Application Note

AM62x DDR Board Design and Layout Guidelines

ABSTRACT

The goal of this application report is to describe how to make the AM62x 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......3
1.1 Board Designs Supported......3
1.2 General Board Layout Guidelines......3
1.3 PCB Stack-Up......4
1.4 Bypass Capacitors......5
1.5 Velocity Compensation......6
2 DDR4 Board Design and Layout Guidance......7
2.1 DDR4 Introduction......7
2.2 DDR4 Device Implementations Supported......7
2.3 DDR4 Interface Schematics......9
2.4 Compatible JEDEC DDR4 Devices......13
2.5 Placement......13
2.6 DDR4 Keepout Region......14
2.7 DBI......14
2.8 VPP......14
2.9 Net Classes......14
2.10 DDR4 Signal Termination......15
2.11 VREF Routing......15
2.12 VTT......15
2.13 POD Interconnect......15
2.14 CK and ADDR_CTRL Topologies and Routing Guidance......15
2.15 Data Group Topologies and Routing Guidance......19
2.16 CK and ADDR_CTRL Routing Specification......20
2.17 Data Group Routing Specification......22
2.18 Bit Swapping......23
3 LPDDR4 Board Design and Layout Guidance......24
3.1 LPDDR4 Introduction......24
3.2 LPDDR4 Device Implementations Supported......24
3.3 LPDDR4 Interface Schematics......24
3.4 Compatible JEDEC LPDDR4 Devices......26
3.5 Placement......26
3.6 LPDDR4 Keepout Region......27
3.7 LPDDR4 DBI......27
3.8 Net Classes......27
3.9 LPDDR4 Signal Termination......27
3.10 LPDDR4 VREF Routing......28
3.11 LPDDR4 VTT......28
3.12 CK0 and ADDR_CTRL Topologies......28
3.13 Data Group Topologies......28
3.14 CK0 and ADDR_CTRL Routing Specification......29
3.15 Data Group Routing Specification......30
3.16 Byte and Bit Swapping......31
4 LPDDR4 Board Design Simulations......32
4.1 Board Model Extraction......32

Copyright © 2024 Texas Instruments Incorporated

4.2 Board-Model Validation.....32
4.3 S-Parameter Inspection.....32
4.4 Time Domain Reflectometry (TDR) Analysis.....33
4.5 System Level Simulation.....33
4.6 Design Example.....40
5 Appendix: AM62x ALW and AMC Package Delays.....45
6 Revision History.....46

Copyright © 2024 Texas Instruments Incorporated

Trademarks

All trademarks are the property of their respective owners.

1 Overview

The AM62x 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 has generic information that is applicable to both DDR4 and LPDDR4, as well as separate sections that are specific to each supported DDR memory type.

Note

To facilitate software configuration of the DDRSS, use the DDR 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:

- Always follow TI's example layouts/EVM designs as close as possible. If concepts or routing strategies are not understood, questions can be posted on E2E.
- All signals need ground reference (strongly suggest on both sides).
- Avoid crossing plane splits in the signal reference planes.
- Use the widest trace that is practical between decoupling capacitors and memory modules.
- Minimize inter-symbol interference (ISI) by keeping impedances matched. This is especially true for any 'T-branch' signals where trace widths are adjusted to match trace impedance.
- 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 (also called VSS) reference for all signals and 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 significant part of PCB-level crosstalk. Dimension and pitch of vias is important. For high speed interfaces, consider GND shielding vias. This via coupling is one factor for recommending data signals be routed on layers closest to processor.
- Via stubs affect signal integrity. Via back-drilling can improve signal integrity, and may be required in some instances.

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

Copyright © 2024 Texas Instruments Incorporated

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 VDDS_DDR reference plane. If a VDDS_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 VDDS_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 singled-ended impedance selected for the PCB specified by PS9 and PS10.

Copyright © 2024 Texas Instruments Incorporated

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.
- Using three-terminal capacitors instead of two terminal capacitors. Three-terminal capacitors provide lower loop inductance, and one three-terminal capacitor could take the place of multiple two-terminal capacitors, further optimizing loop inductance.

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

Table 1-3. High-Speed Bypass Capacitors

Parameter	MIN	TYP	MAX	UNIT
HS bypass capacitor package size (1)		0201	0402	Mils
Distance, HS bypass capacitor to processor being bypassed (2) (3) (4)			150	Mils
Processor HS bypass capacitor count and total capacitance per VDDS_DDR rail(5)	see notes below			
Number of connection vias for each device power/ground ball	1			Vias
Trace length from processor power/ground ball to connection via (2)		35	70	Mils
Distance, HS bypass capacitor to DDR device being bypassed (6)			150	Mils
DDR device HS bypass capacitor count	Refer to DDR manufacturer guidelines			
Number of connection vias for each HS capacitor (7) (8)	2			Vias
Trace length from bypass capacitor to connection via (2) (8)		35	100	Mils
Number of connection vias for each DDR device power/ground ball	1			Vias
Trace length from DDR device power/ground ball to connection via (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 VDDS_DDR balls.
(5) Decoupling capacitor counts and/or capacitor values should be derived from power aware PCB simulations. It is the responsibility of the PCB designer to ensure that any design meets the provided PDN targets.

Copyright © 2024 Texas Instruments Incorporated

(6) Measured from the DDR device power or ground ball to the center of the capacitor package. 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.

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 VDDS_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.

Copyright © 2024 Texas Instruments Incorporated

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 datasheet and the DDR Subsystem (DDRSS) chapter in the AM62x Technical Reference Manual for lists of features and not 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

SDRAM bus width	Number of SDRAM devices	Number of dies in package	BG signals	Number of ranks (CS signals)	Schematic	Max Addressable range
x16	1	1	BG0	1 (CS0_n)	Figure 2-1	4GBytes
x16	1	2	BG[1:0]	1 (CS0n) or 2 (CS[1:0]_n)	Figure 2-2 (VTT optional)	4GBytes (1 rank), 8GBytes (2 ranks)
x8	2	1	BG[1:0]	1 (CS0n) or 2 (CS[1:0]_n)	Figure 2-2 (with VTT)	4GBytes (1 rank), 8GBytes (2 ranks)

Copyright © 2024 Texas Instruments Incorporated

Note

ECC is supported on the DDR4 interface. Unlike traditional ECC interfaces which require dedicated memory pins and devices, ECC is supported inline. The ECC system impact is in interface bandwidth and overall memory density, as ECC data is stored alongside non-ECC data. Max addressable range will be reduced if ECC is enabled. See device TRM for more details.

Copyright © 2024 Texas Instruments Incorporated

2.3 DDR4 Interface Schematics

This section discusses implementations (also called topologies) using 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.

Only implementations with 16-bit data width are supported. Implementations with only 8-bits of data are not supported.

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.

Copyright © 2024 Texas Instruments Incorporated

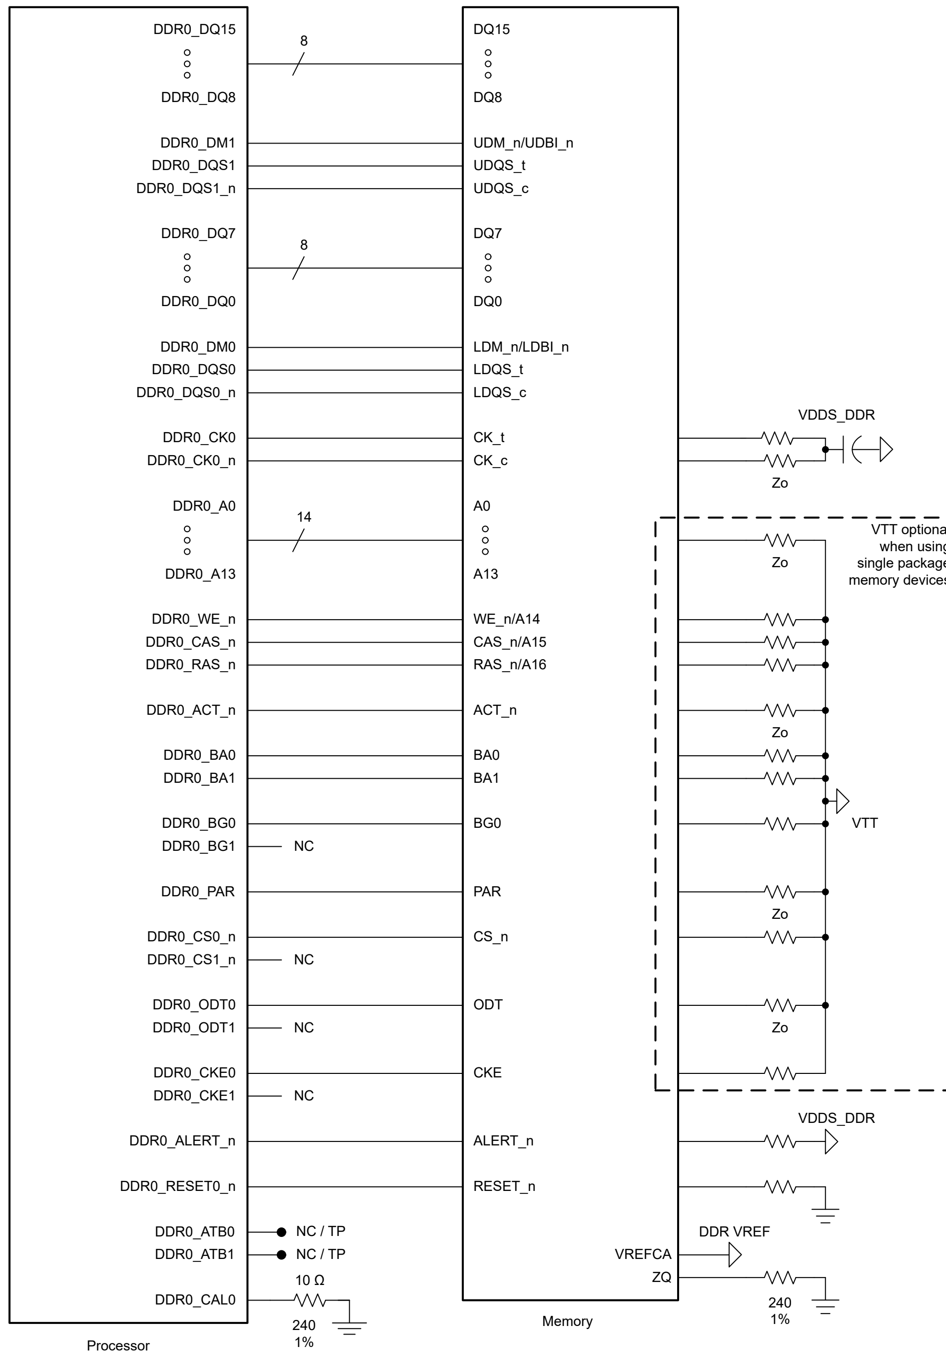

1. When designing with VTT regulator (LDO) which can source and sink current, decoupling capacitors (minimum of one capacitor (1.0uF value) must be used for every two termination resistors) should be used to minimize the effect of VTT supply noise. Refer to the AM64x GP EVM for reference
2. Zo value for resistors is 30-47ohm. Resistor value should closely match trace impedance.
3. VTT is optional on address/control signals when using single package memory devices, but termination as shown on CK0/CK0_n is always required
4. DDR_VREF is supplied by the VTT regulator. When VTT is not used, VREFCA needs to be connected to a voltage divider. Consult the EVM schematic for an example of the voltage divider implementation.

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

Copyright © 2024 Texas Instruments Incorporated

2.3.2 DDR4 Implementation Using 8-Bit SDRAM Devices

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

It is also possible to use memory devices with two dies in one package to achieve higher density designs. The connection would look similar to Figure 2-2, except the two x8 memories are in one package, and thus VTT implementation is not necessary. Ensure to connect BG[1:0] when using these devices, as the internal implementation is still using two x8 memories.

Copyright © 2024 Texas Instruments Incorporated

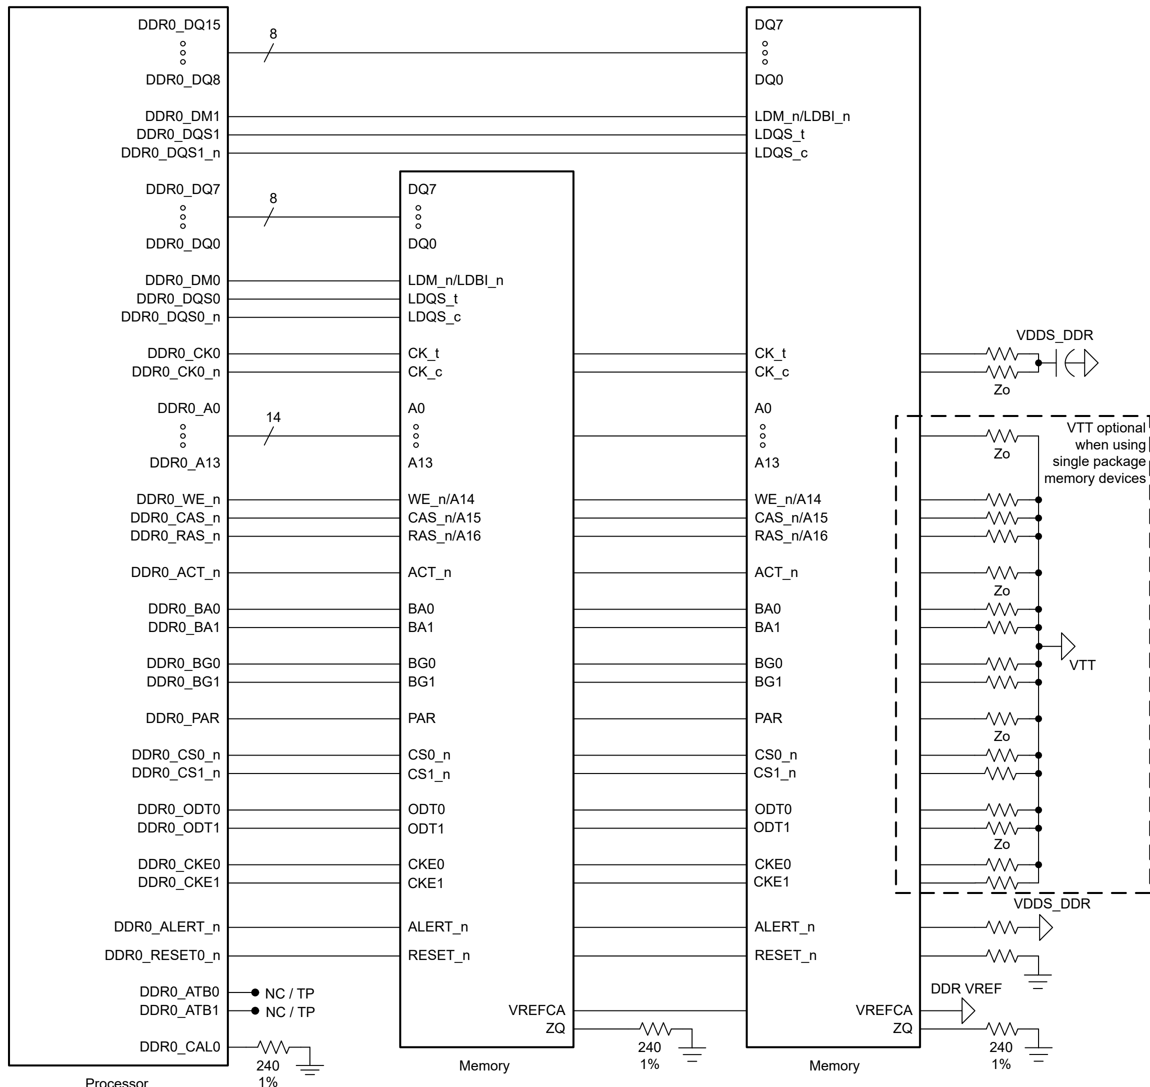

1. When designing with VTT regulator (LDO) which can source and sink current, decoupling capacitors (minimum of one capacitor (1.0uF value) must be used for every two termination resistors) should be used to minimize the effect of VTT supply noise. Refer to the AM64x GP EVM for reference.
2. Zo value for resistors is 30-47ohm. Resistor value should closely match trace impedances.
3. VTT is optional on address/control signals when using single package memory devices, but termination as shown on CK0/CK0_n is always required.
4. DDR_VREF is supplied by the VTT regulator. When VTT is not used, VREFCA needs to be connected to a voltage divider. Consult the EVM schematic for an example of the voltage divider implementation.
5. For Single-Rank designs, CS1_n, ODT1, and CKE1 can be left unconnected
6. Single package memories with two x8 dies do not require VTT.

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

Copyright © 2024 Texas Instruments Incorporated

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 DDR4 Interface Schematics.
(2) For supported data rates, see the device-specific data manual.
(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.

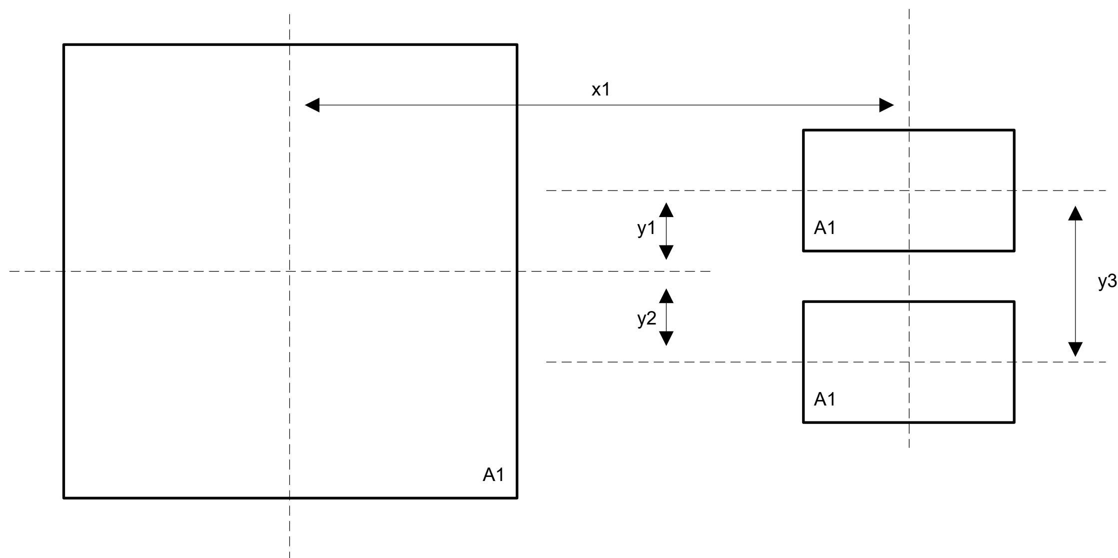

Figure 2-3. DDR4 Placement Specifications

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

Copyright © 2024 Texas Instruments Incorporated

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 VDDS_DDR power plane should exist across the entire keepout region.

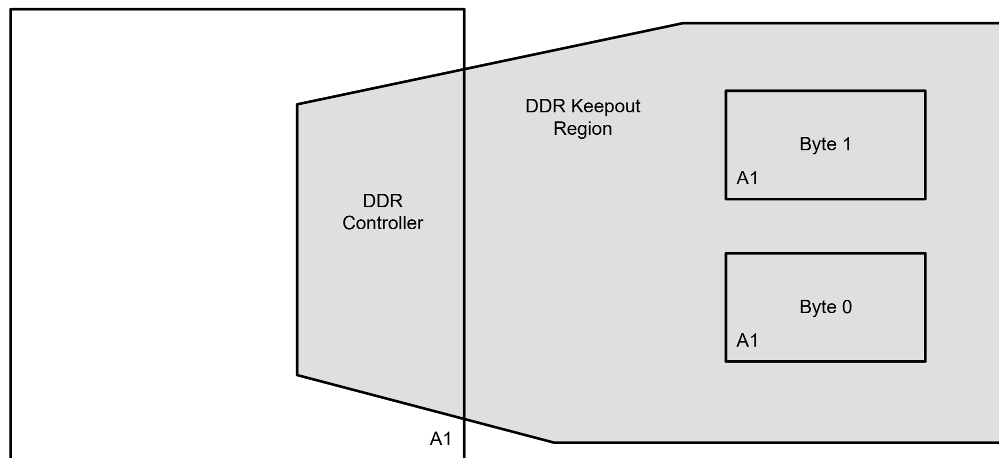

Figure 2-4. DDR4 Keepout Region

2.7 DBI

DBI is recommended to reduce supply/ground noise and to improve the data eye. Therefore the DDR Subsystem Configuration Tool enables Read DBI as a default for DDR4 configurations.

2.8 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.

2.9 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

Copyright © 2024 Texas Instruments Incorporated

2.10 DDR4 Signal Termination

Signal terminators are required on designs with multiple memory devices for the CK and ADDR_CTRL net classes. For designs with a single memory package, VTT termination is optional on the ADDR_CTRL net class (termination is still required for the differential clock signals CK0 and CK0_n). This is shown in the schematic figures in DDR4 Interface Schematics. 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.

2.11 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-μF bypass capacitors near each device connection. Narrowing the VREF trace is allowed to accommodate routing congestion for short lengths near endpoints.

When a VTT power supply is not used, VREFCA should be generated using a voltage divider circuit. Consult the EVM schematics for examples of how the voltage divider circuit is implemented. Ensure high precision resistors (1% tolerance) are using for the voltage divider.

2.12 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 with decoupling capacitors near the terminator resistors. VTT is optional if the design contains only one memory device (ie, all signals are point to point).

2.13 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.

2.14 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.

Copyright © 2024 Texas Instruments Incorporated

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.

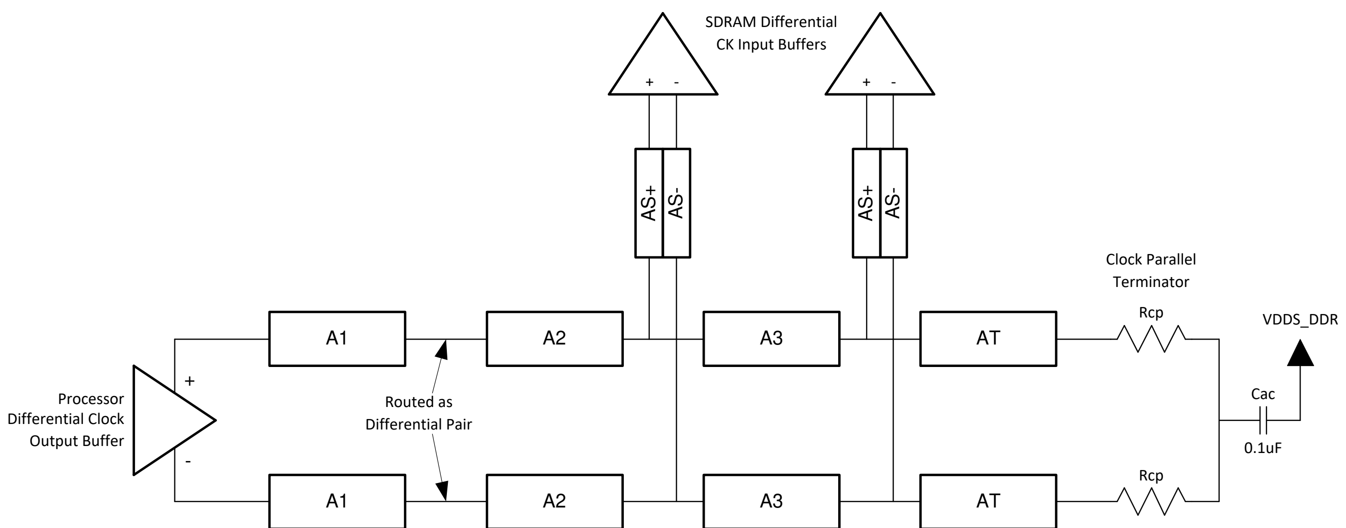

Figure 2-5. CK Topology for Two DDR4 SDRAM Devices

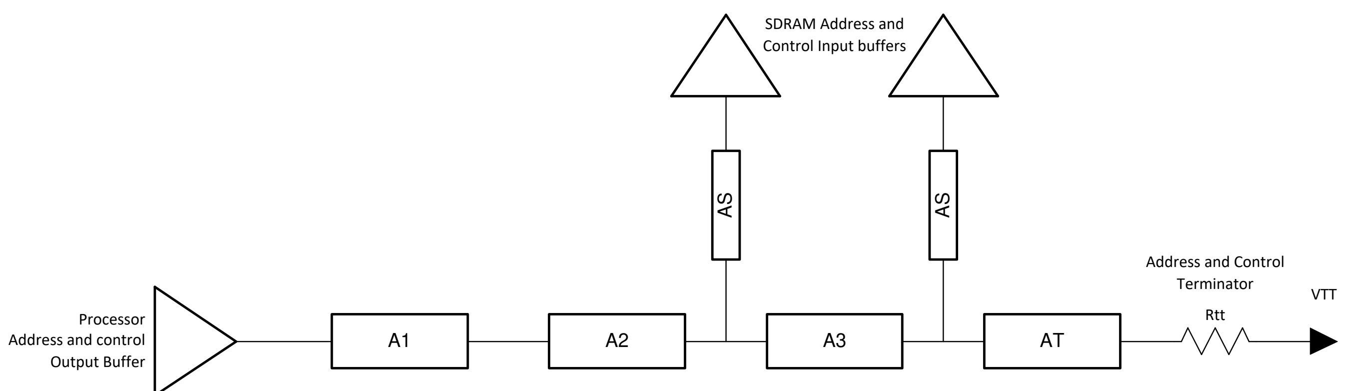

Figure 2-6. ADDR_CTRL Topology for Two DDR4 SDRAM Devices

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.

Copyright © 2024 Texas Instruments Incorporated

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 VDDS_DDR. The routing also includes the routing stubs for both DDR0_CK0 and DDR0_CK0_n at each SDRAM.

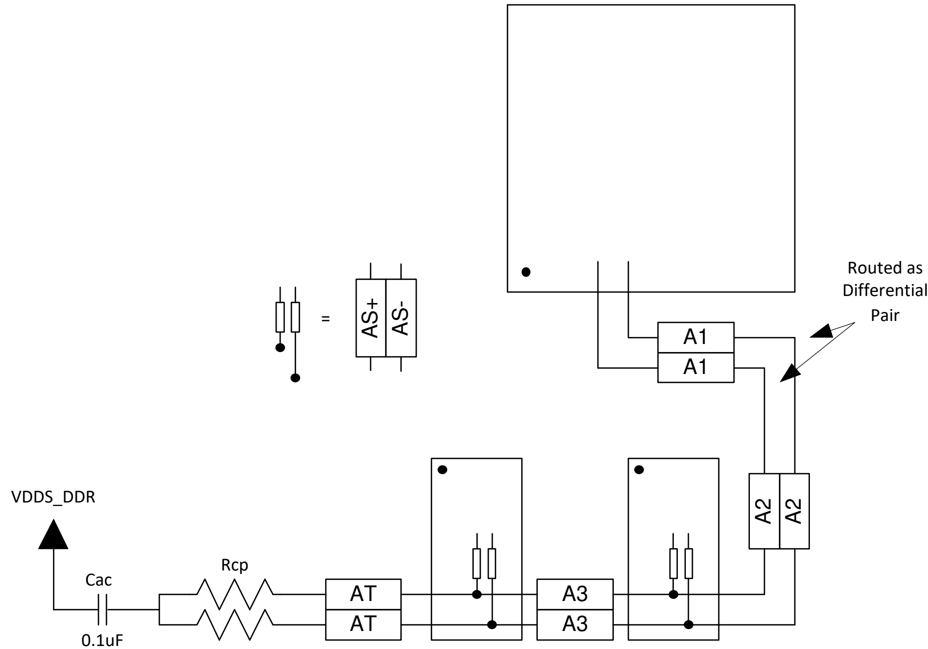

Figure 2-7. CK Routing for Two DDR4 SDRAM Devices

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.

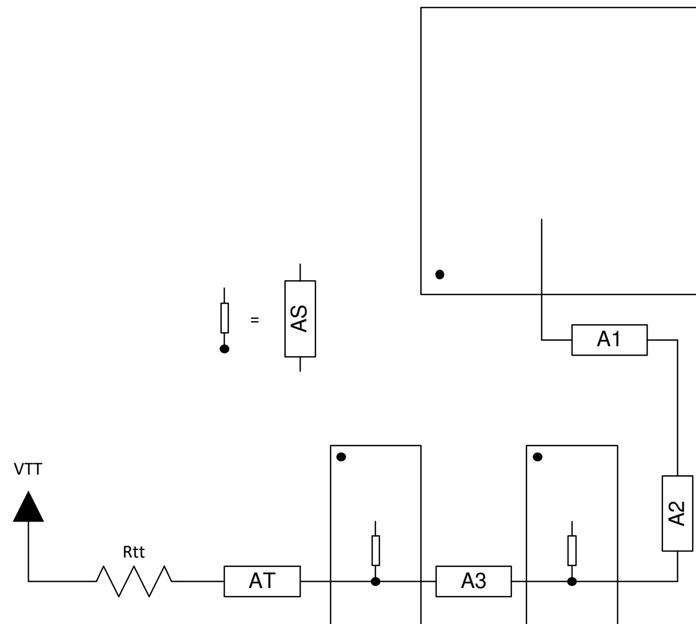

Figure 2-8. ADDR_CTRL Routing for Two DDR4 SDRAM Devices

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

Copyright © 2024 Texas Instruments Incorporated

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.

Copyright © 2024 Texas Instruments Incorporated

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.

2.15 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.

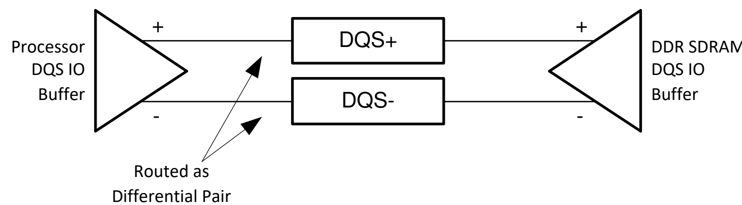

Figure 2-9. DDR4 DQS Topology

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

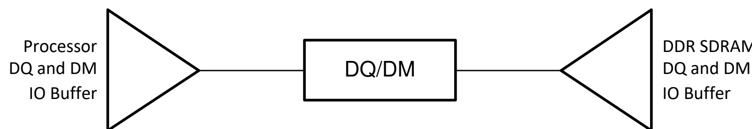

Figure 2-10. DDR4 DQ/DM Topology

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.

Copyright © 2024 Texas Instruments Incorporated

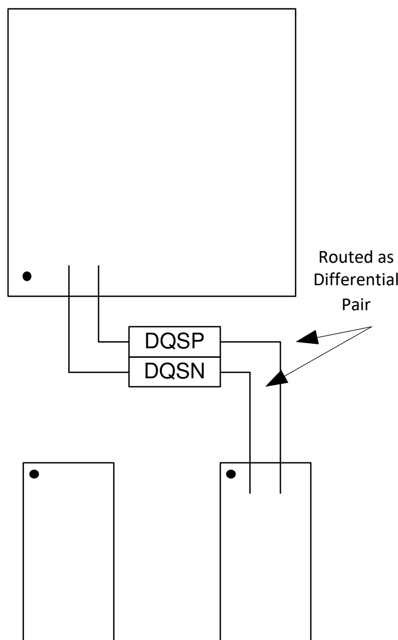

Figure 2-11. DQS Routing to Two DDR4 SDRAM Devices

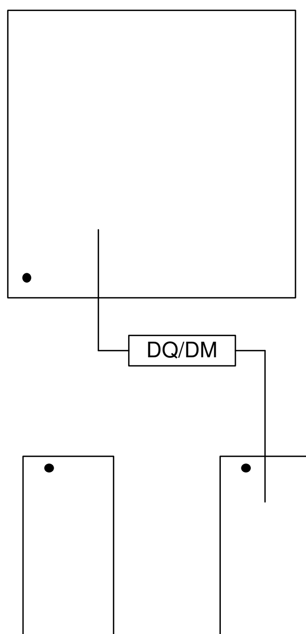

Figure 2-12. DQ/DM Routing to Two DDR4 SDRAM Devices

2.16 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. Consider Z-axis delays (VIAs) with accurate stackup information during analysis.

Copyright © 2024 Texas Instruments Incorporated

2.16.1 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.

2.16.2 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

Copyright © 2024 Texas Instruments Incorporated

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

Number	Parameter	MIN	TYP	MAX	UNIT
17	Via count difference(11)			1(10)	vias
18	Center-to-center CK to other DDR4 trace spacing(5)	4w			
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.

2.17 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. Consider Z-axis delays (VIAs) with accurate stackup information during analysis.

2.17.1 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 and 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.

Copyright © 2024 Texas Instruments Incorporated

2.17.2 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, Zo, 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.

2.18 Bit Swapping

2.18.1 Data Bit Swapping

Data bit swapping is allowed to simplify routing as long as the DQ bits swapped are within the same byte group. This is only possible when not using CRC. Any DQ bits within a byte group can be swapped. The DM and DQS bits must not be swapped with any other signals. Data byte swapping is allowed, as long as all of the associated signals within a byte (DQx, DQSx, and DM) are swapped together. Software configuration changes in the DDR Configuration Tool (*https://dev.ti.com/sysconfig*) are not necessary for normal device functionality when swapping data signals with DDR4

2.18.2 Address and Control Bit Swapping

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

Copyright © 2024 Texas Instruments Incorporated

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. Refer to the datasheet and the DDR Subsystem (DDRSS) chapter in the AM62x Technical Reference Manual for lists of features and not supported features.

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(CS signals)	LPDDR4 Channel Width	DDRSS data width	Schematic	Max Addressable Range
1	1	1	1(CS0_n)	16	16	Figure 3-1	2GBytes

Note

ECC is supported on the LPDDR4 interface. Unlike traditional ECC interfaces which require dedicated memory pins and devices, ECC is supported inline. The ECC system impact is in interface bandwidth and overall memory density, as ECC data is stored alongside non-ECC data. Max addressable range will be reduced if ECC is enabled. See device TRM for more details.

Note

Data bus routing must be point to point between the processor and the memory, and cannot be split on the board. Thus, dual-rank LPDDR4 designs are only possible when using one channel of an LPDDR4 dual channel, dual rank device. If more than 2GBytes is needed, consider using DDR4.

3.3 LPDDR4 Interface Schematics

As stated above, LPDDR4 supports many different implementation topologies but the device 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.

Copyright © 2024 Texas Instruments Incorporated

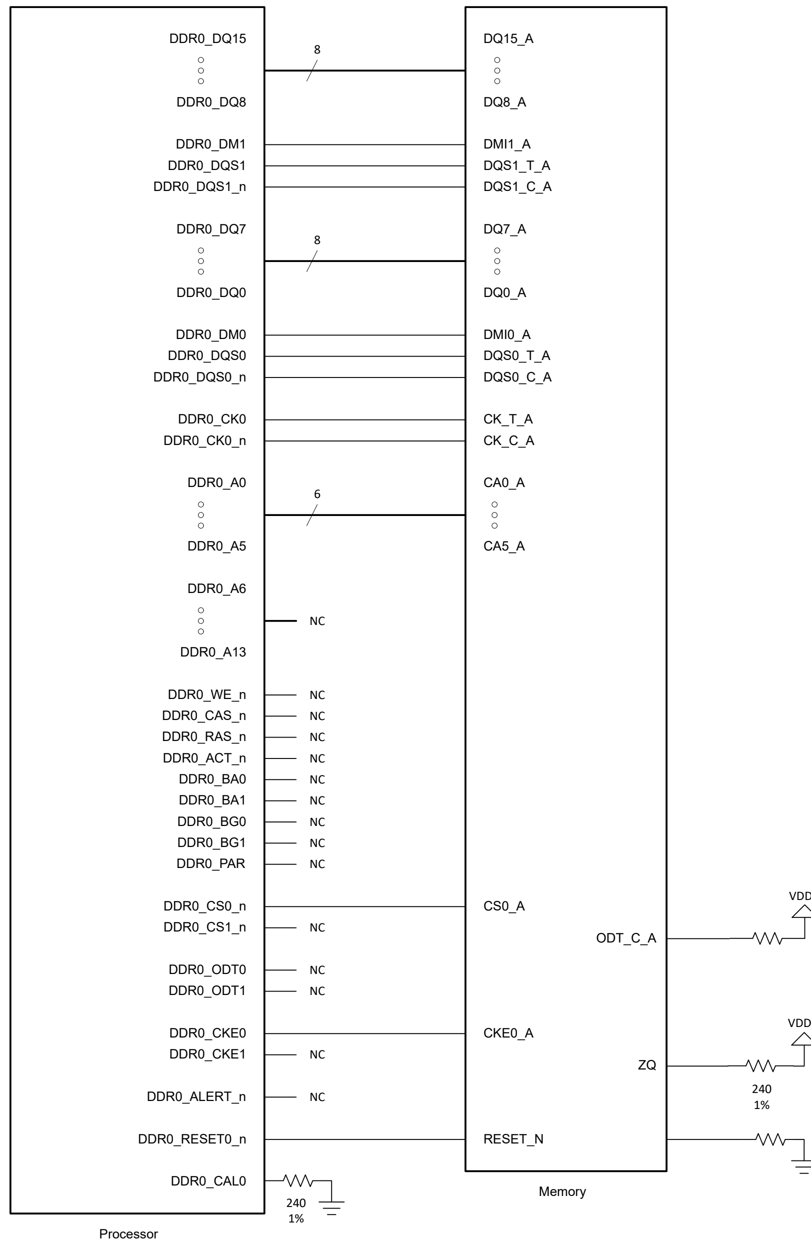

Figure 3-1. 16-Bit, Single-Rank, Single Channel LPDDR4 Implementation

Copyright © 2024 Texas Instruments Incorporated

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.

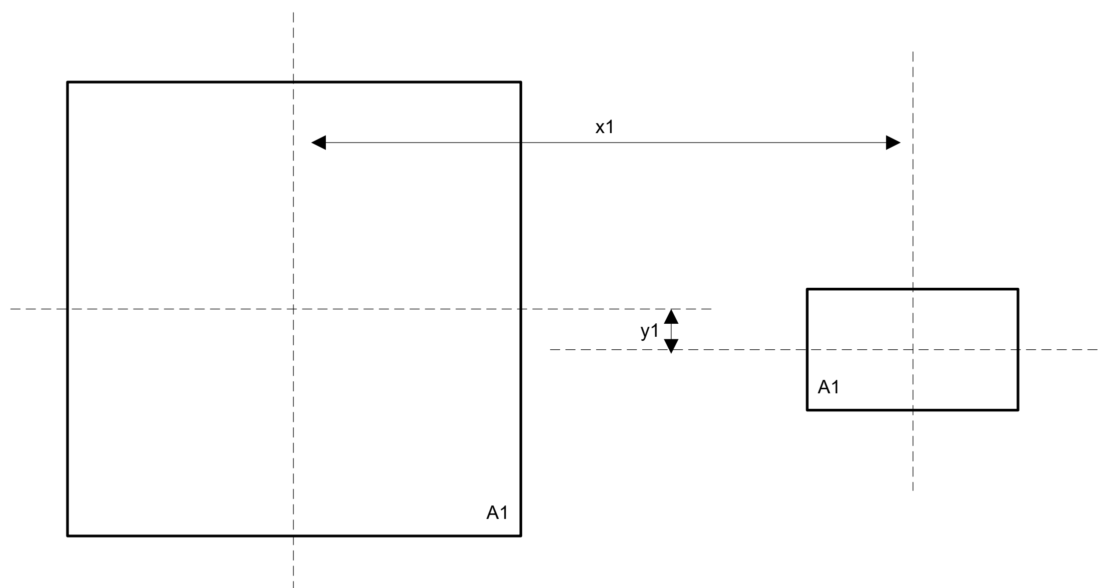

Figure 3-2. LPDDR4 Placement Specification

Table 3-3. LPDDR4 Placement Parameters

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

Copyright © 2024 Texas Instruments Incorporated

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 VDDS_DDR power plane should exist across the entire keepout region.

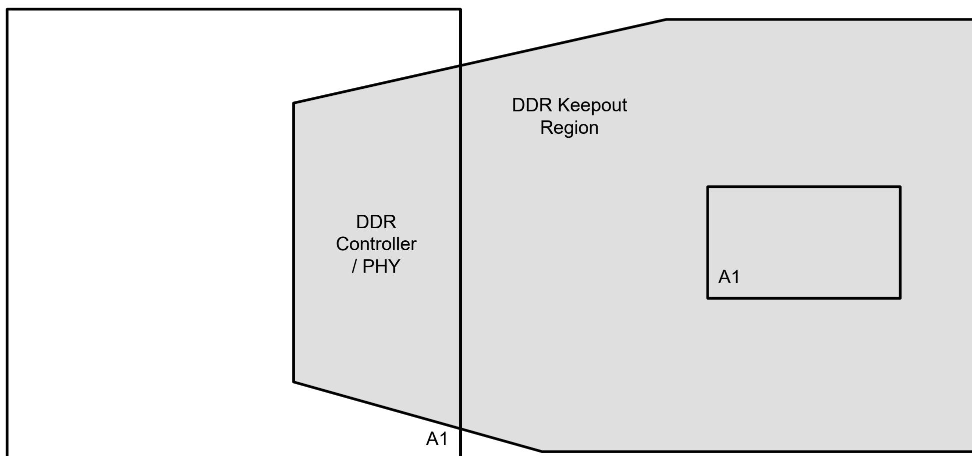

Figure 3-3. LPDDR4 Keepout Region

3.7 LPDDR4 DBI

DBI is recommended to reduce supply/ground noise and to improve the data eye. Therefore the DDR Subsystem Configuration Tool enables Write DBI as a default for LPDDR4 configurations.

3.8 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 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
CK0	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	CK0	DDR0_A[5:0], DDR0_CS0_n, DDR0_CS1_n, DDR0_CKE0, DDR0_CKE1
BYTE0	DQS0	DDR0_DQ[7:0], DDR0_DM0
BYTE1	DQS1	DDR0_DQ[15:8], DDR0_DM1

3.9 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.

Copyright © 2024 Texas Instruments Incorporated

3.10 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.

3.11 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.

3.12 CK0 and ADDR_CTRL Topologies

The CK0 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 CK0 net class requires more care because it runs at a higher transition rate and is differential. The CK0 and ADDR_CTRL topologies are point-to-point.

Figure 3-4 shows the topology of the CK0 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.

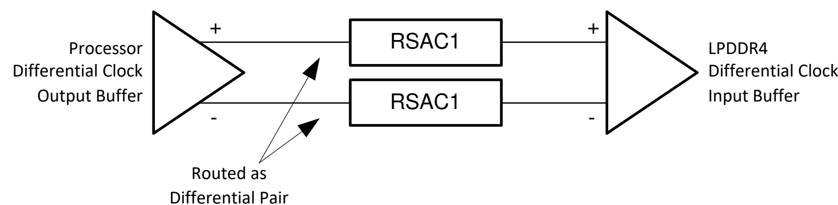

Figure 3-4. LPDDR4 CK0 Topology

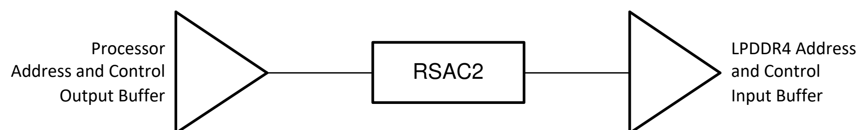

Figure 3-5. LPDDR4 ADDR_CTRL Topology

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 CK0 and ADDR_CTRL routing group topologies. All test and probe access points must be in line without any branches or stubs.

3.13 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 (within ± 250 mils of transition vias). 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 (ie. DQS/DQ/DM) together on the same layer(s) where all nets have the exact same number of vias and the same via barrel length. Microstrip routing can be used to implement DDR routing, but doing so provides lower EMI immunity and signal integrity at high data rates. The designer should evaluate system requirements carefully to determine that the desired product

Copyright © 2024 Texas Instruments Incorporated

requirements can be met. High-speed DQ and DQS/DQSn routing on microstrip layers requires special care and DFM consideration because of more variation in signal propagation. Signals from the entire byte group must be routed together.

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

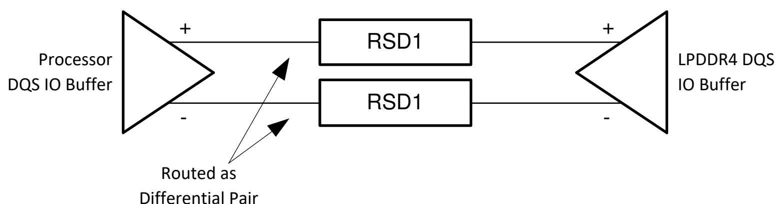

Figure 3-6. LPDDR4 DQS Topology

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

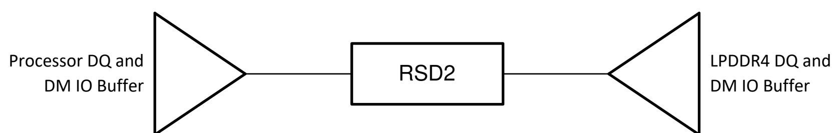

Figure 3-7. LPDDR4 DQ/DM Topology

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.

3.14 CK0 and ADDR_CTRL Routing Specification

Skew within the CK0 and ADDR_CTRL net classes directly reduces setup and hold margin for the ADDR_CTRL nets. Thus, this skew must be controlled. Per-bit deskew capability within the PHY substantially loosens the skew tolerance requirements. The skew budgets in Table 3-6 include total delay from SoC die pad to DRAM pin. (i.e. delay of SOC package + PCB). Package delays are provided in Appendix: AM62x ALW and AMC Package Delays. The designer is free to length match using smaller tolerance than the values shown in the table. 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. Consider Z-axis delays (VIAs) with accurate stackup information during analysis.

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 CK0 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. CK0 and ADDR_CTRL Routing Specifications

Number	Parameter	MIN	TYP	MAX	UNIT
LP4_ACRS1	Propagation delay of net class CK0 (RSAC1)			450(1)	ps
LP4_ACRS2	Propagation delay of net class ADDR_CTRL (RSAC2)			450(1)	ps
LP4_ACRS3	Skew within net class CK0 (Skew of DDR0_CK0 and DDR0_CK0_n) (RSAC1)			0.75(2) (3)	ps
LP4_ACRS6	Skew across ADDR_CTRL and CK0 clock net classes, relative to propagation delay of CK0 net class (RSAC1 - RSAC2)(4)	-312.5(3) (5)		312.5(3) (5)	ps
LP4_ACRS7	Vias per trace			3(1)	vias

Copyright © 2024 Texas Instruments Incorporated

Table 3-6. CK0 and ADDR_CTRL Routing Specifications (continued)

Number	Parameter	MIN	TYP	MAX	UNIT
LP4_ACRS8	VIA Stub Length		20		Mils
LP4_ACRS9	Via count difference			1[(6)]	vias
LP4_ACRS10	Center-to-center CK0 to other LPDDR4 trace spacing	5w[(7)]			
LP4_ACRS11	Center-to-center ADDR_CTRL to other LPDDR4 trace spacing	5w[(7)]			
LP4_ACRS12	Center-to-center ADDR_CTRL to self or other ADDR_CTRL trace spacing	3w[(7)]			
LP4_ACRS13	CK0 center-to-center spacing[(8)]		See note below		
LP4_ACRS14	CK0 spacing to non-DDR net	5w[(7)]			

(1) Max value is based upon conservative signal integrity approach. FR4 material assumed with Dk ~ 3.7 - 3.9 & Df ~ 0.002. This value could be extended only if detailed signal integrity analysis of rise time and fall time confirms desired operation.
(2) Recommendation for PCB layout tool design. Required to be verified by simulation[(9)], confirm JEDEC defined Vix_DQS_ratio (20%) and Vix_CK_ratio (25%) are satisfied, also confirm good eye margins.
(3) Consider the delays from SOC die pad to the DRAM pin (ie. delays of SOC package + delays of PCB upto the DRAM pin. DRAM package delays are omitted). Refer to Appendix: AM62x ALW and AMC Package Delays.
(4) Recommend routing net classes CK0 and ADDR_CTRL on same signal layer for better skew control.
(5) Simulation[(9)] must be performed and the delay report analyzed to ensure delays are within the limit. Delay reports from PCB layout tools use a simplified calculation based on a constant propagation velocity factor. TI recommends initially delay matching in PCB layout tool to a target less than 20% of the limit.
(6) 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.
(7) Center-to-center spacing is allowed to fall to minimum 2w for up to 500 mils of routed length (only near endpoints). Spacing minimums may be relaxed if simulations[(9)] accurately capture crosstalk between neighboring victim and aggressor traces and show good margin. Consider also VIA spacing. Signals with adjacent VIAs near SOC should not also have adjacent VIAs near the DRAM.
(8) P to N spacing set to ensure proper differential impedance. The designer 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. Refer to impedance targets in Table 1-1
(9) Simulation refers to a power-aware IBIS Signal Integrity (SI) simulation. Simulate across process, voltage, and temperature (PVT). Refer to LPDDR4 Board Design Simulations

3.15 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 CK0 clock net class, this skew must be controlled. Per-bit deskew capability within the PHY substantially loosens the skew tolerance requirements. The skew budgets in Table 3-7 include total delay from SoC die pad to DRAM pin. (i.e. delay of SOC package + PCB). Package delays are provided in Appendix: AM62x ALW and AMC Package Delays. The designer is free to length match using smaller tolerance than the values shown in the table. 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. Consider Z-axis delays (VIAs) with accurate stackup information during analysis.

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	TYP	MAX	UNIT
LP4_DRS1	Propagation delay of net class DQSx (RSD1)			450[(1)]	ps
LP4_DRS2	Propagation delay of net class BYTEx (RSD2)			450[(1)]	ps
LP4_DRS3	Difference in propagation delays of CK0 pair and each DQS pair. (RSAC1 - RSD1) [(2)]	0[(3)] [(4)]		3[(3)] [(4)]	tCK
LP4_DRS4	Skew within net class DQSx. Skew of DDR0_DQSx and DDR0_DQSx_n (RSD1)			1.5[(4)] [(6)]	ps
LP4_DRS5	Skew across DQSx and BYTEx net classes. (Skew of RSD1 and RSD2) [(7)]			150[(3)] [(4)]	ps

Copyright © 2024 Texas Instruments Incorporated

Table 3-7. Data Group Routing Specifications (continued)

Number	Parameter	MIN	TYP	MAX	UNIT
LP4_DRS6	Difference in propagation delays of shortest DQ/DM bit in BYTEx and respective DQSx. (RSD2 - RSD1)[8]	-49[3] [4] [5]			ps
LP4_DRS7	Vias Per Trace			2[1]	vias
LP4_DRS8	VIA Stub Length		40		Mils
LP4_DRS9	Via Count Difference			0[9]	vias
LP4_DRS10	RSD1 center-to-center spacing (between different clock net classes)	5w[10]			
LP4_DRS11	RSD1 center-to-center spacing (within clock net class)[11]	See note below			
LP4_DRS12	RSD2 center-to-center spacing (between different signal net classes/ bytes)	5w[10]			
LP4_DRS13	RSD2 center-to-center spacing (to self or within signal net class)	3w[10]			

(1) Max value is based upon conservative signal integrity approach. FR4 material assumed with Dk ~ 3.7 - 3.9 & Df ~ 0.002. This value could be extended only if detailed signal integrity analysis of rise time and fall time confirms desired operation.
(2) Propagation delay of CK0 pair must be greater than propagation delay of each DQS pair.
(3) Simulation[12] must be performed and the delay report analyzed to ensure delays are within the limit. Delay reports from PCB layout tools use a simplified calculation based on a constant propagation velocity factor. TI recommends initially delay matching in PCB layout tool to a target less than 20% of the limit.
(4) Consider the delays from SOC die pad to the DRAM pin (ie. delays of SOC package + delays of PCB upto the DRAM pin. DRAM package delays are omitted). Refer to Appendix: AM62x ALW and AMC Package Delays.
(5) Recommend that the propagation delay of DQS is shorter than all DQx within a byte. If that is not possible, LP4_DRS6 specifies that a DQ can be shorter by at most 49ps
(6) Recommendation for PCB layout tool design. Required to be verified by simulation[12], confirm JEDEC defined Vix_DQS_ratio (20%) and Vix_CK_ratio (25%) are satisfied, also confirm good eye margins.
(7) Skew matching is only done within a byte including DQS. Skew matching across bytes is neither required nor recommended.
(8) Propagation delay of the shortest DQ/DM bit in BYTEx Signal Net Class is recommended to be greater than the the propagation delay of its respective DQSx.
(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 skew maximums are not exceeded.
(10) Center-to-center spacing is allowed to fall to minimum 2w for up to 500 mils of routed length (only near endpoints). Spacing minimums may be relaxed if simulations[12] accurately capture crosstalk between neighboring victim and aggressor traces and show good margin. Consider also VIA spacing. Signals with adjacent VIAs near SOC should not also have adjacent VIAs near the DRAM.
(11) DQS pair spacing is set to ensure proper differential impedance. P to N spacing set to ensure proper differential impedance. The designer 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. Refer to impedance targets in Section 1.3.
(12) Simulation refers to a power-aware IBIS Signal Integrity (SI) simulation. Simulate across process, voltage, and temperature (PVT).

3.16 Byte and Bit Swapping

All address/control signals must be routed from the DDR controller to the LPDDR4 memory as described in the diagrams in LPDDR4 Interface Schematics. Address/control signals cannot be swapped with other signals. Data bit (DQx) and Data Mask (DM) swapping within a byte (for example, swapping D2 with D3) is allowed, but data bit DQx/DM swapping across bytes (for example, swapping D4 and D13) is not allowed.

Swapping byte lanes within a channel (for example, swapping byte 0 and 1) is allowed.When swapping bytes, all of the associated signals of the byte (DQx, DQSx, and DM) must be swapped together.

Use the DDR Subsystem Register Configuration Tool in SysConfig (*https://dev.ti.com/sysconfig*) to describe how the bits are swapped. Check the README link in the tool for detailed instructions.

Copyright © 2024 Texas Instruments Incorporated

4 LPDDR4 Board Design Simulations

This section is intended to provide an overview of the basic system-level board extraction, simulation, and analysis methodologies for high-speed LPDDR4 interfaces (information specific to DDR4 will come in a later revision of this document). This is an essential step to ensure the PCB design meets all the requirements to operate the targeted speeds.

4.1 Board Model Extraction

Note

This section contains various information from J7 devices and board designs. The specific details only apply to J7 board designs, but are provided here as an examples for AM62x board designs. Consult the AM62x datasheet for device specifications.

The board level extraction guidelines listed below are intended to work in any EDA extraction tool and are not tool-specific. It is important to follow the steps outlined in Section 4.2 through Section 4.4 immediately after completing touchstone model extractions. The design should be checked with these steps prior to running IBIS simulations.

1. For DDR extractions, extract power (VDDS_DDR/VDDQ) and signal nets together in a 3D-EM solver.
2. Use wide-band models. It is recommended to extract from DC to at least until 6x the Nyquist frequency (for example, for LPDDR4-3733 extract the model at least until 11.2 GHz).
3. Check the board stack-up for accurate layer thickness and material properties.
    a. It is recommended to use Djordjevic-Sarkar models for the dielectric material definition.
4. Use accurate etch profiles and surface roughness for the signal traces across all layers in the stack-up.
5. If the board layout is cut prior to extraction (to reduce simulation time), define a cut boundary that is at least 0.25 inch away from the signal and power nets.
6. Check the via padstack definitions.
    a. Ensure that the non-functional internal layer pads on signal vias are modeled the same way they would be fabricated.
    b. These non-functional internal layer pads on signal vias are not recommended by TI.
7. Use Spice/S-parameter models (typically available from the vendor) for modeling all passives in the system.

4.2 Board-Model Validation

The extracted board models need to be checked for the following properties:

- Passivity: This ensures that the board model is a passive network and does not generate energy
- Causality: This ensures that the board model obeys the causal relationship (output follows input).

These checks can be performed in any standard EDA simulator or extraction engine.

4.3 S-Parameter Inspection

Once the extracted S-parameters have been verified as causal and passive, the S-parameter plots should be inspected. It is recommended to check for the following:

- Insertion Loss: The single-ended insertion loss is recommended to stay within 0 to 10 dB up to 3 times the Nyquist frequency of operation. For example, if the target frequency is 8 Gbps (4GHz Nyquist), the single-ended insertion loss should stay under 10 dB up to 12 GHz.
- Return Loss: The single-ended return loss is recommended to be less than 15 dB up to 3 times the Nyquist frequency.
- Near and Far end crosstalk (FEXT/NEXT): The FEXT and NEXT are recommended to be under 25 dB for frequencies up to 3 times the Nyquist frequency.

The S-Parameter inspection plots are not pass/fail tests, but rather its more of a guide to check if the design has a reasonable chance of performing a the required level.

Copyright © 2024 Texas Instruments Incorporated

4.4 Time Domain Reflectometry (TDR) Analysis

As a lot of the design fixes are targeted towards maintaining uniform trace impedance, an important analysis method used in assessing the quality of the design is the Time Domain Reflectometry (TDR) analysis. This plots the impedance of a trace as a function of its length, as shown in Figure 4-1.

Figure 4-1. TDR Plot Example With Impedance Mismatch

As shown in Figure 4-1, the TDR plot highlights impedance discontinuities in the trace from one end to the other. This method depends on a reflected waveform from the far-end of the trace. The delay in the plot corresponding to a particular point in the trace actually corresponds to 2 times the distance of that point from the source, owing to the round trip time. This needs to be factored in for assessing the source of impedance discontinuities.

The TDR plot can be generated by reading in the S-parameter models generated by the extraction tool and assessing them in “Time-Domain” mode. A standard EDA simulator such as HyperLynx can perform this function. It is recommended to optimize the design to within a ± 5% deviation from the nominal trace impedance.

The TDR plots are not pass/fail tests, but rather is more of a guide to check if the design has a reasonable chance of performing a the required level.

4.5 System Level Simulation

The methodology for validating the DDR interface is outlined in this section. LPDDR4 interfaces, as defined in the JEDEC specification, uses eye masks defined at a target BER (Bit Error Rate) to determine pass or fail for signal integrity. It is essential to perform channel simulations using IBIS models to generate the signal eye diagrams at the targeted BER. These are introduced for memory interfaces starting from LPDDR4

4.5.1 Simulation Setup

Set up the system-level schematic in the simulator by connecting the SOC IBIS model, board model, power supplies, DRAM package model, and DRAM IBIS model. A typical system-level DDR schematic is shown in Figure 4-2.

Note

Be aware of the DRAM configuration (number of dies in the package, number of ranks, and number of channels) while setting up the system schematic. Be aware the DRAM configuration may also include On-Die Decoupling Circuit.

Copyright © 2024 Texas Instruments Incorporated

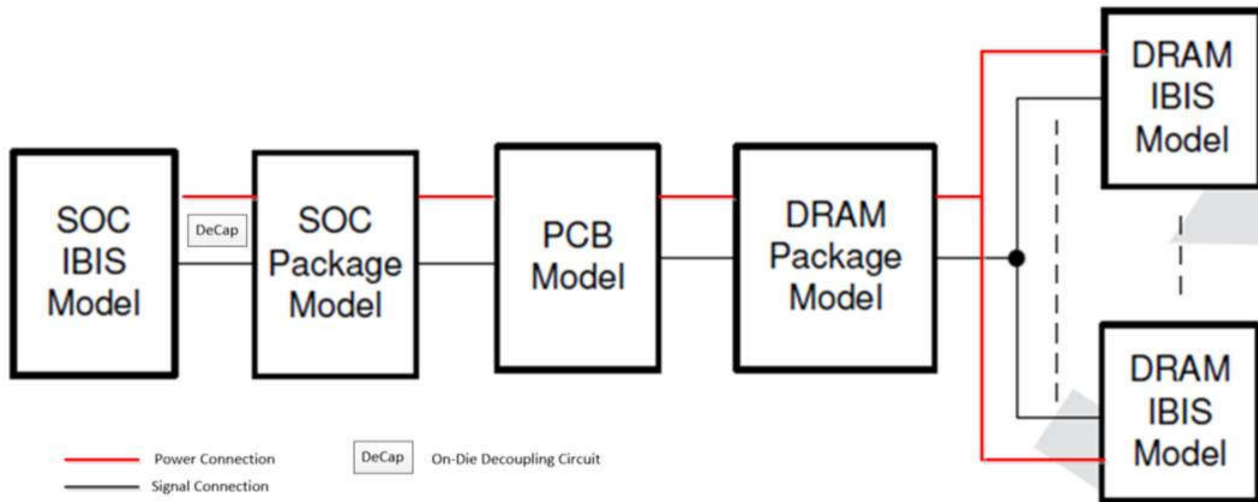

Figure 4-2. Typical System-Level DDR Schematic

- LPDDR4 simulations require power-aware IBIS models for the controller and the memory along with a simulator that supports channel simulations for DDR interfaces.
- SPICE-based transistor-level simulations cannot be used for generating BER signal eyes. Use a simulator that can handle power-aware IBIS simulations and can run channel simulations for the DDR interface.
- IBIS models reduce simulation time with minimal loss in accuracy compared with SPICE-based transistor-level simulations. IBIS models starting from version 5.0 are power-aware models which enables Simultaneous Switching Output (SSO) noise simulations. The TI IBIS model is a power-aware IBIS model.
- Use SPICE models to accurately model the on-die decoupling capacitance on the DDR supply net for both – controller and DRAM. This ensures accurate power noise and Power Supply Induced Jitter (PSIJ) estimation in DDR simulations. The on-die decoupling capacitance information for the DRAM can be obtained from the DRAM vendor.
- Use SPICE or S-parameter files to model the DRAM package. This can be requested from the DRAM vendor. EBD models are not recommended.
- Note that inside the SoC IBIS model, there is a section for the package that contains an RLC matrix for all signal and power nets including DDR. It is recommended to use the SoC IBIS model, not the SOC package S-parameter model. When using SoC IBIS model, be sure to check the "Package Parasitics" (or equivalent parameter in your simulation tools) and use the "Package Model" model type which contains fully coupled L/C information on a per pin basis (denoted in the IBIS file as "[Package Model] am62_pkg").

Copyright © 2024 Texas Instruments Incorporated

- AM62x model for the on-die decoupling capacitance on the DDR supply net:
  - Data

```
*******************************************
* On-die Decoupling circuit for AM62x (DIE_VDDS_DDR to VSS)
*******************************************
* Notes:
* Includes on-die decoupling for all DDR signals
*
* This subcircuit should be added across the AM62x IBIS model's
* DIE_VDDS_DDR and VSS pins
*
*******************************************
* x_decouple DIE_VDDS_DDR vss_die AM62x_ondie_decoupling_alldq
*******************************************
.SUBCKTAM62x_ondie_decoupling_alldq DIE_VDDS_DDR vss_die
Cvddq_c  DIE_VDDS_DDR DIE_VDDS_DDR_c 1.324741e-9
Rvddq_c  vss_die DIE_VDDS_DDR_c 25.0036612e-3
.ENDS
```

  - Address/Command

```
*******************************************
* On-die Decoupling circuit for AM62x (DIE_VDDS_DDR to VSS)
*******************************************
* Notes:
* Includes on-die decoupling for all DDR signals
*
* This subcircuit should be added across the AM62x IBIS model's
* DIE_VDDS_DDR and VSS pins
*
*******************************************
* x_decouple DIE_VDDS_DDR vss_die AM62x_ondie_decoupling_alldq
*******************************************
.SUBCKTAM62x_ondie_decoupling_alldq DIE_VDDS_DDR vss_die
Cvddq_c  DIE_VDDS_DDR DIE_VDDS_DDR_c 4.335517e-9
Rvddq_c  vss_die DIE_VDDS_DDR_c 25.0036612e-3
.ENDS
```


Copyright © 2024 Texas Instruments Incorporated

•

4.5.2 Simulation Parameters

It is important to configure the simulation to exercise the system to real, but worst case parameters.

- Use the worst-case bit pattern to excite the system. The simulator should be able to generate the worst-case bit pattern based on channel characterization.
- Select the controller and DRAM models (sets the drive strength, ODT, VOH levels, and so forth) from the IBIS files which work best for the system.
  - This is typically an iterative process.
  - Every system is unique and the optimal settings for these parameters can vary from system to system.

Table 4-1. Example Data Write ODI/ODT Optimization

Pkg Byte	Board	ODI Ω	ODT Ω	Total EW Margin (ps)	Total EH Margin (mV)B
B3	J7 370HR 10L Ref B3, No BD	40	40	50.28	15.66
B3	J7 370HR 10L Ref B3, No BD	40	48	27.62	11.76
B3	J7 370HR 10L Ref B3, No BD	40	40	33.52	2.92
B3	J7 370HR 10L Ref B3, No BD	48	48	1.54	0.86

- Data bus and address bus ODT and drive strength values can be set independently. As an example, the J7 EVM board (which supports LPDDR4 at similar speeds) used 40-Ω ODT for data read/writes and 80-Ω for CA bus. Drive strength of 40-ohms for data read/write and CA.
  - Data READ Controller model - lpddr4_odt_40, lpddr4_odt_40_diff
  - Data WRITE Controller model - lpddr4_ocd_40p_40n, lpddr4_ocd_40p_40n_diff
  - CA/CLK Controller model - lpddr4_ocd_40p_40n, lpddr4_ocd_40p_40n_diff
- Set up the channel simulation parameters. These typically consists of the data rate, ignore time/bits, minimum number of bits, bit sampling rate, BER floor, number of bits for display, types of BER eyes (voltage and/or timing), and target BER.
  - To determine the minimum number of bits one can run a series of channel simulations with different number of bits. The BER signal eye (and margins) tend to converge after a certain minimum number of bits. This should help determining the minimum number of bits to be used for the system.
  - Run channel simulations to generate the eye diagrams at LBER of -16.
- Run channel simulations with non-ideal power settings at different PVT corners. It is recommended to run the simulations at least at the SSHT and FFLT corners.

Copyright © 2024 Texas Instruments Incorporated

4.5.3 Simulation Targets

Once the simulation successfully completes, generate the DDR analysis reports from the simulation tool. There are several different parameters to be verified, detailed in this section. Each parameter is pass/fail, meaning each must meet the specified target to ensure the design has sufficient margin to operate at the target data rates.

Use the appropriate JEDEC Vref parameters (Vref_min, Vref_max, Vref_step, and Vref_set_tol) and mask parameters (shape, height, width).

4.5.3.1 Eye Quality

The Vix_DQS ratio and Vix_CK ratio for data write and CA bus simulations are to be verified, at the DRAM pin/BGA. Figure 4-3 from the JEDEC specification explains how to measure the Vix ratio, as well as define the ration requirement(s).

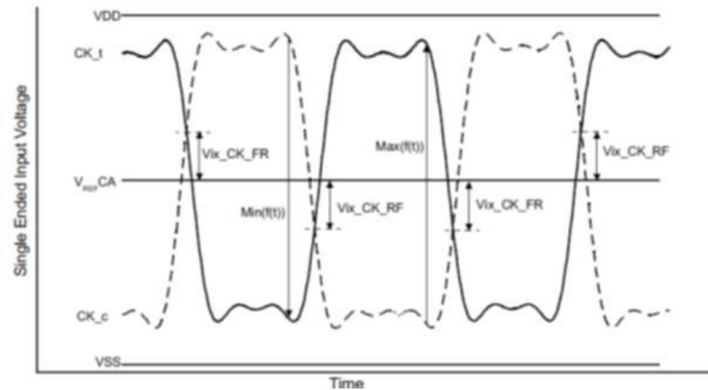

Figure 4-3. Vix_CK and Vix_DQS Ratio for Eye Quality

Copyright © 2024 Texas Instruments Incorporated

Table 4-2. Vix_CK and Vix_DQS Ratio for Eye Quality

Symbol	Data Rate						Unit	Note
	1600/1867		2133/2400/3200		3733/4266			
Vix_CK_ratio	-	25	-	25	-	25	%	(1), (2)
Vix_DQS	-	20	-	20	-	20	%	(1), (2)

(1) Vix_CK_Ratio is defined by the equation: Vix_CK_Ratio = Vix_CK_FR|Min(f(t))|
(2) Vix_CK_Ratio is defined by the equation: Vix_CK_Ratio = Vix_CK_RF|Min(f(t))|

4.5.3.2 Delay Report

The required interconnect delays for DQ, DQS, CA, and CLK are listed in Section 2.16 and Section 2.17. The values listed as 'Typical' are only recommendations. Any minimum/maximum value is a requirement. One key requirement is to ensure the CK delay is greater than any DQS delay. DQSx delays should also be less than the DQ/DM delays in their respective BYTEx. Consider the complete system from SOC die pad, through the PCB, to the pins of the memory package.

4.5.3.3 Mask Report

The minimum jitter and noise margins are to be captured with respect to the eye mask(s). This masks are data rate dependent, and includes:

- Data read eye mask at the SOC die pad for functionality testing
- Data write eye mask (JEDEC spec) at the DRAM pin/BGA for compliance testing
- CA bus eye mask (JEDEC spec) at the DRAM pin/BGA for compliance testing

There should be at least 2 sets of eye diagrams generated by the simulator:

- Vref set to the optimal Vref of the byte offset by the Vref_set_tol in the positive direction (Vref_set_tol is defined in JEDEC spec)
- Vref set to the optimal Vref of the byte offset by the Vref_set_tol in the negative direction

The system-level margins are the worst case noise and jitter margins from all eye diagram measurements listed above (across SSHT and FFLT corners). For all waveforms captured at the DRAM device, margins should be calculated at both the BGA pin and the DRAM pad.

Table 4-3. LPDDR4 Eye Mask Definitions/Requirements

Parameter	Mask Shape	LPDDR4-1600	LPDDR4-3200	LPDDR4-3733
CA eye mask TcIVW	Rectangular (1)	0.3 UI	0.3 UI (1)	(2)
CA eye mask VcIVW	Rectangular (1)	175 mV	155 mV (1)	(2)
Write eye mask TdIVW	Rectangular (1)	0.22 UI	0.25 UI (1)	(2)
Write eye mask VdIVW	Rectangular (1)	140 mV	140 mV (1)	(2)
Read eye mask TdIVW	Diamond	0.42 UI	0.61 UI	0.66 UI
Read eye mask VdIVW	Diamond	140 mV	140 mV	140 mV

(1) Copied from JEDEC specification: Low Power Double Date Rate 4 (LPDDR4).
(2) For details, contact the DRAM vendor.

Copyright © 2024 Texas Instruments Incorporated

Figure 4-4 through Figure 4-6 show the eye mask definitions translated to eye diagrams within captured waveforms.

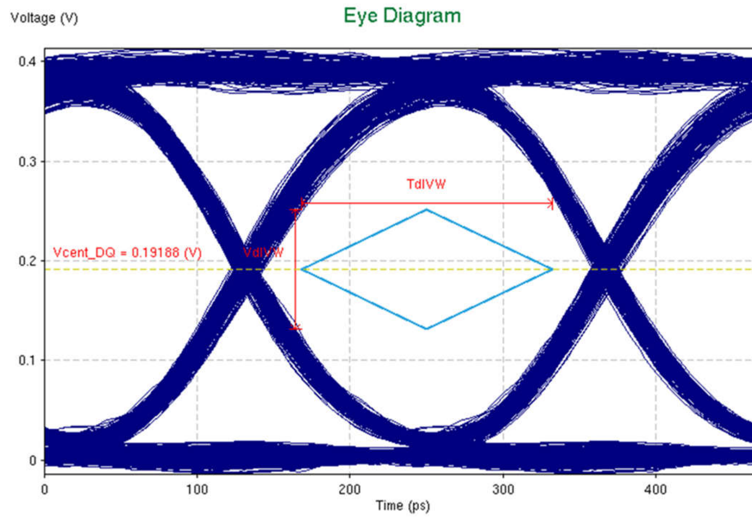

Figure 4-4. Example Simulated LPDDR4-4266 Read Eye With Diamond-Shaped Eye Mask

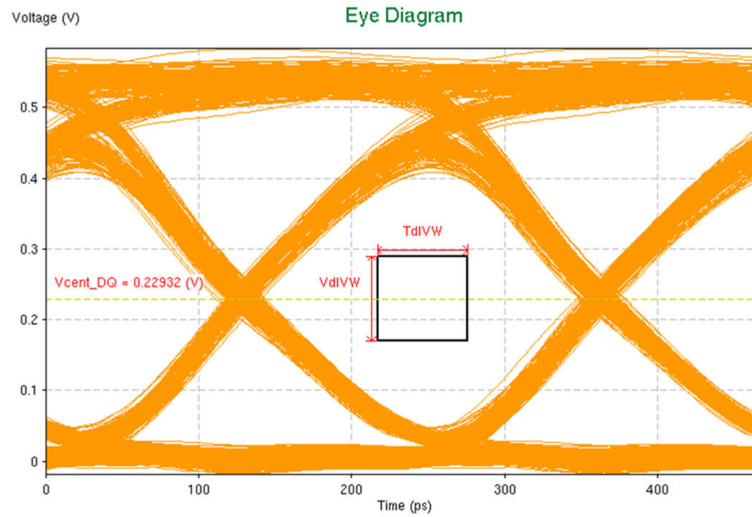

Figure 4-5. Example Simulated LPDDR4-4266 Write Eye With Rectangular JEDEC Eye Mask

Copyright © 2024 Texas Instruments Incorporated

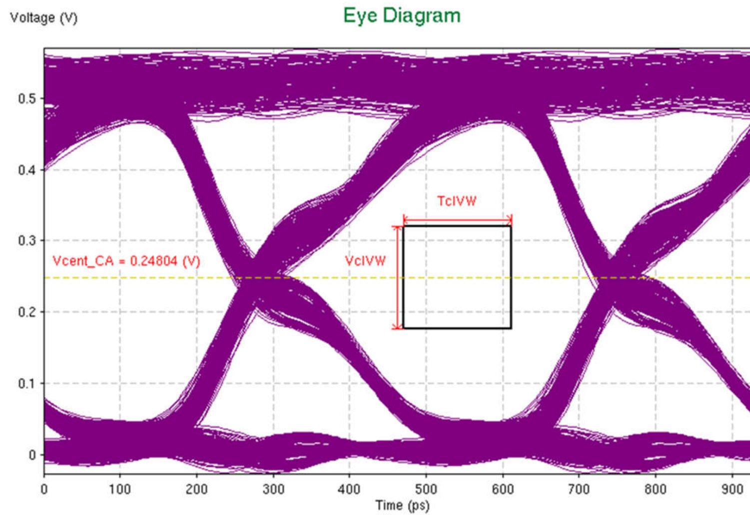

Figure 4-6. Example Simulated LPDDR4-4266 CA Eye With Rectangular JEDEC Eye Mask

4.6 Design Example

4.6.1 Stack-Up

These guidelines recommend a 10- or 12-layer PCB stack-up for full device entitlement. Below are 10- and 12-layer example stack-ups:

- Designs using FR4 products like 370HR are supported, but also recommend higher speed materials like ISOLA I-Speed (or equivalent) for increased margin. IT180A is also another material to help with cost vs. performance tradeoffs
- This example routes data groups on layers 2 and 4. While this minimizes the via travel and therefore reduces via-to-via coupling, but it leaves a longer via stub, which might require back-drill.
- In the 10-layer example, dynamic CA signals are routed on layer 7, and more static control signals routed on layer 9.

Table 4-4. Example 10-layer PCB Stackup for LPDDR4 (J7 EVM)

Layer No	Stackup	Routing Plan Highest Priorities and Layer
	Solder mask	
1	TOP - PWR/SIG	BGA breakouts/VDD_CPU, VDD_CORE and VDD_DDR_1V1
2	PWR/SIG	VDD_CPU and CORE/LPDDR (DBG #3/#1, CAT-Branches)
3	GND	REF
4	PWR/SIG	VDDA_PHYCORE_0V8, VDD_xxx, 0V85/LPDDR (DBG #2/#0)
5	PWR/GND	VDDA_0V8_xxx and GND flood for LPDDR4
6	PWR/GND	VDD_xxx, VDDA_xxx supplies and GND flooded for LPDDR4
7	SIG/PWR	VDD_xxx, VDDA_xxx/LPDDR (Dynamic CA, Trunks)/SERDES
8	GND	REF
9	SIG/PWR	VDD_xxx, VDDA_xxx/LPDDR (static CA)
10	BOTTOM - SIG/PWR	BGA breakouts/Pwr and GND plan segments
	Solder mask	

Copyright © 2024 Texas Instruments Incorporated

Table 4-5. Example 12-layer PCB Stackup for LPDDR4 (AM62Ax LP SK EVM)

Layer No	Stackup	Routing Plan Highest Priorities and Layer
	Solder mask	
1	TOP - PWR/SIG	BGA breakouts, VDD_LPDDR4, GND
2	GND	REF
3	PWR/SIG	VDDA_1V8, GND, LPDDR (DBG #3/#1, CA T-Branches), LVCMOS escape
4	GND	REF
5	SIG/GND	GND, LPDDR (DBG #2/#0), LVCMOS escape
6	PWR/GND	GND (under LPDDR), VDD_CORE, VDDR_CORE, VDDA_1V8, VDDSHVx
7	PWR	DVDD_3V3, DVDD_1V8, VDD1_LPDDR4_1V8
8	PWR	VDD_CORE, VDD_LPDDR4, VDDA_x
9	GND	REF
10	SIG/GND	GND, LPDDR (CA point-to-point, CA Trunks), LVCMOS escape
11	GND	REF
12	BOTTOM - SIG/PWR	GND, decaps, LVCMOS escape
	Solder mask	

Table 4-6. Example 12-layer PCB Stackup for LPDDR4 (AM62Px SK EVM)

Layer No	Stackup	Routing Plan Highest Priorities and Layer
	Solder mask	
1	TOP - PWR/SIG	BGA breakouts, VDD_LPDDR4, GND
2	GND	REF
3	PWR/SIG	VDDA_1V8, GND, LPDDR (DBG #3/#1, CA T-Branches), LVCMOS escape
4	GND	REF
5	SIG/GND	GND, LPDDR (DBG #2/#0), LVCMOS escape
6	GND	REF
7	PWR	VDD_CORE, VDD_LPDDR4, DVDD_3V3
8	PWR/GND	VDD1_LPDDR4_1V8, GND, VDDA_x
9	PWR/GND	GND, VDDR_CORE, VDDA_1V8, DVDD_3V3, DVDD_1V8
10	SIG/GND	GND, LPDDR (CA point-to-point, CA Trunks), LVCMOS escape
11	GND	REF
12	BOTTOM - SIG/PWR	GND, decaps, LVCMOS escape
	Solder mask	

Table 4-7 provides simulation results performed on sample designs, showing the impact of the PCB stackup (material, drill plan, and so forth) on LPDDR4 performance. The results showed that maximum bandwidth could be achieved on a FR4 solution, but required back-drilling. The higher frequency material could achieve same performance without back drill. Note the 8 layer design only achieved 3733, but this was due to other design compromises due to limited layers (solid reference planes, and so forth).

Table 4-7. Example LPDDR4 Performance Impact From J7 EVM Stackup

Design	Material	Layer Count	Via Back Drilling	Maximum LPDDR4 Speed (Mbps) [(1)]
J7 EVM	I-Speed	16	Yes	4266
Ref Board	I-Speed	10	No	4266
Ref Board	370HR	10	Yes	4266
Ref Board	370HR	8	No	3733

(1) These results are for J7 designs. For maximum supported data rates, see the device-specific data sheet.

Copyright © 2024 Texas Instruments Incorporated

4.6.2 Routing

The below examples from a J7 design show the LPDDR4 Clock and CA routing on an example 10-layer PCB design. The clock is routed differentially with target impedance of 70 Ω. For the T-branch to match the impedance of the trace, the impedance needs to be doubled. This can create challenges, as the higher impedances can be difficult to achieve in some PCB stackups. The CA signals are routed targeting 35 Ω, with the T-branch at two times the source impedance.

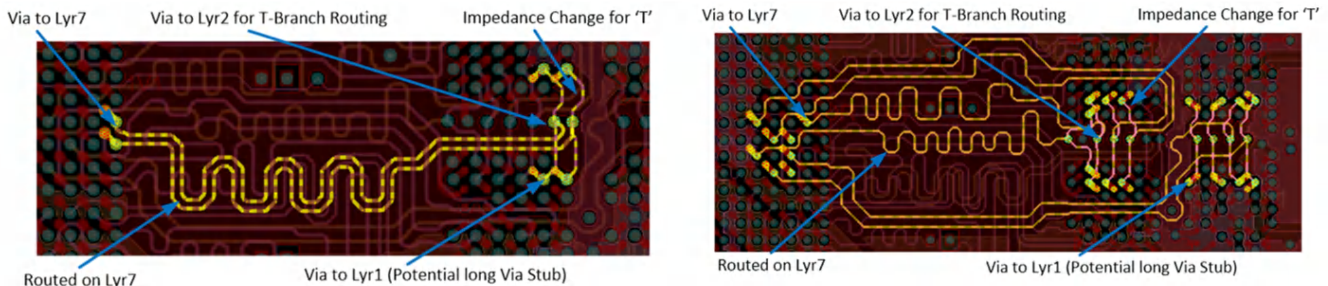

Figure 4-7. Example LPDDR4 Clock and CA Routing

On the same 10 layer reference design, the data groups are routed on layers 2 and 4. The upper layers are used due to the minimum via travel, which minimized the via inductance and via-to-via coupling. Because the data signals are point-to-point, T-branch routing is not required.

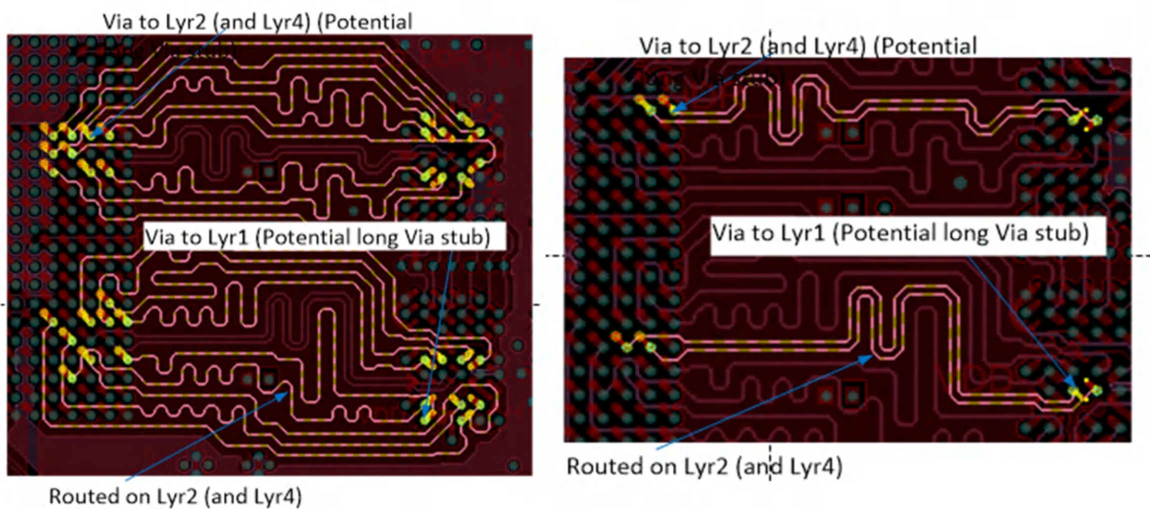

Figure 4-8. Example LPDDR4 Data Byte(s) and DQS(s) Routing

Table 4-8. LPDDR4 Performance Impact on Routing Layer (Read at pad)

Routing Layer	Via Type	Back Drilling	EW Margin (ps)	EH Margin (mV)
L1, L12	PTH	No	-7.10	56.72
L1, L12	PTH	Yes	-4.86	55.71
L1, L3	PTH	No	5.70	40.29
L1, L3	PTH	Yes	8.37	34.54

Table 4-9. LPDDR4 Performance Impact on Routing Layer (Write at pad)

Routing Layer	Via Type	Back Drilling	EW Margin (ps)	EH Margin (mV)
L1, L12	PTH	No	17.42	39.22
L1, L12	PTH	Yes	20.04	41.93
L1, L3	PTH	No	27.66	41.37
L1, L3	PTH	Yes	27.76	48.63

Copyright © 2024 Texas Instruments Incorporated

4.6.3 Model Verification

Before simulating, it is recommended to verify the models. One verification method described is the impedance plot (or impedance scan). The Impedance scans for a 10 layer design are provided.

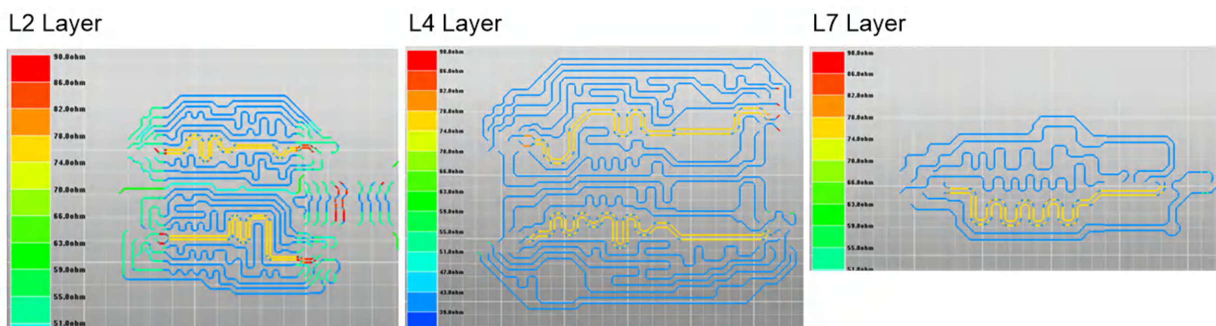

Figure 4-9. Example LPDDR4 Trace Impedance Scan

Table 4-10. Example LPDDR4 Trace Impedance Summary for Data

Layer	DDR Bus	DQ SE Impedance (Ω)	DQS/CLK Difference Impedance (Ω)
L2	B1 and B3	40.9	77.7
L2	CA	51.7	101.4
L4	B0 and B2	41.1	77.7
L7	CA	41.1	77.7

For CK and CA signals, the goal is to have the branch segment equal to two times the impedance of the feed trace. Note its common for the PCB to limit the achievable impedances. Simulations will show you if the compromises are acceptable.

Table 4-11. Example LPDDR4 Trace Impedance Summary for CA

Board	CA Feed Impedance (Ω)	CA Branch Impedance (Ω)	CA Branch Target (Ω)	Impedance Mismatch (Ω)
Initial Design	49.1	59.6	98 (49x2)	19.3
Final Design	41.1	51.7	82 (41x2)	15.3

The simulation results show the improvement made by closer matching the impedances to their targets.

Table 4-12. Example LPDDR4 Simulation Results From Improved Trace Impedance

Board	Total Eye Width Margin (ps)	Total Eye Height Margin (ps)
Initial Design	58.00	14.00
Final Design	124.68	48.08

4.6.4 Simulation Results

The simulation results are provided for the LPDDR4 interface from a 10 layer design. These simulation targets must be met to ensure the design will operate at the desired level of performance.

CA simulations need to be verified at the DRAM pin/BGA. This includes:

- Vix_CK ratio (JEDEC)
- Jitter/noise margins with respect to the eye mask (JEDEC)
- Peak-peak power noise

Copyright © 2024 Texas Instruments Incorporated

At DRAM Pin:

Jitter/Noise Eye Margins

Vix_CA Ratio

Package	Board	Total EW Margin (ps)	Total EH Margin (mV)	Min Rback Margin H (mV)	Min Rback Margin L (mV)	Vix_CK_Ratio (%)
J7 SCK 1207	J7 370HR 10L Ref 121819 ACC, No BD, SIwave	182.04	68.50	89.43	25.50	16.81

Ring-back High/Low Margins

Figure 4-10. LPDDR4 Simulation Results for CA

Data write simulations need to be verified at both the DRAM BGA pin and the DRAM pad. This includes:

- Vix_CK ration (JEDEC)
- Jitter/noise margins with respect to the eye mask (JEDEC)
- Peak-peak power noise

At DRAM Pin:

Jitter/Noise Eye Margins

Vix_DQS Ratio

Pkg Byte	Board	Total EW Margin (ps)	Total EH Margin (mV)	Min Rback margin H (mV)	Min Rback Margin L (mV)	Vix_DQS_Ratio (%)
B0	J7 370HR 10L Ref 121819 B0, With BD	62.32	55.20	68.63	111.83	9.80
B1	J7 370HR 10L Ref 121819 B1, With BD	54.52	94.28	86.46	114.69	7.64
B2	J7 370HR 10L Ref 121819 B2, With BD	53.40	73.96	81.19	106.32	5.55
B3	J7 370HR 10L Ref 121819 B3, With BD	54.86	52.74	34.97	48.55	9.81

Ring-back High/Low Margins

Figure 4-11. LPDDR4 Simulation Results for Write

Data read simulations need to be verified at SOC. This includes:

- Jitter/noise margins with respect to the eye mask
- Peak-peak power noise

Jitter/Noise Eye Margins

Power Noise

Pkg Byte	Board	Total EW Margin (ps)	Total EH Margin (mV)	Min Rback Margin H (mV)	Min Rback Margin L (mV)	P-P VDDQ Noise (mV)
B0	J7 370HR 10L Ref 121819 B0, No BD	17.06	24.00	25.86	26.30	21.33
B1	J7 370HR 10L Ref 121819 B1, No BD	6.72	8.00	28.49	24.44	33.33
B2	J7 370HR 10L Ref 121819 B2, No BD	0.98	FAIL	19.40	17.18	28.97
B3	J7 370HR 10L Ref 121819 B3, No BD	FAIL	FAIL	11.05	20.26	15.87
B0	J7 370HR 10L Ref 121819 B0, With BD	19.02	26.00	27.84	30.10	26.12
B1	J7 370HR 10L Ref 121819 B1, With BD	6.24	8.00	30.79	27.25	39.98
B2	J7 370HR 10L Ref 121819 B2, With BD	5.24	6.00	27.95	28.83	40.24
B3	J7 370HR 10L Ref 121819 B3, With BD	4.60	6.00	35.41	34.52	25.21

Ring-back High/Low Margins

Figure 4-12. LPDDR4 Simulation Results for Read

The simulations results for read includes two sets for data, black and green. The black shows the design failed, as several bytes failed to meet the eye margins. The green is the simulation results of the same design, but with back-drilling the via stubs applied.

Copyright © 2024 Texas Instruments Incorporated

5 Appendix: AM62x ALW and AMC Package Delays

The package delays provided in this appendix are measured from SOC die pad to SOC package pin. The skew limits specified in Table 3-6 and Table 3-7 are measured from SOC die pad to DRAM package pin (including these delays inside the SOC package). The designer can sum these package delays with the PCB delays for each net when checking for complance with the skew limits. Simulations of the propagation delays are then required to confirm the delays satisfy the requirements.

PROCESSOR PIN NAME	AM62x ALW PACKAGE DELAY (ps)	AM62x AMC PACKAGE DELAY (ps)	NET CLASS	DESCRIPTION
DDR0_A0	21.00	28.55	ADDR_CTRL	Used with LPDDR4 and DDR4
DDR0_A1	19.94	26.82	ADDR_CTRL	Used with LPDDR4 and DDR4
DDR0_A2	16.13	26.60	ADDR_CTRL	Used with LPDDR4 and DDR4
DDR0_A3	11.28	26.33	ADDR_CTRL	Used with LPDDR4 and DDR4
DDR0_A4	11.62	26.52	ADDR_CTRL	Used with LPDDR4 and DDR4
DDR0_A5	20.98	26.43	ADDR_CTRL	Used with LPDDR4 and DDR4
DDR0_A6	17.04	33.55	ADDR_CTRL	Used with DDR4
DDR0_A7	19.08	31.69	ADDR_CTRL	Used with DDR4
DDR0_A8	21.55	34.42	ADDR_CTRL	Used with DDR4
DDR0_A9	13.33	34.25	ADDR_CTRL	Used with DDR4
DDR0_A10	10.71	32.94	ADDR_CTRL	Used with DDR4
DDR0_A11	9.00	27.28	ADDR_CTRL	Used with DDR4
DDR0_A12	9.33	29.79	ADDR_CTRL	Used with DDR4
DDR0_A13	23.19	31.52	ADDR_CTRL	Used with DDR4
DDR0_ACT_n	7.98	32.68	ADDR_CTRL	Used with DDR4
DDR0_ALERT_n	17.37	32.54	N/A	Used with DDR4
DDR0_BA0	19.81	26.29	ADDR_CTRL	Used with DDR4
DDR0_BA1	21.85	27.00	ADDR_CTRL	Used with DDR4
DDR0_BG0	15.30	27.61	ADDR_CTRL	Used with DDR4
DDR0_BG1	17.09	26.62	ADDR_CTRL	Used with DDR4
DDR0_CAS_n	11.81	26.42	ADDR_CTRL	Used with DDR4
DDR0_CK0	23.03	34.72	CK0	Used with LPDDR4 and DDR4
DDR0_CK0_n	21.28	33.13	CK0	Used with LPDDR4 and DDR4
DDR0_CKE0	20.94	29.26	ADDR_CTRL	Used with LPDDR4 and DDR4
DDR0_CKE1	13.68	31.10	ADDR_CTRL	Used with LPDDR4 and DDR4
DDR0_CS0_n	7.80	28.47	ADDR_CTRL	Used with LPDDR4 and DDR4
DDR0_CS1_n	18.29	35.18	ADDR_CTRL	Used with LPDDR4 and DDR4
DDR0_DM0	13.95	35.06	BYTE0	Used with LPDDR4 and DDR4
DDR0_DM1	19.07	28.18	BYTE1	Used with LPDDR4 and DDR4
DDR0_DQ0	16.90	37.82	BYTE0	Used with LPDDR4 and DDR4
DDR0_DQ1	14.21	29.20	BYTE0	Used with LPDDR4 and DDR4
DDR0_DQ2	20.40	31.14	BYTE0	Used with LPDDR4 and DDR4
DDR0_DQ3	17.67	28.54	BYTE0	Used with LPDDR4 and DDR4
DDR0_DQ4	23.82	38.78	BYTE0	Used with LPDDR4 and DDR4
DDR0_DQ5	21.95	32.97	BYTE0	Used with LPDDR4 and DDR4
DDR0_DQ6	24.74	35.55	BYTE0	Used with LPDDR4 and DDR4
DDR0_DQ7	24.31	33.64	BYTE0	Used with LPDDR4 and DDR4
DDR0_DQ8	23.28	34.75	BYTE1	Used with LPDDR4 and DDR4
DDR0_DQ9	18.16	32.35	BYTE1	Used with LPDDR4 and DDR4
DDR0_DQ10	19.18	32.72	BYTE1	Used with LPDDR4 and DDR4

Copyright © 2024 Texas Instruments Incorporated

PROCESSOR PIN NAME	AM62x ALW PACKAGE DELAY (ps)	AM62x AMC PACKAGE DELAY (ps)	NET CLASS	DESCRIPTION
DDR0_DQ11	17.78	37.01	BYTE1	Used with LPDDR4 and DDR4
DDR0_DQ12	20.45	35.22	BYTE1	Used with LPDDR4 and DDR4
DDR0_DQ13	16.68	32.03	BYTE1	Used with LPDDR4 and DDR4
DDR0_DQ14	24.67	29.99	BYTE1	Used with LPDDR4 and DDR4
DDR0_DQ15	21.39	26.31	BYTE1	Used with LPDDR4 and DDR4
DDR0_DQS0	27.23	40.21	DQS0	Used with LPDDR4 and DDR4
DDR0_DQS0_n	27.39	40.92	DQS0	Used with LPDDR4 and DDR4
DDR0_DQS1	21.74	39.71	DQS1	Used with LPDDR4 and DDR4
DDR0_DQS1_n	22.68	41.12	DQS1	Used with LPDDR4 and DDR4
DDR0_ODT0	29.40	29.31	ADDR_CTRL	Used with DDR4
DDR0_ODT1	18.45	30.02	ADDR_CTRL	Used with DDR4
DDR0_PAR	25.10	28.63	ADDR_CTRL	Used with DDR4
DDR0_RAS_n	10.64	28.63	ADDR_CTRL	Used with DDR4
DDR0_RESET0_n	31.66	32.77	ADDR_CTRL	Used with LPDDR4 and DDR4
DDR0_WE_n	15.43	31.72	ADDR_CTRL	Used with DDR4

6 Revision History

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

Changes from January 1, 2023 to November 2, 2024 (from Revision A (January 2023) to Revision B (November 2024)) **Page**

- Specified power aware simulation should be used to determine decoupling capacitor count and total values..5
- Clarified DDR4 supported topologies....... 7
- Updated DDR4 Schematics to include info on VTT and dual rank.......9
- Updated info on DDR4 Signal Termination.......15
- Updated VTT as optional for point to point designs.......15
- Updated CK and ADDR_CTRL Routing Limits and Routing Specifications table....... 21
- Updated Data Group Routing Limits and Data Group Routing Specifications tables.......23
- Added support for DDR4 data bit swapping and byte swapping....... 23
- Clarified LPDDR4 supported topologies.......24
- Updated LPDDR4 Data Group Topologies.......28
- Updated LPDDR4 CK and ADDR_CTRL Routing specifications....... 29
- Updated LPDDR4 Data Group Routing Specifications.......30
- Allowing DQ/DM bit swapping and byte swapping....... 31
- Added LPDDR4 Simulation section for AM62x....... 32
- Removed Waveform Quality section (ring-back margins)....... 37
- LPDDR4-3733 Read eye mask VdIVW corrected to 140mV in Table 4-3....... 38
- Add eye masks for LPDDR4-1600....... 38
- Added AM62Px SK EVM stackup into Table 4-6 40
- Corrected Impedance Mismatch calculation in Table 4-11....... 43
- Removed Minimum ring-back margins at high/low levels (JEDEC)....... 43
- Added package delay section.......45

Copyright © 2024 Texas Instruments Incorporated

IMPORTANT NOTICE AND DISCLAIMER

TI PROVIDES TECHNICAL AND RELIABILITY DATA (INCLUDING DATA SHEETS), 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, regulatory 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 or other applicable terms available either on 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.

TI objects to and rejects any additional or different terms you may have proposed.

Mailing Address: Texas Instruments, Post Office Box 655303, Dallas, Texas 75265
Copyright © 2024, Texas Instruments Incorporated