

DDR3 Design Requirements for KeyStone Devices

High-Performance and Multicore Processors

ABSTRACT

This application report provides implementation instructions for the DDR3 interface incorporated in the Texas Instruments (TI) KeyStone series of DSP devices. The DDR3 interface supports 1600 MT/s and lower memory speeds in a variety of topologies. For supported speeds, see the specific-device data manual. This document assumes that you have a familiarization with DRAM implementation concepts and constraints.

Contents

1	Introduction	2
2	Background	2
3	Migrating Designs from DDR2 to DDR3 (Features and Comparisons)	2
4	Prerequisites	6
5	Package Selection	13
6	Physical Design and Implementation	14
7	Simulation and Modeling	34
8	Power	35
9	Disclaimers	37
10	References	37

List of Figures

1	Typical DDR Balanced Line Topology	3
2	Typical DDR3 Fly-By Architecture	4
3	Control and Address Overshoot and Undershoot Requirements	10
4	Data, Clock, Strobe, and Mask Overshoot and Undershoot Requirements	10
5	Signal Overshoot Calculations	11
6	Signal Undershoot Calculations	12
7	DSP-to-SDRAM Connection Example	14
8	DSP to UDIMM Connection	15
9	Interface Topology for Single and Dual Rank	19
10	DDRCLKOUT and DQ/DQS Routing From the DSP to SDRAMs	22
11	Data Group Impedances During Write Cycles on KeyStone I	31
12	Data Group Impedances During Read Cycles on KeyStone I	31
13	Fly-By Impedances on KeyStone I	32
14	Data Group Impedances During Write Cycles on KeyStone II	33
15	Data Group Impedances During Read Cycles on KeyStone II	33
16	Fly-By Impedances on KeyStone II	34
17	DDR3 Simulations	35

List of Tables

1	×8 Width DDR3 SDRAM Possible Configurations Supported	7
2	×16 Width DDR3 SDRAM Possible Configurations Supported	7
3	×32 Width DDR3 SDRAM Possible Configurations Supported	8

4	Discrete SDRAM Configurations with ECC	8
5	DDR3 SDRAM Selection Criteria for KeyStone I and KeyStone II Devices	9
6	OS and US Requirements for Address and Control Lines	9
7	OS and US Requirements for CK, $\overline{\text{CK}}$, DQ, DQS, $\overline{\text{DQS}}$, and DM Lines	10
8	DDR3 Single-Ended Output Levels	12
9	DDR3 Differential Output Levels	12
10	SDRAM Net Class Routing Rules	18
11	Address and Command Line Numeric Routing Rules.....	20
12	Control Line Numeric Routing Rules	20
13	Data Lane Numeric Routing Rules.....	21
14	Data and Data Strobe Byte Lane Grouping	21
15	Clock Lane Numeric Routing Rules	22
16	Maximum Write Leveling Skew Example	25
17	Minimum Write Leveling Skew Example	25
18	Maximum Round Trip Delay Example - Invert Clock Out Enabled and Disabled.....	27

Trademarks

All trademarks are the property of their respective owners.

1 Introduction

The primary goal of this document is to establish a minimum set of requirements necessary to help assure functional success in new application designs for Texas Instruments high performance multiprocessor DSPs incorporating DDR3 memory interfaces.

2 Background

Technological advances in memory architecture in both speed and densities for DDR3 require a different mindset when it comes to application implementation and design compared to the customary and traditional SRAM, DDR, and DDR2 devices.

3 Migrating Designs from DDR2 to DDR3 (Features and Comparisons)

This section is not intended to present a detailed listing of differences between DDR2 and DDR3 designs, but to provide key insight into specific differences that will have a positive impact as customers migrate from a DDR2 to a DDR3 platform (based on the assumption the DDR3 interface is implemented correctly).

3.1 Topologies

In a DDR2 to DDR3 comparison, the single greatest improvement from a topology standpoint is the change from a *Balanced T* to a *Fly-By* architecture. Each architecture is described briefly below.

3.1.1 Balanced Line Topology

In a traditional DDR2 design, a balanced T style topology is typically recommended (if not required) for address and control lines (depending on the number of SDRAMs used). This is generally recommended to balance any delays to each SDRAM device. The general concept of a balanced line topology is not used in DDR3 implementations in favor of fly-by topology, which better accommodates the higher-performance SDRAMs.

Figure 1 shows the general concept of the balanced line topology found in a typical DDR2 design.

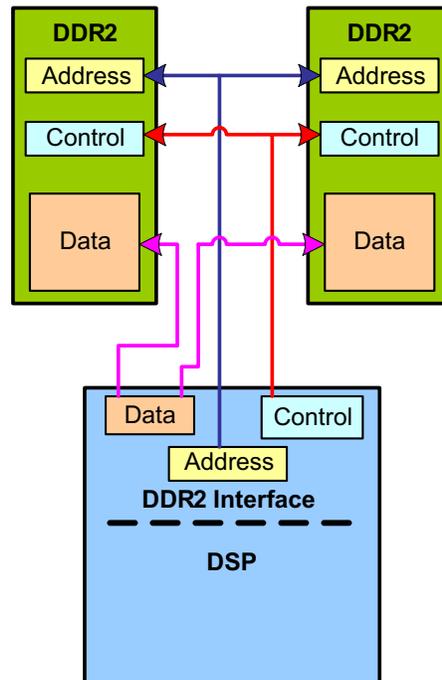


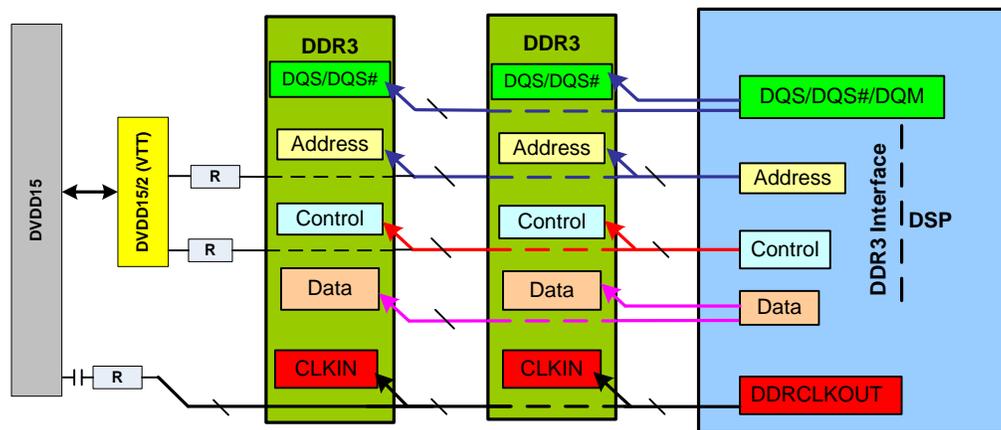
Figure 1. Typical DDR Balanced Line Topology

3.1.1.1 *Balanced Line Topology Issues*

The down side to the use of a balanced T line topology for DDR2 designs is that it may introduce a varying amount of additional skew because of the inclusion of multiple stubs and stub lengths for each individual net. The addition of multiple loads on respective address and control nets limits bandwidth. Skews normally encountered between the address/control and data nets also cause bandwidth limitations.

3.1.2 *Fly-By Topology*

The DDR3 fly-by architecture provides a benefit to layout and routing of control and address signals. In this topology, each respective signal from the DSP DDR3 controller is routed sequentially from one SDRAM to the next, thus eliminating reflections associated with any stub or superfluous traces previously seen in DDR2 designs. Figure 2 shows a typical DDR3 SDRAM fly-by interface in a single rank topology.



- A DVDD15 is a 1.5-V supply rail common to both the SDRAM and DSP.
- B DVDD15/2 or VTT refers to 750 mV.
- C Differential-ended terminations require dedicated discrete components per complementary net.

Figure 2. Typical DDR3 Fly-By Architecture

3.1.2.1 *Balanced Line Topology Issues*

The down side to the use of a fly-by topology for DDR3 designs is the induced delay from the DSP DDR3 controller to the SDRAMs. In fact, the delay is different at each SDRAM. Correction or compensation for the different controller-to-DRAM lengths is handled through read and write leveling. Also keep in mind the following:

- Data nets are point-to-point unless designed as a dual rank implementation.
- Clock nets are point-to-multipoint and end-terminated (differently than control, command, or address nets).
- Control nets are point-to-multipoint and end-terminated to VTT.
- Address/Command nets are point-to-multipoint and are end-terminated to VTT.

3.2 *ECC (Error Correction)*

Error correction has not been supported in previous TI DSPs using JEDEC-compliant SDRAM (DDR or DDR2). Error correction is now supported in TI's new KeyStone DSP processor family. Supporting ECC allows for the automatic detection and correction of single-bit and double-bit errors. ECC software configuration and control is described in detail in the *KeyStone Architecture DDR3 Memory Controller User's Guide* ([SPRUGV8](#)).

3.3 *DDR3 Features and Improvements*

In addition to TI's DDR3 controller now supporting ECC, the DSP DDR3 controller and new DDR3 SDRAMs offer five notable additional features:

- Read leveling
- Write leveling
- Change in pre-fetch size
- ZQ calibration
- A reset pin

These features are described briefly below. Detailed configuration is described in the *KeyStone Architecture DDR3 Memory Controller User's Guide* ([SPRUGV8](#)).

3.3.1 Read Leveling

The memory controller also automatically corrects for delay skew between SDRAMs during read leveling. Read leveling takes advantage of values loaded into the SDRAMs multi-purposed register (MPR). The values loaded into this register are used by the DSP DDR3 controller to calibrate each signal path relative to skew. Each respective SDRAM byte is then internally corrected thus improving performance.

3.3.2 Write Leveling

The memory controller automatically corrects for delay skew between SDRAMs during write leveling. During write leveling, correction for SDRAM skew (the tDQSS, tDSS, and tDSH) is handled using a programmable DQS delay which aligns the timing relationship to the clock and strobe signals. During the write-leveling procedure, the DSP controller delays the DQS until a valid change of state is detected at the SDRAM clock (CK) signal (see [Section 6.3](#) for additional details).

3.3.3 Pre-fetch

The new DDR3 architecture now supports an eight-bit pre-fetch to improve back-to-back accesses (DDR2 allowed only a four-bit pre-fetch).

3.3.4 ZQ Calibration

ZQ calibration is intended to control the on-die termination (ODT) values and output drivers (R_{TT} and R_{ON} respectively) of the SDRAM. ZQ calibration is not a controllable feature from the DSP. It is controlled using a precision ($\leq 1\%$ tolerance) 240- Ω resistor.

The DDR3 SDRAM ZQ calibration cycle is made up of an initial long calibration cycle (ZQ_{CL}) requiring 512 clock cycles to complete (which is why it is typically performed during the initial boot or reset conditions) and a shorter ZQ calibration period.

The subsequent short (ZQ_{CS}) calibration requires only 64 clock cycles and used when the SDRAM is idle. The periodic short calibrations cycles accommodate minor variations in temperature and voltage. The short calibration cycle (ZQ_{CS}) is designed to correct for a minimum 0.5% impedance error within the allotted 64 clock cycles.

See the selected SDRAM data sheet for the maximum ODT, temperature and voltage sensitivity values. The ZQ calibration is intended to help minimize PCB impedance discontinuities between traces and SDRAM drivers.

NOTE: Texas Instruments requires the use of a dedicated ZQ resistor (240- Ω) to be connected to each SDRAM ZQ pin (cannot share pins).

3.3.5 Reset Pin Functionality

The new DDR3 architecture also supports a reset pin. This reset pin is designed to allow the user to clear all data (information) stored the DDR3 SDRAM. The advanced benefit of this feature is that there is no need to reset each control register separately or restart (power down and up again) each individual DDR3 SDRAM. By initiating a reset, the SDRAM will recover in a known good state (if needed).

The reset function of the UDIMM or SDRAM are an active-low ($\overline{\text{RESET}}$) LVCMOS input and referenced to VSS. The SDRAM/UDIMM input pin functions rail-to-rail with a DC HIGH $\geq 0.8 \times VDD$ ($1.5 \text{ V} \times 0.8 = 1.2 \text{ V}$) and DC LOW $\leq 0.2 \times VDD$ ($1.5 \text{ V} \times 0.2 \text{ V} = 0.3 \text{ V}$).

CAUTION

The TI DSP DDR3 controller cannot be held in reset for more than one hour during the initial power-up. Also, the TI DSP DDR3 controller cannot be held in reset for more than 5% of its total power-on hours. Exceeding these limits will cause a gradual reduction in the reliability of the part.

3.3.6 Additional DDR2 to DDR3 Differences

The change from supporting both a single-ended and differential DQS (DDR2) to only a differential DQS (DDR3) improves noise immunity, and allows for longer signal paths without compromising signal integrity.

4 Prerequisites

4.1 High Speed Designs

The goal of this document is to make the DDR3 system implementation and integration easier while reducing the added risk of designing in a high performance interface. It is still expected that the PCB design work (design, layout, and fabrication) is performed, supervised, or reviewed by a highly knowledgeable high-speed PCB designer. This includes a thorough understanding of all high-speed design rules. Specific areas to avoid include ground plane cuts, incorrect spacing, and signal skew mismatches as well as timing violations. The total system should be evaluated for such areas including power, filtering, termination, crosstalk, and EMI.

4.2 JEDEC DDR3 Specification – Compatibility and Familiarity

The DDR3 interface on KeyStone I devices is designed to be compatible with the JEDEC JESD79-3C DDR3 specification. The DDR3 interface on KeyStone II devices is designed to be compatible with the JEDEC JESD79-3E DDR3 specification. It is assumed that the reader is familiar with these specifications and the basic electrical operation of the interface. In addition, several memory manufacturers provide detailed application reports on DDR3 operation.

4.3 Memory Types

Devices from many manufacturers were available at the time this document was prepared. It is recommended that only high-quality DRAM parts that fully comply with the latest DDR3 JEDEC specification and are from well-known manufacturers be used. Incorrect DRAM part or manufacturer selection can cause performance problems that may not be resolvable within the available software or configuration limitations of the DSP.

In addition to supporting the DDR3 SDRAMs, the TI KeyStone family of DSPs also supports SDRAMs including ECC (error correction) or UDIMMs (unbuffered). The KeyStone family of DSPs also supports DDR3L SDRAMs as they are backwards compatible to 1.5 V. DDR3L SDRAMs operating at 1.35 V are not supported at this time.

4.4 Memory Speeds

The KeyStone DSP memory interface currently supports various configurations as specified in the JEDEC DDR3 standard. KeyStone devices support data rates of DDR3 1600 MT/s and lower. For supported data rates, see the device-specific data manual.

4.5 Addressable Memory Space

KeyStone devices can address up to 8GB of contiguous memory. Note that some DDR3 memory interfaces are limited to a smaller memory space such as 2GB. For more information, see the specific device-specific data manual.

4.6 DDR3 SDRAM/UDIMM Memories, Topologies, and Configurations

4.6.1 Topologies

The current DDR3 controller design implementation supports several different DRAM topologies. The following list describes the known available topologies to be used with the TI KeyStone DSPs. For a definitive confirmation on available memory topologies for the DSP before proceeding, see the device-specific data manual.

- x8 (x16 is the minimum supported width, two x8 devices are required)
- x16

- x32 (with and without ECC)
- x64 (with and without ECC)
- UDIMM (unbuffered DIMM) [x72 & x36]

4.6.2 Configurations

The current DDR3 controller design implementation allows multiple DRAM configurations to be used. The following list describes the known usable

DRAM configurations compatible with TI KeyStone family of DSPs. For a definitive confirmation on usable memory configurations and bus widths before proceeding, see the device-specific data manual.

Table 1 through Table 4 are intended to provide a general overview of the possible DDR3 DRAM topologies usable with the KeyStone DSP. In all cases, only JEDEC-compliant (JESD79-3C) SDRAMs are supported. For all of the tables below, the notation * implies possible support, not a plan of record.

Table 1. x8 Width DDR3 SDRAM Possible Configurations Supported

Device Width	Total Memory / Memory Topology	Rank Width	Total Size
x8 SDRAM	1Gb / (16M × 8 × 8) × 2 SDRAMs	x16	256MB
	1Gb / (16M × 8 × 8) × 4 SDRAMs	x32	512MB
	1Gb / (16M × 8 × 8) × 8 SDRAMs	x64	1024MB
	1Gb / (16M × 8 × 8) × 8 SDRAMs × 2 ranks	x64	2048MB
	2Gb / (32M × 8 × 8) × 2 SDRAMs	x16	512MB
	2Gb / (32M × 8 × 8) × 4 SDRAMs	x32	1024MB
	2Gb / (32M × 8 × 8) × 8 SDRAMs	x64	2048MB
	2Gb / (32M × 8 × 8) × 8 SDRAMs × 2 ranks	x64	4096MB
	4Gb / (64M × 8 × 8) × 2 SDRAMs	x16	1024MB
	4Gb / (64M × 8 × 8) × 4 SDRAMs	x32	2048MB
	4Gb / (64M × 8 × 8) × 8 SDRAMs	x64	4096MB
	4Gb / (64M × 8 × 8) × 8 SDRAMs × 2 ranks	x64	8192MB
	8Gb / (128M × 8 × 8) × 2 SDRAMs	x16	2048MB
	8Gb / (128M × 8 × 8) × 4 SDRAMs	x32	4096MB
	8Gb / (128M × 8 × 8) × 8 SDRAMs	x64	8192MB

Table 2. x16 Width DDR3 SDRAM Possible Configurations Supported

Device Width	Total Memory / Memory Topology	Rank Width	Total Size
x16 SDRAM	1Gb / (8M × 16 × 8) × 1 SDRAMs	x16	128MB
	1Gb / (8M × 16 × 8) × 2 SDRAMs	x32	256MB
	1Gb / (8M × 16 × 8) × 4 SDRAMs	x64	512MB
	1Gb / (8M × 16 × 8) × 4 SDRAMs × 2 ranks	x64	1024MB
	2Gb / (16M × 16 × 8) × 1 SDRAM	x16	256MB
	2Gb / (16M × 16 × 8) × 2 SDRAMs	x32	512MB
	2Gb / (16M × 16 × 8) × 4 SDRAMs	x64	1024MB
	2Gb / (16M × 16 × 8) × 4 SDRAMs × 2 ranks	x64	2048MB
	4Gb / (32M × 16 × 8) × 1 SDRAMs	x16	512MB
	4Gb / (32M × 16 × 8) × 2 SDRAMs	x32	1024MB
	4Gb / (32M × 16 × 8) × 4 SDRAMs	x64	2048MB
	4Gb / (32M × 16 × 8) × 4 SDRAMs × 2 ranks	x64	4096MB
	8Gb / (64M × 16 × 8) × 1 SDRAMs	x16	1024MB
	8Gb / (64M × 16 × 8) × 2 SDRAMs	x32	2048MB
	8Gb / (64M × 16 × 8) × 4 SDRAMs	x64	4096MB
	8Gb / (64M × 16 × 8) × 4 SDRAMs × 2 ranks	x64	8192MB

Table 3. x32 Width DDR3 SDRAM Possible Configurations Supported

Device Width	Total Memory / Memory Topology	Rank Width	Total Size
x32 SDRAM	2Gb / (8M x 32 x 8) x 1 SDRAM	x32	256MB*
	2Gb / (8M x 32 x 8) x 2 SDRAMs	x64	512MB*
	2Gb / (8M x 32 x 8) x 1 SDRAM x 2 ranks	x32	512MB*
	2Gb / (8M x 32 x 8) x 2 SDRAMs x 2 ranks	x64	1024MB*
	4Gb / (16M x 32 x 8) x 1 SDRAM	x32	512MB*
	4Gb / (16M x 32 x 8) x 2 SDRAM	x64	1024MB*
	4Gb / (16M x 32 x 8) x 1 SDRAM x 2 ranks	x32	1024MB*
	4Gb / (16M x 32 x 8) x 2 SDRAMs x 2 ranks	x64	2048MB*

In addition to the discrete SDRAM devices configurations listed above, the following ECC configurations are supported.

Table 4. Discrete SDRAM Configurations with ECC

Device Width	Memory rank Topology	Rank Width	Total Size
x8 SDRAM	1Gb (16M x 8 x 8) x 5 SDRAMs	x36	512MB
	1Gb (16M x 8 x 8) x 9 SDRAMs	x72	1024MB
	1Gb (16M x 8 x 8) x 9 SDRAMs x 2 ranks	x72	2048MB
	2Gb (32M x 8 x 8) x 5 SDRAMs	x36	1024MB
	2Gb (32M x 8 x 8) x 9 SDRAMs	x72	2048MB
	2Gb (32M x 8 x 8) x 9 SDRAMs x 2 ranks	x72	4096MB
	4Gb (64M x 8 x 8) x 5 SDRAMs	x36	2048MB
	4Gb (64M x 8 x 8) x 9 SDRAMs	x72	4096MB
	4Gb (64M x 8 x 8) x 9 SDRAMs x 2 ranks	x72	8192MB
	8Gb (64M x 8 x 8) x 5 SDRAMs	x36	4096MB
	8Gb (64M x 8 x 8) x 9 SDRAMs	x72	8192MB
x16 SDRAM ⁽¹⁾	1Gb (8M x 16 x 8) x 3 SDRAMs	x36	256MB
	1Gb (8M x 16 x 8) x 5 SDRAMs	x72	512MB
	1Gb (8M x 16 x 8) x 5 SDRAMs x 2 ranks	x72	1024MB
	2Gb (16M x 16 x 8) x 3 SDRAMs	x36	512MB
	2Gb (16M x 16 x 8) x 5 SDRAMs	x72	1024MB
	2Gb (16M x 16 x 8) x 5 SDRAMs x 2 ranks	x72	2048MB
	4Gb (32M x 16 x 8) x 3 SDRAMs	x36	1024MB
	4Gb (32M x 16 x 8) x 5 SDRAMs	x72	2048MB
	4Gb (32M x 16 x 8) x 5 SDRAMs x 2 ranks	x72	4096MB
	8Gb (64M x 16 x 8) x 3 SDRAMs	x36	2048MB
	8Gb (64M x 16 x 8) x 5 SDRAMs	x72	4096MB
	8Gb (64M x 16 x 8) x 5 SDRAMs x 2 ranks	x72	8192MB

⁽¹⁾ The ECC device can be either x8 or x16 as long as the number of row and column address bits match for all devices in the memory array.

4.6.2.1 Memories – SDRAM Selection Criteria

Table 5 shows the DDR3 SDRAM selection criteria necessary for KeyStone I and KeyStone II devices.

Table 5. DDR3 SDRAM Selection Criteria for KeyStone I and KeyStone II Devices

Description	Min	Max	Unit	Notes
Width	x8	x16	bit	Minimum supported total bus width is x16 (2- x8)
Depth (Density)	512M	8192M	bit	Definition may also include ranked devices (UDIMM)
Data Rate	800	1600	MT/s	
Clock Rate	400	800	MHz	
Temperature Range	0	95	°C	Depends on end-use application
VDD	1.425	1.575	V	
VDDq	1.425	1.575	V	
Latency	5	11		All CAS latencies supported between 5 and 11
JEDEC-compliant DDR3 SDRAM				

DIMM support extends to modules whose SDRAMs fall within the above specifications. DIMM bus width may be x64 or x72 with ECC support. DIMMs must be unbuffered (UDIMMs), and may have up to two ranks. KeyStone I devices do not support address mirroring. KeyStone II devices support UDIMM topologies that have the address bits in the second rank mirrored.

4.7 DRAM Electrical Interface Requirements

This section briefly defines the electrical interface requirements for using JEDEC-compliant DDR3 SDRAM with the KeyStone DSP. Additional information and requirements may exist. Where different or conflicting requirements exist (between applicable standards, SDRAM and DSP data sheets, and this design guide), this design guide should take precedence. Details provided in the following subsection were obtained from TI internal reference material and applicable JEDEC DDR3 SDRAM standards.

4.7.1 Slew

The released JEDEC standard describes in detail skew requirements imposed on the SDRAM, in particular CK, \overline{CK} , DQS, and \overline{DQS} . In order to meet these requirements, loading, SDRAM component selection, and trace routing will have a large impact. For additional information regarding meeting slew rate requirements, see the modeling and simulation section of this guide.

4.7.2 Overshoot and Undershoot Specifications

Overshoot (OS) and undershoot (US) limitations are defined in the following tables and figures. A theoretical instantaneous maximum upper and lower amplitude limit of 400 mV is allowed for overshoots and undershoots (assuming zero time is involved). Because each overshoot and undershoot has a component of time, each applicable signal must be further independently evaluated. Table 6 shows the limitations as listed in the current released standard for all address and control signals. (Measurements are obtained at the pin of the memory DRAM and not the DSP.)

Table 6. OS and US Requirements for Address and Control Lines

Speed	Max pk OS Amplitude (V)	Max pk US Amplitude (V)	Max OS Area Above VDD (V-ns)	Max US Area Below VSS (V-ns)
DDR3-800	0.4	0.4	0.67	0.67
DDR3-1066	0.4	0.4	0.5	0.5
DDR3-1333	0.4	0.4	0.4	0.4
DDR3-1600	0.4	0.4	0.33	0.33

Figure 3 shows the limitations as listed in the current released standard for all address and control pins.

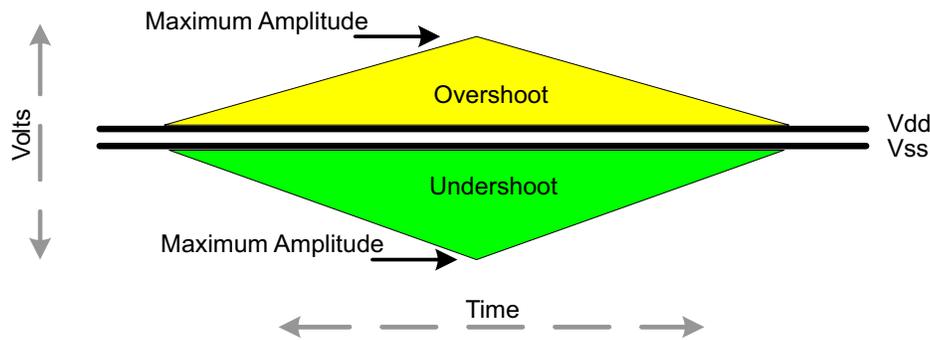


Figure 3. Control and Address Overshoot and Undershoot Requirements

Table 7 shows the limitations as listed in the current released standard for all data, clock, strobe, and mask signals.

Table 7. OS and US Requirements for CK, $\overline{\text{CK}}$, DQ, DQS, $\overline{\text{DQS}}$, and DM Lines

Speed	Max pk OS Amplitude (V)	Max pk US Amplitude (V)	Max OS Area Above VDD (V-ns)	Max US Area Below VSS (V-ns)
DDR3-800	0.4	0.4	0.25	0.25
DDR3-1066	0.4	0.4	0.19	0.19
DDR3-1333	0.4	0.4	0.15	0.15
DDR3-1600	0.4	0.4	0.13	0.13

Figure 4 shows the limitations as listed in the current released standard for all data, clock, strobe, and mask signals.

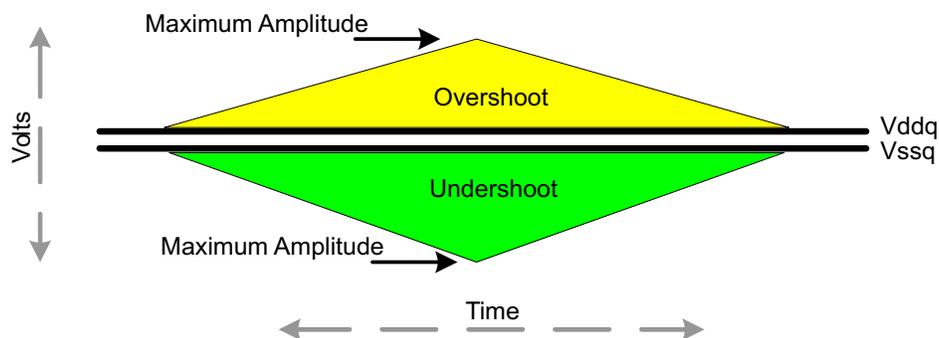


Figure 4. Data, Clock, Strobe, and Mask Overshoot and Undershoot Requirements

4.7.2.1 Overshoot and Undershoot Example Calculations

The following is provided as an example of the necessary steps to calculate the overshoots and undershoots for any waveform.

Assumptions: VDD = 1.5 V; VSS = 0.00 V; Vref = 0.75 V (VDD/2)

Overshoot example (Figure 5)

1. Determine amplitude over VDD
2. Determine the duration of the amplitude
3. Calculate the final value
 - a. OS = Amplitude \times Duration
 - b. OS = 180 mV \times 1 ns
 - c. OS = 0.18 V-ns
4. Compare results to the applicable SDRAM row and column (Section 4.7.2) paying attention to speed grade and signal type differences. If the SDRAM speed was DDR3-800 or DDR3-1066 and this were a data, clock, strobe, or mask net, the overshoot would be acceptable. In the case of data, clock, strobe, or mask DDR3-1333 and DDR3-1600 speed grades, this level of overshoot is not acceptable.

NOTE: This level is acceptable for all speed grades of control and address lines.

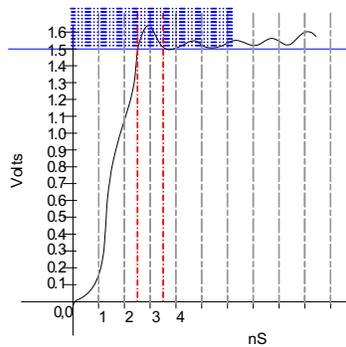


Figure 5. Signal Overshoot Calculations

Undershoot example (Figure 6)

1. Determine amplitude under VSS
2. Determine the duration of the amplitude
3. Calculate the final value
 - a. $US = \text{Amplitude} \times \text{Duration}$
 - b. $US = 215 \text{ mV} \times 1.5 \text{ ns}$
 - c. $US = 0.322 \text{ V}\cdot\text{ns}$
4. Compare results to the applicable SDRAM row and column (Section 4.7.2 paying attention to speed grade and signal type differences). If the SDRAM was DDR3-800 or DDR3-1066 and this were a data net, this overshoot would be acceptable. In the case of DDR3-1333 and DDR3-1600, this level of undershoot is not acceptable. Compare results to the applicable SDRAM row and column (notice speed grade differences). This example undershoot is unacceptable for all speed grades for all data, clock, strobe, and mask signals. It is acceptable for all speed grades of address and control signals.

NOTE: This level is acceptable for all speed grades of control and address lines.

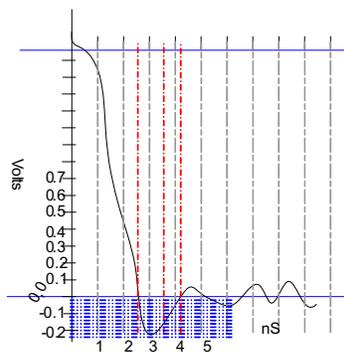


Figure 6. Signal Undershoot Calculations

4.7.3 Typical DDR3 AC and DC Characteristics

Table 8. DDR3 Single-Ended Output Levels

Symbol	Parameter	DDR3-800, 1066, 1333, and 1600	Unit
$V_{OH(DC)}$	DC output high measurement level (for IV curve linearity)	$0.8 \times V_{DDQ}$	V
$V_{OM(DC)}$	DC output mid measurement level (for IV curve linearity)	$0.5 \times V_{DDQ}$	V
$V_{OL(DC)}$	DC output low measurement level (for IV curve linearity)	$0.2 \times V_{DDQ}$	V
$V_{OH(AC)}$	AC output high measurement level (for output SR)	$V_{TT} + 0.1 \times V_{DDQ}$	V
$V_{OL(AC)}$	AC output low measurement level (for output SR)	$V_{TT} - 0.1 \times V_{DDQ}$	V

Table 9. DDR3 Differential Output Levels

Symbol	Parameter	DDR3-800, 1066, 1333, and 1600	Unit
$V_{OH,diff(AC)}$	AC differential output high measurement level (for output SR)	$+ 0.2 \times V_{DDQ}$	V
$V_{OL,diff(AC)}$	AC differential output low measurement level (for output SR)	$- 0.2 \times V_{DDQ}$	V

4.7.4 DDR3 Tolerances and Noise – Reference Signals

Limitations on DC voltage tolerance and AC noise for all reference voltages is well defined in the applicable JEDEC standard (pg. 129 of JESD79-3C). Strict conformity to these limitations is important to help ensure proper functionality of the DDR3 SDRAM interface. The V_{ref} tolerance is $\pm 1\%$ or $VDD/2 \pm 1\%$, which equates to 0.7425 V – 0.7575 V. To achieve this tight tolerance, it is recommended (when using a standard resistor divider network) that better than 1% tolerance components be used. The alternative would be an active reference voltage source. Proper component selection and decoupling is critical. For additional details, see the layout and routing section.

It is important to properly design the reference supply voltage to track VDD/VDDq.

5 Package Selection

Individual SDRAMs are available in four different packages that support the various memory device densities, and widths. The following summarizes and describes the primary differences between package and density for common x8 and x16 SDRAMs.

5.1 Summary

The following packaging subsection summarizes at a high level each DDR3 SDRAM configuration and pinout. For the latest information, see the manufacturers data sheets.

5.1.1 x4 SDRAM

- KeyStone devices do not support DDR3 SDRAMs in an x4 configuration.

5.1.2 x8 SDRAM

- x8 devices are compatible in both the 106-pin and 78-pin packages provided the layout is correct and accounts for the larger package inclusive of the NC (support ball) pin placements.

NOTE: If designed for the larger package, there is no difference in routing required.

- 1Gb and 2Gb devices are pin-compatible in the 106-pin and 78-pin packages
 - The 1Gb, 2Gb, and 4Gb devices are pin-compatible in the 78-pin package
-

NOTE: Other packages not available at the time this document was prepared.

- 1Gb devices, regardless of package, are not compatible between x16 and x8 devices. For additional details, see the device-specific data sheets.
-

NOTE: It is possible to layout for both x16 and x8 devices but this requires additional consideration on switching signals and net attachments, which vary between packages. As a general rule, Texas Instruments does not recommend designing the application board for both topologies [x8 and x16].

- The twin-die version of the x8 SDRAM is not compatible with the monolithic version of the same density.

5.1.3 x16 SDRAM

- 1Gb and 2Gb x16 devices are compatible in the 96-pin package.
-

NOTE: Some pinouts may change across manufacturers – review the data sheets before selecting the DRAMs.

- The twin-die version of the x16 SDRAM are not compatible with the monolithic version of the same density.

5.1.4 x32 SDRAM

- At the time this document was prepared, no x32 devices exist. Texas Instruments cannot ensure functionality or support for x32 SDRAM devices.

5.1.5 x64 SDRAM

- At the time this document was prepared, no x64 devices exist. Texas Instruments cannot ensure functionality or support for x64 SDRAM devices.

6 Physical Design and Implementation

6.1 Electrical Connections

This section discusses the proper electrical interface between JEDEC-compliant DDR3 SDRAMs and UDIMMs to the Texas Instruments KeyStone DSP family DDR3 controller. These sections also provide details for the electrical pin connectivity between the DSP DDR3 interface and DDR3 SDRAM or UDIMM. This is not an all-inclusive listing of available parts as DRAM manufacturers are continuously developing higher density parts or obsoleting others. It should also be noted that not all DSPs will have the same bus width. For details, see the device-specific data sheet. This section assumes both 32-bit and 64-bit wide buses are used.

6.1.1 Pin Connectivity & Unused Pins – SDRAM Examples

Correct DDR3 pin connectivity is vital to help insure the performance and reliability of the DSP/DRAM system. An example set of pin connections from DSP-to-SDRAM are shown in [Figure 7](#).

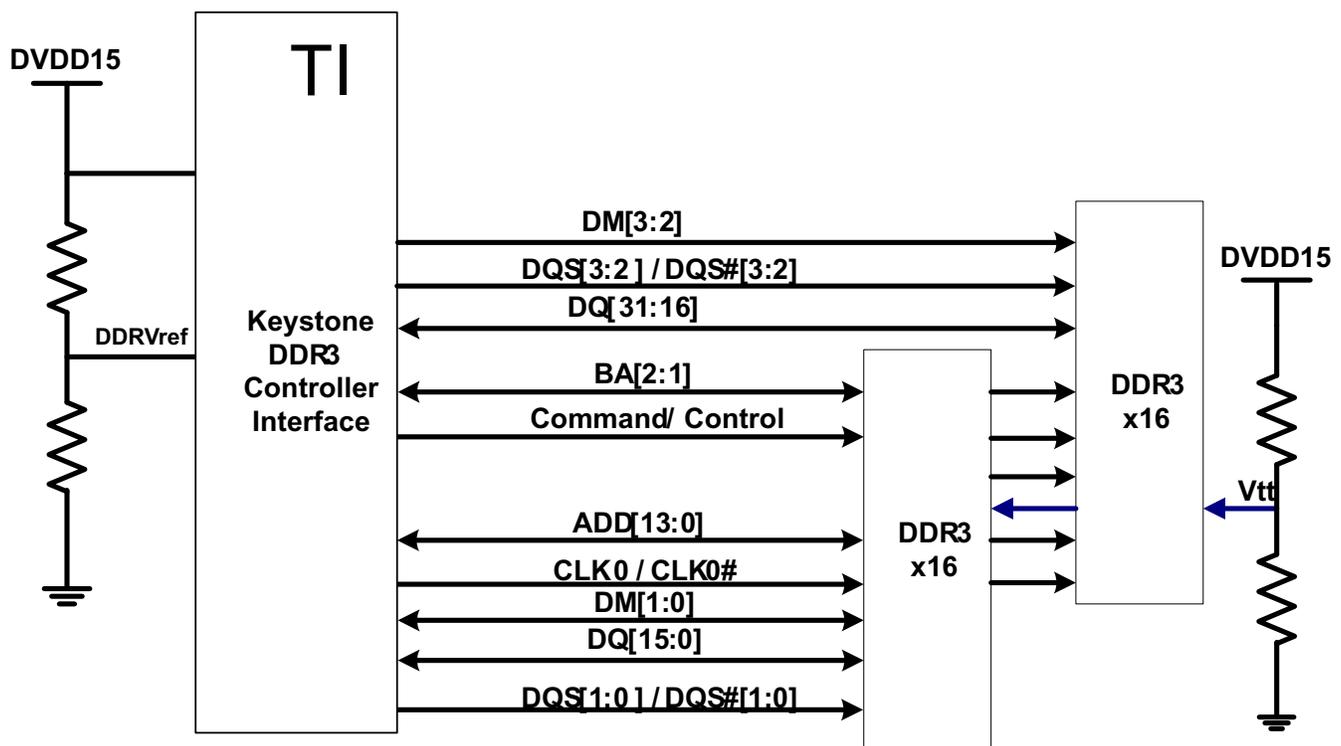


Figure 7. DSP-to-SDRAM Connection Example

Pin nomenclatures are identified in detail in the device-specific DSP and SDRAM data sheets and should be confirmed before releasing the design to layout, PCB fabrication, or production. This section does not include details for non-JEDEC compliant (JESD79-3C) SDRAM components nor does TI recommend the use of any non-compliant JEDEC SDRAMs.

6.1.2 Pin Connectivity – ECC UDIMM & Non-ECC UDIMM Examples

This section defines the recommended configurations when using a standard DDR3 UDIMM (unbuffered DIMM) with and without ECC (Error Detection and Correction). Additional details pertaining to UDIMM implementation can be found in the JEDEC DDR3 UDIMM standard 21C, the JEDEC UDIMM mechanical standard, MO-269, and the JEDEC socket standard, SO-007B (latest revision). For additional details pertaining to routing requirements, see [Section 6.3.2](#).

Correct DDR3 pin connectivity is vital to ensure the performance and reliability of the DSP/UDIMM system. For pin connectivity, see the device-specific data manual.

Pin nomenclatures are identified in detail in the device-specific DSP and UDIMM data sheets and should be confirmed before releasing the design to layout, PCB fabrication, or production. This section does not include specific details for non-JEDEC compliant UDIMMs nor does TI recommend the use of any non-compliant JEDEC UDIMMs.

Figure 8 shows the basic interconnection between the TI KeyStone DSP and device-specific ECC UDIMM.

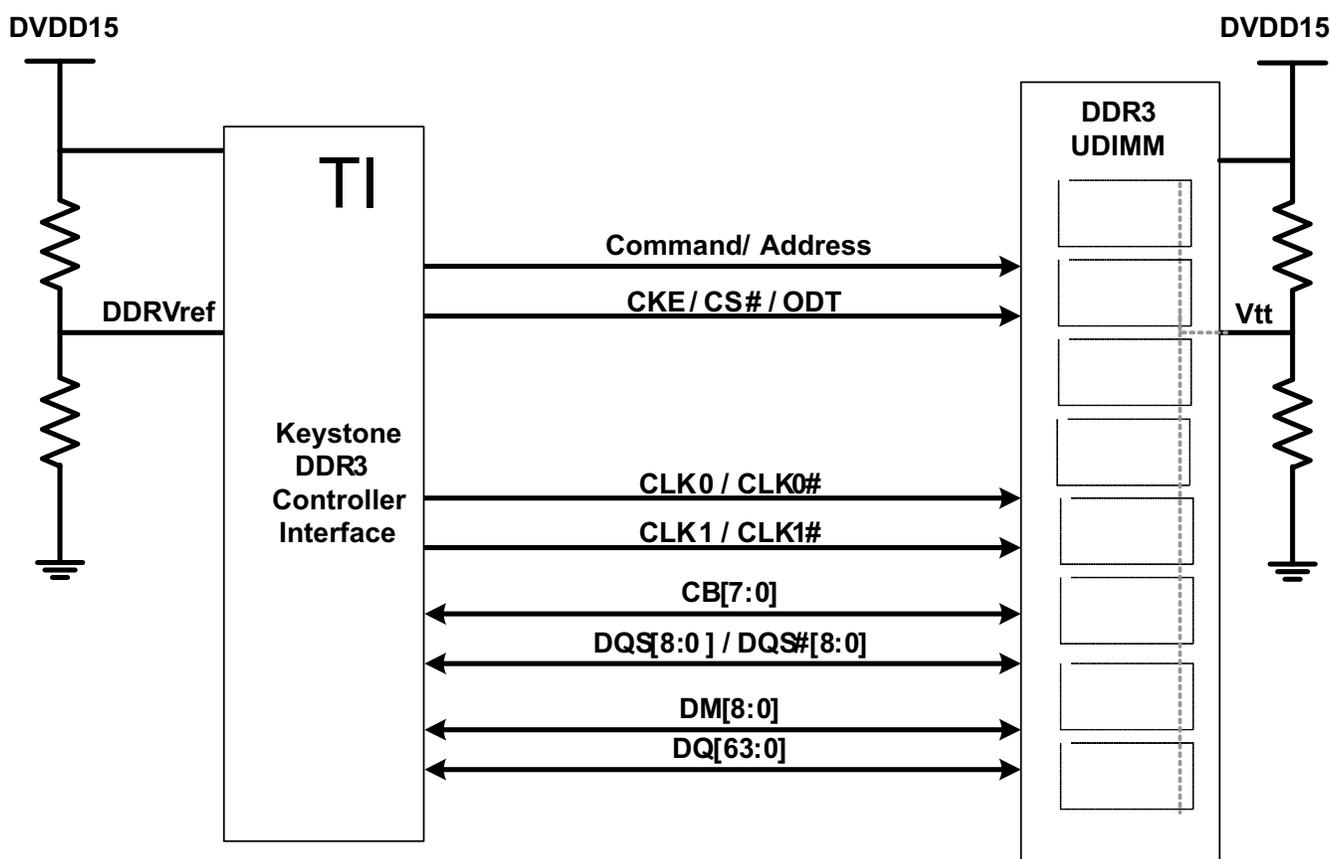


Figure 8. DSP to UDIMM Connection

6.2 Signal Terminations

The following subsection describes the two signal termination methods (leveling and non-leveling) used for DDR3 interfaces, specific termination placement, and the impact of incorrect termination schemes and component values. Terminations are placed at the end of the signal path. In the current fly-by architecture, placing the terminations at the last SDRAM improves the overall signal characteristics, which is an improvement over previous DDR2 SDRAM topology.

6.2.1 External Terminations – When Using Read and Write Leveling

As a rule of thumb, terminations should be applied to all clock, address, and control lines. Although all address, control, and command lines should be end-terminated when using leveling, clocking nets may require different termination values than that used on the command and address nets. Each respective address and command net should be end-terminated using a resistor (in the range of 39 Ω to 42 Ω) and connected to VTT (preferred value is 39 Ω , 1%). VTT is defined as $VDDq/2$ or 0.75 V.

The DDR3 clock nets also must be end-terminated. However, instead of end-terminating the clock nets (DDRCLKOUTP/Nx where x is DDRCLKOUTP/N0 or DDRCLKOUTP/N1, whichever is used), each net should be terminated with a series 39- Ω , 1% resistor to a 0.1- μ F capacitor to DVDD15 (VDDq). [Figure 2](#) (and respective notes) show the required clock termination implementation. All components should be 1% tolerance or better. VTT should be generated using a resistor divider network (1% tolerance or better). For proper operation, the VTT termination must track VDDq/2.

Again, an important point is that the parallel termination should be placed at the last SDRAM in the fly-by or daisy-chained architecture. Each trace to the respective termination should be \leq within 500 mils and the opposite side of the termination resistor should tie directly to the VTT rail.

External terminations may not be required – the only way to determine if your topology requires end terminations is to perform complete simulations inclusive of all topology parasitics.

NOTE: All VTT terminations should be placed at the end of the transmission line (net). Incorrect placement will have an impact on performance and functionality.

6.2.2 External Terminations – When Read and Write Leveling is Not Used

Texas Instruments does not recommend the use of DDR3 without leveling. The use of DDR3 SDRAMs without read and write leveling places an undue burden and constraints on the physical design and topology. In essence, any DDR3 design not making use of read and write leveling becomes nothing more than a DDR2 layout with all of the negative implications associated with DDR2 and none of the benefits of DDR3 (except possibly speed).

6.2.3 Internal Termination – On-Die Terminations

Prior DSPs supporting DDR2 interfaces did not support on-die terminations (ODT), whereas, the KeyStone family of DSPs now supports ODT. One of the primary advantages to using DDR3 is the fact that the data lines no longer require series terminations to optimize signal overshoots and undershoots. The current DDR3 instantiation allows for a wider range of values. In addition, DDR3 now supports dynamic ODT, which has enormous benefits in a complex application board topology. For DDR3, the DSP controller ODT pins (connected to each SDRAM) serve to turn on or off the SDRAM internal termination. The actual ODT functionality of each SDRAM is controlled using the mode registers. For additional information, see the device-specific SDRAM data sheets.

6.2.4 Active Terminations

Active terminations, regardless of configuration, are not required. The combination of TI's new DSP controller and the new DDR3 DRAM architecture in the supported configurations eliminates the need for active terminations.

6.2.5 Passive Terminations

Unlike DDR2, the addition of series passive components is no longer necessary when using DDR3 (with leveling). The only passive terminations to be used (when using DDR3 with leveling) are the end-termination pull-up resistors identified in [Section 8.1.2](#).

6.2.6 Termination Component Selection

All termination components should be 1% tolerance, component size (form factor) will be dependent upon power requirements and parasitics. Texas Instruments recommends 0402 size discrete (passive) components to improve routing, and reduce parasitic inductance.

6.3 Mechanical Layout and Routing Considerations

This subsection defines the basic requirements regarding mechanical layout and PCB routing. These are general guidelines, which, when followed in conjunction with the remainder of sections in this design guide (and good engineering practices), will help provide a functional DDR3 interface. In all cases, it is recommended that the DDR3 interface be modeled to verify functionality and performance.

6.3.1 Routing Considerations – SDRAMs

This subsection describes the specific requirements for discrete SDRAM implementations.

6.3.1.1 Mechanical Layout – SDRAMs

This subsection provides basic information regarding mechanical DDR3 → DSP layout constraints. Included are such topics as routing, stack up, trace lengths, and the use of net classes. Issues not covered (in any great detail) within this document include thermal considerations, part density, pick and place issues, and different packages / footprints.

This document assumes that you have an above average level of understanding regarding mechanical layout and design – including the impact of trace width, spacing, via size, bulk and decoupling capacitance selection, and placement.

6.3.1.2 Stack Up – SDRAMs

Stack up refers to the mechanical layer assembly of the printed circuit board. In all high-speed designs it is good practice to maintain symmetry between the top half of the board and the bottom half of the board. Referencing (sandwiching) signal routing layers containing high speed signals between ground planes reduces EMI problems and provides a constant and controlled impedance transmission path. DDR3 IO power planes can also be used as reference layers for address/command/control signals as long as there are decoupling capacitors distributed across the plane to provide a low-impedance return path to ground. Power and ground planes should ideally be solid planes without breaks. Power planes can have breaks or splits outside of the DDR3 routing region.

Proper stack up must also provide the proper characteristic printed circuit board (PCB) impedance. Most DSP application board systems require a PCB impedance of 50 Ω.

It is recommended that a DDR3 implementation make use of a minimum of four routing layers: two for address/command/control signals and two for data-group signals. Data-group routing should be on layers close to the bottom of the board to minimize stubs. An end-application PCB will likely have a minimum total of 8 layers if sufficient board area is available, but high-performance boards or boards making use of an extended peripheral set may consist of 12 or more layers.

6.3.1.3 Routing Rules – General Overview (SDRAMs)

Several key points to remember when routing any signals on the application board:

- Organize the power, ground, and signal planes so that you eliminate or significantly reduce the number of split or cut planes present in the design (no splits are allowed under any DDR3 routes).
- It is strongly recommended that all SDRAMs are mounted on the top side of the PCB alongside the SoC for single rank designs.
- Apply net classes, for example, group key signals together.
- Maintain an acceptable level of skew across the entire DDR3 interface (by net class).
- Use proper low-pass filtering on the Vref pins.
- Follow the fly-by architecture concept for all address, command, clock, and control lines.
- Increase the size of the decoupling capacitor trace width to as large as possible, keep the stub length as short as possible.
- Center-to-center spacing, including serpentine, must be at least 5 W where W is the trace width. Additional spacing can be added between differential pairs and other routing groups to minimize crosstalk. Spacing of 4 W can be used, but is not appropriate for bus speeds over 1066 MT/s.
- Maintain a common ground reference for all bypass/decoupling capacitors, DSPs, and SDRAMs.

- Take into account the differences in propagation delays between microstrip and stripline nets when evaluating timing constraints. All long routes should be stripline to reduce EMI and timing skew, and any microstrip routed for BGA breakouts should be as short as possible.
- All length-matching is based on an equivalent stripline length. An equivalent stripline length is defined as the length of a stripline trace that will have the same delay as the microstrip portion of the route (see the JEDEC UDIMM specification for more information on velocity compensation).
- There can not be any mid-point vias in the design on any data group net. Extra vias that are located on the data group will negatively impact signal integrity.
- It is strongly recommended that all nets be simulated to assure proper design, performance, and signal integrity.
- Routes along the same path and routing segment must have the same number of vias. Vias can be blind, buried, or HDI microvias for improved SI but are not required for standard data rates. Similarly, back drilling vias is not required for standard data rates but can be used to eliminate via stubs.
- It is strongly recommended that the routing channels between the DSP to SDRAM be dedicated solely to the SDRAM interface and that no other signals be routed in the area. Other signals routed on the same layers must be kept apart from the DDR3 routes. There must be additional separation of the DDR3 nets of at least 6W. In addition, these other traces should not be referenced to the DDR3 IO power planes. If other signals must be routed through this area, they need to be isolated to their own routing planes and shielded from the DDR3 routes by a solid ground plane.

Net classes are an important concept when routing high speed signals that incorporate timing constraints or timing relationships. When routing the DDR3 nets, there are four basic groups (net classes) to consider – [Table 10](#) shows the recommended net classes.

Table 10. SDRAM Net Class Routing Rules

Net Class	Signals
Data	DQS[8:0], \overline{DQS} [8:0], DQ[n:0], CB[7:0], DM[8:0] ^{(1), (2), (3)}
Address/Command	BA[2:0], A[n:0], \overline{RAS} , \overline{CAS} , \overline{WE} ^{(2), (4)}
Control	\overline{CS} , \overline{CKE} , \overline{ODT} ^{(2), (5)}
Clock	DDRCLKOUTP/N[1:0] ⁽⁵⁾

⁽¹⁾ CB[7:0] refer to ECC devices.

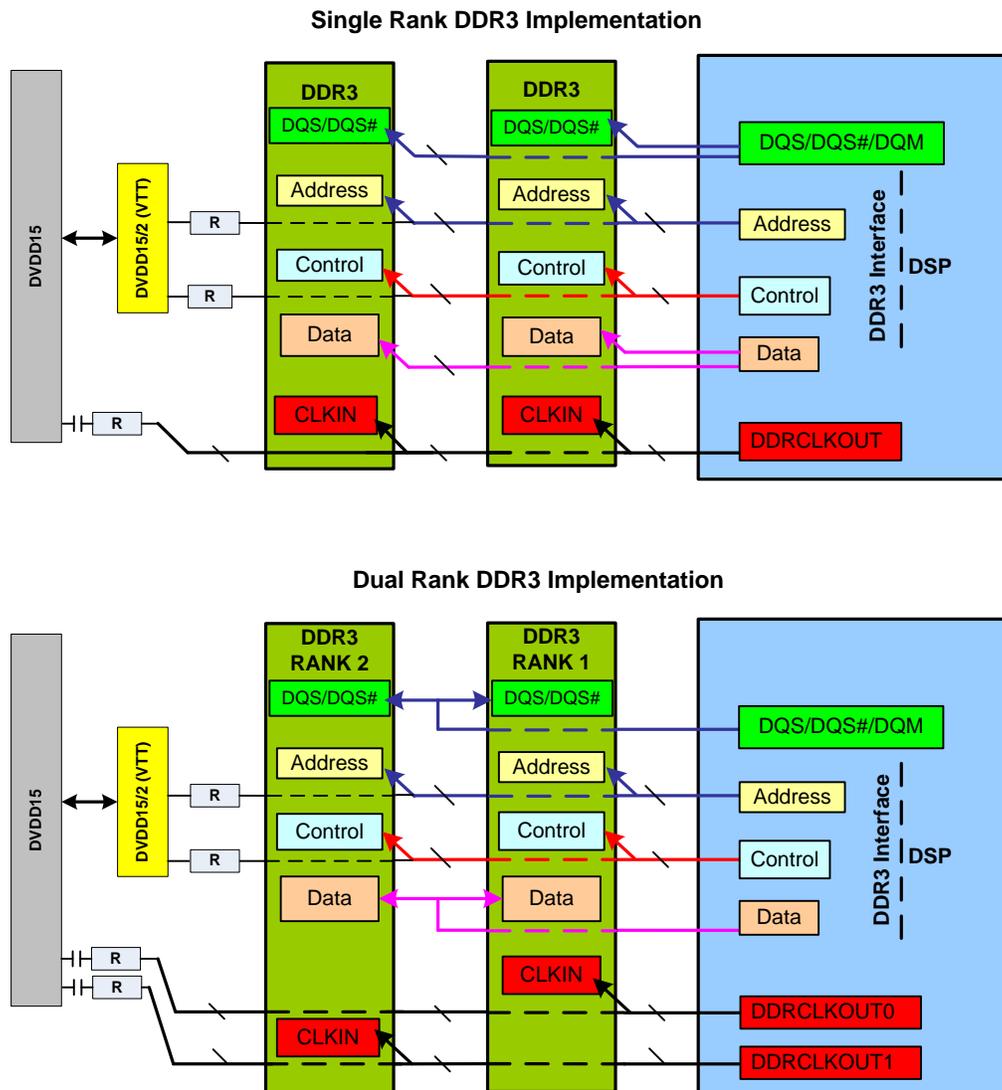
⁽²⁾ n refers to some number of lines and is dependent upon device selected.

⁽³⁾ The data net class will be subdivided into multiple byte lane routing groups each containing a $\overline{DQS}/\overline{DQS}$ pair and the associated DM and 8 data bits.

⁽⁴⁾ Observe relationship between DQ, DQS, \overline{DQS} , DDRCLKOUT, and $\overline{DDRCLKOUT}$.

⁽⁵⁾ Some nets in class may be used only in dual-rank topologies.

Figure 9 shows the typical interface topology when connecting single-rank and dual-rank designs. Note that the dual-rank illustration shows only the connections for byte lane 0, and fly-by connections will still need to include SDRAMs on subsequent byte lanes.



A The CLKOUT, CS, ODT, and CKE pins will be unique to each rank in a dual rank design.

Figure 9. Interface Topology for Single and Dual Rank

6.3.1.4 Routing Rules – Address and Command Lines (SDRAMs)

The following rules must be followed when routing address/command nets in a DDR3 design:

- 50- Ω ($\pm 5\%$) single-ended impedance required.
- All nets in the address and command fly-by groups must route along the same path from the controller to each SDRAM sequentially, and then to the VTT termination.
- All nets in the address and command fly-by groups must be length-matched from the controller to each SDRAM separately within ± 20 mils of the clock along the same route.
- All nets in the address and command fly-by groups must have the same number of vias in each length-matched segment.
- Address lines **cannot** be swapped to simplify routing.
- Address and command fly-by groups must have stubs less than 80 mils and be length-matched within ± 10 mils.
- All nets in the address and command fly-by groups must route adjacent to a solid ground plane or a solid DVDD15 power plane with adequate distributed decoupling to provide high frequency return.
- All nets in the address and command fly-by groups should be routed on close layers to minimize via skew – these are normally close to the center or upper layers of the board.

Table 11 shows the numeric routing rules listed above for address/command lines.

Table 11. Address and Command Line Numeric Routing Rules

Rule Number	Parameter	Value	Unit
1	Net Impedance (single-ended)	50	Ω
2	Skew between fly-by group nets	± 20	mils
3	Stub length	< 80	mils
4	Stub skew	± 10	mils

6.3.1.5 Routing Rules – Control Lines (SDRAMs)

The following rules must be followed when routing control nets in a DDR3 design:

- 50- Ω ($\pm 5\%$) single-ended impedance required.
- All nets in the control fly-by groups must route along the same path from the controller to each SDRAM sequentially, and then to the VTT termination.
- All nets in the control fly-by-groups must be length matched from the controller to each SDRAM separately within ± 20 mils of the clock along the same route.
- All nets in the control fly-by groups must have the same number of vias in each length-matched segment.
- Control fly-by-groups must have stubs less than 80 mils and be length-matched within ± 10 mils.
- All nets in the control fly-by groups must route adjacent to a solid ground plane or a solid DVDD15 power plane with adequate distributed decoupling to provide high frequency return.
- All nets in the control fly-by groups should be routed on close layers to minimize via skew – these are normally close to the center or upper layers of the board.

Table 12 shows the numeric routing rules listed above for control lines.

Table 12. Control Line Numeric Routing Rules

Rule Number	Parameter	Value	Unit
1	Net Impedance (Single-Ended)	50	Ω
2	Skew between fly-by group nets	± 20	mils
3	Stub length	< 80	mils
4	Stub skew	± 10	mils

6.3.1.6 Routing Rules – Data Lines (SDRAMs)

The following rules must be followed when routing data nets in a DDR3 design:

- 100-Ω (±5%) differential impedance required on data strobe (DQS) pairs.
- DQS/ $\overline{\text{DQS}}$ pairs must be routed differentially.
- 50-Ω (±5%) single-ended impedance required on point-to-point routes.
- All nets in each data group must have same number of vias - maximum of two. The number per signal within each byte group must match.
- All nets in a single byte-lane group should be routed on the same layer to eliminate the addition of length-skew from the via barrels.
- All data-group nets must route adjacent to a solid ground plane.
- All data strobe pairs must be length-matched with ±1 mils of each other.
- All nets within a single data group must be length-matched with ±10 mils.
- Data bits within a byte-lane can be swapped to simplify routing.
- Data group nets are routed point-to-point and do not have VTT terminations.

Table 13 shows the numeric routing rules listed above for data lines.

Table 13. Data Lane Numeric Routing Rules

Rule Number	Parameter	Value	Unit
1	Net Impedance (Single-Ended)	50	Ω
2	Net Impedance (Differential)	100	Ω
3	Skew between DQS pairs	±1	mils
4	Skew between data group nets for given byte lane	±10	mils

Table 14 shows the required data signal byte lane groupings.

Table 14. Data and Data Strobe Byte Lane Grouping

Data Group	Data	Data Strobe	Data $\overline{\text{Strobe}}$	Data Mask
BYTE LANE 0	DQ[7:0]	DQS0	$\overline{\text{DQS0}}$	DM0
BYTE LANE 1	DQ[15:8]	DQS1	$\overline{\text{DQS1}}$	DM1
BYTE LANE 2	DQ[23:16]	DQS2	$\overline{\text{DQS2}}$	DM2
BYTE LANE 3	DQ[31:24]	DQS3	$\overline{\text{DQS3}}$	DM3
BYTE LANE 4	DQ[39:32]	DQS4	$\overline{\text{DQS4}}$	DM4
BYTE LANE 5	DQ[47:40]	DQS5	$\overline{\text{DQS5}}$	DM5
BYTE LANE 6	DQ[55:48]	DQS6	$\overline{\text{DQS6}}$	DM6
BYTE LANE 7	DQ[63:56]	DQS7	$\overline{\text{DQS7}}$	DM7
ECC BYTE LANE	CB[7:0]	DQS8	$\overline{\text{DQS8}}$	DM8

6.3.1.7 Routing Rules – Clock Lines (SDRAMs)

The following rules must be followed when routing clock nets in a DDR3 design:

- CK/ $\overline{\text{CK}}$ pairs must be routed differentially.
- 100- Ω ($\pm 5\%$) differential impedance required on all clock pairs.
- Each CK/ $\overline{\text{CK}}$ pair is managed as a separate routing group.
- Each CK/ $\overline{\text{CK}}$ pair must be routed to all SDRAMs within a single rank.
- All nets in the clock fly-by groups must route along the same path from the controller to each SDRAM sequentially and then to an AC VTT termination.
- All clock pairs must be length-matched from the controller to each SDRAM separately within ± 1 mils of each other.
- All nets in the clock fly-by group must have the same number of vias in each length-matched segment.
- Clock pair stubs must be less than 40 mils and length-matched within \pm mil.
- All clock pair nets must route adjacent to a solid ground plane.

Table 15 shows the numeric routing rules listed above for data lines.

Table 15. Clock Lane Numeric Routing Rules

Rule Number	Parameter	Value	Unit
1	Net Impedance (differential)	100	Ω
2	Skew between CK/ $\overline{\text{CK}}$ pairs	± 1	mils
3	Stub length	<40	mils
4	Stub skew	± 1	

Figure 10 shows the required DDRCLKOUT and DQ/ $\overline{\text{DQS}}$ → DQ/ $\overline{\text{DQS}}$ / $\overline{\text{DQS}}$ routing from the DSP to SDRAMs for a single-rank topology.

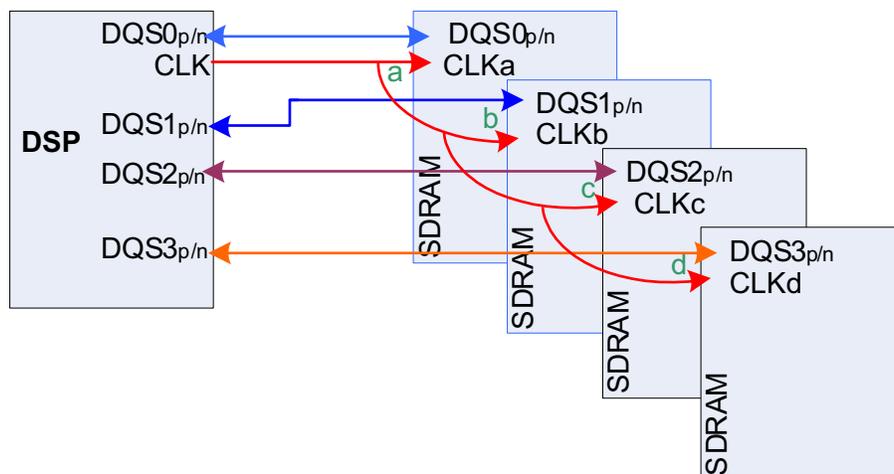


Figure 10. DDRCLKOUT and DQ/ $\overline{\text{DQS}}$ Routing From the DSP to SDRAMs

6.3.1.8 Routing Rules – Power (SDRAMs)

When routing the DDR3 SDRAM Vref voltages, it is necessary to decouple them at the SDRAMs and not at the source. The use of 0.01- μ F and 0.1- μ F ceramic capacitors (0402 or smaller recommended) should be distributed across the Vref power rail with one 0.01- μ F and 0.1- μ F ceramic capacitor located at each Vref pin, and one 0.1- μ F capacitor directly at the source. Traces between the decoupling capacitors and Vref pins should be a minimum of 30 mils (0.762 mm) wide and as short as possible. The Vref pins and interconnection to decoupling capacitors should maintain a minimum of 15 mils (0.381 mm) spacing from all other nets. All Vref nets should be routed on the top layer. Vref pins should be isolated with, or shielded with ground.

When routing the SDRAM VTT power supply, the regulator should be kept close to the VTT pin on the respective SDRAMs. In most cases, a VTT voltage island will be used, and it is recommended that the voltage island be placed on the component side signal layer. There should be a minimum of one 0.1- μ F decoupling capacitor close to each VTT SDRAM pin and a minimum of one 10- μ F to 22- μ F bulk ceramic (low ESR) capacitor on the VTT island. The number of VTT bulk capacitors is based on the size of the island, and topology and loading.

6.3.1.9 Write Leveling Limit Impact on Routing – KeyStone I

The write-leveling process in the DDR3 interface imposes a limit on the maximum and minimum skew between the command delay and the data delay. If these limits are exceeded, the DDR3 interface may fail the write leveling process and data corruption may occur. These limits are sufficiently large that well-controlled topologies will not likely exceed the limits.

The command delay is defined as delay for the clock, command, control, and address group signals from the DSP to a given SDRAM. The data delay is the delay for the data group signals to that same SDRAM. The write-leveling result is effectively the difference, or skew, between these two delays.

The maximum write-leveling skew is the largest difference between the two delays in the topology to a single SDRAM. Likewise, the minimum write-leveling skew is the smallest difference between the two delays in the topology to a single SDRAM.

The write-leveling logic has a theoretical upper limit of 2500 ps. This limit does not scale with SDRAM data rate. The theoretical upper limit equates to two full clock cycles when the clock frequency is 800 MHz for DDR3-1600. It is reduced by half a clock cycle when invert clock out is enabled, as this effectively lengthens the clock by this amount.

The following set of equations provides an approximation of the maximum and minimum write-leveling skew allowed:

- $ddrclkout_{period}$ – period of reference clock the DSP is providing to SDRAM
- $tWLS$ – from JEDEC DDR3 SDRAM specification, write-leveling setup time from rising CK, \overline{CK} crossing to rising DQS, \overline{DQS} crossing
- $tJIT(per, lck)$ – clock period jitter during DLL locking period
- $command_{delay}$ – delay for the clock, command, control and address group signals from the DSP to a given SDRAM
- $data_{delay}$ – delay for the data group signals to that same SDRAM
- $write_leveling_{skew}$ – defined as the value $command_{delay} - data_{delay}$
- $margin$ - additional margin added for preliminary use

Maximum write-leveling skew:

- **Case 1:** Invert clock disabled
 - $write_leveling_{skew} < 2500 \text{ ps} - tWLS - tJIT(per, lck) - margin$
- **Case 2:** Invert clock enabled (adds an additional half-clock period of delay to the command delay term)
 - $command_{delay} + (0.5 * ddrclkout_{period}) - data_{delay} = write_leveling_{skew} < 2500 \text{ ps} - tWLS - tJIT(per, lck) - margin$
 - $write_leveling_{skew} < 2500 \text{ ps} - tWLS - tJIT(per, lck) - (0.5 * ddrclkout_{period}) - margin$

Minimum write-leveling skew:

- **Case 1:** Invert clock disabled
 - $\text{write_leveling}_{\text{skew}} > \text{tWLS} + \text{tJIT}(\text{per}, \text{lck}) + \text{margin}$
- **Case 2:** Invert clock enabled (adds an additional half-clock period of delay to the command delay term)
 - $\text{command}_{\text{delay}} + (0.5 * \text{ddrclkout}_{\text{period}}) - \text{data}_{\text{delay}} = \text{write_leveling}_{\text{skew}} > \text{tWLS} + \text{tJIT}(\text{per}, \text{lck}) + \text{margin}$
 - $\text{write_leveling}_{\text{skew}} > \text{tWLS} + \text{tJIT}(\text{per}, \text{lck}) - (0.5 * \text{ddrclkout}_{\text{period}}) + \text{margin}$

tWLS and tJIT(per, lck) are standard JEDEC DDR3 SDRAM timing parameters that can be obtained from the specific data sheet of the SDRAM chosen.

NOTE: Because this is preliminary guidance, some small margin should be subtracted from these delays to account for additional terms such as multi-rank delay skew. TI currently recommends setting the extra margin term to 100 ps.

Based on the previous equations, the following calculations and summary table shows the write-leveling skew limitations for both invert clock out enabled and disabled given the DDR3-1333 and DDR3-1600 JEDEC SDRAM specification. The first column for each speed-grade category lists the maximum write-leveling skew in picoseconds. The second column for each lists the maximum write-leveling skew in inches, assuming a signal propagation rate of 180 ps/in.

For DDR3-1333:

- $\text{ddrclkout}_{\text{period}} = 1500 \text{ ps}$
- $\text{tWLS} = 195 \text{ ps}$
- $\text{tJIT}(\text{per}, \text{lck}) = \pm 70 \text{ ps}$
- $\text{margin} = 100 \text{ ps}$

Maximum write-leveling skew:

- **Case 1:** Invert clock disabled
 - $\text{write_leveling}_{\text{skew}} < 2500 \text{ ps} - \text{tWLS} - \text{tJIT}(\text{per}, \text{lck}) - \text{margin}$
 - $\text{write_leveling}_{\text{skew}} < 2500 \text{ ps} - 195 \text{ ps} - 70 \text{ ps} - 100 \text{ ps}$
 - $\text{write_leveling}_{\text{skew}} < 2500 \text{ ps} - 195 \text{ ps} - 70 \text{ ps} - 100 \text{ ps}$
 - $\text{write_leveling}_{\text{skew}} < 2135 \text{ ps}$
- **Case 2:** Invert clock enabled (adds an additional half-clock period of delay to the command delay term)
 - $\text{write_leveling}_{\text{skew}} < 2500 \text{ ps} - \text{tWLS} - \text{tJIT}(\text{per}, \text{lck}) - (0.5 * \text{ddrclkout}_{\text{period}}) - \text{margin}$
 - $\text{write_leveling}_{\text{skew}} < 2500 \text{ ps} - 195 \text{ ps} - 70 \text{ ps} - (0.5 * 1500 \text{ ps}) - 100 \text{ ps}$
 - $\text{write_leveling}_{\text{skew}} < 1385 \text{ ps}$

Minimum write-leveling skew:

- **Case 1:** Invert Clock disabled
 - $\text{write_leveling}_{\text{skew}} > \text{tWLS} + \text{tJIT}(\text{per}, \text{lck}) + \text{margin}$
 - $\text{write_leveling}_{\text{skew}} > 195 \text{ ps} + 70 \text{ ps} + 100 \text{ ps}$
 - $\text{write_leveling}_{\text{skew}} > 365 \text{ ps}$
- **Case 2:** Invert clock enabled (adds an additional half-clock period of delay to the command delay term)
 - $\text{write_leveling}_{\text{skew}} > \text{tWLS} + \text{tJIT}(\text{per}, \text{lck}) - (0.5 * \text{ddrclkout}_{\text{period}}) + \text{margin}$
 - $\text{write_leveling}_{\text{skew}} > 195 \text{ ps} + 70 \text{ ps} - (0.5 * 1500 \text{ ps}) + 100 \text{ ps}$
 - $\text{write_leveling}_{\text{skew}} > -385 \text{ ps}$

NOTE: This minimum write-leveling skew calculation with invert clock enabled shows how the invert clock mode can be used to correct a small amount of negative skew between the command and data groups. However, as specified in [Section 6.3.1.7](#), all topologies should be designed for a positive skew between the command delay and data delay to avoid this situation.

For DDR3-1600:

- $ddrclkout_{period} = 1250$ ps
- $tWLS = 165$ ps
- $tJIT(per, lck) = \pm 60$ ps
- margin = 100 ps

Maximum write-leveling skew:

- **Case 1:** Invert clock disabled
 - $write_leveling_{skew} < 2500$ ps - 165 ps - 60 ps - 100 ps
 - $write_leveling_{skew} < 2500$ ps - 165 ps - 60 ps - 100 ps
 - $write_leveling_{skew} < 2175$ ps
- **Case 2:** Invert clock enabled (adds an additional half-clock period of delay to the command delay term)
 - $write_leveling_{skew} < 2500$ ps - 165 ps - 60 ps - (0.5 * 1250 ps) - 100 ps
 - $write_leveling_{skew} < 1610$ ps

Minimum write-leveling skew:

- **Case 1:** Invert clock disabled
 - $write_leveling_{skew} > 165$ ps + 60 ps + 100 ps
 - $write_leveling_{skew} > 325$ ps
- **Case 2:** Invert clock enabled (adds an additional half-clock period of delay to the command delay term)
 - $write_leveling_{skew} > 165$ ps + 60 ps - (0.5 * 1250 ps) + 100 ps
 - $write_leveling_{skew} > -300$ ps

NOTE: This minimum write-leveling skew calculation with invert clock enabled shows how the invert clock mode can be used to correct a small amount of negative skew between the command and data groups. However, as specified in [Section 6.3.1.7](#), all topologies should be designed for a positive skew between the command delay and data delay to avoid this situation.

Table 16. Maximum Write Leveling Skew Example

Speed Grade	Invert Clock Out State			
	Disabled		Enabled	
	Skew in ps	Skew in Inches	Skew in ps	Skew in Inches
DDR3-1333	2135	11.861	1385	7.694
DDR3-1600	2175	12.083	1610	8.944

Table 17. Minimum Write Leveling Skew Example

Speed Grade	Invert Clock Out State			
	Disabled		Enabled	
	Skew in ps	Skew in Inches	Skew in ps	Skew in Inches
DDR3-1333	365	2.027	-385	2.138
DDR3-1600	325	1.805	-300	1.666

NOTE: Because this is preliminary guidance and some small margin should be subtracted or added from these delays to account for additional terms such as multi-rank delay skew, TI recommends that the maximum routing lengths be reduced by 10% and the minimum routing lengths be increased by 10%.

6.3.1.10 Round-Trip Delay Impact on Routing – KeyStone I

The leveling processes in the DDR3 interface impose an upper limit on the maximum round-trip delay. If this limit is exceeded, the DDR3 interface may fail the leveling process and data corruption may occur. This limit is sufficiently large that well-controlled topologies will not likely exceed the limit.

The round-trip delay for a given SDRAM is defined as the sum of two delays. The first is the longest delay for the clock, command, control, and address groups to that SDRAM. The second is the delay for the data group to that same SDRAM. This round-trip delay must be calculated for each byte-lane to each SDRAM device implemented in the DDR3 memory topology, including SDRAM devices on DIMM. All of these individual sums must be below the limit to help ensure robust operation.

Internally, the DDR3 controller logic has a theoretical upper limit of four clock cycles. There are multiple processes that have variation terms that reduce this time window, as listed below:

- **ddrclkout_{period}** – period of reference clock the DSP is providing to SDRAM
- **tDQSCK - DQS to CK skew limit from SDRAM data sheet** – this is stated for each standard speed grade in the JEDEC DDR3 SDRAM standard
- **Invert clock out** – delay of half a clock cycle if enabled – option normally used only in very small memory topologies
- **adjustment_sum** – sum that accounts for all of the leveling errors, buffer delays, and jitter terms associated with the circuitry, and the leveling adjustment – upper limit is defined to be half a clock period plus 225 ps.

The following equation provides an approximation of the maximum round trip delay:

- **Case 1: Invert clock out disabled**
 - $\text{round_trip}_{\text{delay}} < (4 * \text{ddrclkout}_{\text{period}}) - \text{tDQSCK} - \text{adjustment_sum}$
 - $\text{round_trip}_{\text{delay}} < (4 * \text{ddrclkout}_{\text{period}}) - \text{tDQSCK} - 0.5 * \text{ddrclkout}_{\text{period}} - 225 \text{ ps}$
 - $\text{round_trip}_{\text{delay}} < (3.5 * \text{ddrclkout}_{\text{period}}) - \text{tDQSCK} - 225 \text{ ps}$
- **Case 2: Invert clock out enabled (adds an additional half-clock period of delay to the command delay term)**
 - $\text{round_trip}_{\text{delay}} < (4 * \text{ddrclkout}_{\text{period}}) - \text{tDQSCK} - 0.5 * \text{ddrclkout}_{\text{period}} - \text{adjustment_sum}$
 - $\text{round_trip}_{\text{delay}} < (4 * \text{ddrclkout}_{\text{period}}) - \text{tDQSCK} - 0.5 * \text{ddrclkout}_{\text{period}} - 0.5 * \text{drclkoutperiod} - 225 \text{ ps}$
 - $\text{round_trip}_{\text{delay}} < (3 * \text{ddrclkout}_{\text{period}}) - \text{tDQSCK} - 225 \text{ ps}$

Based on the previous equations, the following calculations and summary table show the write leveling skew limitations for both invert clock out enabled and disabled, given the DDR3-1333 and DDR3-1600 JEDEC SDRAM specification. The first column for each speed-grade category lists the maximum write leveling skew in picoseconds. The second column for each lists the maximum write leveling skew in inches assuming a signal propagation rate of 180 ps/in.

For DDR3-1333:

- $ddrclkout_{period} = 1500$ ps
- $tDQSCK = \pm 255$ ps
- margin = 100 ps
- **Case 1:** Invert clock out disabled:
 - $round_trip_{delay} < (3.5 * ddrclkout_{period}) - tDQSCK - 225$ ps
 - $round_trip_{delay} < (3.5 * 1500$ ps) - 255 ps - 225 ps
 - $round_trip_{delay} < 4770$ ps
- **Case 2:** Invert clock out enabled (adds an additional half-clock period of delay to the command delay term):
 - $round_trip_{delay} < (3 * ddrclkout_{period}) - tDQSCK - 225$ ps
 - $round_trip_{delay} < (3 * 1500$ ps) - 255 ps - 225 ps
 - $round_trip_{delay} < 4020$ ps

For DDR3-1600:

- $ddrclkout_{period} = 1250$ ps
- $tDQSCK = \pm 225$ ps
- margin = 100 ps
- **Case 1:** Invert clock out disabled
 - $round_trip_{delay} < (3.5 * ddrclkout_{period}) - tDQSCK - 225$ ps
 - $round_trip_{delay} < (3.5 * 1250$ ps) - 225 ps - 225 ps
 - $round_trip_{delay} < 3925$ ps
- **Case 2:** Invert clock out enabled (adds an additional half-clock period of delay to the command delay term)
 - $round_trip_{delay} < (3 * ddrclkout_{period}) - tDQSCK - 225$ ps
 - $round_trip_{delay} < (3 * 1250$ ps) - 225 ps - 225 ps
 - $round_trip_{delay} < 3300$ ps

Table 18 shows the round trip delay limitations for both invert clock out enabled and disabled. The first column for each lists the maximum round-trip delay in picoseconds. The second column for each lists the maximum routing length in inches assuming a signal propagation rate of 180 ps/in.

Table 18. Maximum Round Trip Delay Example - Invert Clock Out Enabled and Disabled

Speed Grade	Invert Clock Out Disabled		Invert Clock Out Enabled	
DDR3-1333	4770 ps	26.50 in	4020 ps	22.33 in
DDR3-1600	3925 ps	21.81 in	3300 ps	18.33 in

Because this is preliminary guidance and some small margin should be subtracted from these delays to account for additional terms such as multi-rank delay skew, TI recommends that the maximum routing lengths be reduced by 10%.

6.3.1.11 Write Leveling Limit Impact on Routing – KeyStone II

The write-leveling process in the DDR3 interface imposes a limit on the maximum and minimum skew between the command delay and the data delay. If these limits are exceeded, the DDR3 interface may fail the write leveling process and data corruption may occur. These limits are sufficiently large that well-controlled topologies will not likely exceed the limits.

The command delay is defined as delay for the clock, command, control, and address group signals from the DSP to a given SDRAM. The data delay is the delay for the data group signals to that same SDRAM. The write-leveling result is effectively the difference, or skew, between these two delays.

The maximum write-leveling skew is the largest difference between the two delays in the topology to a single SDRAM. Likewise, the minimum write-leveling skew is the smallest difference between the two delays in the topology to a single SDRAM. The write-leveling logic has an upper limit which is the lower of 1.75 clock periods or 2500 ps. For instance, when operating at 1600MT/s, the clock is 800MHz which has a period of 1250ps. The maximum skew is 1.75 clock periods or 2187ps. At a propagation delay of 180ps per inch, this limits the write-leveling routing skew length to just over 12.1 inches.

The minimum write leveling skew occurs when the command delay is less than the data delay. The minimum write-leveling skew is -0.75 times the clock period. At 1600MT/s, this minimum skew is 937ps or about 5.2 inches. The DDR3 PHY implemented on KeyStone II can operate within these limits without the need to invert the clock to the DDR3 SDRAMs.

6.3.1.12 Round-Trip Delay Impact on Routing – KeyStone II

The leveling processes in the DDR3 interface impose an upper limit on the maximum round-trip delay. If this limit is exceeded, the DDR3 interface may fail the leveling process and data corruption may occur. This limit is sufficiently large that well-controlled topologies will not likely exceed the limit.

The round-trip delay for a given SDRAM is defined as the sum of two delays. The first is the delay for the clock, command, control, and address groups to that SDRAM. The second is the delay for the data group to that same SDRAM. This round-trip delay must be calculated for each byte-lane to each SDRAM device implemented in the DDR3 memory topology. All of these individual sums must be below the limit to ensure robust operation.

The DDR3 PHY logic has an upper limit of 6 clock cycles. For instance, when operating at 1600MT/s, the clock is 800MHz which has a period of 1250ps. The maximum round trip delay is 6 clock periods or 7500ps. At a propagation delay of 180ps per inch, this limits the maximum round trip routed length to over 41.6 inches.

6.3.2 Routing Considerations – UDIMMs

6.3.2.1 Mechanical Layout – UDIMMs

When designing the application hardware to use an UDIMM (unbuffered DIMM), there are several issues to take into account that may differ when using individual SDRAMs, including power and route lengths. Key issues including stack-up and power/ground-plane referencing do not differ between the two topologies (SDRAM versus UDIMM).

6.3.2.2 Stack Up – UDIMMs

The board stack-up should not change regardless of whether you are using an UDIMM or individual SDRAMs. General guidelines for stack-up under the SDRAMs section apply. It is worth mentioning that a good-quality connector and proper impedance is critical to minimize signal reflections.

6.3.2.3 Routing Rules – General Overview (UDIMMs)

General routing rules between the KeyStone DSP DDR3 interface and UDIMM are the same as those identified for discrete components (SDRAMs). Note that a single DDR3 interface on any KeyStone DSP cannot currently support multiple DIMMs in a design. Return ground paths and power-plane decoupling also are critical and should be evaluated properly.

All aspects of this application report and especially those pertaining to the use and implementation of UDIMMs assumes that the user has an above average level of understanding regarding mechanical layout and design – including the impact of trace width, spacing, via size, and bulk and decoupling capacitance selection and placement.

6.3.2.4 Routing Rules – Address and Command Lines (UDIMMs)

All address and command nets between the DSP and UDIMM must be referenced to a solid ground or solid power plane — the best option would be a solid ground plane. The entire address bus and/or command nets can be referenced to the DVDD15 power plane as long as sufficient distributed decoupling is located near the endpoints and vias of all routes to provide a low-impedance current return path. All address and command nets routing from the DSP to the UDIMM socket must be adjacent to the referenced plane.

The address bus and command nets between the DSP and UDIMM should be routed away from the data bus and respective nets.

6.3.2.5 Routing Rules – Control Lines (UDIMMs)

All control nets between the DSP and UDIMM must be referenced to a solid ground or solid power plane — the best option would be a solid ground plane. The control nets can be referenced to the DVDD15 power plane as long as sufficient distributed decoupling is located near the endpoints and vias of all control routes to provide a low impedance current return path. All control nets routing from the DSP to the UDIMM socket must be adjacent to the referenced plane.

The control nets between the DSP and UDIMM should be routed away from the data bus and respective nets.

6.3.2.6 Routing Rules – Data Lines (UDIMMs)

All data-group signals must be routed adjacent to a solid ground plane.

If possible, all data-group nets should be routed internally and close to the bottom of the board. All nets in a single byte-group should be routed on the same layer to eliminate addition length skew from the via barrels. All data nets should be skew-matched between data/strobe within the respective byte lane. The byte lanes are identified in the SDRAM routing rules in [Table 14](#).

6.3.2.7 Routing Rules – Clock Lines (UDIMMs)

All DDR3 clocks between the DSP and respective UDIMM must be routed as differential pairs. Each differential clock pair must be length-matched to the address, command, and control signals.

6.3.2.8 Routing Rules – Power (UDIMMs)

For the UDIMM, there exist three separate power supplies, all derived from a common rail. The first is the 1.5-V supply that provides power to all the DDR3 UDIMM I/Os. The second supply is the V_{REF} supply, which must track the VDD15 supply and establishes a reference voltage for the UDIMM. The last supply is the bus termination supply (VTT).

Each of the Vref supplies to the UDIMM (VrefCA and VrefDQ) can originate from a common rail but must be individually decoupled at the UDIMM. V_{REF} must be 50% of the V_{DD}/V_{DDQ} level and meet the tolerances identified in the respective UDIMM data manual. The typical method for establishing each of these reference voltages is through a 1% (or better) resistor divider network. It is important that the V_{REF} (VrefCA and VrefDQ) voltages track the V_{DD}/V_{DDQ} level across all corners (process, temperature, and noise). See the applicable data manual for transient (AC & DC) requirements.

When routing the UDIMM V_{REF} voltages, properly decouple them as close to the socket as possible. The use of 0.01- μ F and 0.1- μ F ceramic capacitors (0402 or smaller recommended) should be distributed across the V_{REF} power rail with one 0.01- μ F and one 0.1- μ F ceramic capacitor located at each V_{REF} pin and one 0.1- μ F capacitor directly at the source. Traces between the decoupling capacitors and V_{REF} pins should be a minimum of 0.030 inch (0.762 mm) wide and as short as possible. The V_{REF} pins and interconnection to decoupling capacitors should maintain a minimum of 0.015 inch (0.381 mm) spacing from all other nets. All V_{REF} nets should be routed on the top layer. V_{REF} pins should be isolated with, or shielded with ground.

The UDIMM termination voltage (VTT) must be at a constant level of 0.750 V and must be capable of sinking a reasonable amount of current while maintaining voltage regulation. VTT must remain stable at all times for the UDIMM to function properly. Issues including noise and crosstalk must be eliminated or reduced to a negligible amount. VTT like V_{REF} must track all variations with respect to VDD/VDDq.

When routing the UDIMM VTT power supply the regulator should be kept as close to the VTT pin on the respective SDRAMs. In most cases a VTT voltage island will be used and it is recommended that the voltage island be placed on the component-side signal layer. There should be a minimum of one 0.1- μ F decoupling capacitor close to each VTT SDRAM pin and a minimum of one 10- μ F to 22- μ F bulk capacitor on the VTT island. The number of VTT bulk capacitors is based on the size of island and topology.

6.3.2.9 Write-Leveling Limit Impact on Routing

The write-leveling process in the DDR3 interface imposes an upper and lower limit on the maximum and minimum skew between the command delay and the data delay. If this limit is exceeded, the DDR3 interface may fail the write leveling process and data corruption may occur. This limit is sufficiently large so that well-controlled topologies will not likely exceed the limit.

The same write-leveling limits in KeyStone I designs outlined for SDRAMs in [Section 6.3.1.9](#) also apply to UDIMMs. Any calculations performed must take both board and UDIMM routing into consideration.

6.4 Timing Considerations

DDR3 requires strict timing relationships between CK (and \overline{CK}) and the address/control lines, and between data and the DQS (and \overline{DQS}) lines. The TI DSP DDR3 interface is designed to comply with the DDR3 JEDEC standard with regards to timing constraints. See the applicable standard when evaluating timing. Additional timing considerations or constraints may be included in the respective DSP or SDRAM data sheets.

6.5 Impedance Considerations

The guidelines below review the impedance considerations for routing and device configuration.

6.5.1 Routing Impedances – KeyStone I Devices

The DDR3 interface on KeyStone I devices has, in general, been verified to operate up to 1333 MT/s with up to nine loads on the fly-by nets (address, control, command and clock). Validation has focused on single-rank memory topologies with only a single load on the data group signals. Topologies with fewer than nine loads should all operate at a 1333 MT/s level of performance or below.

6.5.1.1 Data Group Signals

All data-group signals are point-to-point in the validated topologies. The data-group signals are driven by the KeyStone device on writes and driven by the SDRAM memories during reads. No external resistors are needed on these routes. The receivers in both cases (SDRAMs on writes and KeyStone device on reads) will assert on-die terminations (ODT) at the appropriate times. The following diagrams show the impedances seen on these nets during write and read cycles.

Figure 11 shows the impedances seen on the nets during a write cycle. During writes, the output impedance of the KeyStone I device is approximately $45\ \Omega$. It is recommended that the SDRAM be implemented with a $240\text{-}\Omega$ RZQ resistor and be configured to present an ODT of $RZQ/4$ for an effective termination of $40\ \Omega$.

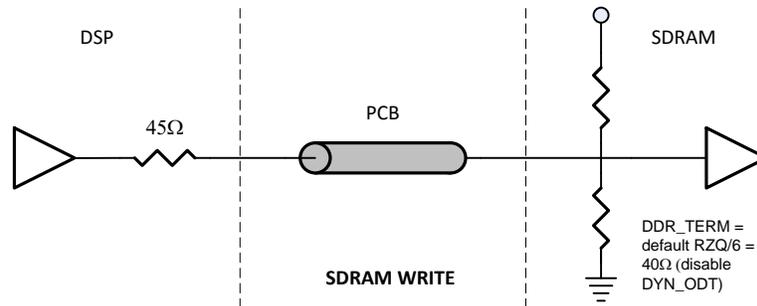


Figure 11. Data Group Impedances During Write Cycles on KeyStone I

Figure 12 shows the impedances seen on the nets during a read cycle. During reads, it is recommended that the SDRAM be configured for an effective drive impedance of $RZQ/7$ or $34\ \Omega$ (assuming RZQ resistor is $240\ \Omega$). The on-die termination (ODT) within the KeyStone I device will have an effective Thevenin impedance of $45\ \Omega$.

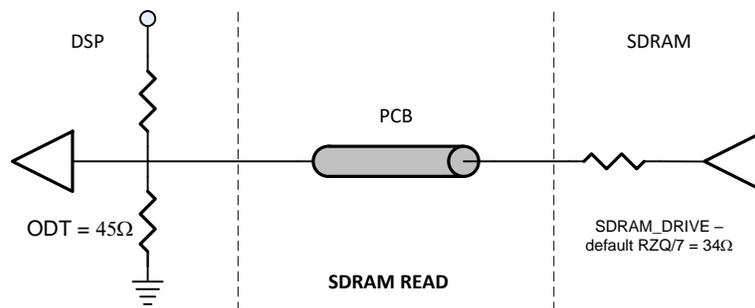


Figure 12. Data Group Impedances During Read Cycles on KeyStone I

6.5.1.2 Fly-By Signals

The fly-by signals include the address, control, command, and clock routing groups. The fly-by signals consist of the fly-by routing from the DSP, stubs at each SDRAM, and terminations after the last SDRAM. The address, control, and command groups will be terminated through a $39.2\text{-}\Omega$ resistor to VTT. The clock pairs will be terminated through $39.2\text{-}\Omega$ resistors to a common node connected to a capacitor that is then connected to VDDQ. The KeyStone I device will present a $45\text{-}\Omega$ output impedance when driving these signals. These relationships are shown in Figure 13.

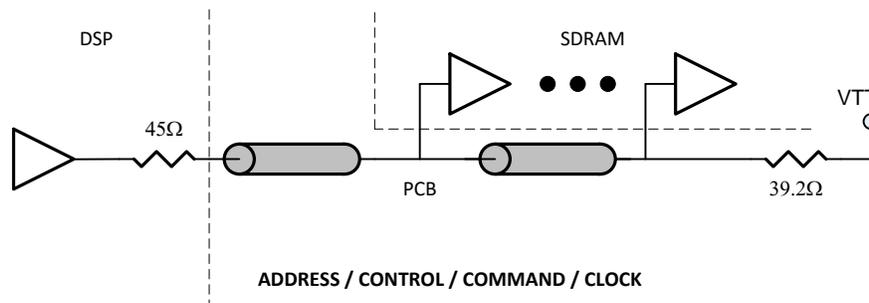


Figure 13. Fly-By Impedances on KeyStone I

6.5.2 Routing Impedances – KeyStone II Devices

The DDR3 interface on KeyStone II devices has, in general, been verified to operate up to 1600 MT/s with up to nine loads on the fly-by nets (address, control, command, and clock). Validation has focused on single-rank memory topologies with only a single load on the data group signals. Topologies with fewer than nine loads should all operate at a 1600 MT/s level of performance or below.

6.5.2.1 Data Group Signals

All data-group signals are point-to-point in the validated topologies. The data-group signals are driven by the KeyStone device on writes and driven by the SDRAM memories during reads. No external resistors are needed on these routes. The receivers in both cases (SDRAMs on writes and KeyStone device on reads) will assert on-die terminations (ODT) at the appropriate times. The following diagrams show the impedances seen on these nets during write and read cycles.

Figure 14 shows the impedances seen on the nets during a write cycle. During writes, the output impedance of the KeyStone II device is approximately $40\ \Omega$. It is recommended that the SDRAM be implemented with a $240\ \Omega$ RZQ resistor and be configured to present an ODT of RZQ/4 for an effective termination of $60\ \Omega$.

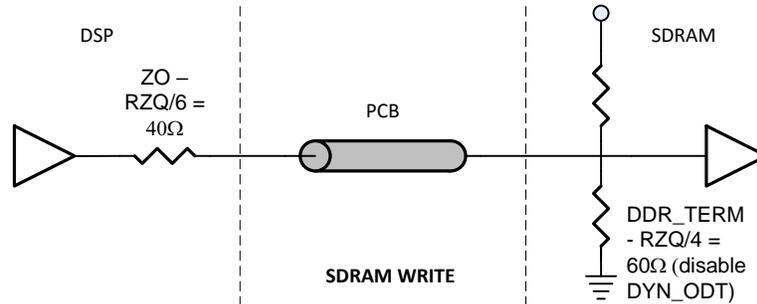


Figure 14. Data Group Impedances During Write Cycles on KeyStone II

Figure 15 shows the impedances seen on the nets during a read cycle. During reads, it is recommended that the SDRAM be configured for an effective drive impedance of RZQ/7 or $34\ \Omega$ (assuming RZQ resistor is $240\ \Omega$). The ODT within the KeyStone II device will have an effective Thevenin impedance of $60\ \Omega$.

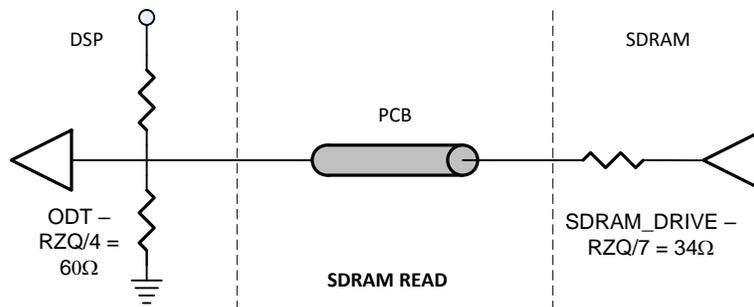


Figure 15. Data Group Impedances During Read Cycles on KeyStone II

6.5.2.2 Fly-By Signals

The fly-by signals include the address, control, command, and clock routing groups. The fly-by signals are composed of the fly-by routing from the DSP, stubs at each SDRAM, and terminations after the last SDRAM. The address, control, and command groups will be terminated through a $39.2\text{-}\Omega$ resistor to VTT. The clock pairs will be terminated through $39.2\text{-}\Omega$ resistors to a common node connected to a capacitor that is then connected to VDDQ. The KeyStone II device can be configured to present either a $40\text{-}\Omega$ output impedance when configured for RZQ/6 (recommended for discrete components) or a $34\text{-}\Omega$ output impedance when configured for RZQ/7 (recommended for UDIMMs). These relationships are shown in Figure 16.

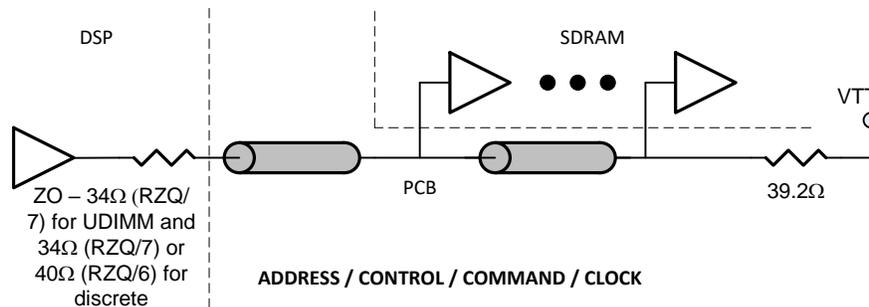


Figure 16. Fly-By Impedances on KeyStone II

6.5.3 Comparison to JEDEC UDIMM Impedance Recommendations

The JEDEC UDIMM specification provides guidance for UDIMM manufacturers including recommendations for routing impedances. These recommendations are similar to the routing impedances discussed above. The JEDEC UDIMM specification recommends $40\text{-}\Omega$ routes on the fly-by nets between the controller and the first SDRAM and then $60\text{-}\Omega$ routes between the SDRAMs and to the terminations. However, footnotes under Table 39 in the JEDEC UDIMM specification state that lightly loaded UDIMMs (like the topologies we expect customers to implement with KeyStone devices) should target $49\text{-}\Omega$ rather than $60\text{-}\Omega$. Differential routes can be constructed from these same single-ended track structures resulting in differential impedances under $100\text{-}\Omega$.

6.6 Switching and Output Considerations

Signal switching and output swings can have a significant if not catastrophic impact on signal integrity and EMI compliance. Routing, SDRAM placement, stack up, and use cases can also attribute to added dynamic switching noise. The concept of dynamic switching noise and induced ground bounce is usually attributed to improper layout and stack up, insufficient decoupling, and improper routing. To minimize the output and switching problems, system modeling is recommended and usually employed.

7 Simulation and Modeling

This section provides a quick overview regarding the importance of simulating and modeling the DSP and DDR3 interfaces.

7.1 Simulation and Modeling

All high-performance interfaces, especially those operating above 300 MHz , should always be modeled. Proper simulation and modeling (which must include a complete application board (PCB) stack up, DSP DDR3 interface, and SDRAMs) is important to verify and confirm component placement, selection, and signal integrity. In a high-performance interface, signal stubs, perturbations, non-monotonic waveforms, inflections, and reflections become significant. Time spent properly modeling the DDR3 interface regardless of the topology or condition selected will pay off in the long run.

Figure 17 shows the impact of proper placement and routing of high speed signals.

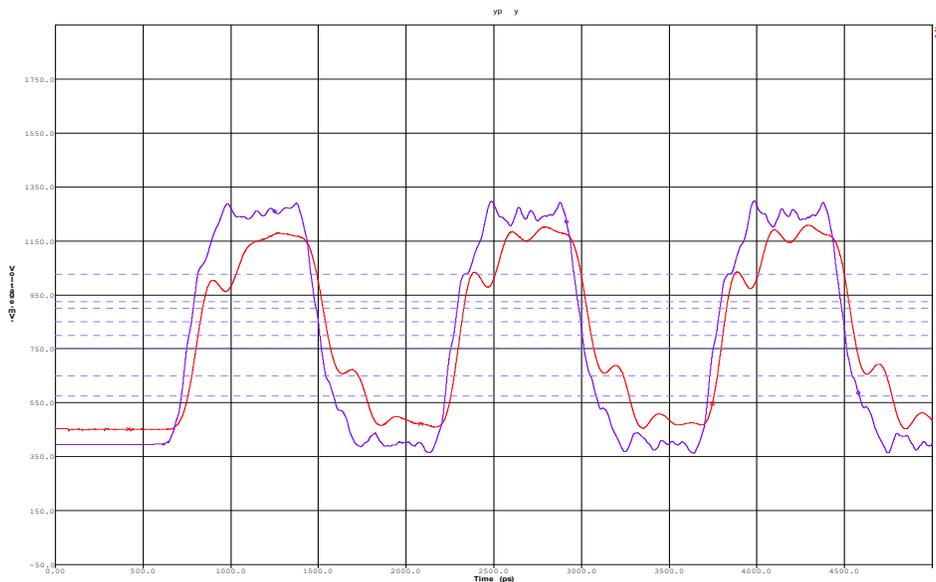


Figure 17. DDR3 Simulations

7.2 Tools

A variety of software tools exists to give customers the ability to model high-performance interfaces. Up until recently, the most common simulation tool was Spice or HSpice. These tools are extremely costly, difficult to use, and vary between types. Other tools including the IBIS standard and protocol were traditionally avoided because of their limitation at high frequencies.

7.3 Models

Different simulation models can be found for the DDR3 SDRAMs and will be available from the SDRAM vendor you have selected. The available models are typically one of the following three formats: Spice, HSpice, or IBIS. Texas Instruments intends to provide IBIS models to our customers.

7.4 TI Commitment

Texas Instruments has taken the initiative to provide IBIS models compliant with the latest 5.0 IBIS standard for the DDR3 interface. The Texas Instruments DSP IBIS model will be correlated to internal Matlab or Spice models, and functionally against timing and performance parameters prior to release (TMS).

8 Power

This section briefly describes the DDR3 relative power, DSP DDR3 interface relative power, power assessment, and power sequencing.

8.1 DDR3 SDRAM Power Requirements

There exists three different power supplies for the SDRAM (1.5-V primary IO supply, Vref, and VTT) and three different power supplies for the DSP (1.1 V, 1.5 V, and Vref).

It is recommended that the KeyStone DSP DDR3 interface and SDRAM share a common 1.5-V supply rail. A common 1.5-V supply rail ($\pm 5\%$ maximum AC/DC tolerance) simplifies the overall design, and reduces the differential between the two devices and minimizes the need for additional layers (power) in the end-use application.

8.1.1 Vref Voltage Requirements

There exist two DDR3 SDRAM Vreference pins: VrefCA and VrefDQ. VrefCA is the reference voltage for all command, address, and control pins. VrefDQ is the reference voltage for the data lines. It is not necessary, but typically recommended, that both Vreference voltages originate from the same supply source. Both Vreference pins must be derived from $VDD/2$ ($VDDq/2$). The recommended Vreference (Vref) implementation is by using a simple resistor divider with 1% or better accuracy. The distance between the source voltage through the divider network and to the decoupled Vreference pins must be short. Each Vreference pin must properly track the $VDD/2$ ($VDDq/2$) variations over voltage, noise, and temperature differences. The pk-to-pk AC and DC noise on the Vreference pins cannot exceed $\pm 2\%$ or 1.5 mV.

8.1.2 VTT Voltage Requirements

The DDR3 SDRAM termination voltage is referred to as VTT and requires a 750-mV supply. The recommended VTT source is a regulator that is capable of sinking a sufficient amount of current while at the same time maintaining a tight voltage tolerance. Like the Vref pins, the distance between the VTT source voltage and SDRAM pin must be short and decoupled properly. The VTT pin must be kept stable and properly track the $VDD/VDDq$ variations over voltage, noise, and temperature differences. The pk-to-pk AC and DC noise on the VTT pin cannot exceed $\pm 2\%$ or 1.5 mV.

8.2 DSP DDR3 Power Requirements

On the DSP, the AVDDA2 pins are designed to supply power to the internal DDR3 clock PLL. These pins must be connected to a clean 1.8-V supply rail. There exist other DSP 1.8-V supplies that can be used provided the voltage tolerance is maintained and a separate filter is used. All DVDD15 pins are designed to supply power to the DSP DDR3 IO buffers. As with the 1.5-V power pins, these must also be connected to a clean 1.5-V supply rail. The DDR3xVREFSSTL pins provides the reference voltage to the DSP DDR3 interface. This supply pin is derived from $VDD/2$ using 1% or better resistors.

For power supply requirements, see the device-specific DSP data sheet and application reports.

8.3 DDR3 Power Estimation

Actual power for each SDRAM is dependent upon many factors. Most DDR3 SDRAM manufacturers have application reports to aid in estimating the DDR3 power. Power can be divided into two categories: active and leakage. All SDRAMs contain some element of each. When designing the power supplies, it is strongly recommended that you take into account both active and leakage power as well as peak in-rush currents. A suitable margin should always be added to the final calculation.

As a general rule of thumb, a single DDR3 SDRAM in a $\times 16$ (256M Byte) configuration running at 1.333 MT/s and assuming 50% reads and 50% writes will consume a maximum total power of 864 mW or 576 mA.

8.4 DSP DDR3 Interface Power Estimation

Excluding the DDR3 SDRAM power discussed above, power estimates for the DSP DDR3 interface can be obtained using the device-specific KeyStone DSP power application report and spreadsheet.

8.5 Sequencing – DDR3 and DSP

The DSP device requires specific power sequencing as well as the DDR3 DRAMs. In all cases, see the device-specific data sheet for verification of power up and power down sequencing requirements.

In all cases, the DRAM cannot be allowed to drive into the DSP until the DSP has been fully powered and sequenced properly.

9 Disclaimers

Comments and proposals identified within this document are presented as guidelines. Texas Instruments can not fully anticipate every possible design permutation, topology, and application variation and therefore requires each end user to perform proper due diligence (using good engineering practices and including proper component selection) when designing for a DDR3 interface. This includes proper simulation and modeling to verify the recommendations provided function in an acceptable manner (in the end-use application).

This application guide is not intended to be the sole source or reference design guide. Many factors not encompassed within this document can force a change in component values, selection, placement, or termination type.

The possible permutations available due to component, layout, and assembly can also induce variations not encompassed in the recommendations provided. Voltage rails, manufacturer parts selection, switching levels should all be verified before committing the design to production.

It is the end user's responsibility to always verify the design, and connectivity between any active or passive component and the targeted DSP. This includes verifying against individual data sheets and application reports.

Simulation and modeling play an important role in system verification and validation. This effort and the need to conduct this level of analysis should not be overlooked in any design – regardless of simplicity or complexity.

This is not a primer for high-speed design. Use of the information provided assumes a strong understanding of electrical and mechanical design requirements in a high-performance DSP application environment.

From time to time there exist DSP or application design changes. It is the end user's responsibility to verify that the appropriate data sheets and application reports are followed.

Specifications continuously change and are updated. Verify the specifications have not changed and that the documents used are correct for the version of silicon you are designing with.

Prior to final design release it is recommended that a full engineering assessment be performed by the end user in order to assure functionality.

10 References

- [KeyStone Architecture DDR3 Memory Controller User's Guide \(SPRUGV8\)](#)
- [Hardware Design Guide for KeyStone I Devices](#)
- [TN-41-04: DDR3 Dynamic On-Die Termination Technical Note \(Micron\)](#)
- [TN-41-06: DDR3 Termination Data Strobe Design Note \(TDQS\) \(Micron\)](#)
- [MO-269D: JEDEC Document: MO \(Module Outline\) {for DDR3}](#)
- [SO-007B: JEDEC Document: SO \(Socket Outline\)](#)
- [TN-42-02: DDR3 ZQ Calibration Design Note \(Micron\)](#)
- [TN-04-54: High-Speed DRAM Controller Design Note \(Micron\)](#)
- [TN-41-01: Calculating Memory System Power for DDR3 Design Note \(Micron\)](#)
- [TN-41-07: DDR3 Power-Up, Initialization, and Reset Design Note \(Micron\)](#)
- [TN-41-08: Design Guide for Two DDR3-1066 UDIMM Systems \(Micron\)](#)
- [JEDEC 21-C: Unbuffered DIMM Design Specification](#)
- [Pub 95 PS-001A: Connector Performance Standards for Outlines of Solid State Related Products – 240 pin DDR3 UDIMM](#)
- [TMS320C6678 Multicore Fixed and Floating-Point Digital Signal Processor Data Sheet](#)

Revision History

NOTE: Page numbers for previous revisions may differ from page numbers in the current version.

Changes from B Revision (January 2015) to C Revision	Page
• Updates were made in Section 4.6.2	7
• Updates were made in Table 1 and Table 3	7
• Added new Section 6.3.1.11	27
• Added new Section 6.3.1.12	28
• Update was made in Section 6.5.1.1	30
• Update was made in Section 6.5.2.1	32

IMPORTANT NOTICE FOR TI DESIGN INFORMATION AND RESOURCES

Texas Instruments Incorporated ("TI") technical, application or other design advice, services or information, including, but not limited to, reference designs and materials relating to evaluation modules, (collectively, "TI Resources") are intended to assist designers who are developing applications that incorporate TI products; by downloading, accessing or using any particular TI Resource in any way, you (individually or, if you are acting on behalf of a company, your company) agree to use it solely for this purpose and subject to the terms of this Notice.

TI's provision of TI Resources does not expand or otherwise alter TI's applicable published warranties or warranty disclaimers for TI products, and no additional obligations or liabilities arise from TI providing such TI Resources. TI reserves the right to make corrections, enhancements, improvements and other changes to its TI Resources.

You understand and agree that you remain responsible for using your independent analysis, evaluation and judgment in designing your applications and that you have full and exclusive responsibility to assure the safety of your applications and compliance of your applications (and of all TI products used in or for your applications) with all applicable regulations, laws and other applicable requirements. You represent that, with respect to your applications, you have all the necessary expertise to create and implement safeguards that (1) anticipate dangerous consequences of failures, (2) monitor failures and their consequences, and (3) lessen the likelihood of failures that might cause harm and take appropriate actions. You agree that prior to using or distributing any applications that include TI products, you will thoroughly test such applications and the functionality of such TI products as used in such applications. TI has not conducted any testing other than that specifically described in the published documentation for a particular TI Resource.

You are authorized to use, copy and modify any individual TI Resource only in connection with the development of applications that include the TI product(s) identified in such TI Resource. NO OTHER LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE TO ANY OTHER TI INTELLECTUAL PROPERTY RIGHT, AND NO LICENSE TO ANY TECHNOLOGY OR INTELLECTUAL PROPERTY RIGHT OF TI OR ANY THIRD PARTY IS GRANTED HEREIN, including but not limited to any patent right, copyright, mask work right, or other intellectual property right relating to any combination, machine, or process in which TI products or services are used. Information regarding or referencing third-party products or services does not constitute a license to use such products or services, or a warranty or endorsement thereof. Use of TI Resources 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.

TI RESOURCES ARE PROVIDED "AS IS" AND WITH ALL FAULTS. TI DISCLAIMS ALL OTHER WARRANTIES OR REPRESENTATIONS, EXPRESS OR IMPLIED, REGARDING TI RESOURCES OR USE THEREOF, INCLUDING BUT NOT LIMITED TO ACCURACY OR COMPLETENESS, TITLE, ANY EPIDEMIC FAILURE WARRANTY AND ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT OF ANY THIRD PARTY INTELLECTUAL PROPERTY RIGHTS.

TI SHALL NOT BE LIABLE FOR AND SHALL NOT DEFEND OR INDEMNIFY YOU AGAINST ANY CLAIM, INCLUDING BUT NOT LIMITED TO ANY INFRINGEMENT CLAIM THAT RELATES TO OR IS BASED ON ANY COMBINATION OF PRODUCTS EVEN IF DESCRIBED IN TI RESOURCES OR OTHERWISE. IN NO EVENT SHALL TI BE LIABLE FOR ANY ACTUAL, DIRECT, SPECIAL, COLLATERAL, INDIRECT, PUNITIVE, INCIDENTAL, CONSEQUENTIAL OR EXEMPLARY DAMAGES IN CONNECTION WITH OR ARISING OUT OF TI RESOURCES OR USE THEREOF, AND REGARDLESS OF WHETHER TI HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

You agree to fully indemnify TI and its representatives against any damages, costs, losses, and/or liabilities arising out of your non-compliance with the terms and provisions of this Notice.

This Notice applies to TI Resources. Additional terms apply to the use and purchase of certain types of materials, TI products and services. These include; without limitation, TI's standard terms for semiconductor products (<http://www.ti.com/sc/docs/stdterms.htm>), [evaluation modules](#), and [samples](http://www.ti.com/sc/docs/sampterm.htm) (<http://www.ti.com/sc/docs/sampterm.htm>).

Mailing Address: Texas Instruments, Post Office Box 655303, Dallas, Texas 75265
Copyright © 2018, Texas Instruments Incorporated