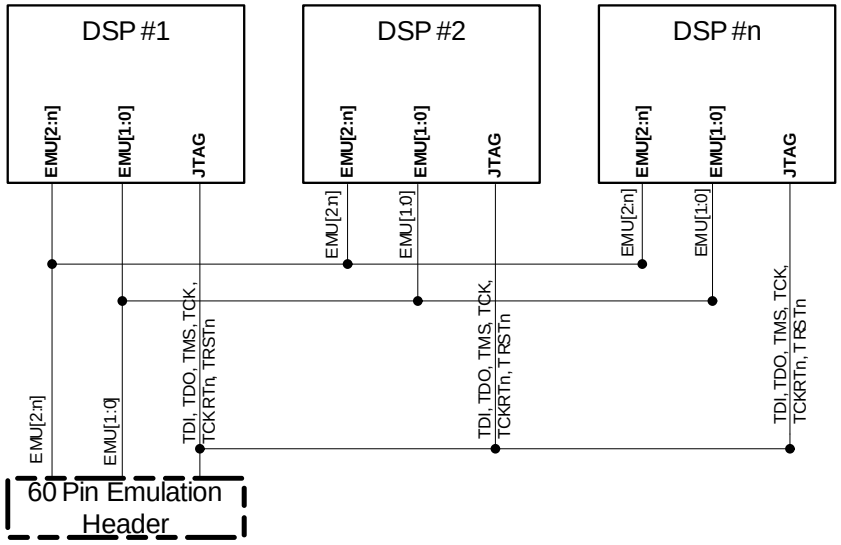

No external pull-up/down resistors are needed because there are internal pull-up/down resistors on all emulation signals.

Although no buffer is shown, for multiple DSP connections, it is recommended to buffer the five standard JTAG signals (TDI, TDO, TMS, TCK, and TRST).

If trace is used, it is *not recommended* to add both a 60-pin header and a 14-pin header because of signal integrity concerns. 60-pin to 14-pin adapters are available to allow connection to emulators that support only the 14-pin connector (although 14-pin emulators do not support trace features).

Some emulators may not support 1.8-V I/O levels. Check emulators intended to operate with the KeyStone I device to verify that they support the proper I/O levels. If 1.8-V levels are not supported, a voltage translator circuit is needed or a voltage converter board may be available. If the KeyStone I device is in a JTAG chain with devices that have a different voltage level than 1.8 V (i.e., 3.3 V), voltage translation is needed.

Figure 20 shows a dual voltage JTAG solution using buffers to perform the voltage translation. The ALVC family for 3.3 V and the AUC family for 1.8 V are both used because they have similar propagation delays.

Figure 20 Emulation Voltage Translation with Buffers

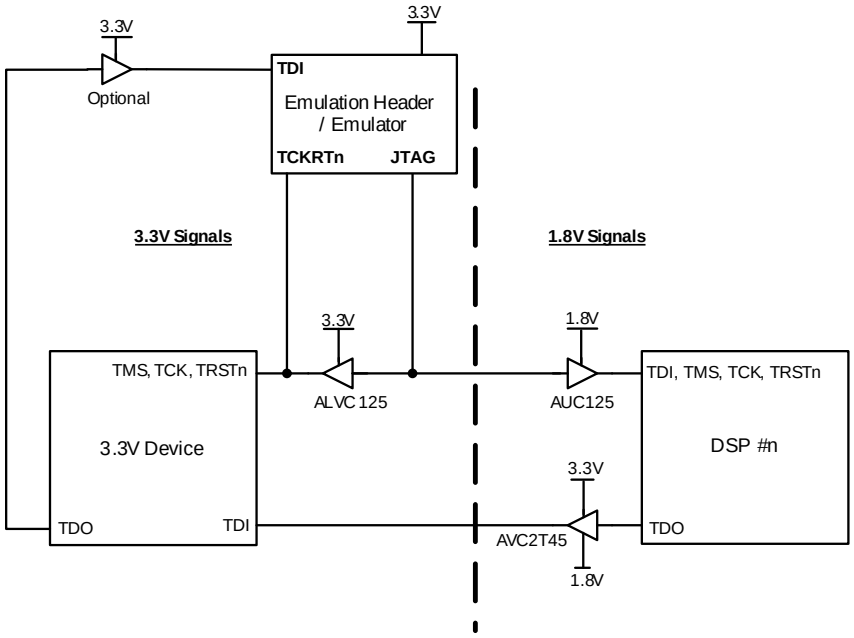

Figure 21 shows an example of using FET switch devices for voltage translation. If EMU0 and EMU1 are connected, use the following approach because these are bidirectional signals.

Figure 21 Emulator Voltage Translation with Switches

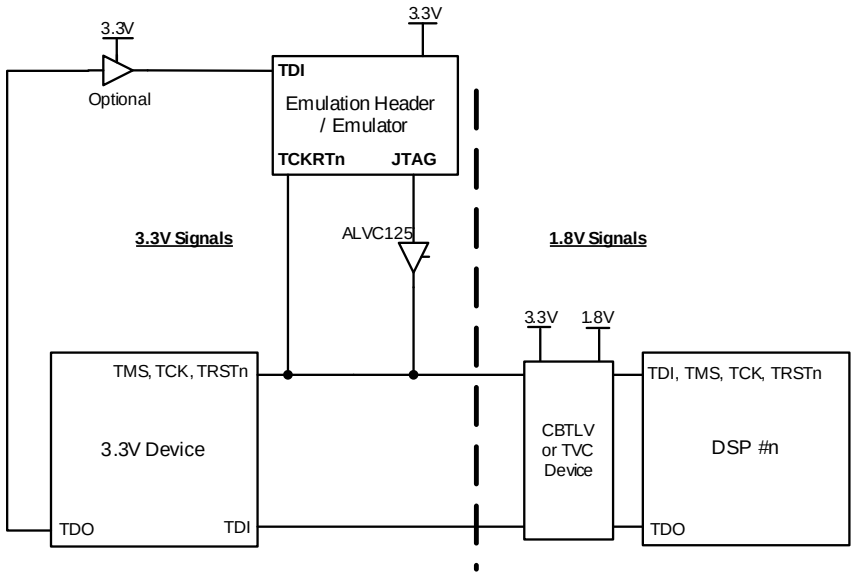

If the trace signals are supported, they may need to have voltage translation as well. Because of the speed of the trace signals, it is recommended that these signals be connected directly to the emulation header.

4.1.3 Unused Emulation Pin Requirement

The TI device contains both a JTAG and emulation pin interface. JTAG is a serial scan mechanism allowing for communication between a debugger and the device. The emulation pins provide for a high-performance output of select data captured from the device and transferred to an emulator and debugger.

If the JTAG and emulation interface is not used, all pins except $\overline{\text{TRST}}$ can be left floating. $\overline{\text{TRST}}$ must be pulled low to ground through a 4.7-kΩ resister.

In the event that the JTAG interface is used and the emulation (or a subset of the emulation pins) interface is not used, the unused emulation pins can be left floating.

See the device-specific data manual and also emulation documentation for additional connectivity requirements.

5 Device Configurations and Initialization

Once the voltage rails and the required clocks are present, the KeyStone I device may be released from reset and initialized.

On KeyStone I devices, boot modes and certain device configuration selections are latched at the rising edge of RESETFULL via specific pins. The configuration and boot mode inputs are multiplexed with general-purpose I/O (GPIO) pins or other pins. Once the level on these pins is latched into the configuration registers, these pins are released to be used for their primary function. In addition to these inputs, some peripheral usage (enabled/disabled) may also be determined by peripheral configuration registers after DSP reset. Most of the peripherals on KeyStone I processors are enabled after reset, however, some are configured as *disabled* by default. The basic information on configuration options, boot modes options, and the use of the power configuration registers can be found in the appropriate KeyStone I data manual.

5.1 Device Reset

The KeyStone I device can be reset in several ways. The methods are defined and described in greater detail in the device-specific data manual. The KeyStone I device incorporates four external reset pins (POR, RESET, LRESET, and RESETFULL) and one reset status pin (RESETSTAT). Additional reset modes are available through internal registers, emulation, and a watchdog timer. Each of the reset pins (POR, RESET, LRESET, and RESETFULL) must have either a high or low logic level applied at all times. They should not be left floating.

Note—See the device-specific data manual for device reset requirements.

5.1.1 POR

POR is an active-low signal that must be asserted during device power-up to reset all internal configuration registers to a known good state. In addition, POR must be asserted (low) on a power-up while the clocks and power planes become stable. Most output signals will be disabled (Hi-Z) while POR is low.

During a POR condition, RESET should be de-asserted before POR is de-asserted on power-up, otherwise the device comes up in the warm reset condition.

Proper POR use is required and must be configured correctly. See the device-specific data manual for additional POR details.

The internal POR circuit is unable to detect inappropriate power supply levels, including power supply brown-outs. It is necessary that all power supplies be at valid levels (stable) prior to the release of POR. The use of power-good logic in the controlling POR circuitry reduces the risk of device degradation due to power failures. Power-good logic can re-assert POR low when a power supply failure or brown-out occurs to prevent over-current conditions.

5.1.2 RESETFULL

RESETFULL is an active-low signal that will reset all internal configuration registers to a known good state. RESETFULL performs all the same functions as POR and is also used to latch the BOOTMODE and configuration pin data. RESETFULL must be asserted (low) on a power-up while the clocks and power planes become stable, and

RESETFULL must remain low until the specified time after POR has been released. For the timing requirements of POR and RESETFULL, see the device-specific data manual. Once POR has been released, the RESETFULL signal may be toggled at any time to place the KeyStone I device into a default state.

Both POR and RESET should be high before RESETFULL is deasserted, otherwise the device may not boot correctly.

The proper use of RESETFULL is required and must follow all timing requirements. See the device-specific data manual for additional RESETFULL details. RESETFULL and POR must be controlled separately. ***The KeyStone I part will not operate correctly if they are tied together.***

If POR is asserted due to an incorrect power supply voltage, then RESETFULL must also be asserted until the prescribed time after POR to assure that the proper configuration is latched into the device.

5.1.3 RESET

RESET is a software-configurable reset function to the device that is available either externally via a pin, or internally through an MMR (memory mapped register). When asserted (active low), RESET functions to reset everything on the device except for the test, emulation, and logic blocks that have reset isolation enabled.

When RESET is active low, all 3-state outputs are placed in a high-impedance (Hi-Z) state. All other outputs are driven to their inactive level. The PLL multiplier and PLL controller settings return to default and must be reprogrammed if required. If necessary, the reset isolation register within the PLL controller can be modified to block this behavior. This setting is mandatory for reset isolation to work. For additional details on reset isolation, see the applicable PLL Controller Specification.

Note—RESET does not latch boot mode and configuration pins. The previous values latched and/or programmed into the Device Status Register remain unchanged.

5.1.4 LRESET

The LRESET signal can be used to reset an individual CorePac or all CorePacs simultaneously. This signal, along with the NMI signal, can be used to affect the state of an individual CorePac without changing the state of the rest of the part. It is always used in conjunction with the CORESEL inputs and the LRESETNMIEN input for proper operation. LRESET does not reset the peripheral devices, change the memory contents, or change the clock alignment. LRESET does not cause RESETSTAT to be driven low.

LRESET must be asserted properly to function correctly. All setup, hold, and pulse-width timing for the proper implementation of LRESET can be found in the KeyStone I device-specific data manual. For proper operation, the following steps should be implemented:

1. Select the CorePac to be reset by driving the code for that CorePac onto the CORESEL input pins. The code values can be found in the device-specific data manual.
2. Drive the LRESET input low. This can be done simultaneously with the CORESEL pins.

3. Once both the CORESEL inputs and the LRESET input have been valid for the specified setup time, drive the LRESETNMIEN input low.
4. Keep the LRESETNMIEN input low for the specified minimum pulse width time.
5. Drive LRESETNMIEN high.
6. Keep the values of CORESEL and LRESET valid for the specified hold time.

5.1.4.1 Unused LRESET Pin

If the LRESET is not used, TI recommends that it be pulled high to DVDD18 (1.8-V) with an external 4.7-kΩ resistor to ensure that the CorePacs are not unintentionally reset. (This is a recommendation, not a requirement. It adds additional noise margin to this critical input.) The LRESETNMIEN and CORESEL inputs are also used in conjunction with the NMI to send an interrupt to one of the CorePacs. If neither LRESET or NMI are used, then it is recommended that the LRESET, NMI, and LRESETNMIEN inputs be pulled high to DVDD18 (1.8-V) with 4.7-kΩ resistors. The CORESEL pins can be left unconnected, relying on their internal resistors to hold them in a steady state.

5.1.5 Reset Implementation Considerations

The sequencing of the POR, RESETFULL, and RESET signals requires some logic to ensure that the signals are of proper length and are properly spaced. POR should be driven by a power management circuit that releases the signal after all power supplies have been enabled and are stable. If any of the power supplies fail during normal operation, the power management circuit should reassert the POR signal. RESETFULL can be used to reset the part through control logic in your design. RESETFULL can be asserted while POR is high but the control logic should monitor the POR signal and reassert RESETFULL if POR is driven low. RESET must be high before POR and RESETFULL are deasserted.

5.1.6 RESETSTAT

The RESETSTAT signal indicates the internal reset state. The RESETSTAT pin is driven low by almost all reset initiators and this pin remains low until the device completes initialization. The only resets that do not cause RESETSTAT to be driven low are initiators of local CorePac reset such as the LRESET. More information is available in the device-specific data manual. Access to the RESETSTAT signal is important for debugging reset problems with the device. This signal should be routed to external logic or to a test point.

5.1.7 EMULATION RESET

The KeyStone I family of devices supports the emulation reset function as part of its standard offering. Emulation reset supports both a hard and soft reset request. The source for this reset is on-chip emulation logic. This reset behavior is the same as hard / soft reset based on the reset requester type. This reset is non-blockable.

5.2 Device Configuration

Several device configuration strapping options are multiplexed with the GPIO pins (see the device-specific data manual for pin assignments). There are dedicated configuration pins, such as CORECLKSEL and DDRSLRATE1:0. The state of these pins is not latched and must be held at the desired state at all times. See the data manual for details on the configuration options.

Device configuration pins should be pulled to ground or to 1.8V (DVDD18) to establish a configuration level.

If the GPIO signals are not used, the internal pull-up (to DVDD18) and pull-down (to ground) resistors can be used to set the input level and external pull-up/down resistors are needed only if the opposite setting is desired. If the GPIO pins are connected to other components or a test point, the internal pull-up/pull-down resistor should not be relied upon. External pull-up and pull-down resistors are recommended for all desired settings. If the pin is pulled in the same direction as the internal pull, then a 4.7-kΩ resistor should be used. If the pin is pulled in the opposite direction as the internal pull, then a 1-kΩ resistor should be used.

If multiple configuration inputs are tied together from one or more KeyStone I devices (only allowed for input-only pins), the external resistor should have a maximum value of 1000/N where N is the number of input pins. If the boot mode or configuration pin is used after the boot mode is captured, then individual resistors should be used.

The phase-locked loop (PLL) multipliers can be set only by device register writes. For details on configuration of the PLL, see the *KeyStone Architecture Phase-Locked Loop (PLL) User Guide* (SPRUGV2) and the data manual.

5.3 Peripheral Configuration

In addition to the device reset configuration, which is covered in the previous section, there are several configuration pins to set. These include: bit ordering (endian), PCI mode, and PCI enable. See the device-specific data manual and GPIO section of this document for correct pin configuration.

Any additional configurations not listed in this or previous sections are done through register accesses. Peripherals that default disabled can be enabled using the peripheral configuration registers. If the boot mode selection specifies a particular interface for boot (SRIO, Ethernet, I^2C), it is automatically enabled and configured. For additional details on peripheral configurations see the Device Configuration section of the data manual.

In some cases, the peripheral operating frequency is dependent on the device core clock frequency and/or boot mode inputs. This should be taken into account when configuring the peripheral.

5.4 Boot Modes

The KeyStone I device contains multiple interfaces that support boot loading. Examples include: EMAC (Ethernet Media Access Controller), SRIO (Serial Rapid IO), PCIe (Peripheral Component Interconnect – express), Emulation, I^2C (Inter-Integrated Circuit), SPI, and HyperLink.

For details regarding boot modes, see the *KeyStone Architecture Bootloader User Guide* (SPRUGY5) and the respective device-specific data manual. Regardless of the boot mode selected, a connected emulator can always reset the device and acquire control.

6 Peripherals Section

This section covers each of the KeyStone I device peripherals/modules (some may not be included, see your specific data manual). This section is intended to be used in addition to the information provided in the device data manual, the peripheral/module guides, and relevant application reports. The four types of documents should be used as follows:

- **Data Manual:** AC Timings, register offsets
- **Module Guide:** Functional Description, Programming Guide
- **Applications Reports:** System level issues
- **This Chapter:** Configuration, system level issues not covered in a separate application report

Each peripheral section includes recommendations on how to handle pins on interfaces that are disabled or for unused pins on interfaces that are enabled. Generally, if internal pull-up or pull-down resistors are included, the pins can be left unconnected. Any pin that is output-only can always be left unconnected. Normally, if internal pull-up and pull-down resistors are not included, pins can still be floated with no functional issues for the device. However, this normally causes additional leakage currents that can be eliminated if external pull-up or pull-down resistors are used. Inputs that are not floating have a leakage current of approximately 100 µA per pin. Leakage current is the same for a high- or low-input (either pull-up or pull-down resistors can be used). When the pins are floating, the leakage can be several milliamps per pin. Connections directly to power or ground can be used only if the pins can be assured to never be configured as outputs and the boundary scan is not run on those pins.

6.1 GPIO/Device Interrupts

Documentation for GPIO/Interrupts:

- *KeyStone Architecture General Purpose Input/Output (GPIO) User Guide* (SPRUGV1)
- KeyStone IBIS Model File (specific to device part number)
- *Using IBIS Models for Timing Analysis* (SPRA839)

6.1.1 Configuration of GPIO/Interrupts

All GPIOs pins are multiplexed with other functions, with several being used for device configuration strapping options. The strapping options are latched by RESETFULL going high. After RESETFULL is high, the pins are available as GPIO pins. Note: All pins must be driven to a logical state during boot and through RESETFULL being released. The GPIO pins are read during a RESETFULL boot function and if required, are available for GPIO functionality after a successful boot.

GPIOs are enabled at power-up and default to inputs.

All GPIOs can be used as interrupts and/or EDMA events to any of the cores.

6.1.2 System Implementation of GPIO/Interrupts

It is recommended that GPIOs used as outputs have a series resistance (22 Ω or 33 Ω being typical values). The value (or need) for the series resistor can be determined by simulating with the IBIS models.

If you want to have a GPIO input default to a particular state (low or high), use an external resistor. A pull-up resistor value of 1 kΩ to 1.8V (DVDD18) is recommended to make sure that it overrides the internal pull-down resistor present on some GPIOs. If this GPIO is also used as a strapping option, the default state needs to also be the desired boot strapping option. The RESETSTAT signal can be used to Hi-Z logic that drives a GPIO boot strapping state during the POR transition.

6.1.3 GPIO Strapping Requirements

Table 18 shows the strapping requirements for one of the KeyStone I family devices (see the data manual for specific pin numbers). Included is the pin-to-function correlation needed to properly configure the device.

Table 18 GPIO Pin Strapping Configurations

			Primary Function		Secondary Function	
Pin	**Name**	**Boot Mode**	**Pull-up**	**Pull-dn**		**Comment**
AJ20	GPIO00	LENDIAN	Little Endian	Big Endian	GPIO (after $\overline{\text{RESETFULL}}$)	
AG18	GPIO01	BOOTMODE00	Boot Device	Boot Device	GPIO (after $\overline{\text{RESETFULL}}$)	See 2.5 for additional details
AD19	GPIO02	BOOTMODE01	Boot Device	Boot Device	GPIO (after $\overline{\text{RESETFULL}}$)	See 2.5 for additional details
AE19	GPIO03	BOOTMODE02	Boot Device	Boot Device	GPIO (after $\overline{\text{RESETFULL}}$)	See 2.5 for additional details
AF18	GPIO04	BOOTMODE03	Device Cfg	Device Cfg	SR ID	Specific device configurations
AE18	GPIO05	BOOTMODE04	Device Cfg	Device Cfg	SR ID	Specific device configurations
AG20	GPIO06	BOOTMODE05	Device Cfg	Device Cfg	GPIO (after $\overline{\text{RESETFULL}}$)	Specific device configurations
AH19	GPIO07	BOOTMODE06	Device Cfg	Device Cfg	GPIO (after $\overline{\text{RESETFULL}}$)	Specific device configurations
AJ19	GPIO08	BOOTMODE07	Device Cfg	Device Cfg	GPIO (after $\overline{\text{RESETFULL}}$)	Specific device configurations
AE21	GPIO09	BOOTMODE08	Device Cfg	Device Cfg	GPIO (after $\overline{\text{RESETFULL}}$)	Specific device configurations
AG19	GPIO10	BOOTMODE09	Device Cfg	Device Cfg	GPIO (after $\overline{\text{RESETFULL}}$)	Specific device configurations
AD20	GPIO11	BOOTMODE10	PLL Multiplier/I²C	PLL Multiplier/I²C	GPIO (after $\overline{\text{RESETFULL}}$)	Specific device configurations
AE20	GPIO12	BOOTMODE11	PLL Multiplier/I²C	PLL Multiplier/I²C	GPIO (after $\overline{\text{RESETFULL}}$)	Specific device configurations
AF21	GPIO13	BOOTMODE12	PLL Multiplier/I²C	PLL Multiplier/I²C	GPIO (after $\overline{\text{RESETFULL}}$)	Specific device configurations
AH20	GPIO14	PCIESSMODE0	Endpt/RootComplex	Endpt/RootComplex	GPIO (after $\overline{\text{RESETFULL}}$)	Specific device configurations
AD21	GPIO15	PCIESSMODE1	Endpt/RootComplex	Endpt/RootComplex	GPIO (after $\overline{\text{RESETFULL}}$)	Specific device configurations
End of Table 18						

Note—See the Boot and Configuration Specification in addition to the data manual for additional details.

6.1.4 Unused GPIO Pin Requirement

All GPIO pins serve multiple functions, so each of the GPIO pins are required to be pulled to a logical state during boot (as defined by the boot method and as listed in the Bootloader Specification). After boot is complete ($\overline{\text{POR}}$ is deasserted), each GPIO pin is released and made available to the user as a general purpose IO. Each GPIO pin (after boot is complete) defaults to its identified state (see the data manual for default pull-up and pull-down conditions) unless otherwise configured. All application hardware must incorporate an appropriate design to permit flexibility between device configuration/boot and functional GPIO requirements.

If an external resistor is used to obtain a specific post boot state for any GPIO pin, use a 1 kΩ to 4.7-kΩ resistor to 1.8V (DVDD18).

Note—If any of the GPIO pins are attached to a trace, test point, another device, or something other than a typical no-connect scenario, a pull-up or pull-down resistor is mandatory. The default internal pull up resistor for GPIO00 and the default internal pull down resistors for GPIO01 through GPIO15 are suitable only if nothing is connected to the pin.

6.2 HyperLink Bus

Documentation for HyperLink:

- *KeyStone Architecture HyperLink User Guide* (SPRUGW8)
- *TMS320C66x DSP CPU and Instruction Set Reference Guide* (SPRUGH7)
- KeyStone IBIS Model File (specific to device part number)

6.2.1 Configuration of HyperLink

The HyperLink interface is a high-performance TI-developed interface intended for high speed data communication. See the HyperLink User Guide for additional configuration details.

6.2.2 System Implementation of HyperLink Interface

See the *SerDes Implementation Guidelines for KeyStone I Devices Application Report* (SPRABC1) for details, guidelines, and routing requirements. This interface is intended to be a device-to-device interconnection. It is not intended for a board-to-board interconnection. Currently, HyperLink is supported only by TI and select third party FPGA IPs.

6.2.3 HyperLink-to-HyperLink Connection

The HyperLink interface bus is intended to be a short-reach high-performance bus supporting a total of four interconnecting lanes between two DSPs operating at up to 10.0 GHz per lane (other configurations may exist). HyperLink data is passed using up to four SerDes lanes. In addition, the HyperLink includes a number of sideband signals used for flow control and power management control. The sideband signals use LVCMOS buffers and operate at speeds up to 166 MHz. Both the SerDes lanes and the sideband signals must be connected for proper operation.

Figure 22 HyperLink Bus DSP-to-DSP Connection

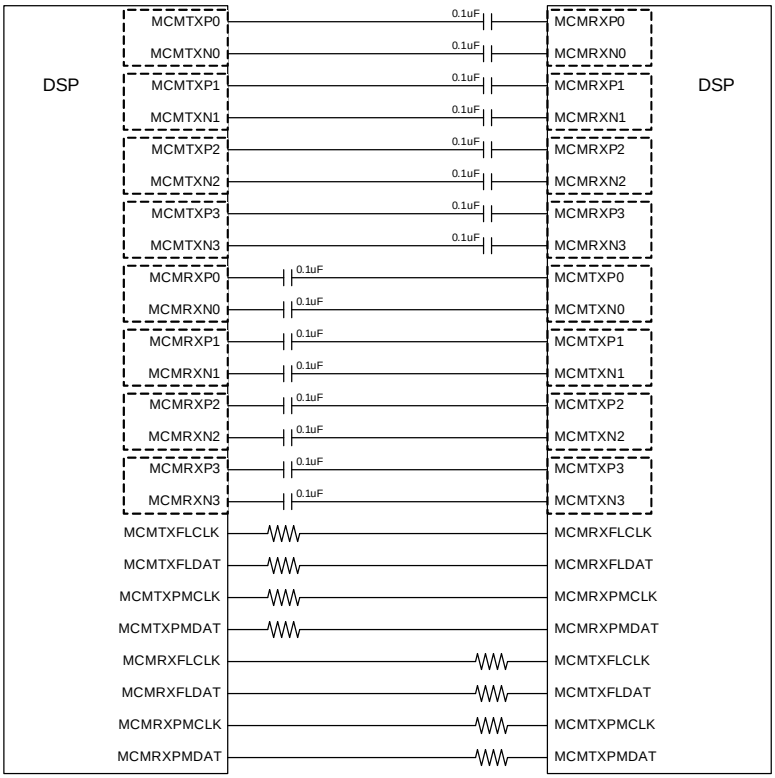

6.2.4 Configuration of HyperLink

If the HyperLink interface is used, a MCMCLKP/N clock must be provided and the SerDes PLL must be configured to generate the desired line rate. The line rate is determined using the MCMCLK frequency, the PLL multiplier, and the rate scale setting. The PLL multiplier value is found in the PLL configuration register and the rate scale setting is found in the transmit and receive configuration registers. The PLL output frequency is equal to the frequency of the MCMCLKP/N multiplied by the PLL multiplier.

PLL Output (MHz) = PLL Multiplier × MCMCLK (MHz)

The VCO frequency range of the HyperLink PLL is 1562.5 MHz to 3125 MHz, so the PLL output frequency must fall within this range.

The highest HyperLink line rates supported by KeyStone I devices is 10000 Mbps. The rate scale is used to translate the output of the SRIO PLL to the line rate using the following formula:

line rate Mbps = PLL Output / rate scale

The HyperLink SerDes interface supports the four rate scales and the associated rate scale value are shown in Table 19.

Table 19 HyperLink Rate Scale Values

Rate Scale	Rate Scale Value	Line Rate (Mbps)
Full	0.25	PLL Output / 0.25
Half	0.5	PLL Output / 0.5
Quarter	1	PLL Output / 1
8th	2	PLL Output / 2

The possible PLL multiplier settings for the three recommended MCMCLKP/N clock frequencies are shown in Table 20. The table includes only the settings for the highest rate. The HyperLink interface can be operated at slower rates. The equations above can be used to generate the settings for lower rates.

Table 20 HyperLink PLL Multiplier Settings

Reference Clock MHz	PLL Multiplier	PLL Output (MHz)	Rate Scale	Line Rate (Mpbs)
156.25	16	2500	0.25	10000
250.00	10	2500	0.25	10000
312.50	8	2500	0.25	10000

6.2.5 Unused HyperLink Pin Requirement

All HyperLink clock and data pins (MCMRXFLCLK, MCMRXFLDAT, MCMTXFLCLK, MCMTXFLDAT, MCMRXPMCLK, MCMRXPMDAT, MCMTXPMCLK, MCMTXPMDAT) can be left floating when the entire HyperLink peripheral is not used. Each unused HyperLink clock and data pin will be pulled low when not in use through internal pull-down resistors.

In the event a partial number of lanes are used (two lanes instead of four), all clock and data pins must be connected. Unused lanes (both transmit and receive) can be left floating.

If the HyperLink interface is not being used, the peripheral should be disabled in the MMR.

If the HyperLink interface is not required, the HyperLink regulator power pin (VDDR1_MCM) must still be connected to the correct supply rail with the appropriate decoupling capacitance applied. The EMI/Noise filter is not required when this interface is not used.

If unused, the primary HyperLink clock input must be configured as indicated in Section 3.3, Figure 13. Pull-up must be to the CVDD core supply.

6.3 Inter-Integrated Circuit (I²C)

Documentation for I²C:

- *KeyStone Architecture Inter-IC Control Bus (I²C) User Guide* (SPRUGV3)
- IBIS Model File for KeyStone Devices (specific to device part number)
- *Using IBIS Models for Timing Analysis Application Report* (SPRA839)
- Philips I²C Specification, Version 2.1

6.3.1 Configuration of I²C

The I²C peripheral powers up enabled. The input clock for the I²C module is SYSCLK, which is further divided down using an internal PLL. There is a prescaler in the I²C module that needs to be set up to reduce this frequency. See the *Bootloader for KeyStone Architecture User's Guide* and *Inter-Integrated Circuit (I²C) for KeyStone Devices User Guide* for additional information and details.

6.3.2 System Implementation of I²C

External pull-up resistors to 1.8 V are needed on the device I²C signals (SCL, SDA). The recommended pull-up resistor value is 4.7 kΩ.

Multiple I²C devices can be connected to the interface, but the speed may need to be reduced (400 kHz is the maximum) if many devices are connected.

The I²C pins are not 2.5 V or 3.3 V tolerant. For connection to 2.5-V or 3.3-V I²C peripherals, the PCA9306 can be used. See the *PCA9306 Dual Bidirectional I²C Bus and SMBus Voltage-Level Translator Data Sheet* (SCPS113) for more information. Section 6.3.3 describes in greater detail the interconnection.

6.3.3 I²C Interface

KeyStone I I²C pins are open-drain IOs, and require pull-up resistors to the DVDD18 rail (same rail as the device 1.8-V supply).

Figure 23 shows the intended configuration inclusive of the PCA9306, DSP, and, as an example, an EEPROM. This configuration shows the connection topology for a single DSP. In the case in which multiple DSPs are used, each I²C interface must include the same configuration.

Figure 23 I²C Voltage Level Translation

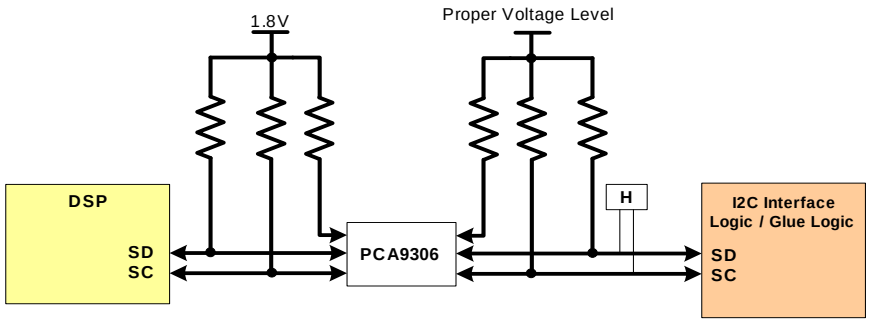

6.3.4 Unused I²C Pin Requirement

If the I²C signals are not used, it is recommended that the SDA and SCL pins be pulled up to 1.8V (DVDD18). These pins can be left floating, but this will result in a slight increase in power consumption due to increased leakage.

6.4 Ethernet

Documentation for EMAC:

- *KeyStone Architecture Gigabit Ethernet (GbE) Switch Subsystem User Guide* (SPRUGV9)
- SGMII Specification (ENG-46158), Version 1.8, dated April 27, 2005 [7]
- *SerDes Implementation Guidelines for KeyStone Devices Application Report* (SPRABC1)

6.4.1 Configuration of EMAC, SGMII and MDIO

The EMAC interface is compliant with the *SGMII Specification (ENG-46158)*, Version 1.8 that specifies LVDS signals. Only the data channels are implemented, so the connected device must support clock recovery and not require a separate clock signal. When the EMAC is enabled, the MDIO interface is enabled.

The MDIO interface (MDCLK, MDIO) uses 1.8-V LVCMOS buffers. The EMAC must be enabled via software before it can be accessed unless the Boot over Ethernet boot mode is selected.

If the EMAC is used, a SRIOSGMIICLKP/N clock must be provided and the SerDes PLL must be configured to generate a 1.25-Gbps line rate. The line rate is determined using the SRIOSGMIICLK frequency, the PLL multiplier, and the rate scale setting. The PLL multiplier value is found in the PLL configuration register and the rate scale setting is found in the transmit and receive configuration registers. The PLL output frequency is equal to the frequency of the SRIOSGMIICLKP/N multiplied by the PLL multiplier.

PLL Output (MHz) = PLL Multiplier x SRIOSGMIICLK (MHz)

The VCO frequency range of the SGMII PLL is 1500 MHz to 3125 MHz, so the PLL output frequency must fall within this range.

The line rate for SGMII is 1250 Mbps. The rate scale is used to translate the output of the SGMII PLL to the line rate using the following formula:

1250 Mbps = PLL Output / rate scale

The SGMII SerDes interface supports the four rate scales and the associated rate scale value are shown in Table 21.

Table 21 SGMII Rate Scale Values

Rate Scale	Rate Scale Value	Line Rate (Mbps)
Full	0.5	PLL Output / 0.5
Half	1	PLL Output / 1
Quarter	2	PLL Output / 2
32nd	16	PLL Output / 16

The possible PLL multiplier settings for the three recommended SRIOSGMIICLKP/N clock frequencies are shown in Table 22. Although the SGMII SerDes share a reference clock with the SRIO SerDes, they have separate PLLs that can be configured with different multipliers.

Table 22 SGMII PLL Multiplier Settings

Reference Clock MHz	PLL Multiplier	PLL Output (MHz)	Rate Scale	Line Rate (Mpbs)
156.25	16	2500	2	1250
250.00	10	2500	2	1250
312.50	8	2500	2	1250

See the data manual, users guide, and application notes for more detailed information.

If EMAC is not used, the SerDes signals and MDIO signals can be left unconnected. If both EMAC and SRIO are not used, the RIOSGMIICLKP/N pins should be terminated as shown in Figure 13. See Section 6.4.4 for additional detail.

6.4.2 System Implementation of SGMII

SGMII uses LVDS signaling. The KeyStone I device uses a CML-based SerDes interface that requires AC coupling to interface to LVDS levels. Texas Instruments recommends the use of 0.1-µF AC coupling capacitors for this purpose. The SerDes receiver includes a 100-Ω internal termination, so an external 100-Ω termination is not needed. Examples of SerDes-to-LVDS connections appear in the following sections.

If the connected SGMII device does not provide common-mode biasing, external components need to be added to bias the LVDS side of the AC-coupling capacitors to the nominal LVDS offset voltage, normally 1.2 V. Additional details regarding biasing can be found in the application note *Clocking Design Guide for KeyStone Devices Application Report (SPRABI4)*.

For information regarding supported topologies and layout guidelines, see the *SerDes Implementation Guidelines for KeyStone I Devices Application Report* (SPRABC1).

The SGMII interface supports hot-swap, where the AC-coupled inputs of the device can be driven without a supply voltage applied.

SRIO/SGMII SerDes power planes and power filtering requirements are covered in Section 7.3.

6.4.3 SGMII MAC to MAC Connection

The SGMII interface can be connected from the KeyStone I device to a PHY or from the KeyStone I device to a MAC, including KeyStone-to-KeyStone device direct connects. An example of the hardware connections is shown in Figure 24. For auto-negotiation purposes, the KeyStone I device can be configured as a master or a slave, or it can be set up for fixed configuration. If the KeyStone I device is connected to another MAC, the electrical compatibility must be evaluated to determine if additional terminations are needed.

Note—Not all devices contain the implementation shown. Your specific KeyStone I device should be properly configured for the correct number of SGMII interfaces shown in the data manual.

Figure 24 SGMII MAC to MAC Connection

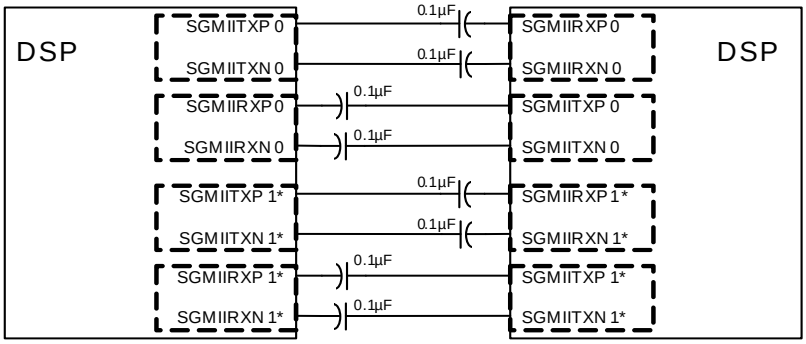

6.4.4 Unused SGMII Pin Requirement

All unused SGMII pins must be left floating. As an additional note, should the SGMII interface not be required, the peripherals should be disabled. In the event only one of the two SGMII interfaces is required, the unused SGMII interface must be disabled through the MMR of the device.

The SGMII shares a common input clock with the SRIO interface. If either of these interfaces is used, the respective power must still be properly connected in accordance with this specification. In addition, if the SGMII interface is used, the appropriate amount of decoupling/bulk capacitance must be included.

If both of the SGMII interfaces are not used, the SGMII regulator power pin (VDDR3_SGMII) must still be connected to the correct supply rail with the appropriate decoupling capacitance applied. The EMI/Noise filter is not required when this interface is not used.

If both SRIO and SGMII are unused, the primary SRIOSGMIICLK clock input must be configured as indicated in Section 3.3, Figure 13. Pull-up must be to the CVDD core supply. (Note: *both interfaces must be unused in order to pull-up the unused clock signal.*)

6.5 Serial RapidIO (SRIO)

Relevant documentation for SRIO:

- *KeyStone Architecture Serial RapidIO (SRIO) User's Guide* (SPRUGW1)
- RapidIO Interconnect Part VI: Physical Layer 1×/4× LP-Serial Specification, Version 2.0.1 [10]
- *SerDes Implementation Guidelines for KeyStone I Devices* (SPRABC1)

6.5.1 Configuration of SRIO

Detailed information on the configuration of the SRIO logic block can be found in the functional operation section of the SRIO user guide.

If the SRIO interface is used, a SRIOSGMIICLKP/N clock must be provided and the SerDes PLL must be configured to generate the desired line rate. The line rate is determined using the SRIOSGMIICLK frequency, the PLL multiplier, and the rate scale setting. The PLL multiplier value is found in the PLL configuration register and the rate scale setting is found in the transmit and receive configuration registers. The PLL output frequency is equal to the frequency of the SRIOSGMIICLKP/N multiplied by the PLL multiplier:

PLL Output (MHz) = PLL Multiplier x SRIOSGMIICLK (MHz)

The VCO frequency range of the SRIO PLL is 1562.5 MHz to 3125 MHz, so the PLL output frequency must fall within this range.

The SRIO line rates supported by KeyStone I devices are 1250 Mbps, 2500 Mbps, 3125 Mbps, and 5000 Mbps. The rate scale is used to translate the output of the SRIO PLL to the line rate using the following formula:

line rate Mbps = PLL Output / rate scale

The SRIO SerDes interface supports the four rate scales and the associated rate scale value shown in Table 23.

Table 23 SRIO Rate Scale Values

Rate Scale	Rate Scale Value	Line Rate (Mbps)
Full	0.25	PLL Output / 0.25
Half	0.5	PLL Output / 0.5
Quarter	1	PLL Output / 1
8th	2	PLL Output / 2
End of Table 23		

The possible PLL multiplier settings for the three recommended SRIOSGMIICLKP/N clock frequencies are shown in Table 24. Although the SRIO SerDes share a reference clock with the SGMII SerDes, they have separate PLLs that can be configured with different multipliers.

Table 24 SRIO PLL Multiplier Settings (Part 1 of 2)

Reference Clock MHz	PLL Multiplier	PLL Output (MHz)	Rate Scale	Line Rate (Mpbs)
156.25	16	2500	2	1250
250.00	10	2500	2	1250
312.50	8	2500	2	1250

Table 24 SRIO PLL Multiplier Settings (Part 2 of 2)

Reference Clock MHz	PLL Multiplier	PLL Output (MHz)	Rate Scale	Line Rate (Mpbs)
156.25	16	2500	1	2500
250.00	10	2500	1	2500
312.50	8	2500	1	2500
156.25	10	1562.5	0.5	3125
156.25	20	3125	1	3125
250.00	12.5	3125	1	3125
312.50	5	1562.5	0.5	3125
312.50	10	3125	1	3125
156.25	16	2500	0.5	5000
250.00	10	2500	0.5	5000
312.50	8	2500	0.5	5000
End of Table 24				

It is possible to configure the KeyStone I device to boot load application code over the SRIO interface. Boot over SRIO is a feature that is selected using boot strapping options. For details on boot strapping options, see the KeyStone I data manual.

If the SRIO peripheral is not used, the SRIO link pins can be left floating and the SerDes links should be left in the disabled state.

The SRIO SerDes ports support hot-swap, in which the AC-coupled inputs of the device can be driven without a supply voltage applied.

If the SRIO peripheral is enabled but only one link is used, the pins of the unused link can be left floating.

6.5.2 System Implementation of SRIO

The Serial RapidIO implementation is compliant with the *RapidIO Interconnect Part VI: Physical Layer 1×/4× LP-Serial Specification*, Version 2.0.1.

For information regarding supported topologies and layout guidelines, see the *SerDes Implementation Guidelines for KeyStone I Devices Application Report* (SPRABC1).

Suggestions on SRIO reference clocking solutions can be found in Section 3.5.

SRIO/SGMII SerDes power planes and power filtering requirements are covered in Section 2.

6.5.3 Unused SRIO Pin Requirement

All unused SRIO lanes must be left floating and the unused lanes properly disabled through the MMR. If the entire SRIO interface is not required, the SRIO peripheral should be disabled.

The SRIO shares a common input clock with the SGMII interface. If either of these interfaces are used, the respective power is still required to be connected in accordance with this specification and the appropriate decoupling/bulk capacitance included.

If the SRIO interface is not required, the SRIO regulator power pin (VDDR4_SRIO) must still be connected to the correct supply rail with the appropriate decoupling capacitance applied. The EMI/Noise filter is not required when this interface is not used.

If both SRIO or SGMII are unused, the primary SRIOSGMIICLK clock input must be configured as indicated in Section 3.3, Figure 13. Pull-up must be to the CVDD core supply.

Note—Both interfaces must be unused in order to pull-up the unused clock signal.

6.6 Peripheral Component Interconnect Express (PCIe)

Relevant documentation for SRIO:

- *KeyStone Architecture Peripheral Component Interconnect Express (PCIe) User Guide* (SPRUGS6)
- PCI Express Base Specification, revision 1.0
- PCI Express Base Specification, Version 2.0.1
- OCP Specification, revision 2.2

6.6.1 PCIe Features Supported

PCI Express is a point-to-point serial signaling protocol incorporating two links (consisting of one TX and one RX differential pair, each). Supported is a single ×2 configuration or a single ×1 configuration (2 separate ×1 links are not supported). TI's PCIe interface supports (depending upon the protocol version) either 2.5 Gbps or 5.0 Gbps data transfer in each direction (less 8b/10b encoding overhead).

The device PCIe interface supports a dual operating mode (either Root complex (RC) or end point (EP)).

6.6.2 Configuration of PCIe

The PCIe mode is controlled by pinstrapping the two MSB GPIO bits (GPIO15:14). Table 25 shows the mode configuration. Specific details and proper configuration of the PCIe (if used) is identified in detail in the PCIe User Guide.

Table 25 PCIe (GPIO) Configuration Table

GPIO15	GPIO14	
PCIESSMODE1	PCIESSMODE0	Selection
0	0	PCIe in End Point Mode
0	1	PCIe in Legacy End Point Mode
1	0	PCIe in Root Complex Mode
1	1	NOT VALID – Do Not Use

If the PCIe interface is used, a PCIECLKP/N clock must be provided and the SerDes PLL must be configured to generate the desired line rate. The line rate is determined using the PCIECLK frequency, the PLL multiplier, and the rate scale setting. The PLL multiplier value is found in the PLL configuration register and the rate scale setting is found in the transmit and receive configuration registers. The PLL output frequency is equal to the frequency of the PCIECLKP/N multiplied by the PLL multiplier:

PLL Output (MHz) = PLL Multiplier x PCIECLK (MHz)

The VCO frequency range of the SGMII PLL is 1500 MHz to 3125 MHz, so the PLL output frequency must fall within this range.

The supported line rates for PCIe are 2500 Mbps and 5000 Mbps. The rate scale is used to translate the output of the PCIE PLL to the line rate using the following formula:

line rate Mbps = PLL Output / rate scale

The PCIE SerDes interface supports the four rate scales and the associated rate scale value shown in Table 26.

Table 26 PCIe Rate Scale Values

Rate Scale	Rate Scale Value	Line Rate (Mbps)
Full	0.5	PLL Output / 0.5
Half	1	PLL Output / 1
Quarter	2	PLL Output / 2
32nd	16	PLL Output / 16
End of Table 26		

The possible PLL multiplier settings for the three recommended PCIECLKP/N clocks frequencies are shown in Table 27.

Table 27 PCIE PLL Multiplier Settings

Reference Clock MHz	PLL Multiplier	PLL Output (MHz)	Rate Scale	Line Rate (Mpbs)
156.25	16	2500	1	2500
250.00	10	2500	1	2500
312.50	8	2500	1	2500
156.25	16	2500	0.5	5000
250.00	10	2500	0.5	5000
312.50	8	2500	0.5	5000
End of Table 27				

See the device-specific data manual, users guide, and application notes for more detailed information.

6.6.3 Unused PCIe Pin Requirement

All unused PCIe lanes must be left floating and the unused lanes properly configured in the MMR. If the entire PCIe interface is not required, the PCIe peripheral should be disabled in the MMR.

If the PCIe interface is not required, the PCIe regulator power pin (VDDR2_PCIe) must still be connected to the correct supply rail with the appropriate decoupling capacitance applied. The EMI/Noise filter is not required when this interface is not used.

If unused, the PCIe clock must be configured as indicated in Section 3.3, Figure 13. Pull-up must be to the CVDD core supply.

6.7 Antenna Interface (AIF)

Relevant documentation for AIF:

- *KeyStone I Architecture Antenna Interface 2 (AIF2) User Guide* (SPRUGV7)
- *SerDes Implementation Guidelines for KeyStone I Devices Application Report* (SPRABC1)
- OBSAI Reference Point 4 Specification, Version 1.1 [1]
- OBSAI Reference Point 3 Specification, Version 4.1 [2]
- OBSAI Reference Point 2 Specification, Version 2.1 [3]
- OBSAI Reference Point 1 Specification, Version 2.1 [4]

- CPRI Specification, Version 4.1 [5]
- Electrical Specification (IEEE-802.3ae-2002), dated 2002 [9]

6.7.1 Configuration of AIF

The AIF defaults to disabled and with the internal memories for the module in a sleep state. The memories, followed by the module, must be enabled by software before use.

There are two protocol modes supported on the AIF interface: OBSAI and CPRI. The mode is selected by software after power-up. All links are the same mode.

The AIF module requires the AIF reference clock (SYSCLKP/N) to drive the SerDes PLLs and requires frame sync timing signals provided by the frame sync module. The frame sync clock provided to the FSM has the following requirements:

- RP1 mode:
  - RP1CLKP/N must be 30.72 MHz (8× UMTS chip rate)
  - RP1FBP/N must provide a UMTS frame boundary signal
- Non-RP1 mode:
  - RADSYNC must be between 1 ms and 10 ms
  - PHYSYNC must provide UMTS frame boundary pulse

For proper operation of the AIF, the SYSCLKP/N (which is the antenna interface SerDes reference clock) and the frame sync clock (either RP1CLKP/N or ALTFSYNCCLK) must be generated from the same clock source and must be assured not to drift relative to each other.

The AIF reference clock and the SerDes PLL multiplier are used to select the link rates. Both CPRI and OBSAI have three supported line rates that run at 2×, 4×, and 8× the base line rate. The SerDes line rates can be operated at half rate, quarter rate, or eighth rate of the PLL output. For that reason, it is suggested that the AIF SerDes PLL be run at the 8× line rate. Each link pair can be configured as half rate (8×), quarter rate (4×), or eighth rate (2×). Table 28 shows the possible clocking variations for the AIF SerDes.

Table 28 AIF SerDes Clocking Options

Protocol Mode	Reference Clock (MHz)	PLL Multiplier	8× (Gbps)	4× (Gbps)	2× (Gbps)
CPRI	122.88	20	4.9152	2.4576	1.2288
	153.60	16	4.9152	2.4576	1.2288
	307.20	8	4.9152	2.4576	1.2288
OBSAI	122.88	25	6.250	3.072	1.536
	153.60	20	6.250	3.072	1.536
	307.20	10	6.250	3.072	1.536
End of Table 28					

To ensure the required data throughput is present, a minimum core frequency of 1 GHz should be used if the antenna interface is configured for OBSAI 8× and 4× links. If an OBSAI 2× link is used, or if CPRI is used, the minimum core frequency is 800 MHz.

The AIF SerDes ports support hot-swap, where the AC-coupled inputs of the device can be driven without a supply voltage applied.

6.7.2 System Implementation of AIF

In OBSAI RP3 mode, the interface is electrically compatible with the *OBSAI RP3 Specification*, Version 4.1.

In CPRI mode, the interface is electrically compatible with the *XAUI Electrical Specification* (IEEE-802.3ae-2002).

For information regarding supported topologies and layout guidelines, see the *SerDes Implementation Guidelines for KeyStone I Devices Application Report* (SPRABC1).

Suggestions on AIF reference clocking solutions can be found in Section 3.2. AIF SerDes power planes and power filtering requirements are covered in Section 2.3.

Table 29 highlights the various uses of the sync timer inputs for the AIF module.

Table 29 AIF2 Timer Module Configuration Options

RAD Timer Sync	PHYTimer Sync	RP1Timer Sync	RP1 Timer Clock	Intended Use
		RP1FBP/N	RP1CLKP/N	RP1 or differential sync & clock
RADSYNC	PHYSYNC			Non-RP1 single-ended sync

6.7.3 Unused AIF Pin Requirement

All unused AIF transmit and receive lanes must be left floating and the unused lanes properly disabled in the MMR. If the entire AIF interface is not required, the AIF peripheral should be disabled in the MMR and all link pins can be left floating. The reference clock should be terminated as shown in Figure 13.

If the AIF interface is not required, the AIF regulator power pins (VDDR5_AIF1 and VDDR6_AIF2) must still be connected to the correct supply rail with the appropriate decoupling capacitance applied. The EMI/Noise filters are not required when this interface is not used.

If the LVDS inputs (RP1CLKP/N and RP1FBP/N) are not used, external connections should be provided to generate a valid logic level. The recommended connections for unused LVDS inputs are shown in Figure 25. The 1-kΩ resistor is used on the N pin to reduce power consumption, and the P pin should be connected to the DVDD18 supply rail.

Figure 25 Unused LVDS Inputs

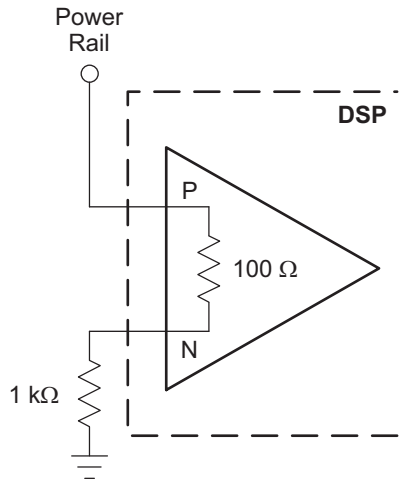

6.7.4 System Implementation of RP1

For an RP1-compliant interface, the frame sync clock and frame sync burst signal defined by the RP1 specification should be connected to RP1CLKP/N and RP1FBP/N, respectively.

LVDS 1:N buffers (such as the SN65LVDS108 or CDCL1810) can be used to connect these signals to multiple devices on a single board, however it is strongly recommended that the new high performance CDCM6208 low-jitter, low-power clock sources be used. The 100-Ω LVDS termination resistor is included in the KeyStone I device LVDS receiver, so an external 100-Ω resistor is not needed. These are standard LVDS inputs, so, unlike the LJCB inputs, these inputs should not be AC coupled. They should be driven directly by an LVDS-compliant driver.

The EXTFRAMEEVENT output is available to generate a frame sync output to other devices based on the frame clock and frame sync inputs. Its switching frequency and offset are programmable in the FSM. This output edge is not aligned with the frame clock input so take care if this output needs to be latched based on this clock.

For information on clocking distribution options for the single-ended frame sync clocks, see Section 10.4.

6.7.5 Unused RP1 Pin Requirement

If the RP1 interface is not used it is recommended that the RP1 clk input pins be connected in accordance with Figure 25. The EXTFRAMEEVENT output pin can be left floating.

6.8 DDR3

Relevant documentation for DDR3:

- *KeyStone Architecture DDR3 Memory Controller User's Guide* (SPRUGV8)
- *DDR3 Design Guide for KeyStone Devices Application Report* (SPRABI1)
- JEDEC JESD79-3C [11]

6.8.1 Configuration of DDR3

The DDR3 peripheral is enabled at power-up.

The DDR3 output clock is derived from the DDR3 PLL that uses DDRCLKP/N as a reference clock. The DDR3 PLL operates at 20× the DDRCLKP/N frequency and the DDR3 output clock is 1/2 of the PLL output clock. For example, a 66.667-MHz reference clock results in a 1,333-MHz PLL output and a DDR3 output clock of 667 MHz for DDR3-1333 support.

6.8.2 System Implementation of DDR3

For information regarding supported topologies and layout guidelines, see the *DDR3 Design Guide for KeyStone Devices Application Report* (SPRABI1).

Suggestions for DDR3 reference clocking solutions can be found in Section 3.2.

6.8.3 Slew Rate Control

The KeyStone I device incorporates two pins necessary to control the DDR3 slew rate. There are four possible slew rate combinations. The slew rate pins (DDRSLRATE1:0) must be pulled low or high (to 1.8V (DVDD18) at all times (they are not latched). Pulling both DDRSLRATE input pins low selects the fastest slew rate. If the DDRSLRATE pins are both pulled high, the resulting slew rate is the slowest. For normal full speed operation, the DDRSLRATE pins should be pulled low. Table 30 shows the possible combinations when controlling the slew rate for the DDR3 interface.

Table 30 Slew Rate Control

Setting	Speed	DDRSLRATE1	DDRSLRATE0
00	Fastest	pull-down	pull-down
01	Fast	pull-down	pull-up
10	Slow	pull-up	pull-down
11	Slowest	pull-up	pull-up

Note—See the DDR3 implementation guide for specific details. Slew rates denoted in the table are relative and are highly dependant upon topologies, component selection, and overall design implementation.

6.8.4 Unused DDR3 Pin Requirement

If the DDR3 peripheral is disabled, all interface signals (including reference clocks) can be left floating and the input buffers are powered down.

All unused DDR3 address pins must be left floating. All unused error correction pins must also be left floating. All remaining control lines must also be left floating when unused. Unused byte lanes must be properly disabled in the appropriate MMR.

If the DDR3 interface is not required, the DDR3 PLL supply (AVDDA2) must still be connected and the ferrite bead installed to minimize noise.

If unused, the primary DDR3 clock input must be configured as indicated in Section 3.3, Figure 13. Pull-up must be to CVDD core supply. The VREFSSTL pin must be connected to the VREF (0.75 V) supply regardless of whether the DDR3 interface is used or not used.

6.9 UART

Relevant documentation for the UART:

- *KeyStone Architecture Universal Asynchronous Receiver/Transmitter (UART) User Guide* (SPRUGP1)
- *TMS320C66x DSP CPU and Instruction Set Reference Guide* (SPRUGH7)
- *TMS320C66x CorePac User Guide* (SPRUGW0)

The UART peripheral performs serial-to-parallel conversion to data received from a peripheral device connecting to the DSP and parallel-to-serial data conversion to data connecting between the DSP and peripheral device.

Figure 26 shows the concept of peripheral-to-DSP connection with autoflow control support.

Figure 26 UART Connections with Autoflow Support

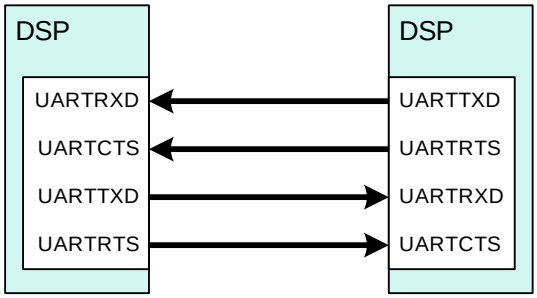

6.9.1 Configuration of the UART

The UART peripheral is based on the industry standard TL16C550 asynchronous communications element, which is a functional upgrade of the TL16C450. Key configuration details are described in the *KeyStone Architecture Universal Asynchronous Receiver/Transmitter (UART) User Guide* (SPRUGP1).

6.9.2 System Implementation of the UART

The UART interface is intended to operate at 1.8 V. Connections between the device UART interface and an external peripheral must be at a 1.8-V level to assure functionality and avoid damage to either the external peripheral or the KeyStone device.

6.9.3 Unused UART Pin Requirements

The UART interface to the device contains four pins (CTS, RTS, TXD, and RXD). These pins (defined with respect to the device) are LVCMOS, which when not used should be connected in the following manner:

Table 31 Unused UART Connections

Signal	Direction	If Unused	Internal Resistor
UARTRTS	Output	Leave Floating	Internal Pull-down
UARTCTS	Input	Leave Floating	Internal Pull-down*
UARTTXD	Output	Leave Floating	Internal Pull-down
UARTRXD	Input	Leave Floating	Internal Pull-down*

Note—*: The above recommendations are based on the peripheral being unused and no traces attached. If traces are attached to any of the device input pins, added pull-up (to DVDD18) or pull-down (to VSS) resistors (4.7 kΩ) are needed.

See the data manual and UART users guide for additional connectivity requirements.

6.10 Serial Port Interface (SPI)

Relevant documentation for the SPI Interface:

- *KeyStone Architecture Serial Peripheral Interface (SPI) User Guide* (SPRUGP2)
- *TMS320C66x CPU and Instruction Set Reference Guide* (SPRUGH7)
- *TMS320C66x CorePac User Guide* (SPRUGW0)

The SPI peripheral is a high-speed synchronous serial input/output port that allows a serial bit stream of programmed length (1 to 32 bits) to be shifted into and out of the device at a programmed bit-transfer rate. The device SPI interface can support multiple SPI slave devices (varies by device type and select/enable pins). The device SPI interface can operate as a master device only.

The SPI is normally used for communication between the device and common external peripherals. Typical applications include an interface to external I/O or peripheral expansion via devices such as shift registers, display drivers, SPI EEPROMs, and analog-to-digital converters.

The DSP must be connected only to SPI-compliant slave devices. Figure 27 shows the concept of a DSP to two-peripheral connection.

Figure 27 SPI Connections

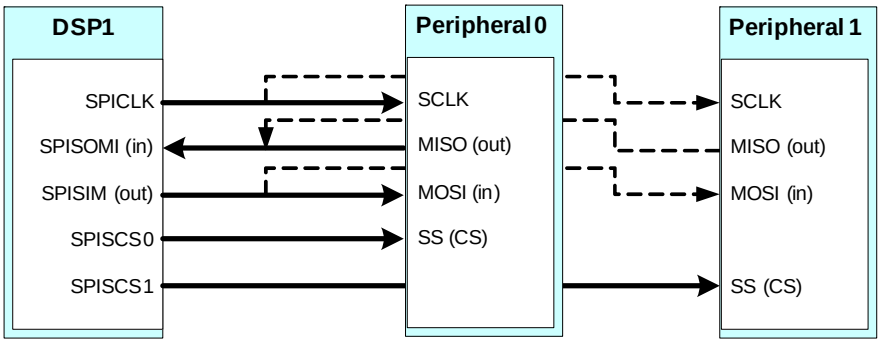

6.10.1 Configuration of the SPI

The SPI peripheral is required to be configured correctly prior to use for a proper bit stream transfer.

6.10.2 System Implementation of the SPI

The SPI interface is intended to operate at 1.8 V. Connection between the device SPI interface and peripheral must be at a 1.8-V level to assure functionality and avoid damage to either the external peripheral or KeyStone device.

Two key controls are necessary to make the SPI to function:

- Chip select control (SPISCS1:0)
- Clock polarity and phase

See the SPI Users Guide for specific and detailed configuration.

6.10.3 Unused SPI Pin Requirements

The SPI interface to the device contains five pins: SPISCS0, SPICS1, SPICLK, SPIDIN, and SPIDOUT. These pins (defined with respect to the device) are 1.8-V LVCMOS, which when not used, should be connected as shown in Table 32.

Table 32 Unused SPI Connections

Signal	Direction	If Unused	Internal Resistor
SPICS0	Output	Leave Floating	Internal Pull-up
SPICS1	Output	Leave Floating	Internal Pull-up
SPICLK	Output	Leave Floating	Internal Pull-down
SPIDOUT	Output	Leave Floating	Internal Pull-down
SPIDIN	Input	Leave Floating	Internal Pull-down*

Note—*: The above recommendations are based on the peripheral being unused and no traces attached. If traces are attached to any of the peripheral input pins, added pull-up (to DVDD18) or pull-down (to VSS) resistors (4.7 kΩ) are needed.

See the data manual and SPI users guide for additional connectivity requirements.

6.11 External Memory Interface 16 (EMIF16)

Documentation for EMIF16:

- *KeyStone Architecture External Memory Interface (EMIF16) User Guide* (SPRUGZ3)
- *KeyStone Architecture Bootloader User Guide* (SPRUGY5)
- BSDL Model for your device
- IBIS Model File for your device
- *Using IBIS Modes for Timing Analysis Application Report* (SPRA839)

6.11.1 Configuration of EMIF16

The EMIF16 peripheral can be disabled using boot strapping options, as defined in the KeyStone I Data Manual and the KeyStone I Bootloader User Guide. If the EMIF16 is enabled via boot strapping, it still needs to be enabled via software after a reset.

The clock to be used for EMIF16 can either be supplied externally to the AECLKIN pin or can be generated internally from the device core clock. The internally-generated clock is called SYSCLK6. After a reset, SYSCLK6 is device core clock / 6. The SYSCLK6 divider can be changed via software register accesses to the PLLDIV6 register. See the *KeyStone Architecture Phase-Locked Loop (PLL) User Guide* (SPRUGV2) for details.

One of the boot modes provides for booting using the EMIF16. In this mode, after reset, the device immediately begins executing from the base address of CS2 in 16-bit asynchronous mode. This would support booting from simple asynchronous memories such as NOR FLASH.

6.11.2 System Implementation of EMIF16

Generally, series resistors should be used on the EMIF16 signals to reduce overshoot and undershoot. They are strongly recommended for any output signal used as a synchronous latch. Acceptable values are 10 Ω, 22 Ω, or 33 Ω. To determine the optimum value, simulations using the IBIS models should be performed to check for signal integrity and AC timings. Significant signal degradation can occur when multiple devices are connected on the EMIF16. Simulations are the best mechanism for determining the best physical topologies and the highest frequency obtainable with that topology.

6.11.3 Unused EMIF16 Pin Requirements

If the EMIF16 peripheral is not used, the EMIF16 inputs can be left unconnected. Internal pull-up and pull-down resistors are included on this interface so leakage will be minimal. If only a portion of the peripheral is used, data pins that are not used can be left floating.

Note—The above recommendations are based on the peripheral being unused and no traces attached. If traces are attached to any of the peripheral input pins, added pull-up (to DVDD18) or pull-down (to VSS) resistors (4.7 kΩ) are needed.

6.12 Telecom Serial Interface Port (TSIP)

Documentation for TSIP:

- *KeyStone Architecture Telecom Serial Interface Port (TSIP) User Guide* (SPRUGY4)
- BSDL Model for your device
- IBIS Model File for your device
- *Using IBIS Models for Timing Analysis Application Report* (SPRA839)

6.12.1 TSIP Configuration

TSIP0 and TSIP1 are independent serial interface ports. Each one has two clock inputs, two frame sync inputs, eight serial data inputs, and eight serial data outputs. All interface timing is derived from the clock and frame sync inputs. Each TSIP module can be individually configured to operate at either 8.192 Mbps, 16.384 Mbps, or 32.768 Mbps. All eight lanes are used at 8.192 Mbps, only the first four lanes are used at 16.384 Mbps, and only the first two lanes are used at 32.768 Mbps.

TSIP outputs are configurable to simplify multiplexing of TDM streams. Unused output timeslots can be configured to drive high, drive low, or be high impedance. Low-speed links can use simple wire-OR multiplexing when unused timeslots are high impedance. Higher speed links can be combined with OR logic when unused timeslots are driven low or AND logic can combine output streams when unused timeslots are driven high.

6.12.2 TSIP System Implementation

Series resistors should be used on clock sources because the TSIP clock inputs are edge-sensitive. These resistors need to be placed near the signal source. This prevents glitches on the clock signal transitions that could potentially disrupt TSIP timing. Most implementations do not need series terminations on the other TSIP control and data lines as long as the traces are kept short and routed cleanly.

The maximum TSIP performance is achievable only when using point-to-point connections. Termination resistors are rarely needed in this topology. In a bused architecture in which several devices are connected to the same TSIP interface, the traces should be routed from device to device in a single, multi-drop route without branches. Then a single termination can be implemented at the end of the bus to reduce over-shoot and under-shoot, which minimizes EMI and crosstalk. This need is even more likely whenever a TSIP bus passes through board-to-board connectors. To determine the optimum value, simulations using the IBIS models should be performed to validate signal integrity and AC timings.

Multiple devices can be connected to a common TSIP bus using TDM mode. The additional loads require a reduction in the operating frequency. Also, the specific routing topology becomes much more significant as additional devices are included. The way to determine the best topology and maximum operating frequency is by performing IBIS simulations.

6.12.3 Unused TSIP Pin Requirements

If any of the TSIP inputs are not used, they can be left floating because all TSIP pins have internal pull-down resistors.

7 I/O Buffers and Termination

This section discusses buffer impedances, I/O timings, terminators, and signaling standards.

7.1 Process, Temperature, Voltage (PTV) Compensated Buffers

The impedance of I/O buffers is affected by process, temperature, and voltage. For the DDR3 interface, these impedance changes can impact performance and make it difficult to meet the JEDEC specifications. For this reason, the KeyStone I device uses PTV-compensated I/O buffers for the DDR3 interface. The PTV compensation works by adjusting internal impedances to nominal values based on an external referenced resistance. This is implemented by connecting a resistor between the PTV pin and VSS. The PTV resistor must be a1%-tolerance 45.3 Ω SMT resistor and connected between the pin and Vss (ground). The resistor should be placed adjacent to the KeyStone device, as close to the pin as possible. For details, see the device-specific data manual and the *DDR3 Design Requirements for KeyStone Devices Application Report* (SPRABI1).

7.2 I/O Timings

The I/O timings in the KeyStone I data manual are provided for, and based on, the tester test load. These timings need to be adjusted based on the actual board topologies. It is highly recommended that timing for all high speed interfaces (including the high-performance SerDes interfaces) on the KeyStone I design be checked using IBIS simulations (at a minimum). Simulating the high performance interfaces (SerDes) will require IBIS 5.0 AMI-compliant models. IBIS models and parasitics incorporated in the IBIS models are based on IBIS loads and not the tester load.

7.3 External Terminators

Series impedance is not always needed but is useful for some interfaces to avoid over-shoot/under-shoot problems. Check the recommendations in the peripherals sections and/or perform IBIS or Matlab® simulations on each interface.

7.4 Signaling Standards

This section discusses signaling standards for various interfaces.

7.4.1 1.8-V LVCMOS

All LVCMOS I/O (input and output or bidirectional) buffers are JEDEC-compliant 1.8-V LVCMOS I/Os as defined in JESD 8-5. Several types of LVCMOS buffers are used in KeyStone I devices, depending mainly on whether an internal pull-up or pull-down resistor is implemented. There also are some differences in the drive strength and impedance for some I/Os. For details on different LVCMOS buffer types, see the KeyStone Data Manual.

These internal pull-up and pull-down resistors can be can be considered a 100 µA current source (with an actual range of 45 µA to 170 µA). This equates to a nominal pull-up/pull-down resistor of 18 kΩ (with a possible range of 10 kΩ to 42 kΩ) The 1.8-V LVCMOS interfaces are ***NOT*** 2.5-V or 3.3-V tolerant, so connections to 2.5-V or 3.3-V CMOS logic require voltage translation.

For input buffers operating at less than 10 MHz, TI's LVC logic family can be operated at 1.8 V and is 3.3-V tolerant. For faster signaling, TI's AUC family is optimized to operate at 1.8 V and is also 3.3-V tolerant. A good option for voltage translation for 1.8 V outputs that need to drive 2.5-V or higher inputs would be the CBT3245A. Some useful TI application reports on voltage translation options are:

- *Flexible Voltage-Level Translation with CBT Family Devices* (SCDA006)
- *Selecting the Right Level-Translation Solution Application Report* (SCEA035)
- Voltage Level Translation Product Portfolio

7.4.2 1.5V SSTL

The KeyStone I DDR3 interface is compatible with the *JEDEC 79-3C DDR3 Specification* [11]. The I/O buffers are optimized for use with direct connections to up to nine DDR3 SDRAMs and can be configured to drive a DDR3 UDIMM module. The DDR3 SSTL interface does not require series terminations on data lines (end terminations may be required). The 1.5-V SSTL DDR3 interface also supports active ODT (on-die-terminations).

7.4.3 SerDes Interfaces

There are many SerDes peripherals on the KeyStone I devices. The SerDes interfaces available include (others may apply): HyperLink, SRIO, PCI, SGMII, and AIF. These serial interfaces all use 8b/10b-encoded links and SerDes macros.

Note—Not all SerDes peripherals are available on all KeyStone I devices. See the device-specific data manual for details.

These interfaces use a clock recovery mechanism so that a separate clock is not needed. Each link is a serial stream with an embedded clock so there are no AC timings or drive strengths as found in the LVCMOS or SSTL interfaces.

There are several programmable settings for each SerDes interface that affect the electrical signaling. The most important of these are: transmitter output amplitude, transmitter de-emphasis, and receiver adaptive equalization. Recommendations for these settings for particular board topologies are provided in *SerDes Implementation Guidelines for KeyStone I Devices Application Report* (SPRABC1). The SerDes interfaces use CML logic. Compatibility to LVDS signals is possible and is described in Section 7.

7.4.4 LJCB Differential Clock Inputs

The reference clock inputs are connected to a Low Jitter Clock Buffer (LJCB) in KeyStone I devices. This buffer is designed to interface with LVDS, LVPECL, and HCSL clock input levels using AC coupling capacitors. The LJCB has internal termination and an internal common mode voltage generator so no external termination is needed between the AC coupling capacitors and the KeyStone I device. The driver selected for your design may require termination on the output. You should review the data manual for the driver to determine what termination is necessary.

Figure 28 LJCB Clock Inputs

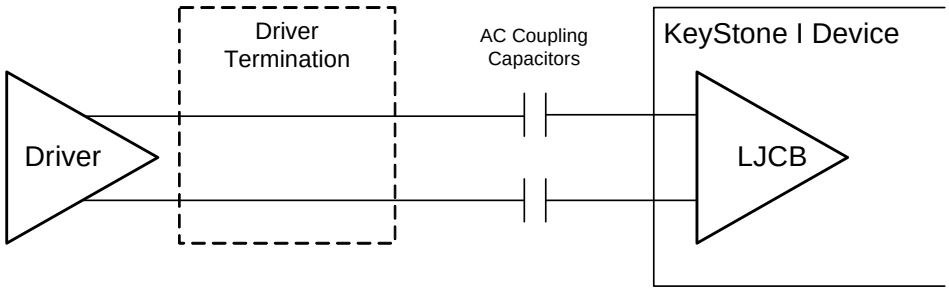

7.5 General Termination Details

The use of terminations in various nets is designed to address one of three major issues:

- Signal overshoot
- Signal undershoot
- Fast slew rates

Depending on the nature of the signal, if an overshoot or undershoot were to occur, the resulting energy induced may be catastrophic to the connected device (or DSP). Many standards – including SDRAM standards – have overshoot and undershoot requirements that must be met to ensure device reliability. The DSP and other devices also incorporate both source and sync current limitations as well as V_{IL}/V_{IH} & V_{OL}/V_{OH} requirements that must be maintained. Many of these established limitations are controlled by using terminations.

Clocking terminations many times require special consideration (AC or DC termination styles). In all cases, proper positioning of the termination and selection of components is critical.

Terminations are also valuable tools when trying to manage reflections. Adding in a series termination on a particular net can relocate the point of reflection outside of the switching region (although this is better managed by physical placement).

It is always recommended that single-ended clock line, data, and control lines be modeled to establish the optimum location on the net.

To determine if a termination should be used (data, control, or single-ended clock traces) the following test should be used:

- If the uni-directional (one way) propagation delay for the trace(s) in question is ≤ 20% of the rise/fall time (whichever is fastest) or ≥ 80% of the rise/fall time (whichever is slowest) then a termination should not be needed.

Further determination of whether or not a termination is required can be ascertained using the second condition:

- A transmission line requires termination (to its intended impedance) if the uni-directional Tpd (propagation delay) for the respective net is ≥ 50% of the respective rise or fall time fall time (fastest of the two).

The most common of terminations are for single-ended nets and follow the basic design shown in Figure 29:

Figure 29 Basic Series Termination

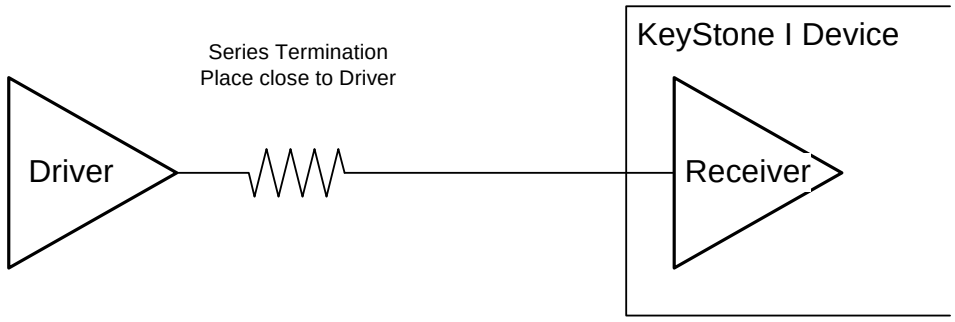

8 Power Saving Modes

The KeyStone I device incorporates advanced power saving methods that can be used when power consumption is a concern. These methods range from disabling particular peripherals to placing the device into a hibernation state. This section describes the low-power methods and hibernation states, and their limitations and configuration methods.

Relevant documentation for Power Savings:

- Device-specific data manual
- *KeyStone Architecture Power Sleep Controller (PSC) User Guide* (SPRUGV4B)

The KeyStone I family of devices has been designed with four basic modes of operation:

- **Off** - Power off, no power supplied to the device at all
- **Active** - Power on, out of reset, and fully functional
- **Standby** - Power on, not in a fully active state, lower power consumption than Active, greater power consumption than in hibernation modes
- **Hibernation** - Two specific operational modes intended to significantly reduce power consumption (includes level 1 & 2).

See the following sections, the device-specific data manual, and respective application reports for additional details on configuration and application of different power saving and operational modes.

8.1 Standby Mode

The KeyStone I family of devices includes a standby mode. This mode of operation is intended to provide maximum readiness from the device. This mode differs from additional power saving modes identified in Section 8.2 in that standby mode does not incorporate all the advanced power savings features associated with other modes.

There are several different combinations of standby mode. Different combinations would take into account such interfaces as SRIO, Ethernet, PCIe, and HyperLink.

During standby mode, the IOs remain in their current state. The IOs are not disabled or placed into reset state.

In standby mode, the status of each peripheral (if available) is indicated in Table 33. See the device-specific data manual for detailed descriptions on power saving modes.

Table 33 Standby Peripheral Status (Part 1 of 2)

Peripheral	Status	Notes
Debug Interface	ON	
PCIe	ON	
DDR3	ON	
SRIO	ON	
AIF	ON	
Accelerator(s)	ON	See data manual for applicable accelerators
HyperLink	ACTIVE	
FFTC	OFF	
Cores	INACTIVE	See application reports and data manual for core status during standby mode

Table 33 Standby Peripheral Status (Part 2 of 2)

Peripheral	Status	Notes
SGMII	ACTIVE	
EMIF	ACTIVE	
End of Table 33		

8.2 Hibernation Modes

The KeyStone I device incorporates several hibernation modes that allow power consumption to be reduced during operation when only partial functionality is required (periods of inactivity). The benefits of using a hibernation mode include faster wake up times (as compared to cold boot) in addition to the typical power savings obtained.

Access to the hibernation modes is obtained through writable registers. See the data manual for the memory map.

Each KeyStone I device incorporates different hibernation modes. See the specific data manual and application reports for additional details.

8.2.1 Hibernation Mode 1

During hibernation mode 1, only the MSMC SRAM memory content is retained. The MSMC MMR is not retained. Prior to entering this specific hibernation mode, the bootcfg MMR PWRSTATECTL must be properly configured and set (see the data manual for the procedure for entering and exiting this hibernation mode).

Exiting hibernation mode 1 requires a chip level reset, which is initiated using an external reset pin, usually the $\overline{\text{RESET}}$. The estimated wake-up time from this hibernation mode is less than 100 ms.

8.2.2 Hibernation Mode 2

During hibernation mode 2, the MSMC SRAM memory contents are lost and only information stored in the DDR3 SDRAM is retained. The MSMC MMR is not retained. After a chip level reset is asserted, the MSMC is reset to its default settings.

Exiting hibernation mode 2 requires a chip level reset, which is initiated using an external reset pin, usually the $\overline{\text{RESET}}$. The estimated wake up time from this hibernation mode is greater than 1.0 second.

Table 34 shows, at a high level, one possible combination of hibernation mode configurations – see the data manual and respective application report for additional information.

Table 34 Hibernation Mode Peripheral Status (Part 1 of 2)

Peripheral	Mode 1 Status	Mode 2 Status	Notes
Debug Interface	OFF	OFF	
PCIe	ON	ON	
DDR3	ON	ON	
MSMC SRAM	ON	OFF	
SRIO	ON	ON	
AIF	ON	ON	Can be multiple configurations

Table 34 Hibernation Mode Peripheral Status (Part 2 of 2)

Peripheral	Mode 1 Status	Mode 2 Status	Notes
Accelerator(s)	ON	OFF	See the data manual for applicable accelerators
HyperLink	OFF	OFF	
FFTC	OFF	OFF	
Cores	INACTIVE	INACTIVE	See the application report and data manual for core status during stand by mode
SGMII	ACTIVE	ACTIVE	Can be multiple configurations
EMIF	ACTIVE	ACTIVE	
End of Table 34			

8.3 General Power Saving & Design Techniques

There exist multiple additional low power techniques that can be implemented in a design (depending on the use case). If the application requires multiple devices and can accommodate partitioning such that certain device can handle the bulk of the workload across defined periods of time, then the use of hibernation modes (identified above) is recommended.

In other use cases, entire sections of a board may be capable of powering down – making the application more "green." However, care must be taken not to drive an IO that does not have power.

Selection of components including pull-up and pull-down resistor values can reduce (or increase) current requirements on the power supplies. Power supply efficiency can be optimized through proper design and implementation.

The following list provides a general overview of typical power saving techniques. Not all techniques apply for all applications. Consult the KeyStone I device data manual for additional information.

- Generally, unused I/O pins should be configured as outputs and driven low to reduce ground bounce.
- Keep ground leads short, preferably tie individual ground pins directly to the respective ground plane.
- Eliminate pull-up resistors (where possible), or use pull-down resistors (where possible).
- The use of multi-layer PCBs that accommodate multiple separate power and ground planes take advantage of the intrinsic capacitance of GND-VCC plane topology and thus (when designed correctly) help to establish a better impedance board while using the intrinsic capacitance inherent in the stack-up design. A better-matched impedance board will prevent over and under shoots and / or reflections.
- Where possible, create synchronous designs that are not affected by momentarily switching pins.
- Keep traces connecting power pins stretching from power pins to a power plane (or island, or a decoupling capacitor) should be as wide and as short as possible. This reduces series inductance, and therefore, reduces transient voltage drops from the power plane to the power pin. Thus, reducing the possibility of ground bounce and coupled noise.

- Use high quality surface-mount low series resistance (ESR) capacitors to minimize the lead inductance. The capacitors should have an ESR value as small as possible. Ceramic capacitor should be used where possible.
- Use separate power planes for analog (PLL) power supplies (where possible).
- All PLL power supply planes should have a ground plane directly next to them in the stackup to minimize all power-generated noise.
- Place analog or digital components only over their respective ground planes.
- Avoid sharing vias. Connect each ground pin to the ground plane individually.
- The use of the recommended filters will aid in isolating the PLL power from the digital power rails.
- Place as many bypass or decoupling capacitors as possible (minimum is the recommended amount), these should be placed between CVDD and ground and be as close as possible. Ideally, the respective CVDD pin should be tied directly to the decoupling capacitor, which, in turn, should be tied directly to ground. This provides the shortest and lowest inductance path possible. If there must be a single trace connecting the decoupling capacitor to ground it should be between the CVDD pin and the decoupling capacitor. If this is the case, then the trace between the device and decoupling capacitor must be as wide as possible.
- Sequence devices appropriately – especially during power up and power down. This will minimize crowbar and peak currents, which, in some cases, could be a 50% increase over the initial design.
- Use of larger diameter vias to connect the capacitor pads between power and ground planes aid in minimizing the overall inductance.

9 Mechanical

9.1 Ball Grid Array (BGA) Layout Guidelines

The BGA footprint and pin escapes can be laid out as defined in *Flip Chip BGA Users Guide* (SPRU811) [15]. If the DDR3 interface is used, there are specific recommendations for the BGA pad and pin escape vias given in *DDR3 Design Requirements for KeyStone Devices Application Report* (SPRABI1).

Given the 0.80 mm-pitch of KeyStone I devices, it is recommended that non-solder-mask-defined (NSMD) PCB pads be used for mounting the device to the board. With the NSMD method, the opening in the solder mask is slightly larger than the copper pad, providing a small clearance between the edge of the solder mask and the pad. (Figure 30).

Figure 30 Non-Solder-Mask Defined (NSMD) PCB Land

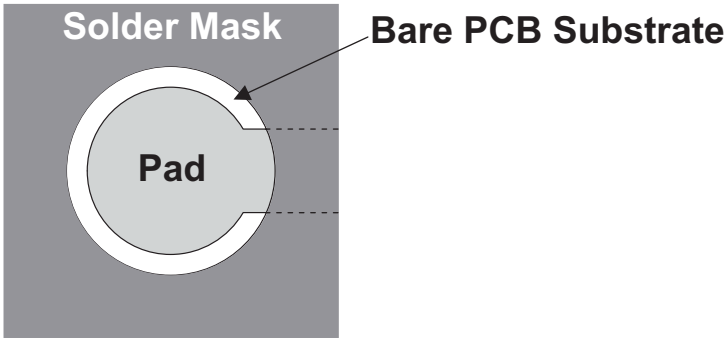

While the size control is dependent on copper etching and is not as accurate as the solder mask defined method, the overall pattern registration is dependent on the copper artwork, which is quite accurate. The trade-off is between accurate dot placement and accurate dot size. NSMD lands are recommended for small-pitch BGA packages because more space is left between the copper lands for signal traces. Dimensioning for the pad and mask are provided in the *Flip Chip BGA Users Guide* (SPRU811).

9.2 Thermal Considerations

Understanding the thermal requirements for the KeyStone I device along with proper implementation of a functional thermal management scheme is necessary to assure continued device performance and reliability. Each application, platform, or system must be designed to ensure that the maximum case temperature of the device never exceeds the allowable limit. See the device-specific data manual for thermal requirements, limitations, and operating ranges. Proper heat dispersion can be obtained using a variety of methods, including heat sinks.

Note—Improper thermal design that results in thermal conditions outside the allowable range will decrease device reliability and may cause premature failure.

9.2.1 Heat Transfer

There exist several methods for meeting the thermal requirements of the device. The two most common methods are conduction and convection.

9.2.1.1 Conduction

Conduction heat transfer refers to the conduction or direct contact between two surfaces resulting in a thermal transfer from one surface (higher temperature) to another (lower temperature). When referring to the cooling of the device in an application/system, this usually involves the direct contact of a heat sink (attached or possibly part of the system enclosure). The heat generated by the device is transferred to an attached heat sink/outer enclosure, thereby reducing the heat (conducting it away) in the device. In all conduction methods, the material used, compression forces, and ambient temperatures affect the transfer rates.

9.2.1.2 Convection

Convection refers to the thermal transfer of heat from one object (e.g., DSP) to the ambient environment typically by means of air movement (e.g., fan). Convection or air movement can be pulled or pushed across the device depending on the application. If a fan is attached to the device the direction is usually to force the air away from the device where as if the system or application board includes external cooling fans, the air direction can be either direction (push or pull). Air movement in any enclosure involves a plenum whether intended or not. It is always best to optimize the air movement to maximize the effects and minimize any excessive back pressure on the fan (to maximize fan life and minimize noise).

9.3 KeyStone I Power Consumption

The KeyStone I processor can consume varying amounts of power depending heavily on usage, implementation, topology, component selection, and process variation. See the device-specific TI power application brief for early power estimations, and the respective power application note and spreadsheet for more detailed configurations and more accurate power figures.

Note—All numbers provided (in product briefs and application reports) assume properly designed power supplies and proper implementation of the recommended variable core (SmartReflex) supply.

9.4 System Thermal Analysis

For an overview on performing a thermal analysis and suggestions on system level thermal solutions, see the *Thermal Design Guide for KeyStone Devices Design Guide* (SPRABI3).

9.5 Mechanical Compression

Mechanical compression, especially where conduction cooling is concerned (Section 9.2.1.1), is important. Excessive compression may improve thermal transfer but also increases the risk of damage to the device or inducing added electrical shorts between BGA balls on the application hardware. Excessive mechanical compression is typically noted when using BGA sockets or interfacing a heat sink/enclosure in direct contact with the device. See the data manual and all relevant application reports regarding mechanical compression and BGA assembly.

The following tables are provided as assistance for applications in which conductive cooling is used and a mechanical compression heat sink is incorporated. The following data assumes an absolute maximum solder ball collapse of 155 µm, 0.8 mm solder ball pitch, 841 pins, and uniform pressure across the entire device lid, a 23 mm × 23 mm lid, and proper device assembly to the PCB.

Table 35 Maximum Device Compression - Leaded Solder Balls

	Maximum Lid Pressure for Leaded Ball - 0.8mm pitch (841 ball package)			
TEMP	5 years	10 years	15 years	20 years
70° C	108.67 psi	90.59 psi	81.57 psi	75.69 psi
90° C	73.20 psi	61.00 psi	54.90 psi	50.83 psi
100° C	61.00 psi	75.69 psi	45.64 psi	42.47 psi

Table 36 Maximum Device Compression - Lead-Free Solder Balls

	Maximum Lid Pressure for Lead Free Ball - 0.8mm pitch (841 ball package)			
TEMP	5 years	10 years	15 years	20 years
70° C	544.47 psi	447.32 psi	397.62 psi	365.99 psi
90° C	442.80 psi	363.73 psi	323.06 psi	298.21 psi
100° C	402.14 psi	329.84 psi	293.70 psi	271.10 psi

Note—The values provided in the above two tables are *Not To Exceed* estimates based on modeling and calculations. Added safety margin should be included to account for variations in PCB planarity, solder mask, and thermal conductive material between the DSP and compression heat sink.

10 Routing Guidelines

This section provides routing guidelines for designs using a KeyStone I device. The layout of the PCB is critical for the proper operation of the device. Some of the information needed to complete a proper layout of KeyStone I interfaces will be found in other documents, specifically the DDR3 interface and the SerDes interfaces. In addition, other components on the board, for example power supplies and Ethernet PHYs, may have additional guidelines in their documentation. It is important to review all the layout guidelines for all the parts in your design.

10.1 Routing Power Planes

Power supply planes should be used to distribute the voltage from a power supply to the associated power supply pins on the Keystone I device and to the bulk and decoupling capacitors. In addition, power planes should be used to route filtered power supply signals such as the AVDDAx and VDDRx. Each of these filtered voltages are typically connected to only one pin on the KeyStone I device, so small cut-out planes may be used to connect the filter to the local decoupling and the power pin.

Each core and I/O supply voltage regulator should be located close to the device (or device array) to minimize inductance and resistance in the power delivery path. Additional requirements include a separate power plane for the core, I/O, and ground rails, all bypassed with high-quality low-ESL/ESR capacitors. See specific power supply data sheets for power plane requirements and support of specific power pads on each power supply.

Because each KeyStone I device has its own AVS supply, the power plane for that supply is usually localized to the area of the device. And because the KeyStone device may draw a significant amount of current, the resistivity of the plane must be kept low to avoid a voltage differential between the power supply and the connections on the device. This voltage differential is minimized by making the copper planes thicker or larger in area. Be sure to consider both the power planes and the ground planes. The resistance of the planes is determined by the following formula:

$$R = \rho * \text{length}/(\text{width} * \text{thickness})$$

where ρ is the resistivity of copper equal to 1.72E-8 Ω-meters.

PCB layer thickness is normally stated in ounces. One ounce of copper is about 0.012 inches or 30.5E-6 meters thick. The width must be de-rated to account for vias and other obstructions. A 50-mm strip of 1-oz copper plane de-rated 50% for vias has a resistance of 0.57 mΩ per inch.

The same calculation should be made for all power planes to ensure a proper current path from the supply to the device is present.

It is highly recommended that 1-oz copper or greater be used for planes used to distribute the supply voltages and for ground planes in KeyStone I designs.

10.1.1 Routing Filtered Power Planes

As stated above, the filtered voltage must be routed with a sufficient current path to avoid voltage droop and to minimize noise coupling. Because the pins for the filtered voltages are generally towards the center of the part, a common mistake is to route this voltage on a 4-mil trace similar to those used for signal traces. Figure 31 shows the bottom layer and one of the inner layers of a PCB layout for a KeyStone I device.

Figure 31 Filtered Voltage Plane

The connections to the ferrite and the local decoupling is highlighted in the lower right and the pin for the KeyStone I device is shown on the upper left. The outline delineates the cut-out on the inner layer creating the small local plane associated with this filtered voltage. The same calculation shown above can be used to calculate the resistivity of the plane.

10.2 Clock Signal Routing

10.2.1 Clock & High Performance Transmission Lines - Microstrip versus Stripline

Clock nets (and most other high-performance nets) are typically designed for either a microstrip or stripline topology. Figure 32 shows the difference between a microstrip and a stripline topology.

Figure 32 Stripline and Microstrip Topologies

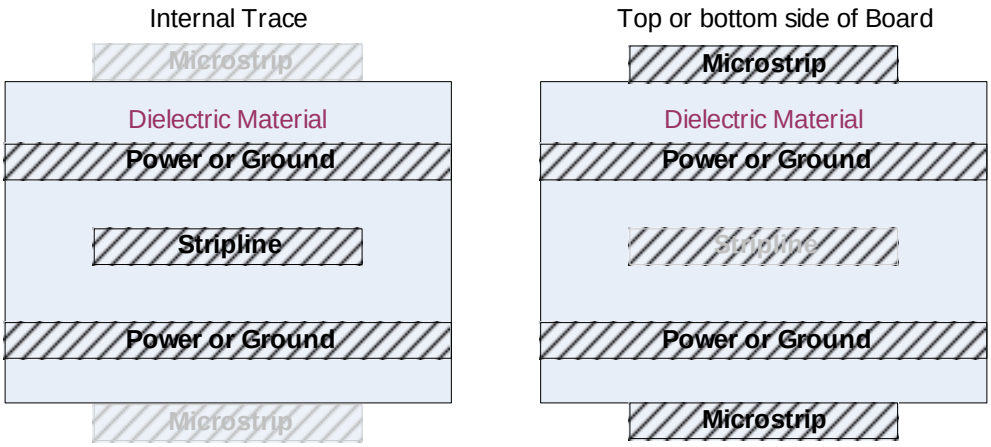

The most obvious difference between these two topologies is the location of the respective net (trace). In a stripline topology (left side) the net or trace is embedded in the PCB and usually between ground or power layers. In the microstrip topology (right side), the net or trace is on an outer layer and adjacent to a power or ground plane.

There are several reasons why one (stripline or microstrip) topology would be selected over the other, key factors determining which topology is used include:

- **Interconnection** - The number of pins on the interconnecting devices may limit the width of traces on the outer layers.
- **Performance** - High speed signals should not be routed on multiple layers, they should generally be routed on the top layers.
- **EMI** - Where emissions or the susceptibility to spurious radiation may impact signal integrity, stripline topologies are recommended.
- **Timing** - Propagation delays differ between microstrip and stripline topologies, the use of either must be comprehended for any and all timing-critical nets.
- **Routing** - The routing of high-speed clock signals is critical to a functional design. Proper return current paths to minimize switching noise is essential. Routing clocks that are > 250 MHz should be done using only a microstrip topology as long as they are < 2 inches (50.8 mm) in length.

A detailed discussion of transmission line routing for clocks and critical signals can be found in Appendix B.

10.2.2 Clock Routing Rules

As a general rule, the following routing guidelines should be followed for all clock signals:

- All high-performance clock signals should be routed on the same layer where possible. If not possible, then the impact of multiple layers on performance, propagation delays, and signal integrity must be taken into account.
- Clock nets, if routed internally, must be captured between two planes. Do not have two clock routing layers adjacent to one another.
- Clock nets involved in synchronous communications must have identical number of vias and be skew-matched within 5% of the total clock period.
- Vias and terminations (if required) must be positioned in locations such that reflections do not impact signal quality or induce an unwanted change of state.
- The target impedance for a single-ended clock net is 50 Ω and the PCB impedance must be ± 5% or 47.5 Ω to 52.5 Ω.
- The target impedance for a differential ended clock nets is 100 Ω and the differential PCB impedance must be ± 5% or 95.0 Ω - 105 Ω.
- All microstrip clock lines must have a ground plane directly adjacent to the clock trace.
- All stripline clock lines must have a ground or power planes directly adjacent (top & bottom) to the clock trace.
- Terminate any single-ended clock signals.
- Confirm whether or not the differential clock sources require terminations.
- A ground plane must be located directly below all clock sources (oscillators, crystals).
- No digital signals must be routed underneath clock sources.
- Clock nets (source to target) must be as short as possible – emphasis on reflections.
- Complementary differential clock nets must be routed with no more then two vias.
- Escape vias from the device or clock source must be of either via-in-pad or dog-bone type design. Dog-bone designed escapes must be short enough to ensure any reflections do not impact signal quality.
- The number of vias used in complementary differential clock pairs should be identical.
- Differential clock signals (complementary nets) must be skew-matched to within a maximum 5 pS of one another.

10.3 General Routing Guidelines

10.3.1 General Trace Width Requirements

Trace widths are largely dependent on frequency and parasitic coupling requirements for adjacent nets within the application design. As a general rule of thumb, and given the complexity of many markets using these high-performance devices, a typical trace width is on the order of 4 mils or 0.1016 mm. Trace widths of this size typically allow for better routing and improved routing densities.

Spacing to other traces is a function of the stack-up impedance as well as induced isolation to minimize (eliminate) crosstalk. Most single-ended nets have a width of 4 mils (0.1016 mm) and are spaced a minimum of 1.5× the trace width from all other traces (2× the trace width is recommended).

10.3.2 General Routing Rules

- Single-ended clock signals should have additional spacing to other nets to avoid crosstalk or coupling.
- Optimal performance should require trace widths to be no closer then 3× the dielectric height (between trace and adjacent ground or power plane).
- The distance between complementary differential nets must be as close as possible (typically 1× the trace width).
- Complementary differential nets must be routed in parallel for the entire length (except escapes) with identical spacing between each complementary net.
- The trace spacing between different differential pairs must be a minimum of 2× the trace width or 2× the spacing between the complementary parasitic differential pairs.
- Do not incorporate split power/ground planes – all planes should be solid.
- Nets routed as power must be of sufficient size to accommodate the intended power plus unintended peak power requirements.
- Keep high-speed signals away from the edges of the PCB. If necessary, *pin* the edges of the PCB to prevent EMI emissions.
- Organize all signals into net classes prior to routing.
- All clock trace routes must be as straight as possible – minimize or eliminate serpentine routing.
- Eliminate all right angle traces. Chamfered traces (if needed) are recommended.
- Determine all constraints prior to routing respective net classes (skew, propagation delays, etc.)
- Optimize all single-ended nets, minimize lengths where possible (as long as they do not violate any net class requirements or violate timing conditions).
- Place decoupling capacitors as close to the target device as possible.
- Place bulk capacitors in close proximity to their respective power supply.
- Use the largest possible vias for connecting decoupling and bulk capacitors to their respective power and ground planes – if traces to connect the decoupling capacitor to planes are used, then they must be short and wide.
- Route single-ended nets perpendicular or orthogonal to other single-ended nets to prevent coupling.
- All transmission line conductors should be adjacent to its reference plane.
- All clocks should be routed on a single layer.
- Nets within a particular net class should have a matched number of vias if used (vias are not recommended if possible).
- Vias and terminations (if required) shall be positioned in locations such that reflections do not impact signal quality or induce an unwanted change of state.
- All microstrip nets shall have a ground plane directly adjacent to the respective trace.

- All stripline nets should have a ground or power plane directly adjacent (top & bottom) to the respective net where possible. If it is not possible to enclose the respective net within a ground/ground or ground/power stack-up, then at least one return path must be provided for proper operation.
- Vias and escape vias to and from the device shall be taken into account with regards to timing, reflections, loading, etc.
- It is always recommended that all net classes be modeled for optimal performance.
- Each power or ground pin should be tied to its respective plane individually.
- AC coupling capacitor as well as DC resistor termination placement shall take into account associated parasitics, coupling capacitance, and multipoint reflections. See Section 10.6 and e Appendix A for additional details.

10.4 DDR3 SDRAM Routing

The routing of the DDR3 SDRAM interface is crucial to the successful function of that interface. A detailed description of the routing requirements for DDR3 PCB layout can be found in the *DDR3 Design Requirements for KeyStone Devices Application Report* (SPRABI1). These requirements must be carefully followed to ensure proper operation. Failure to comply with the requirements documented will result in an interface that is nonfunctional even at lower clock rates.

10.5 SerDes Routing

The routing of each SerDes interface is crucial to the successful function of that interface. A detailed description of the routing requirements for SerDes PCB layout can be found in the *SerDes Implementation Guide for KeyStone I Devices Application Report* (SPRABC1). Each SerDes interface has specific requirements that must be met. This document describes the routing for the AIF, SRIO, SGMII, HyperLink, and PCIe interfaces. Failure to comply with the requirements documented will result in an interface that is nonfunctional even at lower clock rates.

10.6 PCB Material Selection

Proper pcb (printed circuit board) material [also referred to as pwb (printed wiring board) has a large impact on the impedance of the stack up and even more impact on the propagation delays of signals.

Table 40 shows the approximate dielectric constant (Er) differences for common pwb materials used today. Selection of material is paramount to a successful design. Special consideration must be taken if routing signals on different layers where timing is concerned (except where not allowed by this design guide).

10.7 PCB Copper Weight

Proper copper thickness (referred to as planes) is important to the overall impedance, thermal management, and current carrying capabilities of the design. During the examination and review of all application hardware it is important to consider each of these variables. **It is highly recommended that 1 oz copper or greater be used for planes used to distribute the supply voltages and for ground planes in KeyStone I designs.**

As a general rule of thumb, the inner layers or planes of most PCBs are typically 1-oz copper while the outer layers are 0.5-oz copper. Outer layers receive solder masking, which in the finished product, is usually similar to the 1-oz inner copper weight. Table 37 shows basic information regarding copper weight, thickness, and current carrying capabilities (other factors can positively or negatively affect these specifications).

Table 37 Copper Weight and Thickness

Cu Weight	Cu Thickness (mils)	Cu Thickness (mm)
2.0 oz	2.8	0.07112
1.0 oz	1.4	0.03556
½ oz	0.7	0.01778

Figure 33 shows the relationship between trace width, copper thickness (1 oz. shown), and current carrying capability.

Figure 33 Trace Width to Current Specifications

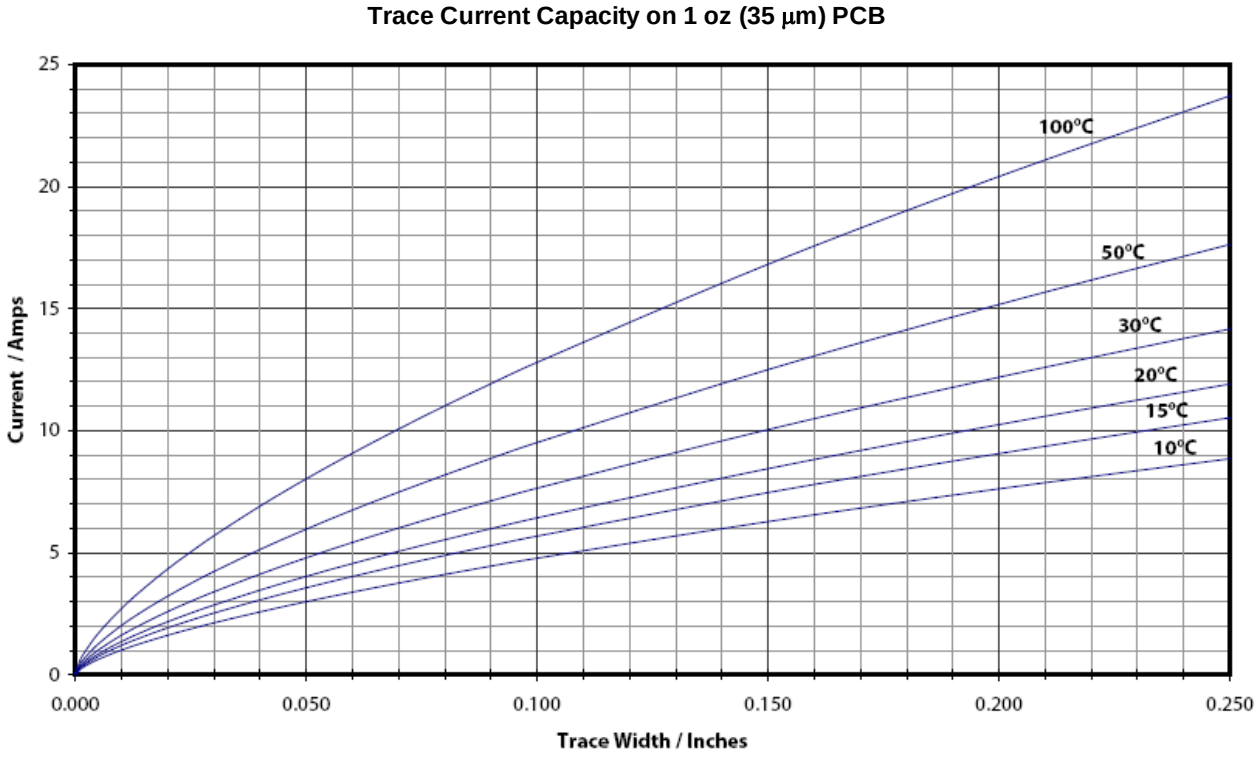

11 Simulation & Modeling

This section contains specific details pertaining to simulation and modeling.

Relevant documentation for Simulation and Modeling:

- KeyStone I data manual
- *Using IBIS Models for Timing Analysis Application Report* (SPRA839)

11.1 IBIS Modeling

Texas Instruments offers a fully-validated standard IBIS model (IBIS 5.0 compliant) inclusive of all SerDes IO buffers. The available model allows for full simulation using industry-standard plug-in tools to such programs as Hyperlynx®. Because of the high-performance nature of many of the KeyStone I IOs, the traditional IBIS model will also include embedded IBIS AMI models.

11.1.1 IBIS AMI Modeling

Given the performance requirements of many of today's input/output buffers, current IBIS models (IBIS 4.x compliant) are unable to properly characterize and produce results for many IOs operating above several hundred megahertz. For this reason, the IBIS standard (4.x and previous) was modified (new revision 5.0 released Aug. 2008) incorporating the new AMI model.

AMI models (Algorithmic Modeling Interface) are black boxes that incorporate a unique parameter definition file that allows for specific data to be interchanged between EDA toolsets and standard models.

Texas Instruments is in the process of developing specific IBIS AMI models for the high-performance SerDes IOs and will be supported and released in the near future.

12 Appendix A

This section contains additional and supporting information to be used during the design and integration of TI's KeyStone I devices.

12.1 Phase Noise Plots

TI recommends specific clock types to be used in order to meet all clocking internal requirements. The following phase noise plots provided are for the CDCE62002 and CDCE62005 alternate clock sources. These clock sources are currently available for purchase and meet the needs of the device.

Figure 34 shows the phase noise data (plot) at 40 MHz. It is representative of both the CDC362002 and CDCE62005. This input frequency would be coupled to the DDR3 clock input of the device (provided the SDRAM were rated for 1.6 Gbps and the PLL multiplier was configured correctly).

Figure 34 Phase Noise Plot - CECE62002/5 @ 40 MHz

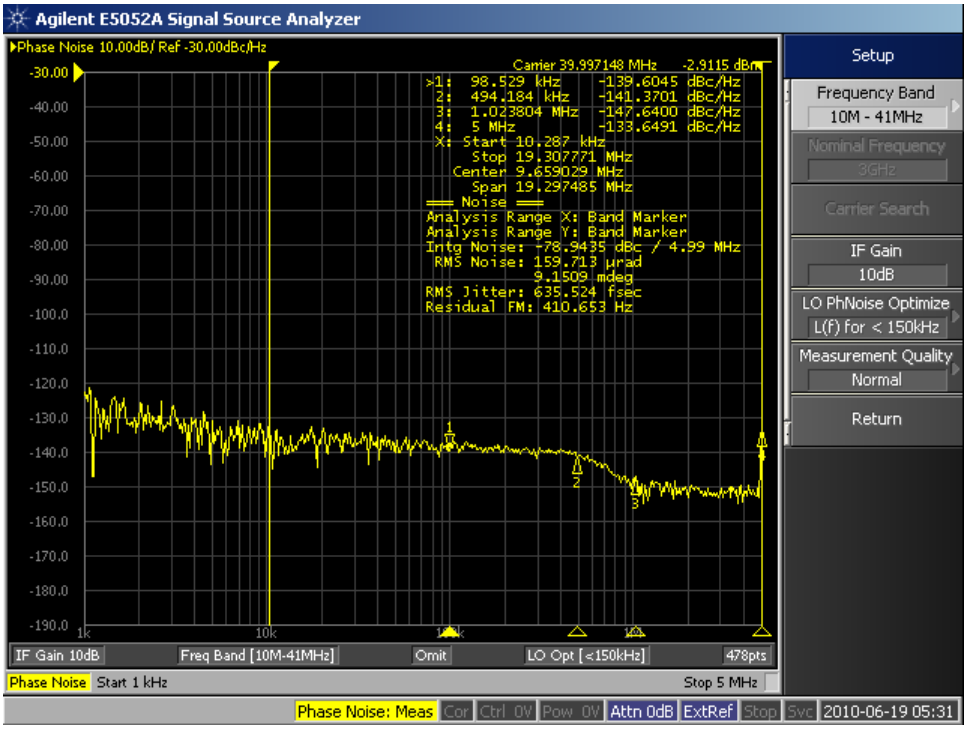

Figure 35 shows the phase noise data (plot) at 66.667 MHz. It is representative of both the CDC362002 and CDCE62005. This input frequency would be coupled to the DDR3 clock input of the device (provided the SDRAM were rated for 1.333 Gbps and the PLL multiplier was configured correctly).

Figure 35 Phase Noise Plot - CECE62002/5 @ 66.667 MHz

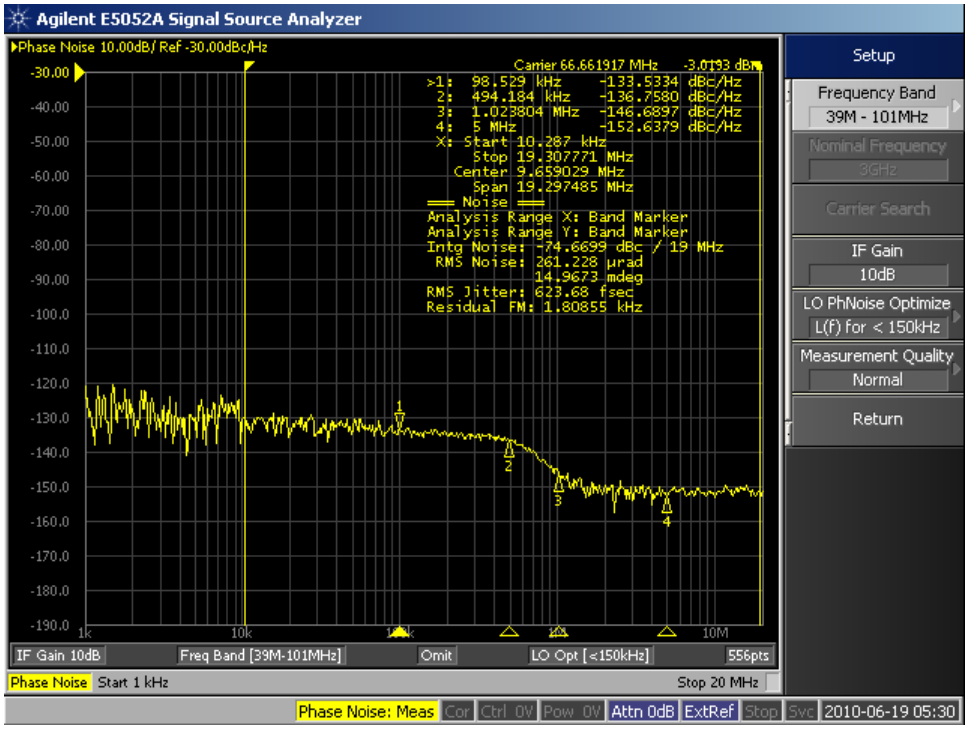

Figure 36 shows the phase noise data (plot) at 122.88 MHz. It is representative of both the CDC362002 and CDCE62005. This input frequency would be coupled to the ALTCORECLK or PASSCLK clock input of the device (provided this was the input frequency selected and the PLL multiplier was configured correctly).

Figure 36 Phase Noise Plot - CECE62002/5 @ 122.88 MHz

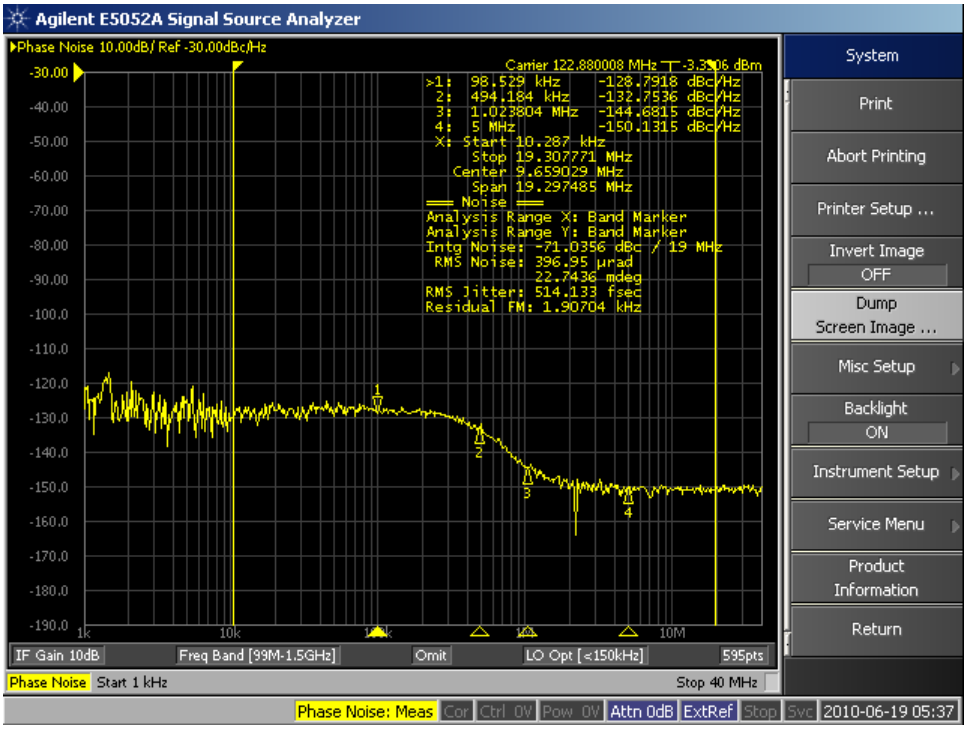

Figure 37 shows the phase noise data (plot) at 153.60 MHz. It is representative of both the CDC362002 and CDCE62005. This input frequency is intended to be coupled to the SYSCLK clock input of the device (provided this was the input frequency selected and the PLL multiplier was configured correctly). When visually overlaying this plot onto the previous SYSCLK plot, we can see that we are below the maximum allowable mask input level to the device.

Figure 37 Phase Noise Plot - CECE62002/5 @ 153.60 MHz

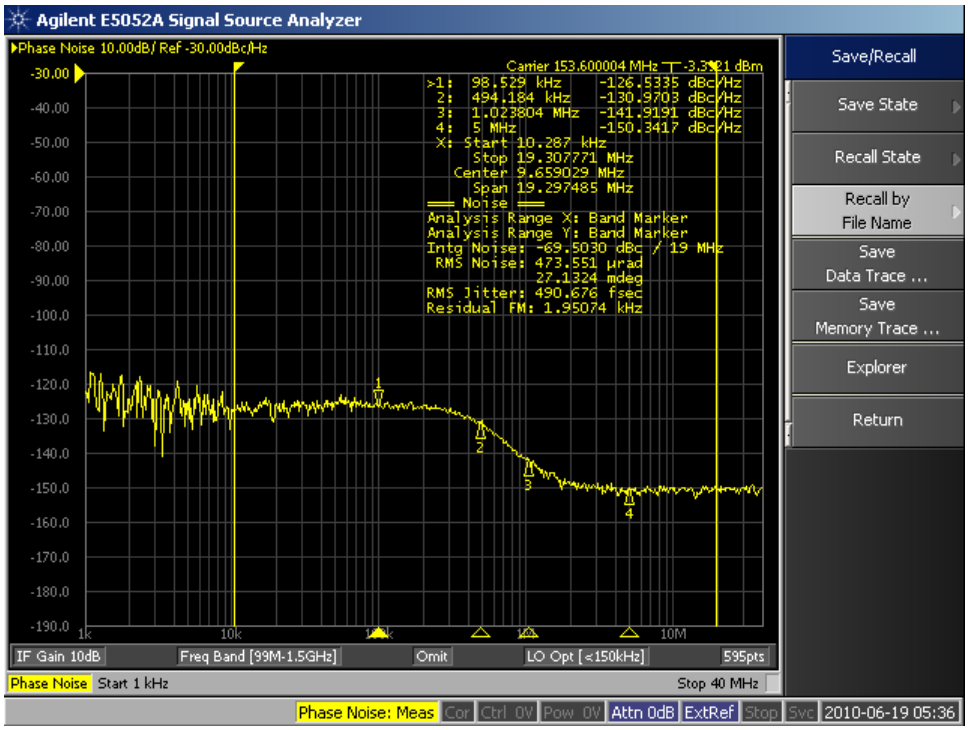

Figure 38 shows the phase noise data (plot) at 156.25 MHz. It is representative of both the CDC362002 and CDCE62005. This input frequency is intended to be coupled to the SRIO_SGMIICLK, PCIe_CLK, or MCMCLK clock input of the device (provided this was the input frequency selected and the PLL multiplier was configured correctly). When visually overlaying this plot onto the previous 156.25 MHz plot, we can see that we are below the maximum allowable mask input level to the device.

Figure 38 Phase Noise Plot - CECE62002/5 @ 156.25 MHz

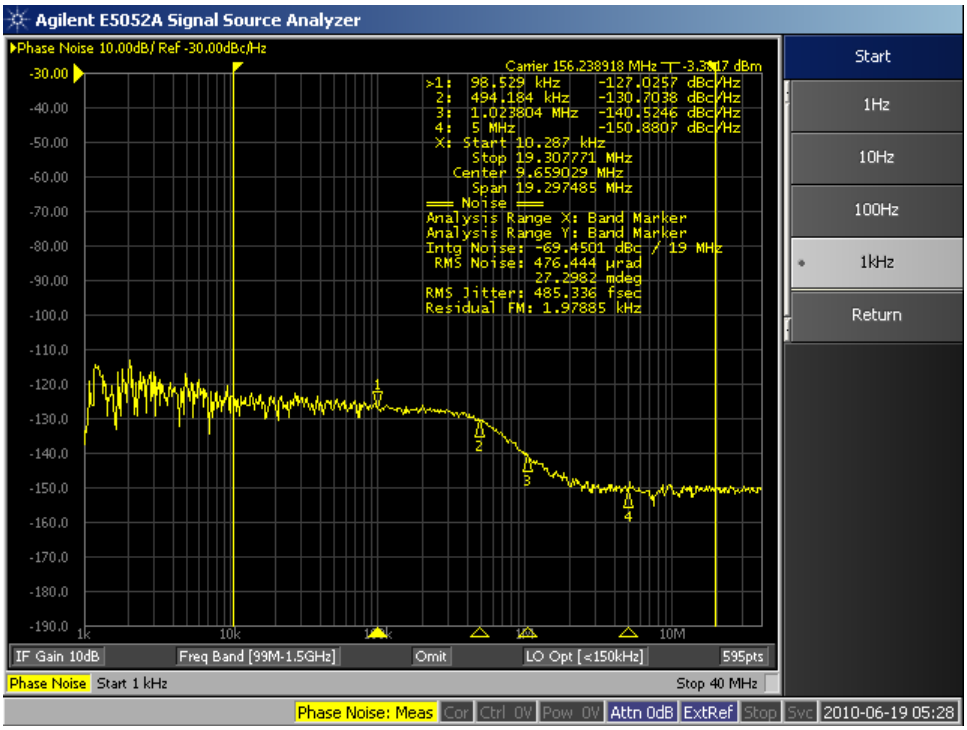

Figure 39 shows the phase noise data (plot) at 312.50 MHz. It is representative of both the CDC362002 and CDCE62005. This input frequency is recommended as the input clock for the SRIO_SGMIICLK, or MCMCLK clock into the device or as an alternate for the PCIe_CLK, ALTCORECLK, PASS_CLK,and DDR3CLK (to increase phase noise margins). Proper device configuration for this input clock frequency is mandatory. When visually overlaying this plot onto the previous 312.50 MHz plot, we can see that we are below the maximum allowable mask input level to the device.

Figure 39 Phase Noise Plot - CECE62002/5 @ 312.50 MHz

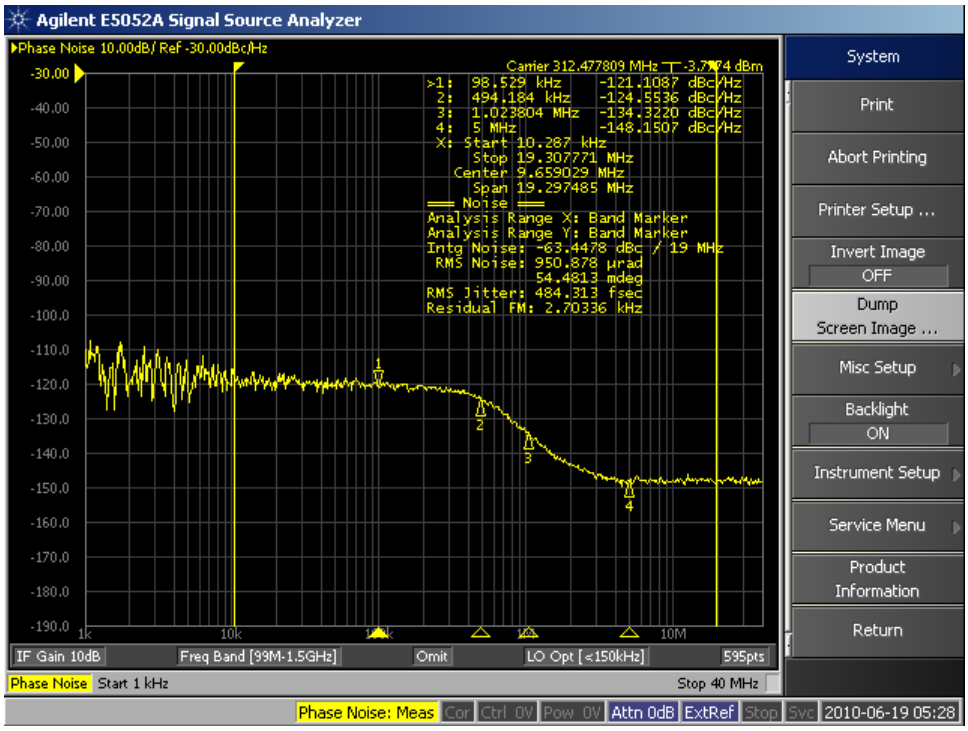

12.2 Reflection Calculations

This section provides designers and engineers with an example of how to calculate, at a first approximation, the impact of a reflection. It does not take into account the magnitude or potential for crosstalk to a parallel signal.

Assumptions:

- Dielectric constant (Er) is 4.1
- Signal is routed internally and externally
- AC coupling capacitor is placed 250 mils (6.35 mm) from the load
- Total net length is 3 inches (76.2mm)
- Frequency = 312.50 MHz or 3.2 ns period
- AC coupling capacitor size is 0402 and mounting pads meet IPC requirements

Calculations:

=85*SQRT((0.475*G13)+0.67) ==> results in the segment Tpd

As a first approximation, each segment was evaluated for Tpd (propagation delay) along a microstrip net. Each signal was further evaluated for the fundamental and secondary reflection (Table 38). The individual segments highlighted in yellow denote which segments are the source of the potential reflection.

For the sake of this example, the period of the reflection is 30 ps. However, this may not always be the case, which is where a simulation is beneficial.

These reflections (highlighted) would occur in the rising or falling edge of the clock periodic wave form. All others (see the Figure 40 timing analysis) would induce reflections, perturbations, or inflections during a logic 0 or logic 1 time frame and depending on magnitude could cause a logic transition, data error, or double clocking.

Reflections can occur at any point where an impedance mismatch occurs. The illustrations provided are only examples and are intended to illustrate the many possible permutations in a simple net that can occur (not all are comprehended in this example). For these reasons, simulation and modeling are always recommended.

Table 38 Reflections in ps (Part 1 of 2)

	t0-t1	t1-t2	t2-t3	t3-t4	t4-t5										
	2.7500	0.0350	0.0180	0.0350	0.2500										
	378.1769	4.81316	2.47534	4.81316	34.37972										
reflection nodes	378.1769	382.99	385.4654	390.2785	424.6582										
	t0-t1	t1-t2	t2-t3	t3-ta2	ta2-ta3	ta3-ta4	ta4-ta5								
	2.7500	0.0350	0.0180	0.0180	0.0180	0.0350	0.2500								
	378.1769	4.81316	**2.47534**	**2.47534**	**2.47534**	**4.81316**	**34.37972**								
reflection nodes	378.1769	382.99	385.4654	387.9407	390.4161	395.2292	429.6089								
	t0-t1	t1-t2	t2-t3	t3-t4	t4-t5	t5-tb4	tb4-tb5								
	2.7500	0.0350	0.0180	0.0350	0.2500	0.2500	0.2500								
	378.1769	4.81316	2.47534	4.81316	34.37972	**34.37972**	**34.37972**								
reflection nodes	378.1769	382.99	385.4654	390.2785	424.6582	459.038	493.4177								
	tb0-tb1	tb1-tc0	tc0-tc1	tc1-tc2	tc2-tc3	tc3-tc4	tc4-tc5								
	2.7500	2.7500	2.7500	0.0350	0.1800	0.0350	0.2500								
	378.1769	**378.1769**	**378.1769**	**4.81316**	**24.7534**	**4.81316**	**34.37972**								
reflection nodes	378.1769	756.3537	1134.531	1139.344	1164.097	1168.91	1203.29								
	tb0-tb1	tb1-tc0	tc0-tc1	tc1-tc2	tc2-tc3	tc3-td2	td2-td3	td3-td4	td4-td5						
	2.7500	2.7500	2.7500	0.0350	0.1800	0.0180	0.0180	0.0350	0.2500						
	378.1769	**378.1769**	**378.1769**	**4.81316**	**24.7534**	**2.47534**	**2.47534**	**4.81316**	**34.37972**						
reflection nodes	378.1769	756.3537	1134.531	1139.344	1164.097	1166.573	1169.048	1173.861	1208.241						
	tb0-tb1	tb1-tc0	tc0-tc1	tc1-tc2	tc2-tc3	tc3-tc4	tc4-tc5	tc5-te4	te4-te5						
	2.7500	2.7500	2.7500	0.0350	0.1800	0.0350	0.2500	0.2500	0.2500						
	378.1769	**378.1769**	**378.1769**	**4.81316**	**24.7534**	**4.81316**	**34.37972**	**34.37972**	**34.37972**						
reflection nodes	378.1769	756.3537	1134.531	1139.344	1164.097	1168.91	1203.29	1237.667	1272.049						
	tb0-tb1	tb1-tc0	tc0-tc1	tc1-tc2	tc2-tc3	tc3-td2	td2-td3	td3-td4	td4-td5	td5-tf4	tf4-tf5				
	2.7500	2.7500	2.7500	0.0350	0.1800	0.0180	0.0180	0.0350	0.2500	0.2500	0.2500				

Table 38 Reflections in ps (Part 2 of 2)

	378.1769	**378.1769**	**378.1769**	**4.81316**	**24.7534**	**2.47534**	**2.47534**	**4.81316**	**34.37972**	**34.37972**	**34.37972**				
reflection nodes	378.1769	756.3537	1134.531	1139.344	1164.097	1166.573	1169.048	1173.861	1208.241	1242.62	1277				
	te0-te1	**te1-te2**	**te2-te3**	**te3-tf4**	**tf4-tg3**	**tg3-tg4**	**tg4-tg5**								
	2.7500	0.0350	0.0180	0.0350	0.0350	0.0350	0.2500								
	378.1769	**4.81316**	**2.47534**	**4.81316**	**4.81316**	**4.81316**	**34.37972**								
reflection nodes	378.1769	382.99	385.4654	390.2785	395.0917	399.9049	434.2846								
	te0-te1	**te1-te2**	**te2-te3**	**te3-tf2**	**tf2-tf1**	**tf1-tf0**	**tf0-tg1**	**tg1-tg2**	**tg2-tg3**	**tg3-th4**	**th4-th5**				
	2.7500	0.0350	0.0180	0.0180	0.0350	2.7500	2.7500	0.0350	0.0180	0.0350	0.2500				
	378.1769	**4.81316**	**2.47534**	**2.47534**	**4.81316**	**378.1769**	**378.1769**	**4.81316**	**2.47534**	**4.81316**	**34.37972**				
reflection nodes	378.1769	382.99	385.4654	387.9407	392.7539	770.9307	1149.108	1153.921	1156.396	1161.209	1195.589				
	te0-te1	**te1-te2**	**te2-te3**	**te3-tf4**	**tf4-tf3**	**tf3-tg2**	**tg2-th1**	**th1-tj2**	**tj2-tj3**	**tj3-tj4**	**tj4-tj5**				
	2.7500	0.0350	0.0180	0.0350	0.0350	0.0180	0.0350	0.0350	0.0180	0.0350	0.2500				
	378.1769	**4.81316**	**2.47534**	**4.81316**	**4.81316**	**2.47534**	**4.81316**	**4.81316**	**2.47534**	**4.81316**	**34.37972**				
reflection nodes	378.1769	382.99	385.4654	390.2785	395.0917	397.567	402.3802	407.1934	409.6687	414.4819	448.8616				
	tk0-tk1	**tk1-tk2**	**tk2-tk3**	**tk3-tf4**	**tk4-tk5**	**tk5-tl4**	**tl4-tl3**	**tl3-tl2**	**tl2-tl1**	**tl1-tl0**	**tl0-tm1**	**tm1-tm2**	**tm2-tm3**	**tm3-tm4**	**tm4-tm5**
	2.7500	0.0350	0.0180	0.0350	0.2500	0.2500	0.0350	0.0180	0.0350	2.7500	2.7500	**0.035**	**0.018**	**0.035**	**0.25**
	378.1769	**4.81316**	**2.47534**	**4.81316**	**34.37972**	**34.37972**	**4.81316**	**2.47534**	**4.81316**	**378.1769**	**378.1769**	**4.81316**	**2.47534**	**4.81316**	**34.37972**
reflection nodes	378.1769	382.99	385.4654	390.2785	424.6582	459.038	463.8511	466.3265	471.1396	849.3165	1227.493	1232.307	1234.782	1239.595	1273.975
End of Table 38															

Figure 40 Reflections Timing Analysis

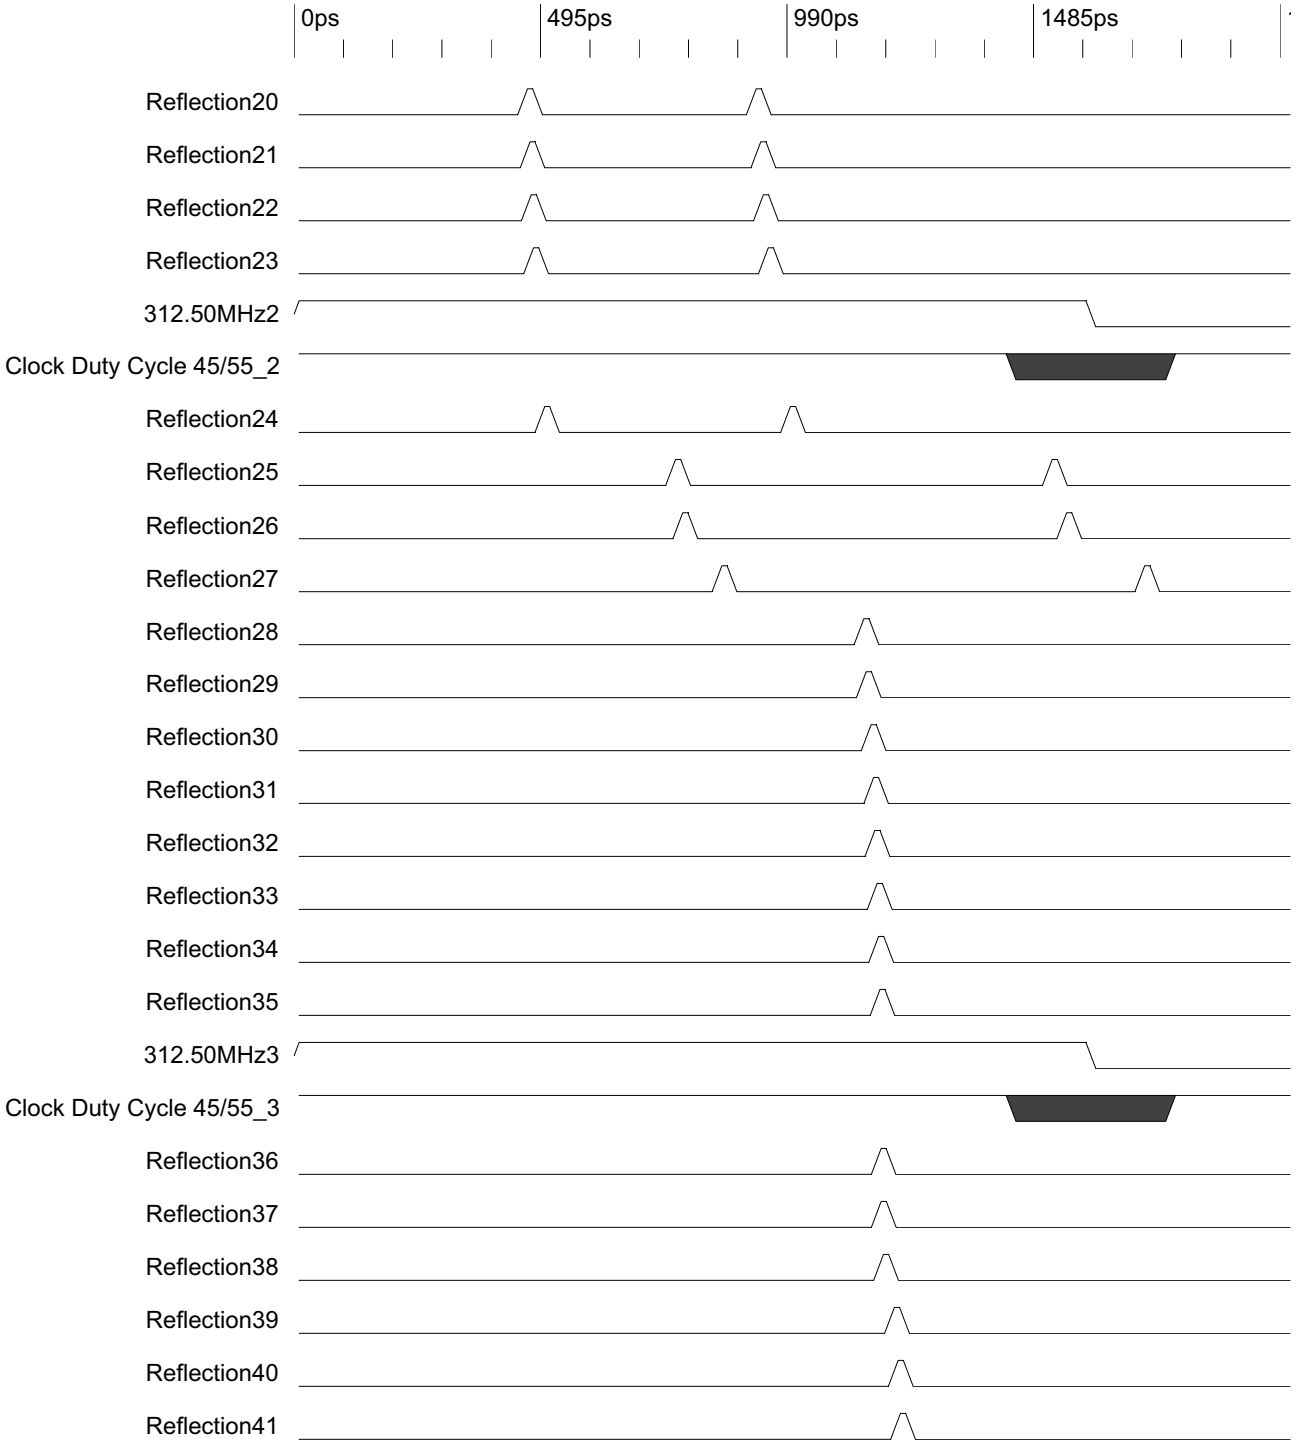

Table 39 Possible First 50 Reflection Combinations (Part 1 of 2)

	Initial Reflections			Secondary Reflections	
#	Rising Edge	Falling Edge		Rising Edge	Falling Edge
1	0.3782	0.4082		0.7564	0.7864
2	0.383	0.413		0.766	0.796
3	0.3855	0.4155		0.7709	0.8009
4	0.3879	0.4179		0.7759	0.8059
5	0.3903	0.4203		0.7806	0.8106
6	0.3904	0.4204		0.7808	0.8108
7	0.3928	0.4228		0.7855	0.8155
8	0.3951	0.4251		0.7902	0.8202
9	0.3952	0.4252		0.7905	0.8205
10	0.3976	0.4276		0.7951	0.8251
11	0.3999	0.4299		0.7998	0.8298
12	0.4024	0.4324		0.8048	0.8348
13	0.4072	0.4372		0.8144	0.8444
14	0.4097	0.4397		0.8193	0.8493
15	0.4145	0.4445		0.829	0.859
16	0.4247	0.4547		0.8493	0.8793
17	0.4296	0.4596		0.8592	0.8892
18	0.4343	0.4643		0.8686	0.8986
19	0.4489	0.4789		0.8977	0.9277
20	0.459	0.489		0.9181	0.9481
21	0.4639	0.4939		0.9277	0.9577
22	0.4663	0.4963		0.9327	0.9627
23	0.4711	0.5011		0.9423	0.9723
24	0.4934	0.5234		0.9868	1.0168
25	0.7564	0.7864		1.5127	1.5427
26	0.7709	0.8009		1.5419	1.5719
27	0.8493	0.8793		1.6986	1.7286
28	1.1345	1.1645		2.2691	2.2991
29	1.1393	1.1693		2.2787	2.3087
30	1.1491	1.1791		2.2982	2.3282
31	1.1539	1.1839		2.3078	2.3378
32	1.1564	1.1864		2.3128	2.3428
33	1.1612	1.1912		2.3224	2.3524
34	1.1641	1.1941		2.3282	2.3582
35	1.1666	1.1966		2.3331	2.3631
36	1.1689	1.1989		2.3378	2.3678
37	1.169	1.199		2.3381	2.3681
38	1.1739	1.2039		2.3477	2.3777
39	1.1956	1.2256		2.3912	2.4212
40	1.2033	1.2333		2.4066	2.4366
41	1.2082	1.2382		2.4165	2.4465
42	1.2275	1.2575		2.455	2.485
43	1.2323	1.2623		2.4646	2.4946

Table 39 Possible First 50 Reflection Combinations (Part 2 of 2)

	Initial Reflections			Secondary Reflections	
#	Rising Edge	Falling Edge		Rising Edge	Falling Edge
44	1.2348	1.2648		2.4696	2.4996
45	1.2377	1.2677		2.4753	2.5053
46	1.2396	1.2696		2.4792	2.5092
47	1.2426	1.2726		2.4852	2.5152
48	1.272	1.302		2.5441	2.5741
49	1.274	1.304		2.5479	2.5779
50	1.277	1.307		2.554	2.584
End of Table 39					

13 Appendix B

This section provides detailed information on transmission lines and routing clocks.

13.1 Overview

High-frequency signaling requires the use of transmission lines where signals require minimal distortion, crosstalk, and EMI radiation. They are used for clocks, differential SerDes signaling, and a variety of other high-speed signaling. Microstrip and stripline are two popular types of electrical transmission lines that can be fabricated using printed circuit board technology. Clock and most other high performance signaling generally use one of these transmission line types.

An industry standards body, the IPC, has created standard IPC-2141A (2004) *Design Guide for High-Speed Controlled Impedance Circuit Boards* dealing with their construction.

Before proceeding, a bit of background information is helpful. The use of the term *ground plane* implies a large area, low-impedance reference plane. In practice, it may be either a ground plane or a power plane, both of which are assumed to be at zero AC potential.

It should be understood that there are numerous equations, all represented as describing transmission line design and behavior. It is a safe to question the accuracy of these equations and assume none of them are perfect, being only good approximations, with accuracy depending upon specifics. The best known and most widely quoted equations are those in the IPC 2141A standard, but even these come with application caveats. The source of the formulas herein is the above mentioned standard.

It is helpful to know that there are a number of transmission line geometries:

- Microstrip
- Embedded microstrip
- Symmetric stripline
- Asymmetric stripline
- Wire microstrip
- Wire stripline
- Edge-coupled microstrip
- Edge-coupled stripline
- Broadside coupled stripline

These transmission line geometries are shown in Figure 41. This document briefly covers only microstrip and symmetric stripline.

Figure 41 Transmission Line Geometries

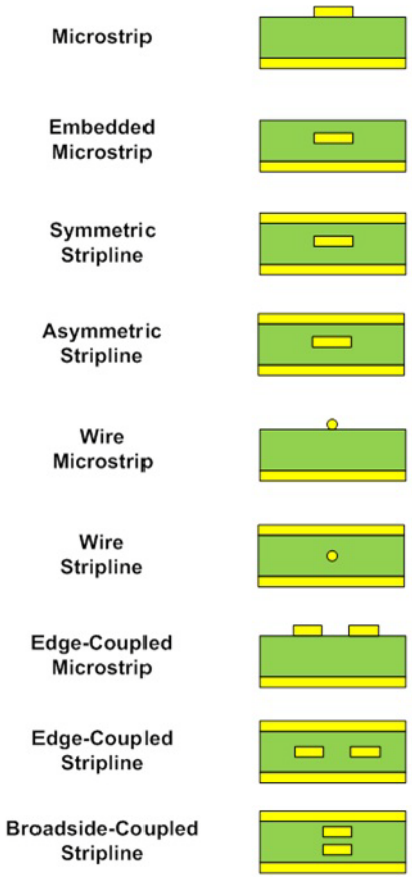

There are multiple impedance calculators available to designers, many of which comprehend these geometries. Two such calculators are available at www.eeweb.com/toolbox/microstrip-impedance/ and at www.technick.net/public/code/.

13.2 Contrasting Microstrip and Stripline Transmission Lines

A simplified side-by-side view of microstrip and stripline transmission lines is shown in Figure 42.

Figure 42 Transmission Line Geometries

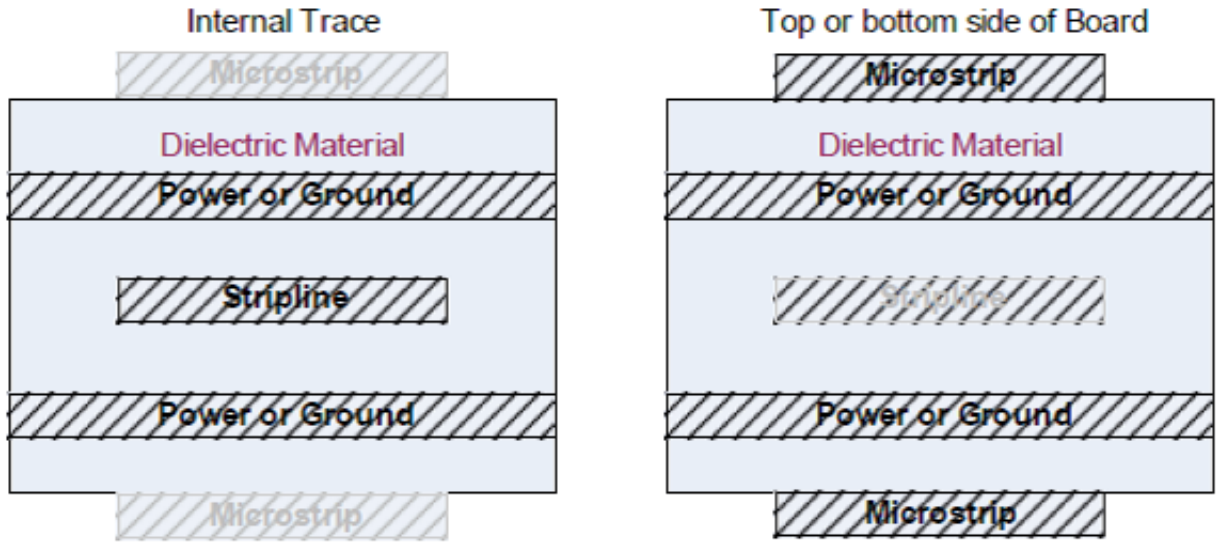

A number of physical and electrical factors govern whether signaling should use microstrip or stripline topology. The key factors determining which topology is used include:

- **BGA escape paths** - A routing escape path may be on an internal or external layer.
- **Timing/performance** - Microstrip propagation delays are smaller than those of stripline.
- **EMI** - Susceptibility to and emission of radiated noise is less with stripline
- **Loss characteristics** - Stripline is less lossy than microstrip

13.2.1 Microstrip Transmission Lines Construction

A cross section of microstrip transmission line is shown in Figure 43. This geometry is known as a surface microstrip, or simply, microstrip. It is formed by a dielectric (D) (typically air) above a conductor (C), with another dielectric substrate (B) separating the conductor from a ground plane (A).

Figure 43 Microstrip Transmission Line Construction

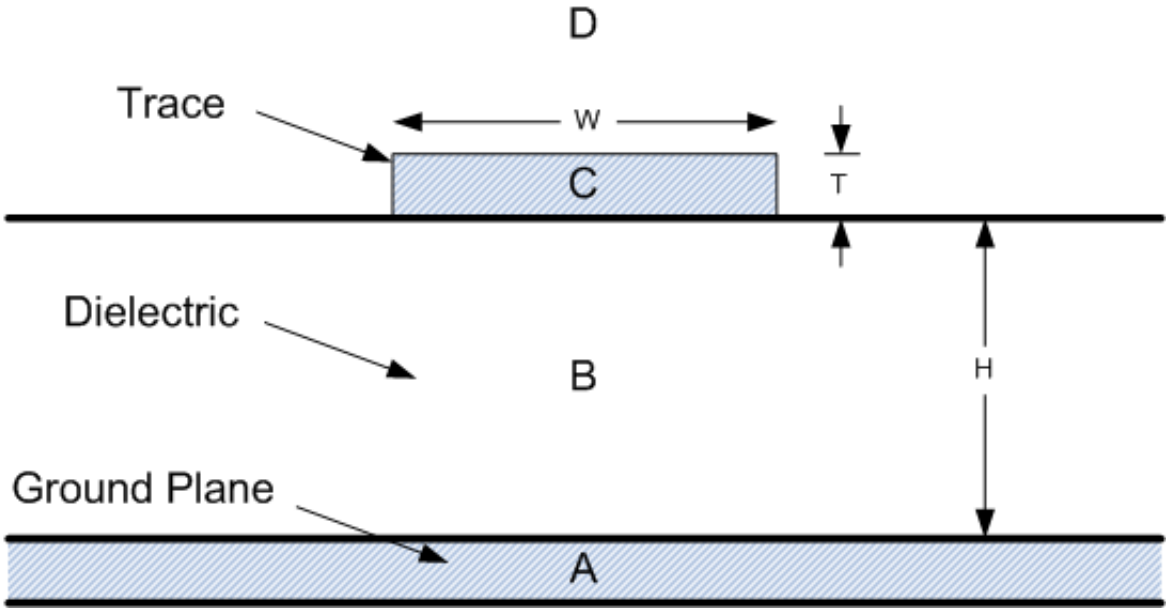

The characteristic impedance of a microstrip transmission line depends on a number of factors including, the width (W) and thickness (T) of the trace (C), and the thickness (H) of dielectric.

For a given PCB laminate and copper weight, note that all parameters will be predetermined except for W, the width of the signal trace. The trace width is then varied to create the desired characteristic impedance. Printed circuit boards commonly use 50 Ω and 100 Ω characteristic impedances.

Equation 1 can be used to calculate the characteristic impedance of a signal trace of width W and thickness T, separated by a dielectric with a dielectric constant ε_r and thickness H (all measurements are in common dimensions (mils) from a constant potential plane).

Figure 44 Transmission Line Equation 1

$$Z_o(\Omega) = \frac{87}{\sqrt{\varepsilon_r + 1.41}} \ln\left[\frac{5.98H}{(0.8W + T)}\right] \quad \text{Equation 1}$$

The characteristic capacitance of a microstrip transmission line can be calculated in terms of pF/inch as shown in Equation 2.

Figure 45 Transmission Line Equation 2

$$C_o(\text{pF/in}) = \frac{0.67(\varepsilon_r + 1.41)}{\ln[5.98H/(0.8W + T)]} \quad \text{Equation 2}$$

For example, a microstrip transmission line constructed with a 1-ounce (T=1.4) copper-clad 10-mil (H) FR-4 ($\varepsilon_r = 4.0$) where the trace width (W) is 20-mils wide results in an impedance of about 50 Ω.

For the 75-Ω video standard, the W would be about 8.3 mils. W may easily be adjusted to create other impedances. The above equations are from the IPC standards. These equations are most accurate between 50 Ω and 100 Ω, but is less accurate for lower or higher impedances.

A useful guideline pertaining to microstrip PCB impedance is that for a dielectric constant of 4.0 (FR-4), a W/H ratio of 2/1 results in an impedance of close to 50 Ω (as in the first example, with W = 20 mils). This is within 5% of the value predicted by Equation 1 (46 Ω) and generally consistent with the accuracy (>5%) of observed measurements.

13.2.2 Microstrip Propagation Delay

The propagation delay of a microstrip transmission line is somewhere between the speed of an electromagnetic waves in the substrate, and the speed of electromagnetic waves in air. This is due to the electromagnetic wave partial propagation in the dielectric substrate, and partial propagation in the air above it — propagation delay being a hybrid of the two.

The one way propagation delay of the microstrip transmission line can be calculated using Equation 3 to obtain ns/ft or Equation 4 to obtain ps/in.

Figure 46 Transmission Line Equation 3 & 4

$$t_{pd}(\text{ns/ft}) = 1.017\sqrt{0.475\varepsilon_r + 0.67} \quad \text{Equation 3}$$

$$t_{pd}(\text{ps/in}) = 85\sqrt{0.475\varepsilon_r + 0.67} \quad \text{Equation 4}$$

13.2.3 Symmetric Stripline Transmission Lines

A stripline transmission line is constructed with a signal trace sandwiched between two parallel ground planes as shown in Figure 47. It is similar to coaxial cable as it is non-dispersive and has no cutoff frequency. When compared to microstrip transmission lines, it provides better isolation between adjacent traces and propagation of radiated RF emissions, but has slower propagation speeds.

Figure 47 Stripline Transmission Line Construction

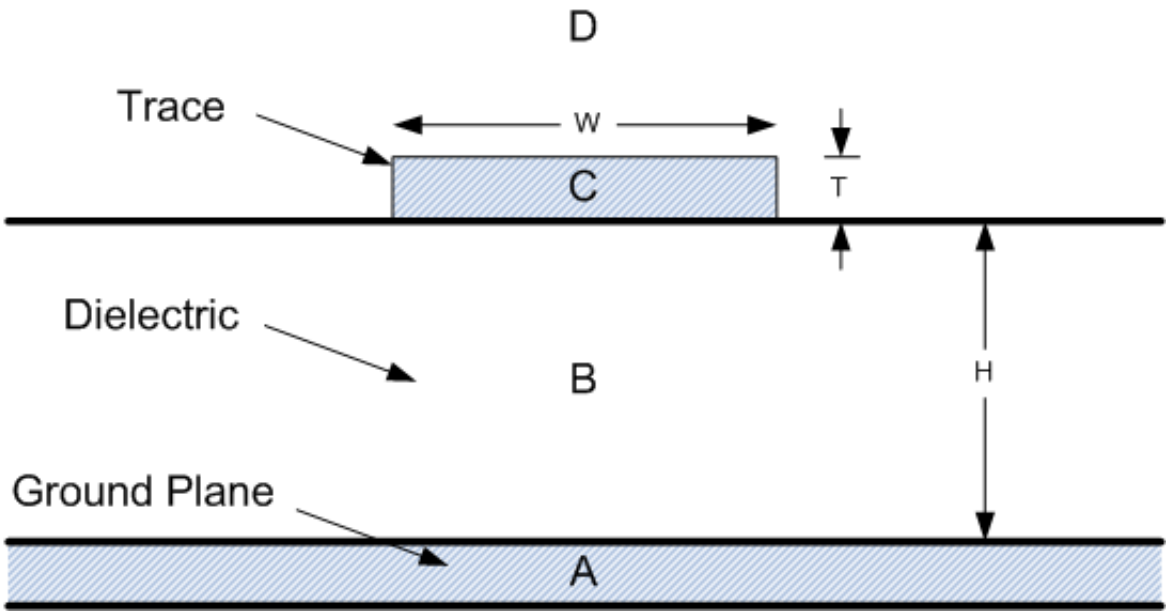

13.2.4 Characteristic Impedance

The width and thickness (W and T) of the trace, the thickness of the dielectric separating the planes (H), and their relative permittivity (ε_r) determine the transmission line's characteristic impedance. In the general case, the trace need not be equally spaced between the ground planes and the dielectric material may be different above and below the central conductor. In this case the transmission line is an asymmetric stripline transmission line.

Equation 7 defines Z_O of a symmetric (H = H1 = H2) symmetric stripline transmission line. The accuracy of the equation is typically on the order of 6%.

Figure 48 Transmission Line Equation 7

$$Z_o(\Omega) = \frac{60}{\sqrt{\varepsilon_r}} \ln\left[\frac{1.9(S)}{(0.8W + T)}\right]. \qquad \text{Equation 7}$$

A useful ballpark estimate for constructing a 50-Ω symmetric stripline transmission line where ε_r= 4.0 (FR4) is to make the S/W ratio 2 to 2.2. Recognize that this guideline is only an approximation as it neglects the value of T. It is important to note that it may be difficult to create some transmission lines of a high impedance while maintaining a trace width suitable for high yield PCB manufacturing.

The characteristic capacitance of a symmetric stripline transmission line in pf/in can be calculated using Equation 8.

Figure 49 Transmission Line Equation 8

$$C_o\,(pF/in) = \frac{1.41(\varepsilon_r)}{\ln[3.81H/(0.8W + T)]} \qquad \text{Equation 8}$$

13.2.5 Stripline Propagation Delay

The propagation delay of a symmetric stripline transmission line in ns/ft and ps/inch can be calculated using Equations 9 and 10, respectively. For a PCB dielectric constant of 4.0 (FR-4), a symmetric stripline's delay is 2 ns/ft, or 170 ps/inch.

Figure 50 Transmission Line Equation 9 & 10

$$t_{pd}\,(ns/ft) = 1.017\sqrt{\varepsilon_r} \qquad \text{Equation 9}$$

$$t_{pd}\,(ps/in) = 85\sqrt{\varepsilon_r} \qquad \text{Equation 10}$$

Table 40 Signal Delay Through Various Materials

			Trace Length in Inches Versus Delay in ps									
Er	**Name**	**µ Strip**	**1″**	**2″**	**3″**	**4″**	**5″**	**6″**	**7″**	**8″**	**9″**	**10″**
1	Air	90.9540	90.95	181.91	272.86	363.82	454.77	545.72	636.68	727.63	818.59	909.54
2.2	PTFE/Glass	111.3143	111.31	222.63	333.94	445.26	556.57	667.89	779.20	890.51	1001.83	1113.14
2.9	Ceramic	121.6273	121.63	243.25	364.88	486.51	608.14	729.76	851.39	973.02	1094.65	1216.27
3.5	GETEK	129.8165	129.82	259.63	389.45	519.27	649.08	778.90	908.72	1038.53	1168.35	1298.17
4	FR-4	136.2654	136.27	272.53	408.80	545.06	681.33	817.59	953.86	1090.12	1226.39	1362.65
4.1	FR-4	137.5189	137.52	275.04	412.56	550.08	687.59	825.11	962.63	1100.15	1237.67	1375.19
4.2	FR-4	138.761	138.76	277.52	416.28	555.04	693.81	832.57	971.33	1110.09	1248.85	1387.61
4.3	FR-4	139.9922	139.99	279.98	419.98	559.97	699.96	839.95	979.95	1119.94	1259.93	1399.92
4.4	FR-4	141.2126	141.21	282.43	423.64	564.85	706.06	847.28	988.49	1129.70	1270.91	1412.13
4.5	FR-4	142.4226	142.42	284.85	427.27	569.69	712.11	854.54	996.96	1139.38	1281.80	1424.23
End of Table 40												

13.2.6 Changes in Transmission Line Direction (Bends)

Microstrip transmission lines used with PCBs do not always follow a straight line from point A to Point B. The transmission line often turns one or more times between these points. Abrupt turns in the transmission line are not permitted as this will cause a significant portion of the wave energy to be reflected back towards the wave's source, with only part of the wave energy transmitted past the bend. This effect is generally best minimized by curving the path of the strip in an arc of a radius at least 3× the strip-width [8]. Another more commonly used method is to use a mitred bend as this consumes a smaller routing area, where the mitered portion cuts-away a fraction of the diagonal between the inner and outer corners of the un-mitred bend as shown in Figure 47.

Figure 51 Mitered Bend

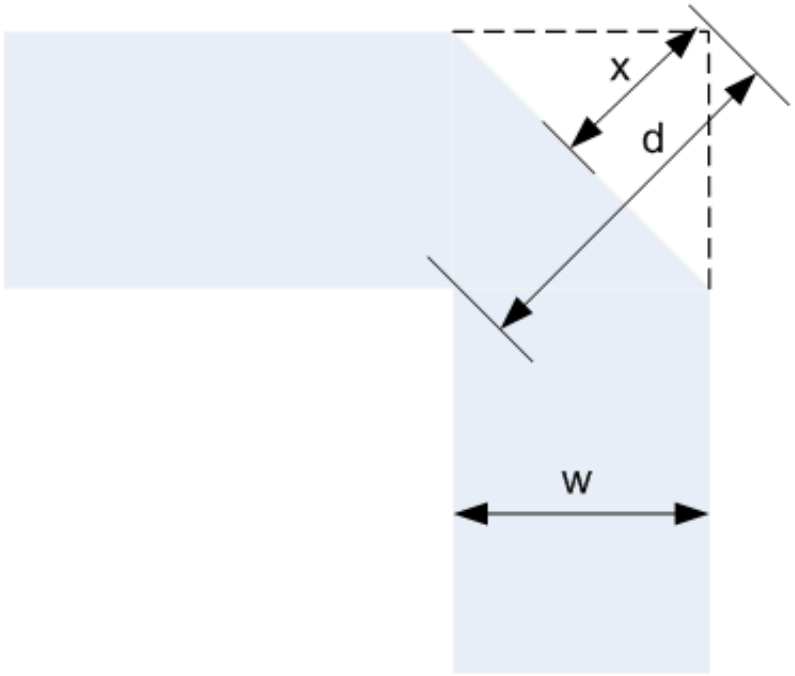

Douville and James have experimentally determined the optimum mitre for a wide range of microstrip geometries. They have presented evidence for cases where 0.25 <= W/H <= 2.75 and 2.5 <= ε_r <= 25. Their findings indicate that the optimum percentage mitre can be represent by Equation 11 where M is the percentage of miter.

Figure 52 Transmission Line Equation 11

$$M = 100\frac{x}{d}\% = (52 + 65e^{-\frac{27}{20}\frac{w}{h}})\% \quad \text{Equation 11}$$

They report a Voltage Standing Wave Ratio (VSWR of better than 1.1 (i.e., a return better than 26 dB) for any percentage mitre of D within 4% produced by the formula.

It is important to note that for a minimum W/H ratio of 0.25, the percentage mitre is 96%, so that the strip is very nearly cut through. This means that care must be taken to understand the effect the percentage of miter has on the manufacturability of the PCB. At some point, the manufacturing tolerances may limit the use of mitering and the use of appropriate trace radii may be required. Note that the propagation delay is affected by changes in trace direction as the electrical length is somewhat shorter than the physical path-length of the strip for both the curved and mitred bends.

14 References

1. OBSAI Reference Point 4 Specification, Version 1.1
2. OBSAI Reference Point 3 Specification, Version 4.1
3. OBSAI Reference Point 2 Specification, Version 2.1
4. OBSAI Reference Point 1 Specification, Version 2.1
5. CPRI Specification, Version 4.1 CPRI v4.1
6. JEDEC SSTL_18 Specification (EAI/JESD8-15a), dated September 2003
7. SGMII Specification (ENG-46158), Version 1.8, dated April 27, 2005
8. LVDS Electrical Specification (IEEE-1596.3-1996), dated 1996
9. XAUI Electrical Specification (IEEE-802.3ae-2002), dated 2002
10. RapidIO Interconnect Part VI: Physical Layer 1×/4× LP-Serial Specification, Version 2.0.1, dated March 2008
11. JEDEC DDR3 SDRAM Specification (EAI/JESD79-3C)
12. Boundary Scan Test Specification (IEEE-1149.1)
13. AC Coupled Net Test Specification (IEEE-1149.6)
14. I2C Bus Specification (9398 393 40011), Version 2.1, dated January, 2000
15. Flip Chip BGA Users Guide (SPRU811)
16. Emulation and Trace Headers Technical Reference Manual (SPRU655)
17. JEDEC 2.5 V ± 0.2 V (Normal Range), and 1.8 V to 2.7 V (Wide Range) Power Supply Voltage and Interface Standard for Nonterminated Digital Integrated Circuit Specification (EAI/JESD8-5), dated October 1995 http://www.jedec.org/download/search/JESD8-5.pdf
18. JEDEC Power Supply Voltage and Interface Standard Specification (EAI/JESD8-7), dated February 1997 http://www.jedec.org/download/search/JESD8-7.pdf
19. DDR3 Design Guide for KeyStone Devices (SPRABI1)
20. Communications Infrastructure KeyStone Data Manual
21. Serial RapidIO (SRIO) for KeyStone Devices User's Guide (SPRUGW1)
22. SerDes Implementation Guidelines for KeyStone Devices (SPRABC1)
23. DDR3 Memory Controller for KeyStone Devices User's Guide (SPRUGV8)
24. Ethernet Media Access Controller (EMAC) for KeyStone Devices User's Guide (SPRUGV9)
25. Inter-Integrated Circuit (I2C) for KeyStone Devices User's Guide (SPRUGV3)
26. General-Purpose Input/Output (GPIO) for KeyStone Devices User's Guide (SPRUGV1)
27. 64-Bit Timer (Timer64) for KeyStone Devices User's Guide (SPRUGV5)
28. Using IBIS Modes for Timing Analysis (SPRA839)
29. C66x CPU and Instruction Set Reference Guide (SPRUGH7)
30. Phase-Locked Loop (PLL) Controller for KeyStone Devices User's Guide (SPRUGV2)
31. Antenna Interface 2 (AIF2) for KeyStone Devices User's Guide (SPRUGV7)
32. KeyStone Architecture Bootloader User Guide (SPRUGY5)
33. Filtering Techniques: Isolating Analog and Digital Power Supplies in TI's PLL-Based CDC Devices (SCAA048)
34. UCD9244 Quad PWM System Controller (http://www.ti.com/product/ucd9244)

35. CDCx706/x906 Termination and Signal Integrity Guidelines (SCAA080)
36. Voltage Translation Between 3.3-V, 2.5-V, 1.8-V, and 1.5-V Logic Standards With the TI AVCA164245 and AVCB164245 Dual-Supply Bus-Translating Transceivers Application Report(SCEA030)
37. Selecting the Right Level-Translation Solution Application Report (SCEA035)
38. PCA9306 Dual Bidirectional I²C Bus and SMBus Voltage-Level Translator Data Sheet (SCPS113)
39. Thermal Design Guide for KeyStone Devices (SPRABI3)
40. AC-Coupling Between Differential LVPECL, LVDS, HSTL, and CML (SCAA059)
41. DC-Coupling Between Differential LVPECL, LVDS, HSTL, and CM (SCAA062)
42. TMS320C6678/74/72, TMS320TCI6608 Power Consumption (SPRABI5)
43. DDR3 Design Guide for KeyStone Devices (SPRABI1)
44. Clocking Design Guide for KeyStone Devices (SPRABI4)
45. UCD9222 Dual Digital PWM System Controller (http://www.ti.com/product/ucd9222)
46. UCD7242 - Digital Dual Synchronous-Buck Power Driver (SLUS962)
47. CDCE62002 - Four Output Clock Generator/Jitter Cleaner With Integrated Dual VCOs (SCAS882)
48. CDCE62005 - Five/Ten Output Clock Generator/Jitter Cleaner With Integrated Dual VCOs (SCAS862)
49. Common Trace Transmission Problems and Solutions (SPRAAK6)
50. Using xdsprobe with the XDS560 and XDS510 (SPRA758A)

15 Revision History

Table 41 shows the change history for this document.

Table 41 Document Revision History

Revision	Date	Comments
SPRABI2C	August 2013	• Major rewrite of most sections.
SPRABI2B	March 2012	• Added the Random Jitter section. (Page 1-32) • In the Clocking Requirements table, added CORECLKp and CORECLKn. (Page 1-37) • In the Device Reset section, reset modes information updated. (Page 1-53) • In the PORz section, noted that: Most output signals will be disabled (high-impedance) while POR is low. (Page 1-53) • In the PORz section, noted that: Power good logic can re-assert POR low when a power supply failure or brown-out occurs to prevent over-current conditions. (Page 1-53) • In the RESETz section, noted that RESETz does not latch bootstrapping. (Page 1-54) • In the RESETz section, noted that RESETz is a software configurable function. (Page 1-54) • In the RESETz section, noted that: When RESETz is active low, all 3-state outputs are placed in a high-impedance state. (Page 1-54) • In the Slew Rate Control table, corrected the "Setting" column. (Page 1-74) • Renamed the "Specific SerDes Input Clock Requirements" section to "Reference Clock Jitter Requirements" and updated the content. (Page 1-31) • Renamed the "Specific SerDes Input Clock Requirements" section to "Reference Clock Jitter Requirements" and updated the content. (Page 1-30) • Reworked the Boot Modes section. (Page 1-56) • Reworked the LRESETz section. (Page 1-54) • Reworked the Reset Implementation Considerations section. (Page 1-55) • Reworked the RESETSTATz section. (Page 1-55) • Statement added for multiple configuration inputs. (Page 1-56) • The table Maximum Device Compression - Leaded Solder Balls was updated. (Page 1-90) • The table Maximum Device Compression - Lead-Free Solder Balls was updated. (Page 1-90) • Updated the bulleted information points in the Clock Signal Routing section. (Page 1-93) • Updated references to HyperLink in the text that referred to it as HyperBridge or MCM. (Page 1-56) • Corrected the bits in the Setting column of the Slew Rate Control table. (Page 1-74) • Rewrite of the LRESETz section. (Page 1-54) • Updated Jitter information. (Page 1-31) • Updated the ‘‘Managing Unused Clock Inputs’’ table. (Page 1-36) • Updated the ‘‘Maximum Device Compression - Lead-Free Solder Balls’’ table. (Page 1-90) • Added CORECLK to the Clocking Requirements table. (Page 1-40) • Added CORECLK to the Clocking Requirements table. (Page 1-31) • Added CORECLK to the Clocking Requirements table. (Page 1-30) • Added CORECLK to the Clocking Requirements table. (Page 1-30) • Added CORECLK to the Clocking Requirements table. (Page 1-29) • Added CORECLK to the Clocking Requirements table. (Page 1-37)
SPRABI2A	November 2011	• All references to CDCM6212 removed. (Page 1-43) • Change made to pull TRST to ground instead of pulling it high. (Page 1-52) • Managing Unused Clock Inputs added. (Page 1-36) • Updated the Thermal Considerations (formerly Thermal Issues) section. (Page 1-88)
SPRABI2	November 2010	Initial release.

IMPORTANT NOTICE

Texas Instruments Incorporated and its subsidiaries (TI) reserve the right to make corrections, enhancements, improvements and other changes to its semiconductor products and services per JESD46, latest issue, and to discontinue any product or service per JESD48, latest issue. Buyers should obtain the latest relevant information before placing orders and should verify that such information is current and complete. All semiconductor products (also referred to herein as “components”) are sold subject to TI’s terms and conditions of sale supplied at the time of order acknowledgment.

TI warrants performance of its components to the specifications applicable at the time of sale, in accordance with the warranty in TI’s terms and conditions of sale of semiconductor products. Testing and other quality control techniques are used to the extent TI deems necessary to support this warranty. Except where mandated by applicable law, testing of all parameters of each component is not necessarily performed.

TI assumes no liability for applications assistance or the design of Buyers’ products. Buyers are responsible for their products and applications using TI components. To minimize the risks associated with Buyers’ products and applications, Buyers should provide adequate design and operating safeguards.

TI does not warrant or represent that any license, either express or implied, is granted under any patent right, copyright, mask work right, or other intellectual property right relating to any combination, machine, or process in which TI components or services are used. Information published by TI regarding third-party products or services does not constitute a license to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property of the third party, or a license from TI under the patents or other intellectual property of TI.

Reproduction of significant portions of TI information in TI data books or data sheets is permissible only if reproduction is without alteration and is accompanied by all associated warranties, conditions, limitations, and notices. TI is not responsible or liable for such altered documentation. Information of third parties may be subject to additional restrictions.

Resale of TI components or services with statements different from or beyond the parameters stated by TI for that component or service voids all express and any implied warranties for the associated TI component or service and is an unfair and deceptive business practice. TI is not responsible or liable for any such statements.

Buyer acknowledges and agrees that it is solely responsible for compliance with all legal, regulatory and safety-related requirements concerning its products, and any use of TI components in its applications, notwithstanding any applications-related information or support that may be provided by TI. Buyer represents and agrees that it has all the necessary expertise to create and implement safeguards which anticipate dangerous consequences of failures, monitor failures and their consequences, lessen the likelihood of failures that might cause harm and take appropriate remedial actions. Buyer will fully indemnify TI and its representatives against any damages arising out of the use of any TI components in safety-critical applications.

In some cases, TI components may be promoted specifically to facilitate safety-related applications. With such components, TI’s goal is to help enable customers to design and create their own end-product solutions that meet applicable functional safety standards and requirements. Nonetheless, such components are subject to these terms.

No TI components are authorized for use in FDA Class III (or similar life-critical medical equipment) unless authorized officers of the parties have executed a special agreement specifically governing such use.

Only those TI components which TI has specifically designated as military grade or “enhanced plastic” are designed and intended for use in military/aerospace applications or environments. Buyer acknowledges and agrees that any military or aerospace use of TI components which have ***not*** been so designated is solely at the Buyer's risk, and that Buyer is solely responsible for compliance with all legal and regulatory requirements in connection with such use.

TI has specifically designated certain components as meeting ISO/TS16949 requirements, mainly for automotive use. In any case of use of non-designated products, TI will not be responsible for any failure to meet ISO/TS16949.

Products		**Applications**	
Audio	www.ti.com/audio	Automotive and Transportation	www.ti.com/automotive
Amplifiers	amplifier.ti.com	Communications and Telecom	www.ti.com/communications
Data Converters	dataconverter.ti.com	Computers and Peripherals	www.ti.com/computers
DLP® Products	www.dlp.com	Consumer Electronics	www.ti.com/consumer-apps
DSP	dsp.ti.com	Energy and Lighting	www.ti.com/energy
Clocks and Timers	www.ti.com/clocks	Industrial	www.ti.com/industrial
Interface	interface.ti.com	Medical	www.ti.com/medical
Logic	logic.ti.com	Security	www.ti.com/security
Power Mgmt	power.ti.com	Space, Avionics and Defense	www.ti.com/space-avionics-defense
Microcontrollers	microcontroller.ti.com	Video and Imaging	www.ti.com/video
RFID	www.ti-rfid.com		
OMAP Applications Processors	www.ti.com/omap	**TI E2E Community**	e2e.ti.com
Wireless Connectivity	www.ti.com/wirelessconnectivity		

Mailing Address: Texas Instruments, Post Office Box 655303, Dallas, Texas 75265