



ABSTRACT

The goal of this document is to describe how to make the AM64x\AM243x DDR system implementation straightforward for all designers. The requirements have been distilled down to a set of layout and routing rules that allow designers to successfully implement a robust design for the topologies TI supports.

Table of Contents

1 Overview	2
1.1 Board Designs Supported	2
1.2 General Board Layout Guidelines	2
1.3 PCB Stack-Up	3
1.4 Bypass Capacitors	4
1.5 Velocity Compensation	5
2 DDR4 Board Design and Layout Guidance	6
2.1 DDR4 Introduction	6
2.2 DDR4 Device Implementations Supported	6
2.3 DDR4 Interface Schematics	7
2.4 Compatible JEDEC DDR4 Devices	10
2.5 Placement	10
2.6 DDR4 Keepout Region	11
2.7 VPP	11
2.8 Net Classes	11
2.9 DDR4 Signal Termination	12
2.10 VREF Routing	12
2.11 VTT	12
2.12 POD Interconnect	12
2.13 CK and ADDR_CTRL Topologies and Routing Guidance	13
2.14 Data Group Topologies and Routing Guidance	16
2.15 CK and ADDR_CTRL Routing Specification	17
2.16 Data Group Routing Specification	19
2.17 Bit Swapping	20
3 LPDDR4 Board Design and Layout Guidance	21
3.1 LPDDR4 Introduction	21
3.2 LPDDR4 Device Implementations Supported	21
3.3 LPDDR4 Interface Schematics	22
3.4 Compatible JEDEC LPDDR4 Devices	23
3.5 Placement	23
3.6 LPDDR4 Keepout Region	24
3.7 Net Classes	24
3.8 LPDDR4 Signal Termination	24
3.9 LPDDR4 VREF Routing	25
3.10 LPDDR4 VTT	25
3.11 CK and ADDR_CTRL Topologies	25
3.12 Data Group Topologies	25
3.13 CK and ADDR_CTRL Routing Specification	27
3.14 Data Group Routing Specification	28
3.15 Channel, Byte, and Bit Swapping	28
4 Revision History	28

Trademarks

All trademarks are the property of their respective owners.

1 Overview

The AM64x and AM243x processors support two different types of DDR memories: DDR4 and LPDDR4. This allows customer board designs to be implemented with the memory type that best meets their target market at the lowest possible DDR SDRAM cost. This document is divided into three sections. The first section contains material applicable to board designs containing either of the DDR SDRAM memory types. This is followed by sections containing information specific to each of the DDR memory types.

Note

To facilitate software configuration of the DDRSS, use the DDR Subsystem Configuration Tool in SysConfig (<https://dev.ti.com/sysconfig>)

1.1 Board Designs Supported

The goal of this document is to make the DDR system implementation straightforward for all designers. Requirements have been distilled down to a set of layout and routing rules that allow designers to successfully implement a robust design for the topologies that TI supports. At this time, TI does not provide timing parameters for the processor's DDR PHY interface.

It is still expected that the PCB design work (design, layout, and fabrication) be performed and reviewed by a highly knowledgeable high-speed PCB designer. Problems such as impedance discontinuities when signals cross a split in a reference plane can be detected visually by those with the proper experience.

TI only supports board designs using DDR4 and LPDDR4 memory that follow the guidelines in this document. These guidelines are based on well-known transmission line properties for copper traces routed over a solid reference plane. Declaring insufficient PCB space does not allow routing guidelines to be discounted.

1.2 General Board Layout Guidelines

To ensure good signaling performance, the following general board design guidelines must be followed:

- Avoid crossing plane splits in the signal reference planes.
- Some signals require a ground (also called VSS) reference plane to obtain the needed signal integrity. Some may even need it on both sides.
- Use the widest trace that is practical between decoupling capacitors and memory modules.
- Minimize inter-symbol interference (ISI) by keeping impedances matched.
- Minimize crosstalk by isolating sensitive signals, such as strobes and clocks, and by using a proper PCB stack-up.
- Avoid return path discontinuities by adding vias or capacitors whenever signals change layers and reference planes.
- Minimize reference voltage noise through proper isolation and proper use of decoupling capacitors on the reference input pins on the SDRAMs.
- Keep the signal routing stub lengths as short as possible.
- Add additional spacing for clock and strobe nets to minimize crosstalk.
- Maintain a common ground (VSS) reference for all bypass and decoupling capacitors.
- Consider the differences in propagation delays between microstrip and stripline nets when evaluating timing constraints.
- Via-to-via coupling can be a significant part of PCB-level crosstalk. GND shielding vias may need to be inserted between adjacent signal vias.
- Via stubs affect signal integrity. Via back-drilling may be required in some instances to improve signal integrity.

For more information, see the [High-Speed Interface Layout Guidelines](#). It provides additional general guidance for successful routing of high-speed signals.

1.3 PCB Stack-Up

The minimum stack-up for routing the DDR interface is a six-layer stack up. However, this can only be accomplished on a board with routing room with large keep-out areas. Additional layers are required if:

- The PCB layout area for the DDR Interface is restricted, which limits the area available to spread out the signals to minimize crosstalk.
- Other circuitry must exist in the same area, but on layers isolated from the DDR routing.
- Additional planes layers are needed to enhance the power supply routing or to improve EMI shielding.

Board designs that are relatively dense require 10 or more layers to properly allow the DDR routing to be implemented such that all rules are met.

DDR signals with the highest frequency content (such as data or clock) must be routed adjacent to a solid VSS reference plane. Signals with lower frequency content (such as address) can be routed adjacent to either a solid VSS or a solid VDDSD_DDR reference plane. If a VDDSD_DDR reference plane is used, bypass capacitors must be implemented near both ends of every route to provide a low-inductance, AC path to ground for these routes. Similarly, when multiple VSS reference planes exist in the DDR routing area, stitching vias must be implemented nearby wherever vias transfer signals to a different VSS reference plane. This is required to maintain a low-inductance return current path.

It is strongly recommended all DDR signals be routed as strip-line. Some PCB stack-ups implement signal routing on two adjacent layers. This is acceptable only as long as the routing on these layers is perpendicular and does not allow for broad-side coupling. Severe crosstalk occurs on any trace routed parallel to another trace on an adjacent layer, even for a short distance. Also, DDR signal routing on two adjacent layers is only allowed when implementing offset stripline routing, where the distance between the adjacent routing layers is more than 3x the distance from the traces to their adjacent reference plane.

Table 1-1. PCB Stack-up Specifications

Number	Parameter	MIN	TYP	MAX	UNIT
PS1	PCB routing plus plane layers	6			
PS2	Signal routing layers	3			
PS3	Full VSS reference layers under DDR routing region ⁽¹⁾	1			
PS4	Full VDDSD_DDR power reference layers under the DDR routing region ⁽¹⁾	1			
PS5	Number of reference plane cuts allowed within DDR routing region ⁽²⁾			0	
PS6	Number of layers between DDR routing layer and reference plane ⁽³⁾			0	
PS7	PCB routing feature size		4		Mils
PS8	PCB trace width, w		4		Mils
PS9	Single-ended impedance		40		Ω
PS10	Differential impedance		80		Ω
PS11	Impedance control ⁽⁴⁾	Z-10%	Z	Z+10%	Ω

- (1) Ground reference layers are preferred over power reference layers. Return signal vias need to be near layer transitions. When using power reference layers, include bypass caps to accommodate reference layer return current, as the trace routes switch routing layers.
- (2) No traces should cross reference plane cuts within the DDR routing region. High-speed signal traces crossing reference plane cuts create large return current paths, which can lead to excessive crosstalk and EMI radiation. Beware of reference plane voids caused by via antipads, as these also cause discontinuities in the return current path.
- (3) Reference planes are to be directly adjacent to the signal layer, to minimize the size of the return current loop.
- (4) Z is the nominal single-ended impedance selected for the PCB specified by PS9 and PS10.

1.4 Bypass Capacitors

1.4.1 Bulk Bypass Capacitors

Bulk bypass capacitors are required for moderate speed bypassing of the DDR SDRAMs and other circuitry. [Table 1-2](#) contains the minimum numbers and capacitance required for the bulk bypass capacitors. [Table 1-2](#) only covers the bypass needs of the SoC's DDR PHY. Additional bulk bypass capacitance may be needed for other circuitry. For any additional decoupling requirements for the SDRAM devices, see the manufacturer's data sheet

Table 1-2. Bulk Bypass Capacitors

Number	Parameter	MIN ⁽²⁾	MAX	UNIT
1	VDDS_DDR bulk bypass capacitor count ⁽¹⁾	1		Devices
2	VDDS_DDR bulk bypass total capacitance	22		μF

- (1) These capacitors should be placed near the devices they are bypassing, but preference should be given to the placement of the high-speed (HS) bypass capacitors and DDR signal routing.
- (2) The capacitor recommendations in this guide reflect only the needs of this processor. For determining the appropriate decoupling capacitor arrangement for the memory device itself, see the memory vendor's guidelines.

1.4.2 High-Speed Bypass Capacitors

High-speed (HS) bypass capacitors are critical for proper DDR interface operation. It is particularly important to minimize the parasitic series inductance of the HS bypass capacitors to VDDS_DDR and the associated ground connections. [Table 1-3](#) contains the specification for the HS bypass capacitors and for the power connections on the PCB. Generally speaking, TI recommends:

- Fitting as many HS bypass capacitors as possible.
- Minimizing the distance from the bypass capacitor to the pins and balls being bypassed.
- Using the smallest physical sized ceramic capacitors possible with the highest capacitance readily available.
- Connecting the bypass capacitor pads to their vias using the widest traces possible and using the largest via hole size possible.
- Minimizing via sharing. Note the limits on via sharing shown in [Table 1-3](#).

For any additional SDRAM requirements, see the manufacturer's data sheet.

Table 1-3. High-Speed Bypass Capacitors

Number	Parameter	MIN	TYP	MAX	UNIT
1	HS bypass capacitor package size ⁽¹⁾		0201	0402	10 Mils
2	Distance, HS bypass capacitor to processor being bypassed ⁽²⁾ ⁽³⁾ ⁽⁴⁾			400	Mils
3	Processor HS bypass capacitor count per VDDS_DDR rail	See PDN Guide ⁽⁹⁾			Devices
4	Processor HS bypass capacitor total capacitance per VDDS_DDR rail	See PDN Guide ⁽⁹⁾			μF
5	Number of connection vias for each device power/ground ball	1			Vias
6	Trace length from processor power/ground ball to connection via ⁽²⁾		35	70	Mils
7	Distance, HS bypass capacitor to DDR device being bypassed ⁽⁵⁾			150	Mils
8	DDR device HS bypass capacitor count ⁽⁶⁾	12			Devices
9	DDR device HS bypass capacitor total capacitance ⁽⁶⁾	0.85			μF
10	Number of connection vias for each HS capacitor ⁽⁷⁾ ⁽⁸⁾	2			Vias
11	Trace length from bypass capacitor to connection via ⁽²⁾ ⁽⁸⁾		35	100	Mils
12	Number of connection vias for each DDR device power/ground ball	1			Vias
13	Trace length from DDR device power/ground ball to connection via ⁽²⁾ ⁽²⁾		35	60	Mils

- (1) LxW, 10-mil units, that is, a 0402 is a 40x20-mil surface-mount capacitor.
- (2) Closer/shorter is preferable.
- (3) Measured from the nearest processor power or ground ball to the center of the capacitor package.

- (4) Three of these capacitors should be located underneath the processor, among the cluster of VDDSD_DDR balls.
- (5) Measured from the DDR device power or ground ball to the center of the capacitor package. Refer to the guidance from the SDRAM manufacturer.
- (6) Per DDR device. Refer to the guidance from the SDRAM manufacturer.
- (7) An additional HS bypass capacitor can share the connection vias only if it is mounted on the opposite side of the board. No sharing of vias is permitted on the same side of the board.
- (8) An HS bypass capacitor may share a via with a DDR device mounted on the same side of the PCB. A wide trace should be used for the connection, and the length from the capacitor pad to the DDR device pad should be less than 150 mils.
- (9) The capacitor recommendations in this guide reflect only the needs of this processor. See the memory vendor's guidelines for determining the appropriate decoupling capacitor arrangement for the memory device itself.

1.4.3 Return Current Bypass Capacitors

Use additional bypass capacitors if the return current reference plane changes due to DDR signals hopping from one signal layer to another, resulting in the reference plane changing from VDDSD_DDR to VSS. The bypass capacitor here provides a path for the return current to hop planes along with the signal. Use as many of these return current bypass capacitors as possible – up to one per signal via. Because these are returns for signal current, the via size for these bypass capacitors can be the smaller via used for signal routing.

1.5 Velocity Compensation

Because portions of the DDR signal traces are microstrip (top and bottom layers) while the majority of the trace segment length is stripline (internal layers), and because there is a wide variation in the proportion of track length routed as microstrip or stripline, the length/delay matching process should include a mechanism for compensating for the velocity delta between these two types of PCB interconnects. A compensation factor of 1.1 has been specified for this purpose by JEDEC. All microstrip segment lengths are to be divided by 1.1 before summation into the length matching equation. The resulting compensated length is termed the 'stripline equivalent length'. While some amount of residual velocity mismatch skew remains in the design, the process is a substantial improvement over simple length matching.

2 DDR4 Board Design and Layout Guidance

2.1 DDR4 Introduction

DDR4 board designs are similar to DDR3 board designs. Fly-by routing is required just as it is with DDR3, and thus leveling is required. To achieve higher data rates with DDR4, there are several enhancements added to the interface specification that must be accommodated by both the SDRAM and the processor's interface (PHY). The enhancements that affect the board interconnect and layout are listed below:

- Addition of ACT_n pin – This pin provides signaling to allow the pins previously called Command pins (RAS_n, CAS_n and WE_n) to be used as additional address pins. These pins behave as row address pins when ACT_n is low and as command pins when ACT_n is high. This is valid only when CS_n is low.
- Removal of one BA (Bank Address) pin and addition of 2 BG (Bank Group) pins – This adds flexibility with accesses similar to DDR3, but with 16 banks bundled in four bank groups of four banks each. This results in additional timing parameters, because adjacent accesses within a bank group are faster than adjacent accesses to another bank group. Successive accesses to locations within a single bank are the fastest option.
- Addition of PAR (Parity) and ALERT_n pins (use is optional) – The PAR pin supplies parity monitoring for the command and address pins using even parity from the controller to the SDRAM. ALERT_n is the indicator (open-drain output) from the SDRAMs that indicate when a parity error has been detected.
- Change to POD termination – Pseudo-Open Drain (POD) output buffers are implemented rather than traditional SSTL push-pull outputs. This allows the data bit termination, ODT, to go to the I/O power rail, VDDQ, rather than to the mid-level voltage, VTT. Power consumption may be reduced, because only driving a bit low draws current.
- Addition of DBI – Data bus invert (DBI) is a feature that allows the data bus to be inverted whenever more than half of the bits are zero. This feature may reduce active power and enhance the data signal integrity when coupled with POD termination.
- Addition of a VPP power input – The VPP power supply (2.5 V) provides power to the internal word line logic. This voltage increase allows the SDRAM to reduce overall power consumption.
- Separation of data VREF from address/control VREF – The data reference voltage, VREFDQ, is now internally generated both within the SDRAM and within the PHY. It can be programmed to various levels to provide the optimum sampling threshold. The optimum threshold varies based on the ODT impedance chosen, the drive strength, and the PCB track impedance. The address/control reference voltage, VREFCA, is a mid-level reference voltage, the same as it is on DDR3.

Note

These features may not be supported on all devices. Refer to the device-specific documentation for supported features.

2.2 DDR4 Device Implementations Supported

There are several possible combinations of SDRAM devices supported by the DDR4 EMIF. [Table 2-1](#) lists the supported device combinations. The SDRAMs used in each combination must be identical: that is, they must have the same part number.

Table 2-1. Supported DDR4 SDRAM Combinations

Number of DDR4 SDRAMs	DDR4 SDRAM Width (bits)	DDR4 EMIF Width (bits)
1	16	16
2	8	16

2.3 DDR4 Interface Schematics

This section discusses implementations (also called topologies) using single-rank x16 and x8 SDRAM devices. This section does not discuss recommendations for implementations that support low-power operation, such as when the SDRAM is held in self-refresh and the processor is powered off. It also does not discuss the DDR-less implementations. These options are under study and may be supported in future versions of this document.

2.3.1 DDR4 Implementation Using 16-Bit SDRAM Devices

The DDR4 interface schematics vary, depending upon the width of the DDR4 SDRAM devices used and the width of the EMIF bus implemented. General connectivity is straightforward and consistent between the implementations. 16-bit SDRAM devices look like two 8-bit devices. [Figure 2-1](#) shows the schematic connections for a 16-bit interface using a single x16 SDRAM.

When not using one of the byte lanes on the processor, the proper method of handling the unused pins is to tie off the unused DDR_DQSxP pins to ground through a 1k- Ω resistor and to tie off the unused DDR_DQSxN pins to the VDDS_DDR supply, also referred to as the I/O supply VDDQ, through a 1k- Ω resistor. This must be done for each byte not used. Although these signals have internal pullups and pulldowns, external pullups and pulldowns provide additional protection against external electrical noise causing activity on the signals.

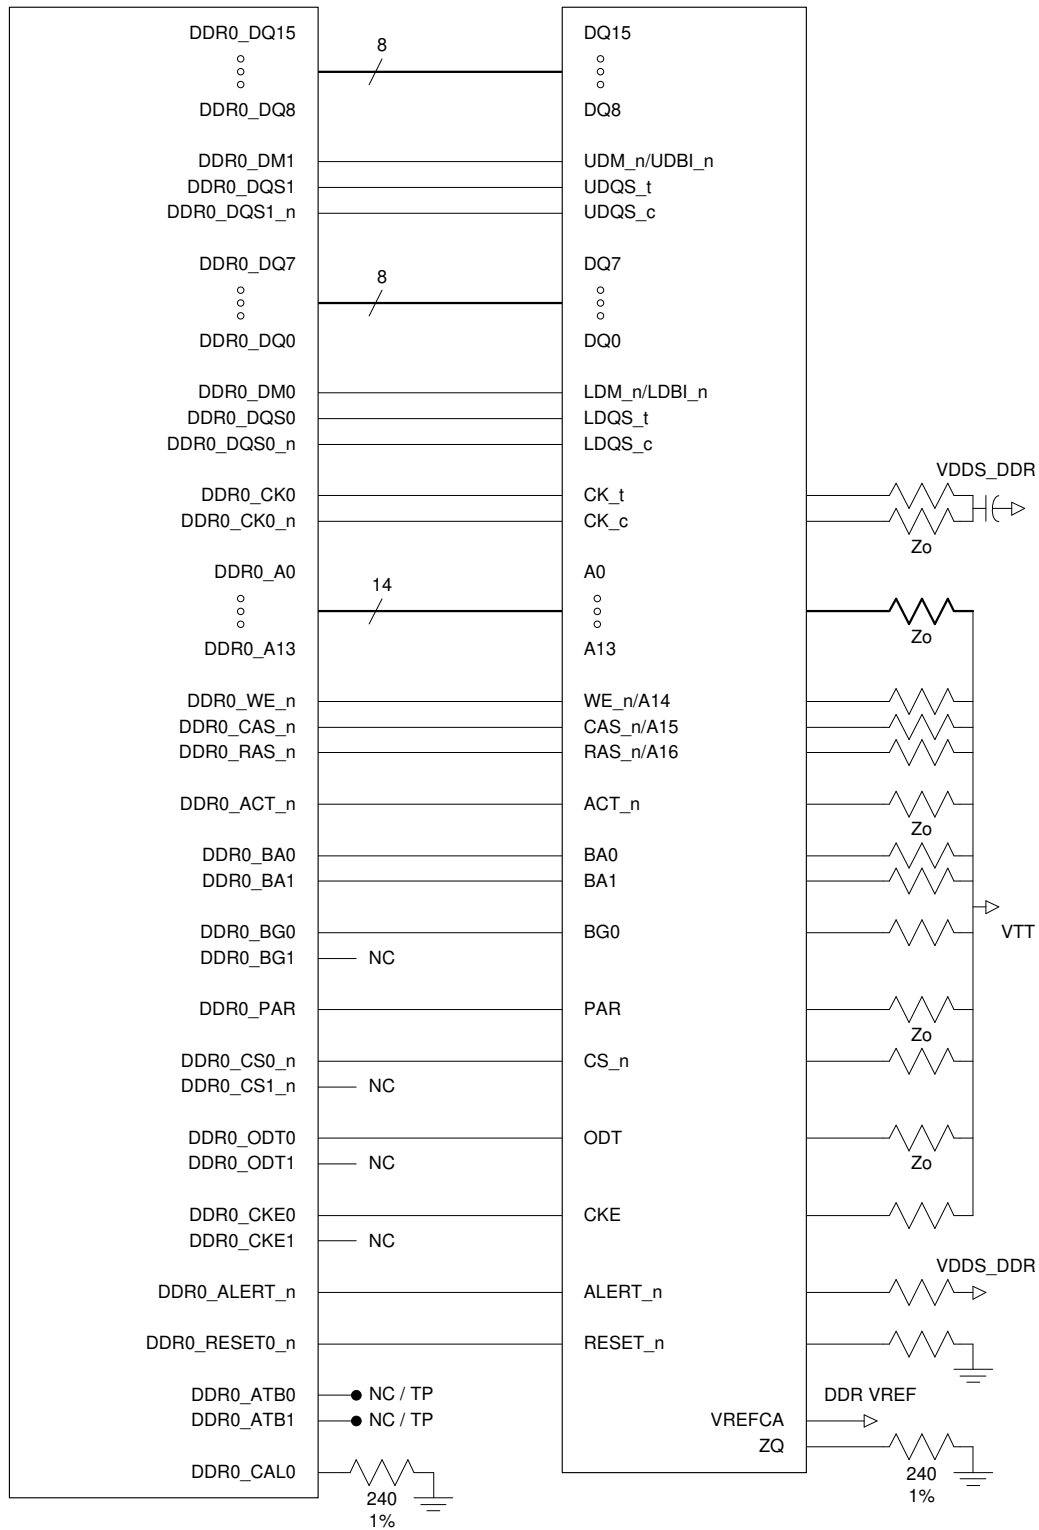


Figure 2-1. 16-Bit, Single-Rank DDR4 Implementation Using x16 SDRAM

2.3.2 DDR4 Implementation Using 8-Bit SDRAM Devices

Figure 2-2 shows the schematic connections for a 16-bit interface using x8 devices.

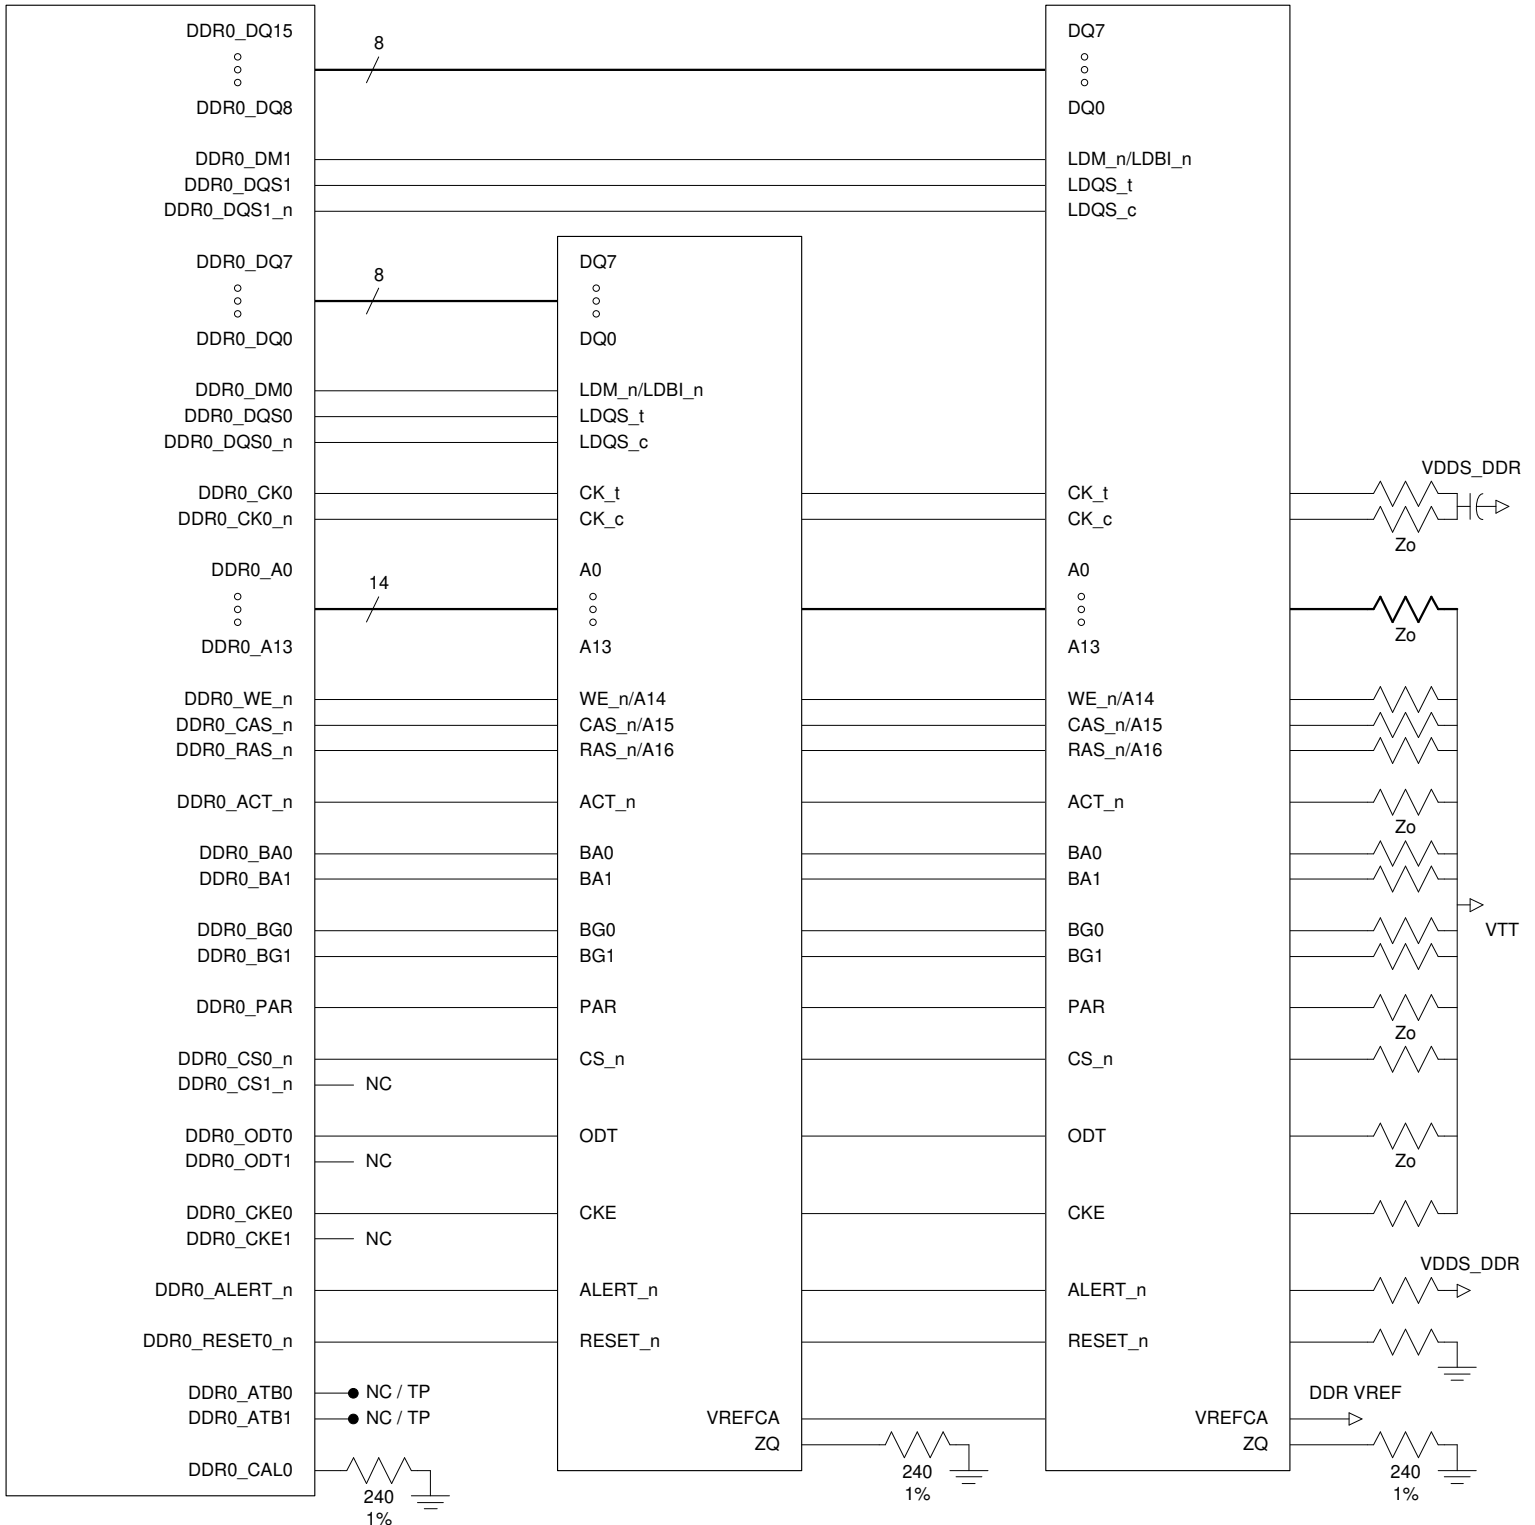


Figure 2-2. 16-Bit, Single-Rank DDR4 Implementation Using x8 SDRAMs

2.4 Compatible JEDEC DDR4 Devices

Table 2-2 shows the parameters of the JEDEC DDR4 devices compatible with this interface. Generally, the DDR4 interface is compatible with all JEDEC-compliant DDR4 SDRAM devices in x8 or x16 widths.

Table 2-2. Compatible JEDEC DDR4 Devices

Number	Parameter	MIN	MAX	UNIT
1	JEDEC DDR4 data rate ⁽²⁾ ⁽³⁾		1600	MT/s
2	JEDEC DDR4 device bit width	x8	x16	Bits
3	JEDEC DDR4 device count ⁽¹⁾	1	2	Devices

- (1) For valid DDR4 device configurations and device counts, see Figure 2-1 and Figure 2-2.
- (2) Refer to the device data manual for supported data rates.
- (3) SDRAMs in faster speed grades can be used provided they are properly configured to operate at the supported data rates. Faster speed grade SDRAMs may have faster edge rates, which may affect signal integrity. SDRAMs with faster speed grades must be validated on the target board design.

2.5 Placement

Figure 2-3 shows the required placement for the processor and the DDR4 devices. The dimensions for this figure are defined in Table 2-3. The placement does not restrict the side of the PCB on which the devices are mounted. The ultimate purpose of the placement is to limit the maximum trace lengths and allow for proper routing space.

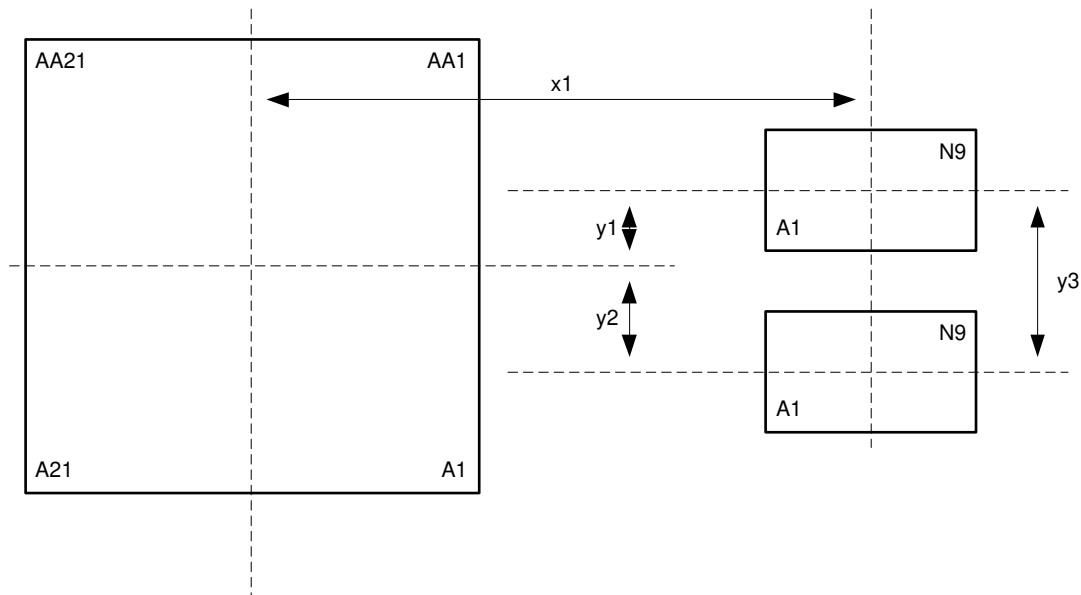


Figure 2-3. DDR4 Placement Specifications

Table 2-3. Placement Parameters

Number	Parameter	MIN	MAX	UNIT
1	x1		2000	Mils
2	y1		500	Mils
3	y2		1000	Mils
4	y3		750	Mils

2.6 DDR4 Keepout Region

The region of the PCB used for DDR4 circuitry must be isolated from other signals. The DDR4 keepout region is defined for this purpose and is shown in Figure 2-4. The size of this region varies with the placement and DDR routing. Non-DDR4 signals should not be routed on the DDR signal layers within the DDR4 keepout region. Non-DDR4 signals may be routed in this region only if they are routed on other layers separated from the DDR signal layers by a ground layer. No breaks are allowed in the reference ground layers in this region. In addition, a solid VDDSD_DDR power plane should exist across the entire keepout region.

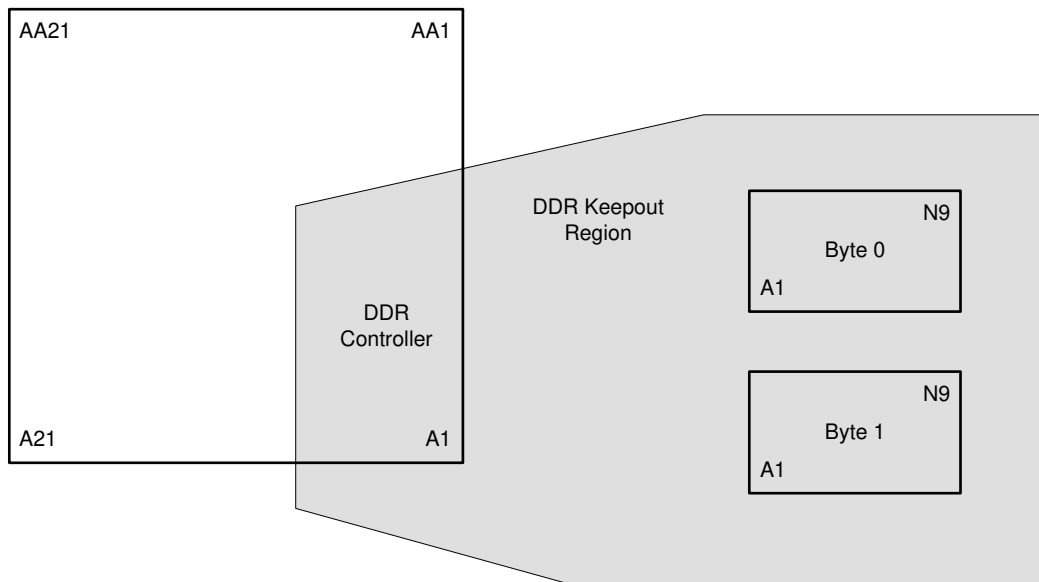


Figure 2-4. DDR4 Keepout Region

2.7 VPP

VPP is a new supply input on DDR4 SDRAMs. This supply must provide an average of less than 5 mA in active and standby modes and 10 to 20 mA during refresh. There is not a constant current draw during refresh. The VPP power supply and decoupling capacitors must be able to supply short bursts of current up to 60 mA during this time.

2.8 Net Classes

Routing rules are applied to signals in groups called net classes. Each net class contains signals with the same routing requirements. This simplifies the implementation and compliance of these routes. Table 2-4 lists the clock net classes for the DDR4 interface. Table 2-5 lists the signal net classes, and associated clock net classes, for signals in the DDR4 interface. These net classes are then linked to the termination and routing rules that follow.

Table 2-4. Clock Net Class Definitions

Clock Net Class	Processor Pin Names
CK	DDR0_CK0 / DDR0_CK0_n
DQS0	DDR0_DQS0 / DDR0_DQS0_n
DQS1	DDR0_DQS1 / DDR0_DQS1_n

Table 2-5. Signal Net Class Definitions

Signal Net Class	Associated Clock Net Class	Processor Pin Names
ADDR_CTRL	CK	DDR0_A[13:0], DDR0_WE_n, DDR0_CAS_n, DDR0_RAS_n, DDR0_ACT_n, DDR0_BA0, DDR0_BA1, DDR0_BG0, DDR0_BG1, DDR0_PAR, DDR0_CS0_n, DDR0_CS1_n, DDR0_ODT0, DDR0_ODT1, DDR0_CKE0, DDR0_CKE1
BYTE0	DQS0	DDR0_DQ[7:0], DDR0_DM0
BYTE1	DQS1	DDR0_DQ[15:8], DDR0_DM1

2.9 DDR4 Signal Termination

Signal terminators are required for the CK and ADDR_CTRL net classes. This is shown in [Figure 2-1](#) and [Figure 2-2](#). The data group nets are terminated by ODT in the processor and SDRAM memories, and thus the data group PCB traces must be unterminated. Detailed termination specifications are covered in the routing rules in the following sections.

2.10 VREF Routing

JEDEC defines two reference voltages that are used with DDR4 memory interfaces. These are VREFDQ and VREFCA. VREFDQ is the reference voltage used for the data group nets during reads and writes. VREFCA is the reference voltage used for command and address inputs to the SDRAMs. DDR4 SDRAMs generate their own VREFDQ internally. Similarly, the processor's DDR4 PHY generates its own VREFDQ internally. The VREFCA reference voltage must be generated on the board and propagated to all of the SDRAMs. VREFCA is intended to be 50% of the DDR4 power supply voltage and is typically generated with the DDR4 VTT power supply. It should be routed as a nominal 20-mil wide trace with 0.1- μ F bypass capacitors near each device connection. Narrowing the VREF trace is allowed to accommodate routing congestion for short lengths near endpoints.

2.11 VTT

As with VREFCA, the nominal value of the VTT supply is 50% of the DDR4 supply voltage. Unlike VREFCA, the VTT supply is expected to source and sink current; specifically the termination current for the ADDR_CTRL net class Thevenin terminators. VTT is needed at the end of the address and control bus and it should be routed as a power sub-plane. VTT must be bypassed near the terminator resistors.

2.12 POD Interconnect

Prior to DDR4, the output buffers were push-pull CMOS buffers. They would sink current when driving low and source current when driving high. They were then terminated to a mid-level Thevenin resistance to obtain optimum power transfer and signal integrity. Unfortunately, this resulted in current flowing, and power being dissipated, whenever the buffers were enabled at either high or low. Pseudo Open Drain (POD) is a connection type where the termination at the load, ODT, is only connected to VDDQ. POD connections only consume power when driving low, thus reducing power. In DDR4, both the PHY (for reads) and SDRAM (for writes) provide these terminations to VDDQ internally on all of the data group pins.

Signals look different on connections using POD terminations as compared to previous DDR connections, where the data group signals went from VSS to VDDQ and sampling was based on a mid-level reference voltage. The high level is still at VDDQ. However, the low level is now calculated based on the drive impedance and the ODT resistance. If they are both set to 50 Ω , the low-level voltage is now at VDDQ/2. That then requires a sampling voltage half way between those voltages, or $3/4 * VDDQ$, for optimum performance.

2.13 CK and ADDR_CTRL Topologies and Routing Guidance

The CK and ADDR_CTRL net classes are routed similarly, and are length matched from the DDR PHY in the processor to each SDRAM to minimize skew between them. The CK net class requires more care because it runs at a higher transition rate and is differential.

The CK and ADDR_CTRL net classes are routed in a ‘fly-by’ implementation. This means that the CK and ADDR_CTRL net classes are routed as a multi-drop bus from the DDR controller in the processor sequentially to each SDRAM, and each signal has a termination at the end. To complete this routing, a small stub trace exists on each net at each SDRAM. These stubs must be short and approximately the same length to manage the reflections. The ADDR_CTRL net class is length matched to the CK net class, at each SDRAM, so that the ADDR_CTRL signals are properly sampled at each SDRAM.

Note

Fly-by routing is required for DDR4 layouts. Balanced-T routing, previously used for DDR2 layouts, is not supported.

Section 2.2 discussed that there are multiple possible memory topologies, or implementations, ranging from a single x16 SDRAM up to a maximum of two x8 SDRAMs. Regardless of the number of SDRAMs implemented, the routing requirements must be followed. TI recommends that all SDRAMs be implemented on the same side of the board, preferably on the same side of the board as the processor. It is possible to implement the SDRAMs on both sides of the board, but the routing complexity and the number of PCB layers required is significantly increased.

Figure 2-5 shows the topology of the CK net class, and Figure 2-6 shows the topology for the corresponding ADDR_CTRL net class. The fly-by routes have been broken into segments to simplify the length matching analysis. Care must be taken to avoid excessive length error accumulation with this method.

Segments A1 and A2 comprise the lead-in section. Segment AT is the track to the termination at the end of the net. Segments A3 are the routed track between the stubs that branch off to each SDRAM. For topologies with fewer SDRAMs, remove an A3 segment for each SDRAM not present. Length matching requirements for the routing segments are detailed in Table 2-6.

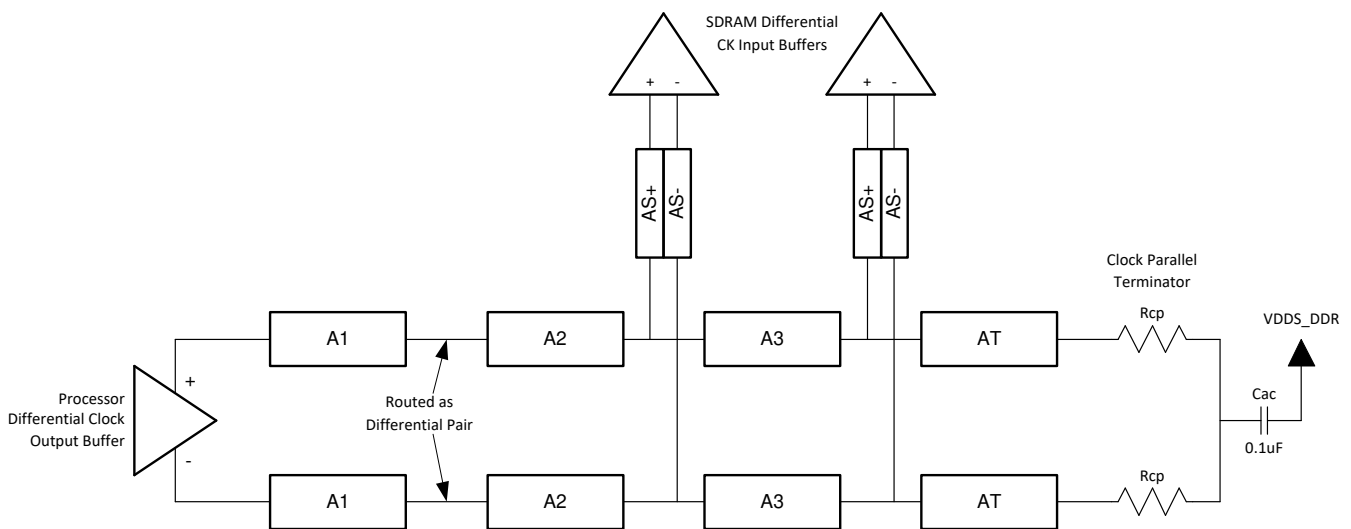


Figure 2-5. CK Topology for Two DDR4 SDRAM Devices

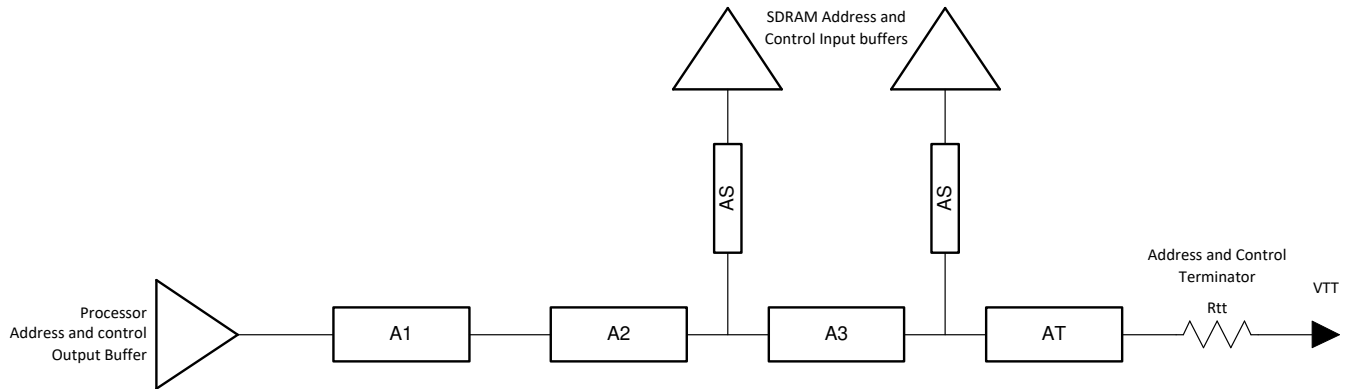


Figure 2-6. ADDR_CTRL Topology for Two DDR4 SDRAM Devices

The previous figures show the circuit topology such that the track lengths can be managed and the routed track length matching rules can be followed. The next two figures again show the routing for the CK and ADDR_CTRL routing groups depicted from the perspective of tracks routed on the PCB.

Figure 2-7 shows the CK group routing for two SDRAM devices. The fly-by routing is made clear in this figure. The DDR0_CK0 and DDR0_CK0_n tracks (the CK routing group) are routed as a differential pair from the processor to the SDRAM at the end that will contain BYTE0 data. This differential pair routing then proceeds to the other SDRAM and ends with the AC termination to VDDSDR. The routing also includes the routing stubs for both DDR0_CK0 and DDR0_CK0_n at each SDRAM.




Figure 2-7. CK Routing for Two DDR4 SDRAM Devices

Figure 2-8 shows the ADDR_CTRL routing for two SDRAM devices. These are also routed in a fly-by manner along the same path because the ADDR_CTRL routing group is length-matched to the CK routing group.




Figure 2-8. ADDR_CTRL Routing for Two DDR4 SDRAM Devices

The absolute order is not significant. The fly-by routing that starts at the processor can also route down to the SDRAM containing the last byte of data (or whichever SDRAM that is opposite in the row from the one containing the BYTE0 data). The fly-by routing then proceeds to the other SDRAM as discussed above, until it routes to VTT through the Rtt termination after the BYTE0 SDRAM.

Minimize layer transitions during routing. If a layer transition is necessary, it is preferable to transition to a layer using the same reference plane. If this cannot be accommodated, ensure there are nearby stitching vias to allow the return currents to transition between reference planes when both reference planes are ground or VDDSD_DDR. Alternately, ensure there are nearby bypass capacitors to allow the return currents to transition between reference planes when one of the reference planes is ground and the other is VDDSD_DDR. This must occur at every reference plane transition. The goal is to minimize the size of the return current path thus minimizing the inductance in this path. Lack of these stitching vias or capacitors results in impedance discontinuities in the signal path that increase crosstalk and signal distortion.

2.14 Data Group Topologies and Routing Guidance

Regardless of the number of DDR4 devices implemented, the data line topology is always point-to-point. Minimize layer transitions during routing. If a layer transition is necessary, it is better to transition to a layer using the same reference plane. If this cannot be accommodated, ensure there are nearby ground vias to allow the return currents to transition between reference planes. The goal is to provide a low inductance path for the return current. Also, to optimize the length matching, TI recommends routing all nets within a single data routing group on one layer where all have the exact same number of vias and the same via barrel length.

DQSP and DQSN lines are point-to-point signals routed as a differential pair. [Figure 2-9](#) shows the DQS connection topology.




Figure 2-9. DDR4 DQS Topology

DQ and DM lines are point-to-point signals routed singled-ended. [Figure 2-10](#) shows the DQ and DM connection topology.




Figure 2-10. DDR4 DQ/DM Topology

Similar to the figures above for the CK and ADDR_CTRL routes, [Figure 2-11](#) and [Figure 2-12](#) show an example of the PCB routes for a DQS routing group and the associated data routing group nets.

The routing example shows DQS0P and DQS0N, which are routed as a differential pair from the processor to the SDRAM that contains Byte 0. This is implemented as a point-to-point routed differential pair without any board terminations. There are no stubs allowed on these nets of any kind. All test access probes must be in line without any branches or stubs. Similar DQS pair routing exists from the processor to each SDRAM for the byte lanes implemented.

[Figure 2-12](#) shows a routing example for a single net in the Byte 0 routing group. The DQ and DM nets are routed single-ended and are also point-to-point without any stubs or board terminations. Point-to-point routes exist for each of the DQ and DM nets implemented.

The DQ and DM nets are routed along the same path as the DQSP and DQSN pair for that byte lane, so that they can be length matched to the DQS pair.

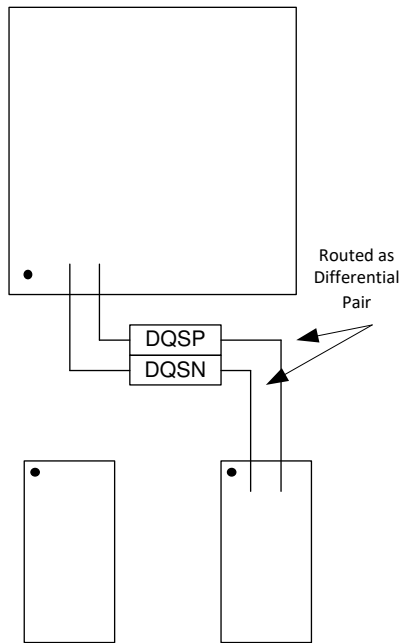


Figure 2-11. DQS Routing to Two DDR4 SDRAM Devices




Figure 2-12. DQ/DM Routing to Two DDR4 SDRAM Devices

2.15 CK and ADDR_CTRL Routing Specification

Skew within the CK and ADDR_CTRL net classes directly reduces setup and hold margin for the ADDR_CTRL nets. Thus, this skew must be controlled. Routed PCB track has a delay proportional to its length. Thus, the delay skew must be managed through matching the lengths of the routed tracks within a defined group of signals. The only way to practically match lengths on a PCB is to lengthen the shorter traces up to the length of the longest net in the net class and its associated clock pair, DDR0_CK0 and DDR0_CK0_n.

2.15.1 CACLM - Clock Address Control Longest Manhattan Distance

A metric to establish a maximum length is Manhattan distance. The Manhattan distance between two points on a PCB is the length between the points when connecting them only with horizontal or vertical track segments. A

reasonable limit to the trace route length is to its Manhattan distance plus some margin. CACLM is this limit and it is defined as the Clock Address Control Longest Manhattan distance.

Given the clock and address pin locations on the processor and the DDR4 memories, the maximum possible Manhattan distance can be determined given the placement of these parts. It is from this distance that this rule-of-thumb limit on the lengths of the routed track for the CK and ADDR_CTRL routing groups is determined.

It is likely that the longest CK and ADDR_CTRL Manhattan distance will be for Address Input A13 on the DDR4 SDRAM device, because it is at the farthest corner in the placement. Assuming A13 is the longest, calculate CACLM as the sum of CACLMY(A13) + CACLMX(A13) + 300 mils. The extra 300 mils allows for routing past the first DDR4 SDRAM and returning up to reach pin A13. Use this as a guideline for the upper limit to the length of the routed traces from the processor to the first SDRAM.

2.15.2 CK and ADDR_CTRL Routing Limits

Table 2-6 lists the limits for the individual segments that comprise the routing from the processor to the SDRAM. These segment lengths coincide with the CK and ADDR_CTRL topology diagram shown previously in Figure 2-5 and Figure 2-6. By matching the length for the same segments of all signals in a routing group, the signal delay skews are controlled.

Recall that the CK and ADDR_CTRL nets route along the same path for each segment. This simplifies the length matching. The skew limits for the CK group compare the length of DDR0_CK0P to the length of DDR0_CK0N. Then the skew limits for the ADDR_CTRL group nets are compared to the CK group nets.

Most PCB layout tools can be configured to generate reports to assist with this validation. If this cannot be generated automatically, this must be generated and verified manually.

Table 2-6 also lists skew limits for the full routes from the processor to each SDRAM. This must be checked in addition to the skew limits in the individual sections to verify that there is not accumulating error in the layout.

To use length matching (in mils) instead of time delay (in ps), multiply the time delay (in ps) limit by 5. The microstrip routes propagate faster than stripline routes. A standard practice when using length matching is to divide the microstrip length by 1.1 to achieve a compensated length to normalize the microstrip length with the stripline length and to align with the delay limits provided. This is called velocity compensation (see Section 1.5).

Table 2-6. CK and ADDR_CTRL Routing Specifications

Number	Parameter	MIN	TYP	MAX	UNIT
1	A1+A2 length			500 ⁽¹⁾	ps ⁽¹²⁾
2	A1+A2 skew ADDR_CTRL to CK ⁽⁴⁾			3	ps
4	A3 skew ADDR_CTRL to CK ⁽⁴⁾			3	ps
3	A3 length			125	ps
5	A1+A2 skew DDR0_CK0 to DDR0_CK0_n			0.4	ps
6	A3 skew DDR0_CK0 to DDR0_CK0_n			0.4	ps
7	AS length		5 ⁽¹⁾	17	ps
8	AS skew		1.3 ⁽¹⁾	3	ps
9	AS+/AS- length		5	17	ps
10	AS+/AS- skew			0.4	ps
11	AT length ⁽³⁾		75		ps
12	AT skew ADDR_CTRL to CK ⁽⁴⁾		14		ps
13	AT skew DDR0_CK0 to DDR0_CK0_n			0.4	ps
14	Total DDR0_CK0 to DDR0_CK0_n skew from processor to each SDRAM ⁽²⁾			0.8	ps
15	Total CK to ADDR_CTRL skew from processor to each SDRAM ⁽²⁾			4	ps
16	Vias per trace ⁽¹¹⁾			3 ⁽¹⁾	vias
17	Via count difference ⁽¹¹⁾			1 ⁽¹⁰⁾	vias
18	Center-to-center CK to other DDR4 trace spacing ⁽⁵⁾	4w			

Table 2-6. CK and ADDR_CTRL Routing Specifications (continued)

Number	Parameter	MIN	TYP	MAX	UNIT
19	Center-to-center ADDR_CTRL to other DDR4 trace spacing ⁽⁵⁾	4w			
20	Center-to-center ADDR_CTRL to other ADDR_CTRL trace spacing ⁽⁵⁾	3w			
21	CK center-to-center spacing ^{(6) (7)}	See notes below			
22	CK spacing to other net ⁽⁵⁾	4w			
23	Rcp ⁽⁸⁾	Zo-1	Zo	Zo+1	Ω
24	Rtt ^{(8) (9)}	Zo-5	Zo	Zo+5	Ω

- (1) Max value is based upon conservative signal integrity approach. This value could be extended only if detailed signal integrity analysis of rise time and fall time confirms desired operation.
- (2) This is the combined length from the processor to the SDRAM. It must be computed for each SDRAM to ensure that the segment matching does not result in accumulated error. For the first SDRAM, it is A1 + A2 + AS, computed for each signal. For the 2nd SDRAM, it is A1 + A2 + A3 + AS, computed for each signal.
- (3) While this length can be increased for convenience, its length should be minimized.
- (4) ADDR_CTRL net class relative to its CK net class.
- (5) Center-to-center spacing is allowed to fall to minimum 2w for up to 500 mils of routed length (only near endpoints).
- (6) CK spacing set to ensure proper differential impedance.
- (7) The user must control the impedance so that inadvertent impedance mismatches are not created. Generally speaking, center-to-center spacing should be either 2w or slightly larger than 2w to achieve a differential impedance equal to twice the single-ended impedance, Zo, on that layer.
- (8) Source termination (series resistor at driver) is specifically not allowed.
- (9) Termination values should be uniform across the net class.
- (10) Via count difference may increase by 1 only if accurate 3-D modeling of the signal flight times – including accurately modeled signal propagation through vias – has been applied to ensure all segment skew maximums are not exceeded.
- (11) Count vias individually from processor to each SDRAM.
- (12) PCB track length shown as ps is a normalized representation of length. 1 ps can be equated to 5 mils as a simple transformation. This is stripline equivalent length where velocity compensation must be used for all segments routed as microstrip track.

2.16 Data Group Routing Specification

Skew within the DQS and DQ/DM net classes directly reduces setup and hold margin for the DQ and DM nets. Thus, this skew must be controlled. Routed PCB track has a delay proportional to its length. Thus, the length skew must be managed through matching the lengths of the routed tracks within a defined group of signals. The only way to practically match lengths on a PCB is to lengthen the shorter traces up to the length of the longest net in the net class and its associated clock pair, DQSP, and DQSN.

2.16.1 DQLM - DQ Longest Manhattan Distance

As with CK and ADDR_CTRL, a reasonable trace route length is to within a percentage of its Manhattan distance. DQLMn is defined as DQ Longest Manhattan distance n, where n is the byte number. For a 16-bit interface, there are two DQLMs, DQLM0 and DQLM1.

Note

It is not required nor recommended to match the lengths across all byte lanes. Length matching is only required within each byte.

Given the DQS, DQ, and DM pin locations on the processor and the DDR4 memories, the maximum possible Manhattan distance can be determined given the placement. It is from this distance that an upper limit on the lengths of the transmission lines for the data bus can be established. Unlike the CACLM, there is no margin added to the DQLMn limits. These limits are simply the sum of the horizontal and vertical distances for the longest pin to pin route for that byte group.

2.16.2 Data Group Routing Limits

Table 2-7 contains the routing specifications for DQS, DQ, and DM routing groups. Each byte lane is routed and matched independently.

To use length matching (in mils) instead of time delay (in ps), multiply the time delay (in ps) limit by 5. The microstrip routes propagate faster than stripline routes. A standard practice when using length matching is to divide the microstrip length by 1.1, to achieve a compensated length to normalize the microstrip length with the stripline length and to align with the delay limits provided (see Section 1.5).

Table 2-7. Data Group Routing Specifications

Number	Parameter	MIN	MAX	UNIT
DRS31	BYTE0 length		500	ps ⁽¹⁰⁾
DRS32	BYTE1 length		500	ps
DRS36	DQSn+ to DQSn- skew		0.4	ps
DRS37	DQSn to DQn skew ^{(2) (3)}		2	ps
DRS38	Vias per trace		2 ⁽¹⁾	vias
DRS39	Via count difference		0 ⁽⁹⁾	vias
DRS310	Center-to-center BYTEn to other DDR4 trace spacing ⁽⁵⁾	4		w ⁽⁴⁾
DRS311	Center-to-center DQn to other DQn trace spacing ⁽⁶⁾	3		w ⁽⁴⁾
DRS312	DQSn center-to-center spacing ^{(7) (8)}	See notes below		
DRS313	DQSn center-to-center spacing to other net	4		w ⁽⁴⁾

- (1) Max value is based upon conservative signal integrity approach. This value could be extended only if detailed signal integrity analysis of rise time and fall time confirms desired operation.
- (2) Length matching is only done within a byte. Length matching across bytes is neither required nor recommended.
- (3) Each DQS pair is length matched to its associated byte.
- (4) Center-to-center spacing is allowed to fall to minimum 2w for up to 500 mils of routed length (only near endpoints).
- (5) Other DDR4 trace spacing means other DDR4 net classes not within the byte.
- (6) This applies to spacing within the net classes of a byte.
- (7) DQS pair spacing is set to ensure proper differential impedance.
- (8) The user must control the impedance so that inadvertent impedance mismatches are not created. Generally speaking, center-to-center spacing should be either 2w or slightly larger than 2w to achieve a differential impedance equal to twice the single-ended impedance, Z_o , on that layer.
- (9) Via count difference may increase by 1 only if accurate 3-D modeling of the signal flight times – including accurately modeled signal propagation through vias – has been applied to ensure DQn skew and DQSn to DQn skew maximums are not exceeded.
- (10) PCB track length shown as ps is a normalized representation of length. 1 ps can be equated to 5 mils as a simple transformation. This is stripline equivalent length where velocity compensation must be used for all segments routed as microstrip track.

2.17 Bit Swapping

2.17.1 Data Bit Swapping

Data bit swapping is allowed to simplify routing as long as the bits swapped are within the same byte group. This is only possible when not using CRC. However, the prime bit, the lowest numbered bit in each byte, must be connected to the corresponding bit on the SDRAM without swapping. That is bit 0 and bit 8. Also, the DM and DQS bits must not be swapped.

2.17.2 Address and Control Bit Swapping

Bit swapping of the address or control bits is not allowed, as this breaks functionality.

3 LPDDR4 Board Design and Layout Guidance

3.1 LPDDR4 Introduction

LPDDR4 is an SDRAM device specification governed by the JEDEC standard JESD209-4, *Low Power Double Data Rate 4 (LPDDR4)*. This standard strives to reduce power and improve signal integrity by implementing a lower voltage I/O power rail, employing ODT on the Command/Address bus, and reducing the overall width of the Command/Address bus, among other features. Unlike other DDR types, LPDDR4 has been organized into 16-bit channels.

The following sections detail the routing specification and layout guidelines for an LPDDR4 interface.

3.2 LPDDR4 Device Implementations Supported

LPDDR4 supports many different implementation topologies. However, the devices only support a single 16-bit channel for LPDDR4. SDRAMs with additional channels and/or dies can be implemented but the additional channels/dies will be unconnected and not used. [Table 3-1](#) lists the only supported LPDDR4 device combination.

Table 3-1. Supported LPDDR4 SDRAM Combinations

LPDDR4 SDRAM Count	Channels	Die	Ranks	LPDDR4 Channel Width	DDRSS Data Width
1	1	1	1	16 bits	16 bits

3.3 LPDDR4 Interface Schematics

As stated above, LPDDR4 supports many different implementation topologies but the devices only supports a single 16-bit channel for LPDDR4. Figure 3-1 illustrates the supported 16-bit, single-rank, single-channel LPDDR4 implementation. SDRAMs with additional channels and/or dies can be used but the additional channels/dies will be unconnected and not used.




Figure 3-1. 16-Bit, Single-Rank, Single Channel LPDDR4 Implementation

3.4 Compatible JEDEC LPDDR4 Devices

Table 3-2 shows the parameters of the JEDEC LPDDR4 devices compatible with this interface.

Table 3-2. Compatible JEDEC LPDDR4 Devices

Number	Parameter	MIN	MAX	UNIT
1	Data Rate ^{(1) (2)}		1600	MT/s
2	Channel Bit Width	x16	x16	Bits
3	Channels	1	1	-
4	Ranks	1	1	-
5	Die	1	1	-
6	Device Count	1	1	-

- (1) Refer to the device data manual for supported data rates.
- (2) SDRAMs in faster speed grades can be used, provided they are properly configured to operate at the supported data rates. Faster speed grade SDRAMs may have faster edge rates, which may affect signal integrity. SDRAMs with faster speed grades must be validated on the target board design.

3.5 Placement

Figure 3-2 shows the required placement for the processor and the LPDDR4 device. The dimensions for this figure are defined in Table 3-3. The placement does not restrict the side of the PCB on which the devices are mounted. The ultimate purpose of the placement is to limit the maximum trace lengths and allow for proper routing space.

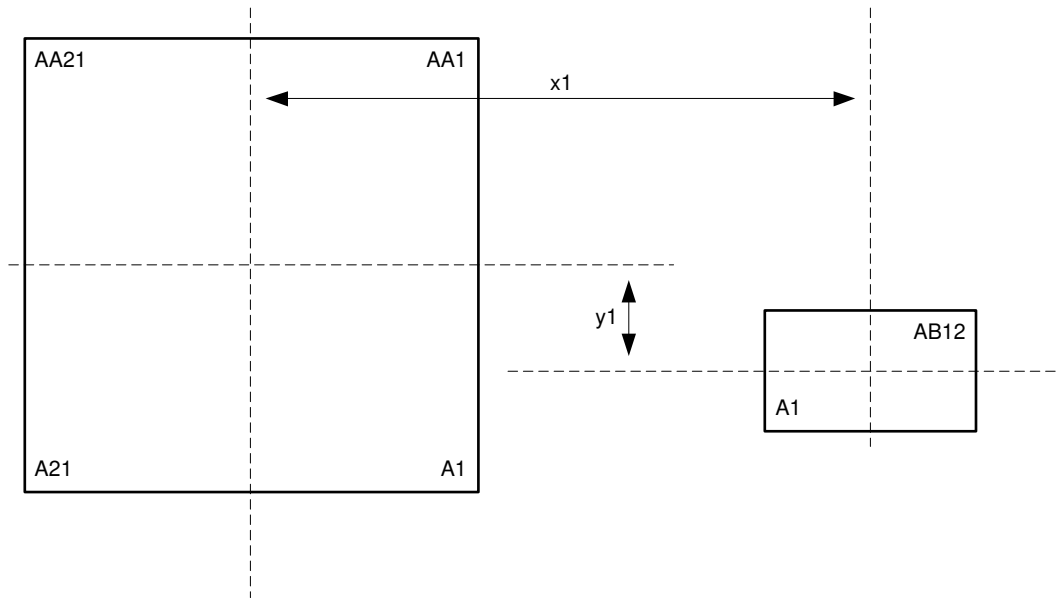


Figure 3-2. LPDDR4 Placement Specification

Table 3-3. LPDDR4 Placement Parameters

Number	Parameter	MIN	MAX	UNIT
1	x1		2000	Mils
2	y1		1000	Mils

3.6 LPDDR4 Keepout Region

The region of the PCB used for LPDDR4 circuitry must be isolated from other signals. The LPDDR4 keepout region is defined for this purpose and is shown in [Figure 3-3](#). The size of this region varies with the placement and DDR routing. Non-LPDDR4 signals should not be routed on the DDR signal layers within the LPDDR4 keepout region. Non-LPDDR4 signals may be routed in this region only if they are routed on other layers separated from the DDR signal layers by a ground layer. No breaks are allowed in the reference ground layers in this region. In addition, a solid VDDSD_DDR power plane should exist across the entire keepout region.

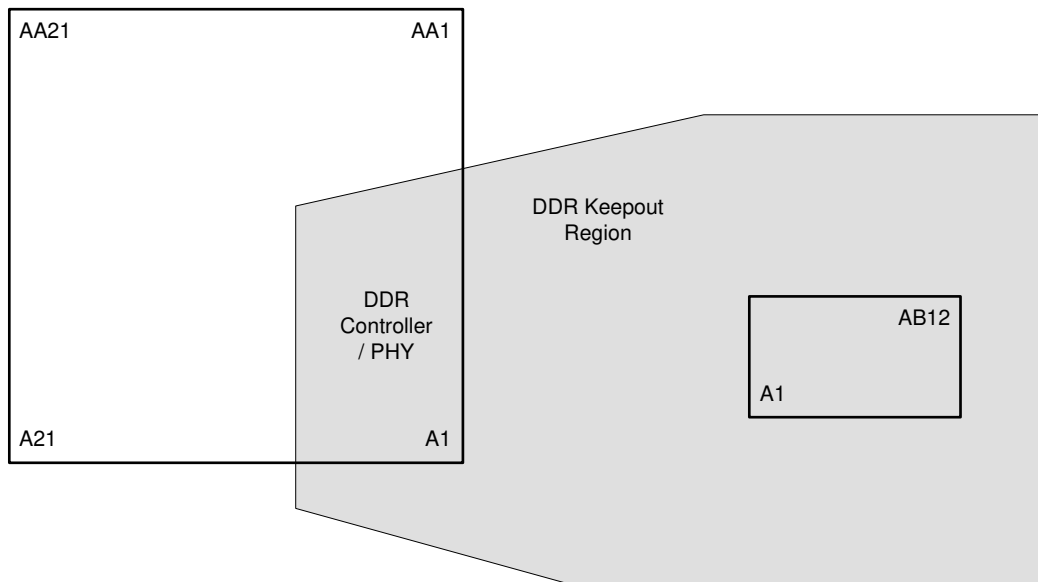


Figure 3-3. LPDDR4 Keepout Region

3.7 Net Classes

Routing rules are applied to signals in groups called net classes. Each net class contains signals with the same routing requirements. This simplifies the implementation and compliance of these routes. [Table 3-4](#) lists the clock net classes for the LPDDR4 interface. [Table 3-5](#) lists the signal net classes, and associated clock net classes, for signals in the LPDDR4 interface. These net classes are then linked to the termination and routing rules that follow.

Table 3-4. Clock Net Class Definitions

Clock Net Class	Processor Pin Names
CK	DDR0_CK0 / DDR0_CK0_n
DQS0	DDR0_DQS0 / DDR0_DQS0_n
DQS1	DDR0_DQS1 / DDR0_DQS1_n

Table 3-5. Signal Net Class Definitions

Signal Net Class	Associated Clock Net Class	Processor Pin Names
ADDR_CTRL	CK	DDR0_A[5:0], DDR0_CS0, DDR0_CKE0
BYTE0	DQS0	DDR0_DQ[7:0], DDR0_DM0
BYTE1	DQS1	DDR0_DQ[15:8], DDR0_DM1

3.8 LPDDR4 Signal Termination

LPDDR4 memories have software configurable on-die termination for the data group nets. The DDR subsystem also contains software configurable on-die termination for the address / control group nets. Thus, termination is not required on any DDR signals for an LPDDR4 configuration.

3.9 LPDDR4 VREF Routing

LPDDR4 memories generate their own VREFCA and VREFDQ internally for the address / command bus and data bus, respectively. Similarly, the DDR PHY also provides its own reference voltage for the data group nets during reads. Thus unlike DDR3 and DDR4, VREF does not need to be generated on the board, and there is no required VREF routing for an LPDDR4 configuration.

3.10 LPDDR4 VTT

Unlike DDR3 and DDR4, there is no required termination on the PCB of the address/control bus of an LPDDR4 configuration. All termination is handled internally (on-die). Thus, VTT does not apply for LPDDR4.

3.11 CK and ADDR_CTRL Topologies

The CK and ADDR_CTRL net classes are routed similarly, and are length matched from the DDR controller in the processor to the LPDDR4 SDRAM to minimize skew between the signals and ensure that the ADDR_CTRL signals are properly sampled at the SDRAM. The CK net class requires more care because it runs at a higher transition rate and is differential. The CK and ADDR_CTRL topologies are point-to-point.

Figure 3-4 shows the topology of the CK net class, and Figure 3-5 shows the topology for the corresponding ADDR_CTRL net classes. Length matching requirements for the routing segments are detailed in Table 3-6.




Figure 3-4. LPDDR4 CK Topology




Figure 3-5. LPDDR4 ADDR_CTRL Topology

Minimize layer transitions during routing. If a layer transition is necessary, it is preferable to transition to a layer using the same reference plane. If this cannot be accommodated, ensure there are nearby stitching vias to allow the return currents to transition between reference planes when both reference planes are ground or VDDS_DDR. Alternately, ensure there are nearby bypass capacitors to allow the return currents to transition between reference planes when one of the reference planes is ground and the other is VDDS_DDR. This must occur at every reference plane transition. The goal is to minimize the size of the return current path thus minimizing the inductance in this path. Lack of these stitching vias or capacitors results in impedance discontinuities in the signal path that increase crosstalk and signal distortion.

There are no stubs or terminations allowed on the nets of the CK and ADDR_CTRL routing group topologies. All test and probe access points must be in line without any branches or stubs.

3.12 Data Group Topologies

The data line topology is always point-to-point for LPDDR4 implementations, and is separated into two different byte routing groups. Minimize layer transitions during routing. If a layer transition is necessary, it is better to transition to a layer using the same reference plane. If this cannot be accommodated, ensure there are nearby ground vias to allow the return currents to transition between reference planes. The goal is to provide a low inductance path for the return current. To optimize the length matching, TI recommends routing all nets within a single data routing group on one layer where all nets have the exact same number of vias and the same via barrel length.

DQSP and DQSN lines are point-to-point signals routed as a differential pair. Figure 3-6 illustrates the DQSP/N connection topology.




Figure 3-6. LPDDR4 DQS Topology

DQ and DM lines are point-to-point signals routed as single-ended. [Figure 3-7](#) illustrates the DQ and DM connection topology.




Figure 3-7. LPDDR4 DQ/DM Topology

There are no stubs or termination allowed on the nets of the data group topologies. All test and probe access points must be in line without any branches or stubs.

3.13 CK and ADDR_CTRL Routing Specification

Skew within the CK and ADDR_CTRL net classes directly reduces setup and hold margin for the ADDR_CTRL nets. Thus, this skew must be controlled. The routed PCB track has a delay proportional to its length. Thus, the delay skew must be managed through matching the lengths of the routed tracks within a defined group of signals. The only way to practically match lengths on a PCB is to lengthen the shorter traces up to the length of the longest net in the net class and its associated clock.

Table 3-6 lists the limits for the individual segments that comprise the routing from the processor to the SDRAM. These segment lengths coincide with the CK and ADDR_CTRL topology diagram shown previously in Figure 3-4 and Figure 3-5. By matching the length for the same segments of all signals in a routing group, the signal delay skews are controlled. Most PCB layout tools can be configured to generate reports to assist with this validation. If this cannot be generated automatically, this must be generated and verified manually.

Table 3-6. CK and ADDR_CTRL Routing Specifications

Number	Parameter	MIN	MAX	UNIT
LP4_ACRS1	Propagation delay of net class CK, RSAC1		500 ⁽¹⁾	ps
LP4_ACRS2	Propagation delay of net class ADDR_CTRL, RSAC2		500 ⁽¹⁾	ps
LP4_ACRS3	Skew within net class CK (DDR0_CK0 to DDR0_CK0_n Skew)		0.4	ps
LP4_ACRS4	Skew across net class ADDR_CTRL (RSAC2)		3	ps
LP4_ACRS5	Skew across ADDR_CTRL net class and associated CK clock net class (RSAC1 to RSAC2)		3	ps
LP4_ACRS6	Vias per trace		3 ⁽¹⁾	vias
LP4_ACRS7	Via count difference		1 ⁽²⁾	vias
LP4_ACRS8	Center-to-center CK to other LPDDR4 trace spacing ⁽³⁾	4w		
LP4_ACRS9	Center-to-center ADDR_CTRL to other LPDDR4 trace spacing ⁽³⁾	4w		
LP4_ACRS10	Center-to-center ADDR_CTRL to other ADDR_CTRL trace spacing ⁽³⁾	3w		
LP4_ACRS11	CK center-to-center spacing ^{(4) (5)}		See notes below	
LP4_ACRS12	CK spacing to other net ⁽³⁾	4w		

- (1) Max value is based upon conservative signal integrity approach. This value could be extended only if detailed signal integrity analysis of rise time and fall time confirms desired operation.
- (2) Via count difference may increase by 1 only if accurate 3-D modeling of the signal flight times – including accurately modeled signal propagation through vias – has been applied to ensure all segment skew maximums are not exceeded.
- (3) Center-to-center spacing is allowed to fall to minimum 2w for up to 500 mils of routed length (only near endpoints).
- (4) CK spacing set to ensure proper differential impedance.
- (5) The user must control the impedance so that inadvertent impedance mismatches are not created. Generally speaking, center-to center spacing should be either 2w or slightly larger than 2w to achieve a differential impedance equal to twice the single-ended impedance, Zo, on that layer.

3.14 Data Group Routing Specification

Skew within the Byte signal net class directly reduces the setup and hold margin for the DQ and DM nets. Thus as with the ADDR_CTRL signal net class and associated CK clock net class, this skew must be controlled. The routed PCB track has a delay proportional to its length. Thus, the length skew must be managed through matching the lengths of the routed tracks within a defined group of signals. The only way to practically match lengths on a PCB is to lengthen the shorter traces up to the length of the longest net in the net class and its associated clock.

Note

It is not required nor recommended to match the lengths across all byte lanes. Length matching is only required within each byte.

Table 3-7 contains the routing specifications for the Byte0 and Byte1 routing groups. Each signal net class and its associated clock net class is routed and matched independently.

Table 3-7. Data Group Routing Specifications

Number	Parameter	MIN	MAX	UNIT
LP4_DRS1	Propagation delay of net class DQSx (RSD1)		500	ps
LP4_DRS2	Propagation delay of net class BYTE _x (RSD2)		500	ps
LP4_DRS3	Skew within net class DQSx (DDR0_DQSx to DDR0_DQSx_n skew)		0.4	ps
LP4_DRS4	Skew across net class DQSx and BYTE _x (RSD1 to RSD2 skew) ^{(1) (2)}		2	ps
LP4_DRS5	Skew within net class BYTE _x (DQ/DM to DQ/DM skew) ⁽¹⁾		2	ps
LP4_DRS6	Vias Per Trace		2 ⁽⁴⁾	vias
LP4_DRS7	Via Count Difference		0 ⁽³⁾	vias
LP4_DRS8	RSD1 center-to-center spacing (between clock net class) ⁽⁵⁾	4w		
LP4_DRS9	RSD1 center-to-center spacing (within clock net class) ^{(6) (7)}	See notes below		
LP4_DRS10	RSD2 center-to-center spacing (between signal net class) ⁽⁵⁾	4w		
LP4_DRS11	RSD2 center-to-center spacing (within signal net class) ⁽⁵⁾	3w		

- (1) Length matching is only done within a byte. Length matching across bytes is neither required nor recommended.
- (2) Each DQS pair is length matched to its associated byte.
- (3) Via count difference may increase by 1 only if accurate 3-D modeling of the signal flight times – including accurately modeled signal propagation through vias – has been applied to ensure DQ_n skew and DQS_n to DQ_n skew maximums are not exceeded.
- (4) Max value is based upon conservative signal integrity approach. This value could be extended only if detailed signal integrity analysis of rise time and fall time confirms desired operation.
- (5) Center-to-center spacing is allowed to fall to minimum 2w for up to 500 mils of routed length (only near endpoints).
- (6) DQS pair spacing is set to ensure proper differential impedance.
- (7) The user must control the impedance so that inadvertent impedance mismatches are not created. Generally speaking, center-to-center spacing should be either 2w or slightly larger than 2w to achieve a differential impedance equal to twice the single-ended impedance, Z_o, on that layer.

3.15 Channel, Byte, and Bit Swapping

All signals, including data and address / control, must be routed 1 to 1 from the DDR controller to the LPDDR4 memory. Byte swapping across channels or within a channel is not allowed. Similarly, data bit swapping across byte lanes or within a byte is also not allowed.

4 Revision History

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

DATE	REVISION	NOTES
	*	Initial Release

Changes from January 1, 2021 to July 30, 2021 (from Revision * (January 2021) to Revision A (July 2021))
Page

• Added Velocity Compensation section.....	5
• Updated DDR4 Interface Schematics section.....	7
• Updated 32-Bit, Single-Rank DDR4 Implementation With ECC Using x16 SDRAMs image.....	7
• Updated 32-Bit, Single-Rank DDR4 Implementation With ECC Using x8 SDRAMs image.....	9
• Updated VREF Routing section.....	12
• Updated Data Group Topologies and Routing Guidance section.....	16
• Updated CK and ADDR_CTRL Routing Limits section.....	18
• Updated CK and ADDR_CTRL Routing Specifications table.....	18
• Updated Data Group Routing Limits section.....	20
• Updated Data Group Routing Specifications table.....	20

IMPORTANT NOTICE AND DISCLAIMER

TI PROVIDES TECHNICAL AND RELIABILITY DATA (INCLUDING DATASHEETS), DESIGN RESOURCES (INCLUDING REFERENCE DESIGNS), APPLICATION OR OTHER DESIGN ADVICE, WEB TOOLS, SAFETY INFORMATION, AND OTHER RESOURCES "AS IS" AND WITH ALL FAULTS, AND DISCLAIMS ALL WARRANTIES, EXPRESS AND IMPLIED, INCLUDING WITHOUT LIMITATION ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT OF THIRD PARTY INTELLECTUAL PROPERTY RIGHTS.

These resources are intended for skilled developers designing with TI products. You are solely responsible for (1) selecting the appropriate TI products for your application, (2) designing, validating and testing your application, and (3) ensuring your application meets applicable standards, and any other safety, security, or other requirements. These resources are subject to change without notice. TI grants you permission to use these resources only for development of an application that uses the TI products described in the resource. Other reproduction and display of these resources is prohibited. No license is granted to any other TI intellectual property right or to any third party intellectual property right. TI disclaims responsibility for, and you will fully indemnify TI and its representatives against, any claims, damages, costs, losses, and liabilities arising out of your use of these resources.

TI's products are provided subject to TI's Terms of Sale (<https://www.ti.com/legal/termsofsale.html>) or other applicable terms available either on [ti.com](https://www.ti.com) or provided in conjunction with such TI products. TI's provision of these resources does not expand or otherwise alter TI's applicable warranties or warranty disclaimers for TI products.

Mailing Address: Texas Instruments, Post Office Box 655303, Dallas, Texas 75265
Copyright © 2021, Texas Instruments Incorporated