88AP510

High Performance SoC with Integrated CPU, 2D/3D Graphics Processor, and High-Definition Video Decoder

Hardware Design Guide

Doc. No. MV-S301669-U0, Rev. E
July 5, 2011, Preliminary
Document Classification: Proprietary Information

Document Conventions

	**Note:** Provides related information or information of special importance.
	**Caution:** Indicates potential damage to hardware or software, or loss of data.
	**Warning:** Indicates a risk of personal injury.

Document Status

Doc Status: Preliminary	Technical Publication: 0.xx

For more information, visit our website at: www.marvell.com

Revision History

Table 1: Revision History

Revision	Date	Comments
Rev. D.1	July 5, 2011	Revised Release
1. Added R3/R4 resistor labels to figures in Section 3, 32-bit SDRAM DDR2 Interface, on page 34 and Section 4, 32-bit SDRAM DDR3 Interface, on page 50.		
Rev. D	June 6, 2011	Revised Release
2. Revised the Implementation Notes sections in Section 2, Generic Guidelines for SERDES Interfaces, on page 16. 3. Added a reference to copper surface roughness in Section 2.1, Insertion Loss and Loss Budget, on page 16. 4. Revised Section 2.2, Inter-Symbol Interference (ISI), on page 19. 5. Revised much of Section 2.4, Crosstalk, on page 21. 6. Added the section Coplanar coupling on a two-layer PCB on page 23 and Section 2.4.4, 3-Dimensional Structures with High-Crosstalk Susceptibility, on page 23. 7. Revised Section 2.5, Return Path Continuity, on page 25, Section 2.6, Target Routing Impedance, on page 27, and Section 2.7, Capacitive Discontinuities, on page 28. 8. Changed BGA to BGA/QFP in Section 2.9, Via Structures, on page 30. 9. Revise Section 2.9.1, Removing Unused Pads in a Through-hole Via Construction, on page 30 and Section 2.9.2, Effect of the Via Stub, on page 31. 10. Changed the differential impedance tolerance from 5% to 10% in all of the topology tables in Section 3, 32-bit SDRAM DDR2 Interface, on page 34 and Section 4, 32-bit SDRAM DDR3 Interface, on page 50. 11. Added new routing notes to Section 3.2, Interface Signals Layout Guidelines, on page 37 and Section 4.2, Interface Signals Layout Guidelines, on page 55. 12. In Section 4, 32-bit SDRAM DDR3 Interface, on page 50: • Changed the maximum operating frequency to 500 MHz. • Changed the clock skew min values to -8 inches in all of the Routing Constraints tables. • Added a note that R3 and R4 must be placed within 500 mils of the device after the Routing Constraint tables. 13. Added VDDO_M pull-up/down information to: • Figure 29, 2 x 16-bit Wide Memory Topology—Address and Control in Daisy Chain, on page 56 • Figure 31, 4 x 16-bit Wide Memory Topology for Data, on page 60 • Figure 33, 4 x 8-bit Wide Memory Topology for Data, on page 63 • Figure 34, 8 x 8-bit Wide Memory Topology for Data, on page 66 14. Added a maximum value of 200 nF for C1 in Table 16, System Board with Existing Connector, on page 74, Table 17, Add-in Card, on page 75, and Table 18, Same-Board Connection, on page 76. 15. In Table 34, Routing Constraints for VGA Connection, on page 104: • Changed the TL[x] internal-group separation (crosstalk) value to max -20 dB. • Changed the TL[x] to TL[y] cross-group separation (crosstalk) value to max -26 dB. 16. Added Section 16.3, Polygon Shape Considerations, on page 115. 17. Added Section 16.7, Bulk Capacitors, on page 119. 18. Added the following note to Section 17.3, System Power Topology Implementation, on page 124: **NOTE:** When using the 88AP510 with a CPU frequency of 1 GHz, the CPU power supply recommended operating condition is stricter. To meet those specifications, Marvell® recommends using the Marvell® 88PH845-T184 regulator.		

Table 1: Revision History (Continued)

Revision	Date	Comments
Rev. C	August 16, 2010	Revised Release

1. In Section 2.6.1, Implementation Notes, on page 27, added a bullet:
   - The dielectric constant varies depending on the frequency, causing matched impedance to increase with frequency. To meet the matched impedance at Fb/2, take into consideration Er frequency dependency and the target.
2. Added Section 2.9.3, Via Structures Crosstalk, on page 31.
3. In Section 3.2, Interface Signals Layout Guidelines, on page 37, added a bullet:
   - The DDR CKE signals must have a 4.7 kohm pulldown resistor in designs, where VTT termination on CKE is not required.
4. Added the Reset signal to Table 8, Signal Groups, on page 50.
5. Added the Reset signal to the following figures:
   - Figure 25, 2 x 16-bit Wide Memory Devices Connected to the Controller, on page 51
   - Figure 26, 4 x 16-bit Wide Memory Devices with Two Chip Selects Connected to the Controller, on page 52
   - Figure 27, 4 x 8-bit Wide Memory Devices Connected to the Controller, on page 53
   - Figure 28, 8 x 8-bit Wide Memory Devices Connected to the Controller, on page 54
6. In Section 4.2, Interface Signals Layout Guidelines, on page 55, added a bullet:
   - The DDR CKE signals must have a 4.7 kohm pulldown resistor.
7. Added the following note after Figure 30, 4 x 16-bit Wide Memory Topology for Clock/Address/Command/Control Implemented as Daisy Chain, on page 59:

NOTE: Multiple clocks can be used to ease the routing. If multiple output clocks are available in the specific Marvell® device, be sure that the routing does not violate the clock's timing constraints.

8. Added Section 4.1.5, Connectivity with Two Chip Select 8 x 8-bit Wide Memory, on page 54 and Section 4.2.4, Routing Topologies—Two Chip Selects 8 x 8-bit Wide Memory, on page 66.
9. Revised Table 13, ODT/RTT Setting Matrix for a Two-Chip-Select Topology, on page 69.
10. Revised the way that the PCIe input clock is presented in Table 14, Signal Groups, on page 70.
11. Revised the figures Figure 38, Topology with an Existing Connector and/or an Add-in Card, on page 73 and Figure 39, Topology with a Same-Board Connection, on page 76. Also, after the tables following these figures revised the note:
    High-speed, low-ESL, low-ESR capacitors are recommended to:
    It is essential to use high-speed, low-ESL, low-ESR capacitors.
12. Changed the (TL1n+TL2n)-(TL1m+TL2m), (TL3n+TL4n)-(TL3m+TL4m) length max value to 5 inches in Table 18, Same-Board Connection, on page 76.
13. Revised Figure 49, Vbus Connectivity when Configured as a Host, on page 93.
14. Revised Figure 59, VGA Connector Topology, on page 104.
15. In Table 34, Routing Constraints for VGA Connection, on page 104:
    - Added TL4 length value.
    - Added C and FB – RGB signal filters information.
    - Added TL4 to control to data skew equation.
    - Added TL4 to single-ended impedance.
16. Added Section 12.3, Analog Signal Reference Plane and Return Path Considerations, on page 105.
17. In Section 17, Power Management Unit (PMU) Board Design Guidelines, on page 121, the following sections were added:
    - Added Section 17.5.1, PCI Express Power Saving, on page 131.
    - Added Section 17.5.2, USB Power Saving, on page 139.
    - Added Section 17.5.3, Ethernet Power Saving, on page 139.
    - Added Section 17.5.4, SATA Power Saving, on page 141.
    - Added Section 17.5.5, Clock Gating for Unused Interfaces, on page 142.
18. Added CKE_MASK and M_RESET_MASK to Table 39, Power Management Signals, on page 123.
19. Added system consideration notes at the end of Section 17.4.1, CPU DVFS (Dynamic Voltage and Frequency Scaling), on page 126.

Table 1: Revision History (Continued)

Revision	Date	Comments
20. Changed the period for the BootROM driving the signal low to 300 us for MPP in Section 17.4.5, Special Consideration for DDR Reset and CKE During Standby, on page 129. 21. Revised Figure 78, CKE and Reset Implementation for Standby Support, on page 130. 22. Revised Figure 82, Device Clock Topology, on page 143. 23. Added Appendix C, LCD Clock Tree Architecture Appendix, on page 165.		
Rev. B	January 3, 2010	Revised Release
1. Change the product number (PN) to 88AP510. 2. Revised Section 1.2, Related Documentation, on page 14. 3. Added a new note for the R1, R2 parameters in the routing constraint tables in Section 3, 32-bit SDRAM DDR2 Interface, on page 34 and Section 4, 32-bit SDRAM DDR3 Interface, on page 50: Place a bypass capacitor between VDDIO and GND, and close to the Thevenin terminations. The ratio of capacitors to terminations is 1:3. 4. Added note about the maximum frequency that the interface connectivity and routing guidelines apply to in Section 3.1, Interface Connectivity, on page 34 and Section 4.1, Interface Connectivity, on page 50. 5. Added a note about clock routing guidelines after Figure 37, PCI Express Interface Connectivity and Reference Clock Supplied by Marvell® Device, on page 72. 6. Revised note 7 following the table In Table 16, System Board with Existing Connector, on page 74. 7. Revised the description of parameter *TL1[x] to TL1[y] separation*, and revised note 11 in Table 18, Same-Board Connection, on page 76. 8. Revised the description of parameter *TL1[x] to TL1[y] separation*, and revised the minimum value for parameter *TL1[x]* or *TL3[x] to TL[x] lane separation* from other traces from 4 to 5.5 signal width in Table 25, Connection to a Standard SATA II 1-Meter Cable, on page 87. 9. Revised the description of parameter *TL1[x] to TL1[y] separation*, and revised the minimum value for parameter *TL1[x] or TL3[x] to TL[x] lane separation* from other traces from 4 to 5.5 signal width and revised note 9 following Table 26, End-to-End Connection Parameters for a Well-Defined Connection Cable or Board, on page 89. 10. Revised Section 10, Sony/Philips Digital Interface (S/PDIF), on page 97 to indicate that the device only supports output data on this interface. 11. Revised Figure 64, JTAG Connection, on page 112. 12. Added CPU and Core power down (Standby mode) information to Table 38, 88AP510 Power Modes, on page 122. 13. Updated Figure 82, Device Clock Topology, on page 143. 14. Added Section 11, Liquid Crystal Display (LCD) Interface Guidelines, on page 99 and Section 13, NAND Flash Interface Guidelines, on page 107. 15. In Section 16.4.1, Implementation Notes, on page 116: • Added a third bullet in the filter implementation notes. • Changed the value from 10–100 nF to 100 nF in the second bullet of the decoupling implementation notes. 16. In Table 37, Definition of Analog Power Filter Symbols, on page 117, revised the value of Cx from 10–100 nF to 100 nF and revised note 2. 17. Completely revised Section 17, Power Management Unit (PMU) Board Design Guidelines, on page 121. 18. Added Appendix A, Hardware Abstraction Layer Power Management Procedures Appendix, on page 148.		
Rev. A	June 8, 2009	Initial release.

Table of Contents

List of Tables

List of Figures

1 Introduction

This hardware design guide is provided by Marvell® for board designs that use the 88AP510 device. It provides designers with guidelines and design rule recommendations.

If the guidelines listed in this document are not followed, it is important to perform thorough signal integrity and timing simulations for the design. This will ensure proper signal integrity and timing. Any deviation from the guidelines should be simulated.

For full hardware specifications, refer to the *88AP510 Hardware Specifications* document.

1.1 Relevant Devices

This document is relevant for the 88AP510 device.

In this document, the 88AP510 is often referred to as the “device”.

1.2 Related Documentation

The following documents contain additional information related to the 88AP510. For the latest revision, contact a Marvell representative.

- *88AP510 Hardware Specifications*, Doc. No. MV-S105141-U0
- *88AP510 Functional Specifications*, Doc. No. MV-S105142-U0
- *88SV581x-V6 PJ4 ARM v6 Sheeva™ CPU Core Datasheet,* Doc No. MV-S105190-00[1]
- *88SV581x-V7 PJ4 ARM v7 Sheeva™ CPU Core Datasheet,* Doc No. MV-S105190-01[1]
- *AN-216, EMI Guidelines*, Doc. No. MV-S300935-00[1]
- *Green Ethernet with Marvell® Alaska® PHYs and Integrated Switches,* Doc No. MV-S301313-00D[1]
- *PCI Express Base Specification.1.1, Copyright © 2002-2005 PCI-SIG*

See the Marvell Extranet website for the latest product documentation.

1.3 Acronyms

Table 2 lists the definition for acronyms used throughout this document.

Table 2: Acronyms

Acronym	Definition
PMU	Power Management Unit
HAL	Hardware Abstraction Layer
PM	Power Management
GPIO	General Purpose Input/Output pin
DFS	Dynamic Frequency Scaling
DVFS	Dynamic Voltage Frequency Scaling
PLL	Phase Locked Loop

1. This document is a Marvell® proprietary, confidential document, requiring an NDA and can be downloaded from the Marvell Extranet.

Table 2: Acronyms (Continued)

Acronym	Definition
RTC	Real Time Clock
SRAM	Synchronous Random Access Memory

2 Generic Guidelines for SERDES Interfaces

This section presents the generic design and routing guidelines for differential pairs of serial interfaces. While the interface protocols vary, many of the electrical aspects are similar in most high-speed serial differential interfaces.

This generic section refers to both BGA/QFP packages and some of information may be irrelevant to the package type used in a specific device. Some references to the package type can refer to other chips that are on the board.

This section addresses the following issues that must be dealt with to ensure a robust, high-speed serial interface board design:

- Insertion Loss and Loss Budget
- Inter-Symbol Interference (ISI)
- Placement of Devices and Connectors on the Board
- Crosstalk
- Return Path Continuity
- Target Routing Impedance
- Capacitive Discontinuities
- Selecting the Appropriate Components
- Trace Symmetry and Matching—Mode Conversion
- Via Structures

Every sub-section includes two parts. The first part provides theoretical background and motivation, and the second part contains the actual design recommendations.

Note

If there is a miscorrelation between the data presented in this *Generic Guidelines for SERDES Interfaces* section and data presented in a specific interface design guidelines section, the rule to comply to is the data presented in the specific interface section.

2.1 Insertion Loss and Loss Budget

In this section, the term insertion loss is referred to as the overall signal quality degradation factors (not just to skin effect, dielectric loss, and surface roughness).

The insertion loss versus frequency curve describes the attenuation of the signal per frequency, in terms of amplitude and phase. It is generally acceptable to measure the insertion amplitude loss curve in units of dB - Insertion loss [db] = 20 log (V far-end / V near-end). For high-speed serial interfaces, the differential insertion loss is usually measured, as it describes the attenuation of a differential signal.

Most high-speed serial differential electrical specifications define the maximum allowed insertion loss of specific sections of the system. This definition is referred to as insertion loss budget and is usually defined in terms of differential insertion loss at predefined frequency points. The initial parameter to check, when designing a high-speed serial interface board, is the specific interface loss budget.

The insertion loss and loss budget refer to the degradation of a high-speed signal propagating through a printed circuit board trace or a stacking cable. The loss is mainly affected by:

- Material and geometrical dimensions effects:
  - Skin effect
  - Dielectric loss
  - Copper surface roughness
- Other signal integrity effects: Impedance mismatches and discontinuities

To allow an end-to-end reliable link operating at a desired bit error rate (BER), the receiver sensitivity and jitter tolerance characteristics requirements can be determined by the following—all of which will determine the eye pattern at the receiver:

- The transmitter eye characteristics
- The interconnect loss characteristics
- Crosstalk characteristics: The crosstalk effect for loss budget calculation is explained in Section 2.4, Crosstalk, on page 21.
- Impedance mismatches, including manufacturing tolerance
- Other disturbances such as RFI

It is important to have a smooth insertion-loss curve shape over the desired frequency range. Notch-like behavior should be avoided (see Figure 1). Notch-like behavior indicates a major phase shift of the insertion loss at the specific frequency. Phase changes between different signal symbols can be a major contributor to jitter.

Figure 1: Insertion Loss Curve

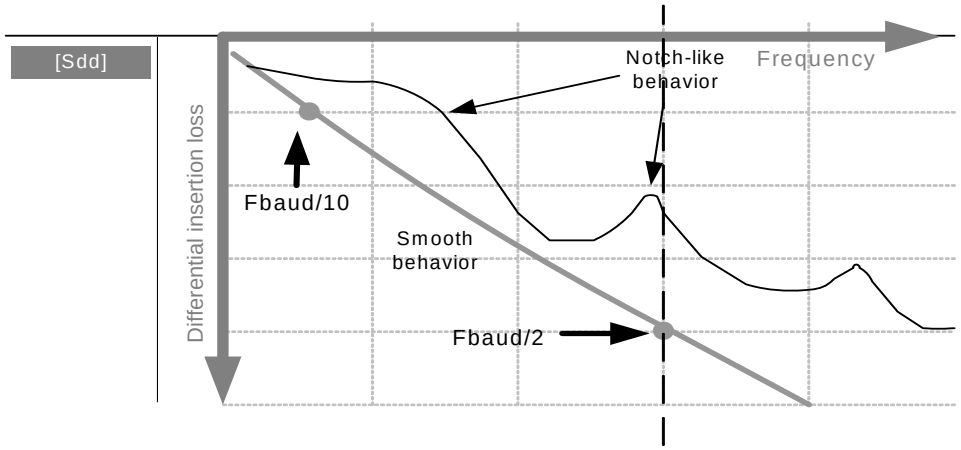

2.1.1 Implementation Notes

- Examine the specific interface loss budget definition. Calculate or simulate the end-to-end loss and take into account:
  - Insertion loss of board traces.
  - Insertion loss of any component the signal travels through, including connectors, AC coupling capacitors, and cables.
- Properly select trace geometrical dimensions and choose the suitable material to assure achieving a predefined amount of end-to-end loss.

- Wider traces are characterized by lower skin effect. Final manufactured trace width may be narrower than the planned and may influence skin effect. Monitor the actual manufactured trace width. Recalculate the loss using the actual manufactured width and verify that it meets the insertion-loss budget at the predefined frequencies.
- The µ-strip (outer layer) is characterized by lower dielectric loss compared to the strip line (inner layer) and usually requires fewer vias. However, the µ-strip is usually characterized by higher copper surface roughness and high-impedance fluctuations. Consider using core material for µ-strips.
- Monitor the material loss coefficient and the dielectric constant parameter during manufacturing. Verify that there are a consistent loss coefficient and dielectric constant over material batches.
- Verify that the insertion-loss curve shape is smooth.
  - Avoid stubs on high-speed serial interface traces and limit the maximal via stub length.
  - Minimize the number of vias along a high-speed serial link differential pair.
  - Minimize discontinuities caused by trace impedance and parasitic capacitance.

Note

For Electromagnetic interference (EMI) issues related to using µ-strip routing, see *AN-216, EMI Guidelines*, Doc. No. MV-S300935-00.

2.2 Inter-Symbol Interference (ISI)

ISI is defined as the difference in [dB] between the insertion loss magnitude at the following two frequencies, examining the first harmonic (see Figure 2):

Frequency of longest symbol or longest repetitive bit sequence: 0.5*Fbaud/(Maximal amount of consecutive similar bits of the data), for example, for the 8–10 bit data scrambling protocol, the lowest frequency is Fbaud/10, and for longer repeating bit sequences the frequency may be as low as Fbaud/Bit_Sequence_Length. Check the specific standard specifications or the specific design guidelines section for the scrambling protocol.

Frequency of shortest symbol: 0.5*FBaud

Note The eye pattern is also affected by harmonics higher than the first harmonic.

Figure 2: Typical Trace Differential Amplitude Insertion Loss and ISI

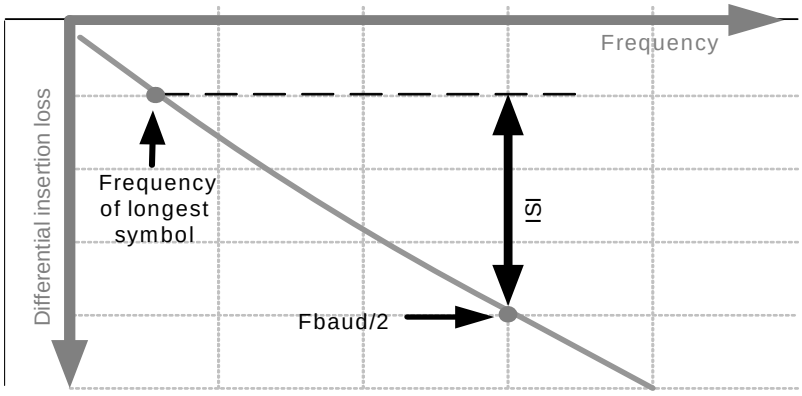

Due to the difference in attenuation magnitude of different frequencies and/or the difference in group delay (phase), data symbols of different lengths endure different amplitude degradation. The different symbol amplitudes caused by the ISI are a major cause of jitter, as well as, vertical eye opening degradation (see Figure 3 and Figure 4).

Figure 3: Difference in Attenuation Due To ISI - Time Domain Influence

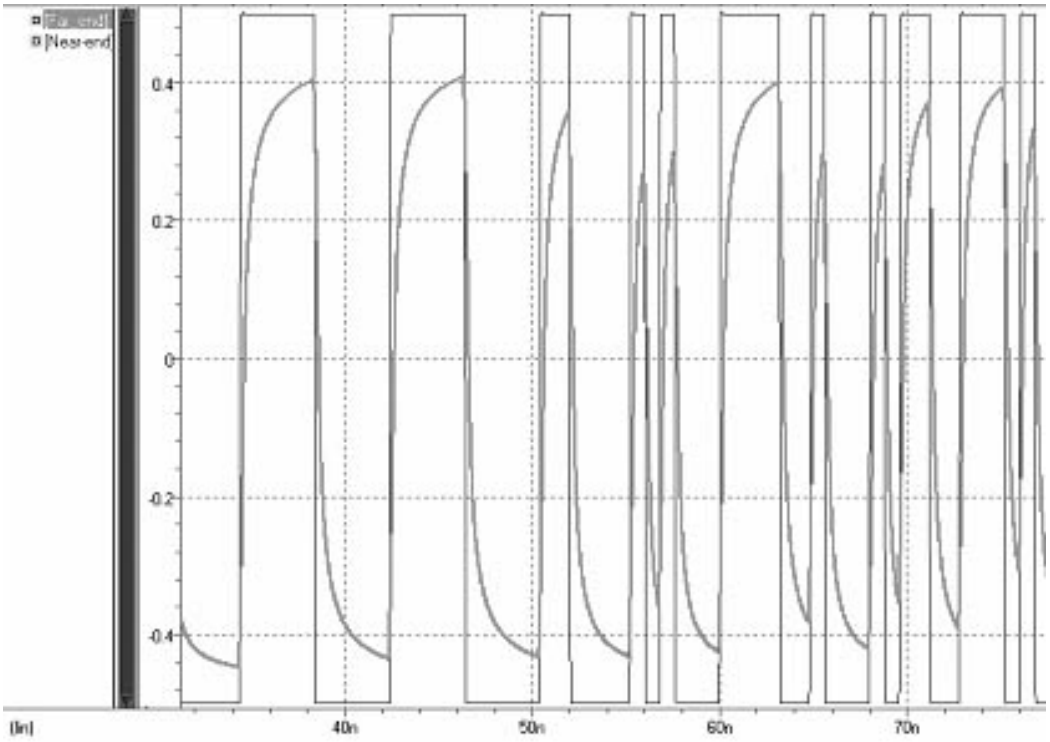

Figure 4: ISI Originated Jitter - Eye Pattern

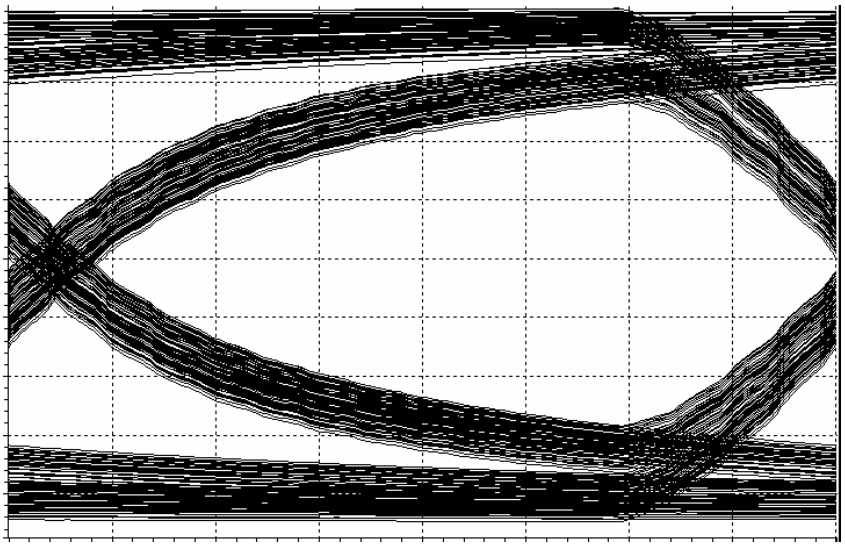

Every high-speed serial interface has a maximal allowed ISI budget. The ISI allowed amount can be a function of the receiver jitter tolerance, as well as, the availability of ISI cancellation equalization methods as pre-emphasis on the transmitter side and feedforward equalizer (FFE), such as, continuous time linear equalizer (CTLE) and/or decision feedback equalizer (DFE) on the receiver side.

2.2.1 Implementation Notes

- Calculate/simulate the interconnect insertion loss and verify that it meets both the allowed ISI amount and the maximal allowed loss at a certain frequency.

  Example: In the PCI Express design guidelines, 6.6 dB is allowed for ISI loss. Calculating the maximal ISI contribution of 0.33 dB/inch for a specific trace geometry and material results in a maximal trace length of 20 inches.
- To meet the required trace impedance, as well as, to minimize discontinuities, properly select the traces' geometrical dimensions and material loss characteristics to comply with these requirements. Note that vias along the trace may introduce loss and discontinuities that need to be subtracted from the total loss.

2.3 Placement of Devices and Connectors on the Board

Placement of devices and connectors on the board can considerably influence the performance of the system.

Placement of the high-speed serial interface device can determine the trace length, as well as, influence the proximity to noise sources.

2.3.1 Implementation Notes

When starting a new design, follow the placement guidelines described below. Prior to routing, place the high-speed serial interface devices and other layout sensitive components on the board. Refer to the following constraints:

- Estimate the expected trace length. Calculate/simulate/extract its loss and compare to the insertion loss budget.
- In cases where a high-speed serial interface is destined to travel through a connector to a backplane or a cable (both stacking or as part of an interface definition) consider placing the high-speed serial interface devices close to the connector, to minimize the loss spent for board trace attenuation.
  NOTE: Take into account that connectors are capacitive in nature (see Section 2.7, Capacitive Discontinuities).
- Consider the location of noise sources (such as DC/DC converters, magnetic devices, and wide parallel interface buses) relative to high-speed serial interface devices. Usually, the differential noise induced by the noise sources should be -46 dB relative to its own amplitude.
  NOTE: The Rx amplitude value near the receiver can be very low, and there may be additional total power sum crosstalk requirements for the specific interface.
- Sketch the path of the high-speed serial differential interfaces. Verify that they have continuous reference planes. Routing through a high pin-count through-hole device/BGA device with a via matrix creates a Swiss-cheese-like reference plane. Place the devices accordingly to provide a continuous reference plane.
- For most designs, Marvell® recommends routing the differential interfaces referenced to GND to allow the shortest, lowest inductance return path.

2.4 Crosstalk

High-speed serial interfaces can have low differential amplitudes (mostly at the receiving side). Therefore they are susceptible to induced noise originating from crosstalk. Crosstalk can induce

differential and common mode noise. Common mode noise can mainly harm EMI compatibility, while differential crosstalk-induced noise can increase the amount of jitter as well as the eye pattern opening and harm the interface performance.

The terminology used in this specification defines:

- **Aggressors:** The signals inducing the crosstalk noise
- **Victim:** The differential high-speed serial pair exposed to crosstalk

Consider the following, while performing analysis to minimize the amount of crosstalk:

- Near-end crosstalk
- Far-end crosstalk
- Stack-up with higher probability of crosstalk issues
- Return path discontinuities
- 3-dimensional structures with a high susceptibility to crosstalk

2.4.1 Near-end and Far-end Crosstalk

Near-end crosstalk needs to be considered mainly between transmitting aggressor signals (that are characterized by short rise time and high amplitudes) and receiving victim signals. To minimize near-end crosstalk, perform analysis and increase the coupling distance between traces, as required to achieve the desired level of crosstalk.

Far-end crosstalk is accumulative over trace overlapping distance. To minimize far-end crosstalk, minimize the coupling length of traces. Internal layers (homogenous strip-lines), when designed properly, can dramatically decrease far-end crosstalk.

2.4.2 Stack-up with Higher Probability of Crosstalk Issues

Consider the following two principles of induced board crosstalk:

Coplanar coupling: Produced by having the victim geometrically close to any aggressor relative to the return path reference plane. Layers characterized by high-dielectric distance between the signal layer and the closest reference plane may require increased separation from the closest coplanar aggressor. Routing high-speed serial differential interface traces on these layers should be examined and the minimal distance to the aggressor calculated according to the desired amount of crosstalk.

Tandem coupling: Stack-up layer order of: (Reference plane - Signal - Signal - Reference plane) is susceptible to tandem coupling.

The geometric dimensions H1, H2, H3 and the separation determine the amount of crosstalk induced onto the victim pair of traces from the aggressor pair of traces (see Figure 5). Increase H2 and the separation as much as required to achieve the desired level of crosstalk.

Figure 5: Stack-up Cross-section with a High Probability of Crosstalk

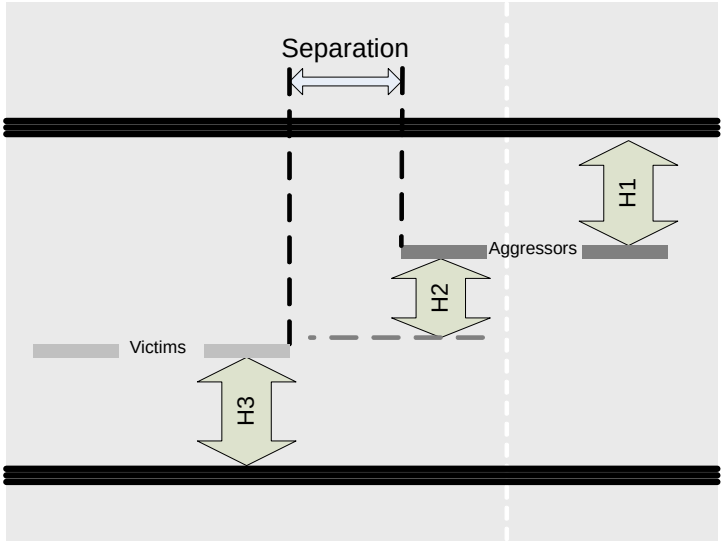

The aggressor pair can induce crosstalk on the victim pair, depending on the geometry and the dielectric characteristics.

Coplanar coupling on a two-layer PCB

Routing differential traces on a two-layer PCB can have a complicated return current path creating electromagnetic coupling between adjacent traces.

When using a 2-layer board, a return path plane must be created on the same layer (coplanar) as the SERDES traces. Provide a coplanar return current path by adding coplanar GND strips, as well as GND shielding between sensitive high-speed SERDES traces and any possible aggressor signal.

2.4.3 Return Path Discontinuity

Where return path discontinuity occurs, the return path of a signal may form a big loop to the closest return path closure point. Signals residing inside the current closure loop may suffer high crosstalk.

Return path discontinuity may occur where a signal changes routing layer and reference plane, or in cases where the reference plane is non-continuous for any possible reason.

2.4.4 3-Dimensional Structures with High-Crosstalk Susceptibility

Structures that form a non-homogeneous electromagnetic medium, such as vias, have a high susceptibility to electromagnetic induced crosstalk. Vias can also induce a high amount of crosstalk.

2.4.5 Implementation Notes

- Calculate/simulate the amount of near-end differential crosstalk, Tx to Rx, and ensure that it is below -46 dB, and/or meets the required standard.
- Calculate/simulate the amount of far-end differential crosstalk and ensures that it is meets the required standard.
- Look for coplanar coupling, especially in cases of a distant reference plane, or two-layer PCB routing.
- Calculate/simulate cross-layer coupling between tandem traces.
  - Verify that the crosstalk value meets the far-end and near-end specified value.
  - Marvell® recommends using adjacent signal layers to route perpendicular direction traces.
  - It is essential to define the minimum separation that will ensure the desired amount of crosstalk.
- Look for possible return path discontinuities.
- Provide a means of supplying return path closure (vias/capacitors)
- Verify that no victim high-speed serial pair resides within the return current path closure loop.
- Maintain a differential-via-pair to differential-via-pair separation of at least 1.5*Via_Length, or place GND shielding vias between the two pairs.

2.5 Return Path Continuity

To ensure signal quality, it is vitally important to supply high-speed differential serial interface traces with a continuous return current path.

Pay special attention to locations where the signal changes routing layer and therefore may change reference layer.

Routing over non-continuous planes increases return current path inductance. This increases the exposure to crosstalk with other signals, as well as, degrades the rise time and can increase ISI. Crosstalk in this case may induce both differential and common mode noise. Return path discontinuity increases the amount of common mode noise. Marvell® recommends routing high-speed differential serial interface traces referenced to a continuous ground plane.

2.5.1 Routing High-speed Serial Interface Traces Referenced to a Non-continuous Reference Plane

If this routing cannot be avoided, it is possible to moderate the effect of the discontinuity by adding bypass capacitors near the planes' cross-points, between the two reference planes. It is essential to supply enough bypass capacitors and calculate their equivalent series inductance (including via inductance).

Using capacitors does not completely overcome the effect of the return path discontinuity, therefore it is not an optimal solution.

Figure 6: Non-continuous Reference Plane Return Current Path Influence

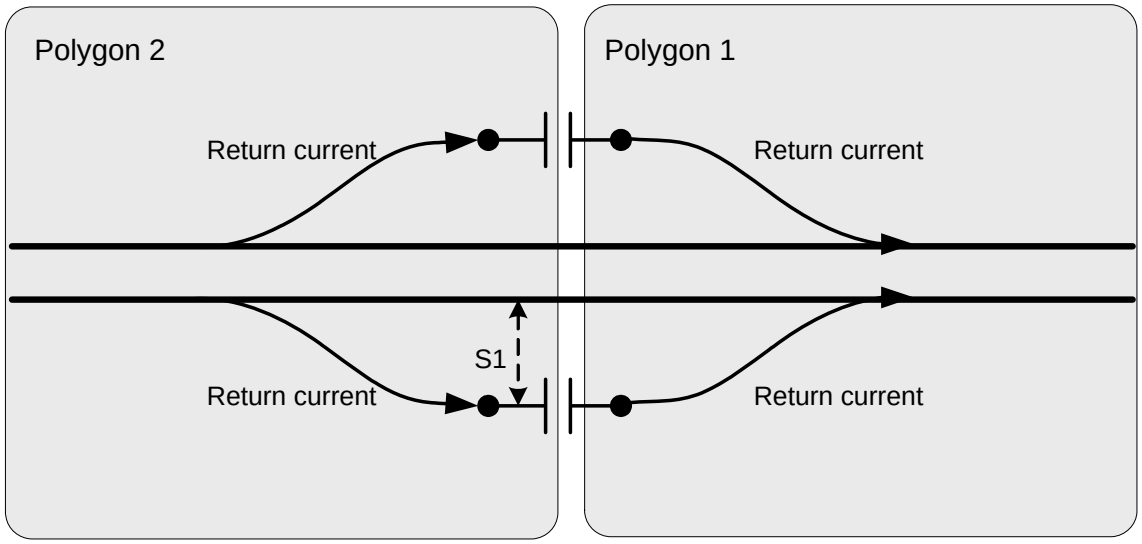

Figure 7: Layer Transfer

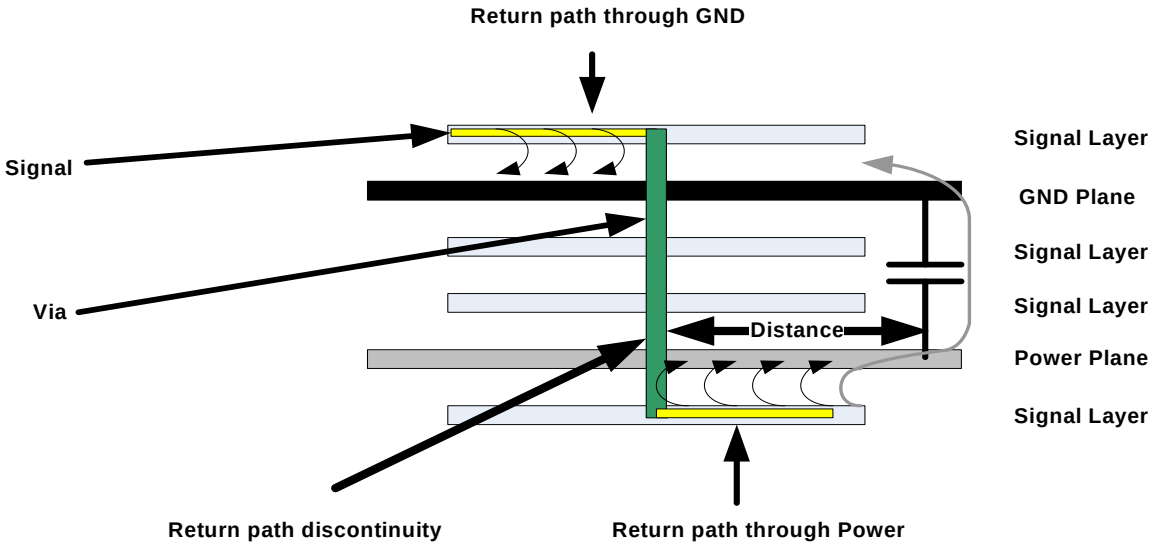

2.5.2 Implementation Notes

- Route high-speed traces referenced to continuous GND planes.
- Avoid routing over non-continuous planes.
- As much as possible, avoid changing the routing layer.
- Provide the means of supplying return path closure (vias) in cases of reference layer change (see Section 2.9, Via Structures, on page 30).
- Only if changing reference from one PWR/GND to another cannot be avoided:
  - Add bypass capacitors in cases of reference change.
  - Place these capacitors between the two reference planes, as close as possible to the discontinuity location (reference plane change point / polygon edge).
  - Marvell® recommends using a 100 nF capacitor, characterized by low equivalent serial inductance value (ESL) and low equivalent serial resistance value (ESR).
  - The inner-pair skew at the reference change point should be less than 50 mil (if specified, take into account the package trace length skew).
  - Where the reference plane is continuous, but other than ground, supply the capacitors near both ends of the reference plane.
  - If the power reference plane accompanies the signal up to a device, supply the capacitors as close as possible to the device pins and as close as possible to the far-end point of reference discontinuity (may be near a connector or a peer device).
  - Verify a maximum loop inductance of 1 nH through the 100 nF capacitor referenced to the return path reference change point (including via inductance, capacitor ESL, and plane loop inductance) per high-speed differential pair of traces.
  - In the case of traces routing to a connector, take into account that you will need to supply return path capacitors up to 100 mil from the connector and verify the loop inductance.

2.6 Target Routing Impedance

The main consideration of the SERDES traces impedance matching is to achieve the differential impedance stated in the specific interface guidelines section {usually 100 ohm (50 ohm odd mode impedance)}.

Loosely coupling the differential SERDES routing traces allows single-ended impedance matching of these traces to be close to 50 ohm. Tightly coupling the differential SERDES traces provides higher durability to crosstalk aggressors.

Sharing the same signal layer for SERDES signals and high-speed single-ended traces may cause a target-matched impedance conflict. SERDES is usually routed using wide traces. The differential impedance target of usually 100 ohm dictates that the single-ended impedance be above 50 ohm. Usually trace width is a point of interest for monitoring skin effect loss, while high-speed single-ended traces usually need a target impedance of 50 ohm single-ended at a routing trace width of 5 mils (to allow easier wide bus routing). When this conflict occurs, the main idea is to compromise the impedance matching of the interface that has a greater timing margin, while still complying with each interface's AC and DC specifications. Any compromise in the high-speed serial interface's differential impedance (have it lower than 100 ohm) must be deducted from the available trace loss.

Changing the differential impedance along the trace causes reflections and effects loss, impacting the loss budget.

Neighboring traces/planes may create a coplanar effect on trace target impedance. Inner-pair separation has a strong effect on differential impedance.

- Matching of high-speed differential traces should be done taking into account the Er (DK) at Fbaud/2. It is recommended that the impedance matching at Fb/2 not be higher than that recommended by the interface (usually 100 ohm).
- Matching traces to impedance lower than 100 ohm may be advisable where there are many capacitive discontinuities along the traces (see Section 2.7, Capacitive Discontinuities, on page 28).

2.6.1 Implementation Notes

- Keep clearance from the plane voids. Six times the height above the reference plane is recommended. Four times the height is the minimum (see Figure 8).

Figure 8: Reference Plane Clearance

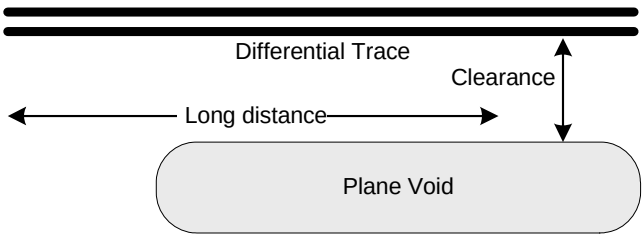

- The dielectric constant varies depending on the frequency, causing matched impedance to increase with frequency. To meet the matched impedance at Fb/2, take into consideration Er frequency dependency and the target.
- Avoid changing the trace width along the trace. If (for mechanical or geometrical reasons) this cannot be avoided (for example during a BGA escape section) Marvell® recommends calculating the new width's characteristic impedance and design this section to match the desired differential impedance, using the specific section's trace width.

- Keep adequate separation from coplanar traces/polygons. Keeping a separation equal to at least three times the height from the nearest reference plane is recommended.

2.7 Capacitive Discontinuities

Capacitive discontinuities are most common in differential high-speed serial trace routing topologies. Capacitive discontinuities may be formed by any combination of a pad and a nearby reference plane forming a plate capacitor. Long via stubs can also behave as capacitive discontinuities. Capacitive discontinuities can considerably degrade insertion loss curve behavior by having a major effect on the matched impedance of the trace.

Another main factor to address is the distance between two different capacitive discontinuities. Any two capacitive discontinuities together form an insertion loss low-pass filter. The distance between the two capacitive discontinuities and their value will determine the cut-off frequency of the filter. This phenomenon is commonly referred to as multiple-reflection loss.

2.7.1 Implementation Notes

Minimize capacitive discontinuities values by looking for any plate capacitor that can be formed between any component pad and the reference plane:

- Form a void in the adjacent reference plane, underneath the AC coupling capacitors' pads, SMT connector pads, QFP lead frame, and BGA outer perimeter pads. Verify that forming voids do not introduce coupling between signals through the void or due to the effect of the void.
- Avoid forming return path discontinuities when reference plane voids become one with via anti-pads.

Figure 9: Voids Underneath AC Coupling Capacitors' Pads

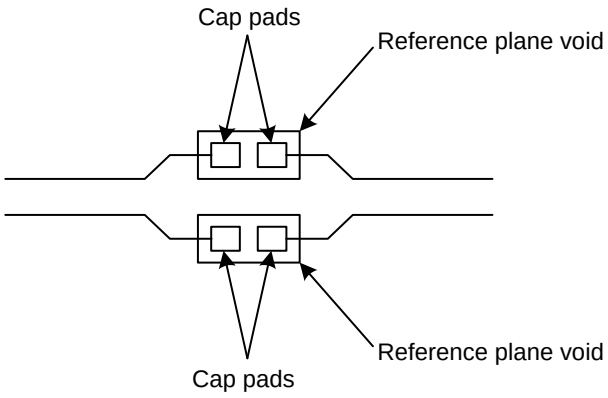

Figure 10: Voids Underneath BGA Device Ball

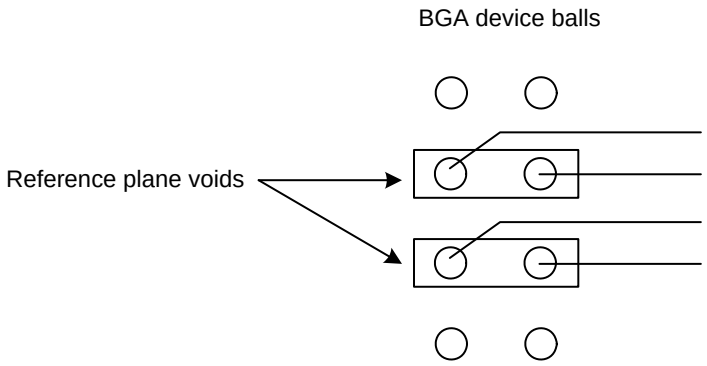

Note For low pitch devices, such as 0.65 mm pitch BGA devices, the reference plane voids should have the same size and shape as the ball pads. This will avoid forming slots in the reference plane.

Figure 11: Voids Underneath SMT Connector/QFP Lead Frame Pads

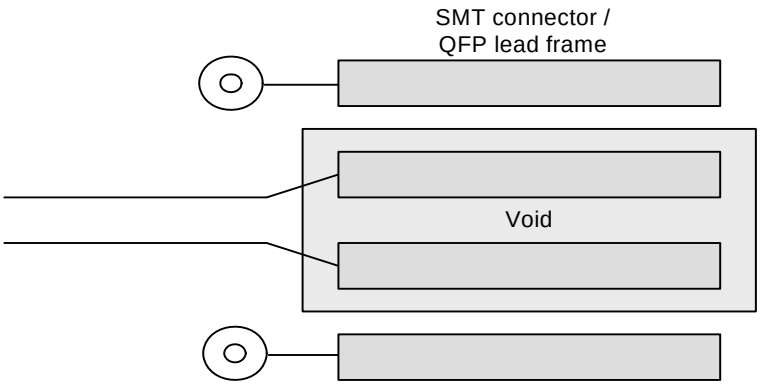

- Look at via structures:
  - Form rectangular via anti-pads to reduce via capacitance.
  - Minimize via stub length (see Section 2.9, Via Structures, on page 30).
  - Keep in mind that through hole devices' pads have large vias. Look at the specific device's recommendations for routing high-speed serial differential traces.
- Restrictions on the distance between any two capacitive discontinuities:
  - Avoid having capacitive discontinuities at a distance greater than 0.5 inches and less than 3 inches from each other.
  - In cases of low board trace losses, the minimal distance between any two capacitive discontinuities should not be less than 6 inches.

2.8 Trace Symmetry and Matching—Mode Conversion

Any non-symmetrical effect on a differential pair of traces creates a common mode signal. This effect is called mode conversion and it increases EMI emission and jitter.

2.8.1 Implementation Notes

- Match each differential pair per segment.
- Place components on a symmetrical location, along a differential pair (for them to have a symmetrical effect on both traces).
- Avoid having any non-symmetrical effect on traces that can be created from proximity to vias, traces or polygons.

2.9 Via Structures

To route signal traces between different signal layers of a multi-layer, stack-up PCB, the board designer must use a via structure. Vias must be used to connect the inner strip-line trace to the outer placed BGA/QFP pins or to a connector.

The following via structures vary in cost, as well as performance:

- Through-hole vias
- Blind vias
- Buried vias
- Back-drilled (counter-bored) vias

The following via guidelines discuss only through-hole via.

Figure 12: Via Structures

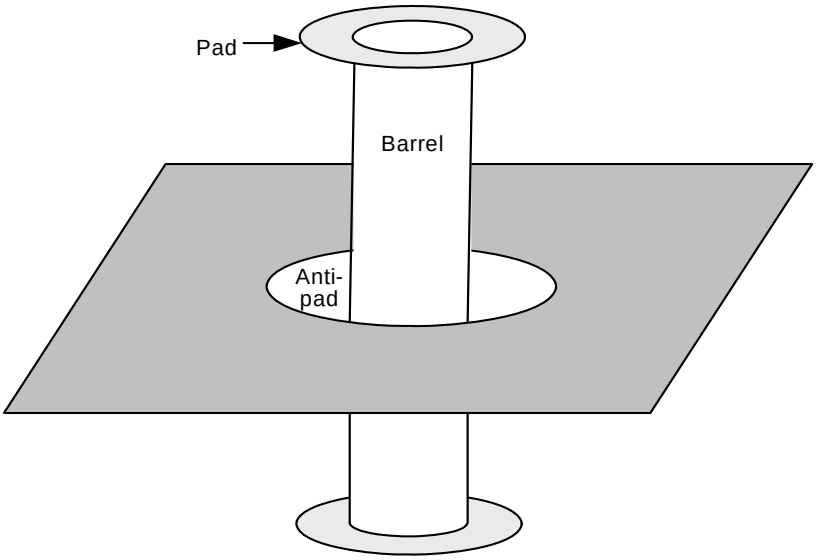

2.9.1 Removing Unused Pads in a Through-hole Via Construction

Through-hole vias are manufactured by drilling the PCB from side to side and plating the inner side of the via drill (barrel). The via can have round pads on all layers of the PCB, or only on desired

layers. The via is separated from the ground and power layers (if it is a signal via) by a clearance on those layers called anti-pad.

The via pad is used to connect the signal trace to the via on both layers; the layer on which the signal enters the via, and the layer on which the signal exits the via. It is common practice and recommended, with high-speed signals to remove the unused pads (all inner pads that are not used to route the signal) to reduce the via's parasitic capacitance. However, removing the unused pads mechanically weakens the via structure. Therefore, before removing any unused pads, consult with the printed circuit board manufacturer.

2.9.2 Effect of the Via Stub

Through-hole vias are cylindrical conductors placed throughout the height of the PCB. Pads are placed in such a way that connecting the via to any inner layer is possible, as necessary. A via can connect a trace located on the component side micro-strip to a trace located on the print side micro-strip of the board. In that case, the signal travels all the way through the via from one side of the board to the other.

Additional possibilities of using vias to connect signal traces are connecting a micro-strip to a strip-line, a strip-line to another strip-line layer. In these cases, a portion of the via is not used to transport the signal. This portion acts as an open electric circuit stub. The longer the stub, the bigger is its effect. To minimize the unwanted effect of the via stubs, Marvell® recommends utilizing the largest possible portion of the via for signal transfer (using the via to connect traces).This minimizes the length of the via stub (see Figure 13). Where only a portion of the via is utilized for signal transfer, consider using counter boring (back-drilling) of the via stub.

Figure 13: Via Stubs

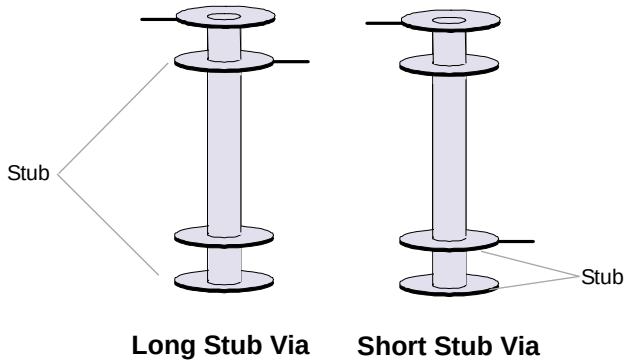

2.9.3 Via Structures Crosstalk

Board via structures can cause a high amount of crosstalk, especially in high-speed SERDES.

It is recommended to extract the crosstalk of the vias. A separation between the Tx and Rx vias of at least 1.5 time the thickness of the board, may help reduce crosstalk.

2.9.4 Implementation Notes

- Use the same number of vias for both traces of a differential pair.
- Marvell® recommends placing return path stitching vias (vias connecting the reference planes) close to signal vias, according to the recommended via structure (see Section 2.9.4.1).

Figure 14: Differential Via Structure

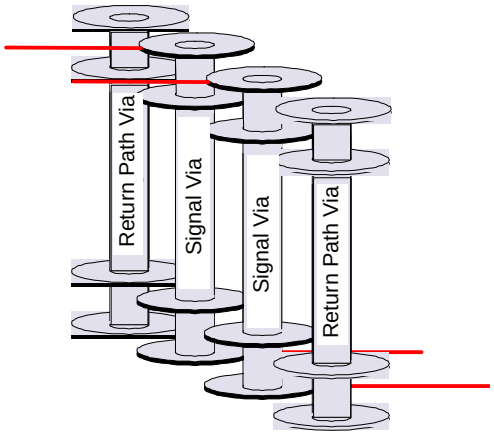

2.9.4.1 Recommended Via Structure

- Marvell® recommends designing signal vias for high-speed SERDES differential traces in a GND-Signal-Signal-GND, four-via in a row construction.
- Both stitching return path vias must maintain the same distance from the closest signal via.
- Optimized dimensions:
  - Via drill diameter: 12 mil
  - Via pad diameter: 22 mil
  - Via anti-pad diameter: 45 mil
  - Signal via to signal via center to center pitch: 40 mil
  - Signal via to GND via center to center pitch: 40 mil

Figure 15: Recommended Via Structure Dimensions

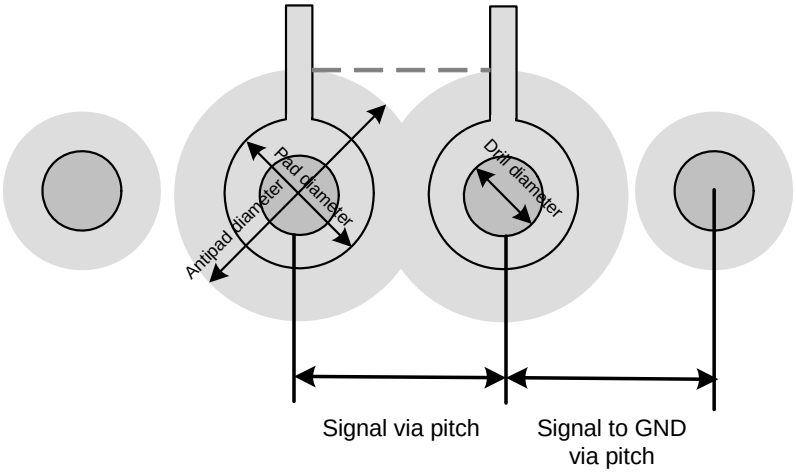

2.10 Selecting the Appropriate Components

The selection of a specific electrical/electromechanical component along a high-speed serial differential interface can have a major influence on signal quality.

The component is part of the signal path, therefore it influences the insertion loss curve. Thus care should be taken when choosing connectors, AC coupling capacitors, sockets, etc.

2.10.1 Implementation Notes

- It is preferable to use high-speed capacitors characterized by low equivalent serial inductance value (ESL) and low equivalent serial resistance value (ESR).
- During the component selection phase, pay adequate attention to the specific device's influence on the parameters specified in the introduction subsection (see introduction to Section 2.7, Capacitive Discontinuities, on page 28).
- Read the manufacturer’s design guidelines for the components, regarding high-speed serial interface routing, and follow them strictly.
- It is usually preferable to use SMT devices over through-hole devices, due to lower capacitance discontinuity and the ease of implementation of a reduction in its value.

3 32-bit SDRAM DDR2 Interface

The design guidelines presented in this section are relevant for on-board SDRAM devices using a single memory bank (one chip select) and a dual memory bank (two chip selects). The configurations that are listed below are:

- Two memory devices, each with 16-bit-wide on-board routing
- Four memory devices, each with 8-bit-wide on-board routing, using one chip select signal
- Four memory devices, each with 16-bit-wide on-board routing, using two chip select signals

Note For SDRAM power management guidelines, see Section 17, Power Management Unit (PMU) Board Design Guidelines, on page 121.

3.1 Interface Connectivity

Note The interface connectivity and routing guidelines apply to an SDRAM interface with a maximum frequency of 400 MHz (DDR 800).

3.1.1 Signal Groups

Table 3 lists the signal groups and indicates which signals are differential pairs.

Table 3: Signal Groups

Group Name	Sub-Group Name	Marvell® Device Signal	Comments
Byte[0]	Data[7:0], DM[0]	DQ[7:0], DM[0]	
	Data Strobe[0]	DQS[0]/DQSn[0]	Differential pair
Byte[1]	Data[15:8], DM[1]	DQ[15:8], DM[1]	
	Data Strobe[1]	DQS[1]/DQSn[1]	Differential pair
Byte[2]	Data[23:16], DM[2]	DQ[23:16], DM[2]	
	Data Strobe[2]	DQS[2]/DQSn[2]	Differential pair
Byte[3]	Data[31:24], DM[3]	DQ[31:24], DM[3]	
	Data Strobe[3]	DQS[3]/DQSn[3]	Differential pair

Table 3: Signal Groups (Continued)

Group Name	Sub-Group Name	Marvell® Device Signal	Comments
Address and Control	Clock	CLK_OUT[1:0]/CLK_OUTn[1:0]	Differential pair
	Address/Command	BA[2:0], A[15:0], WEn, RASn, CASn	The Address bus may have fewer than 16 signals in some devices.
	Control[0]	CSn[0], CKE[0], ODT[0]	
	Control[1]	CSn[1], CKE[1], ODT[1]	

3.1.2 Connectivity with One Chip Select 2 x 16-bit Wide Memory

Figure 16 depicts the connection for two 16-bit wide on-board memory devices.

Figure 16: 2 x 16-bit Wide Memory Devices Connected to the Controller

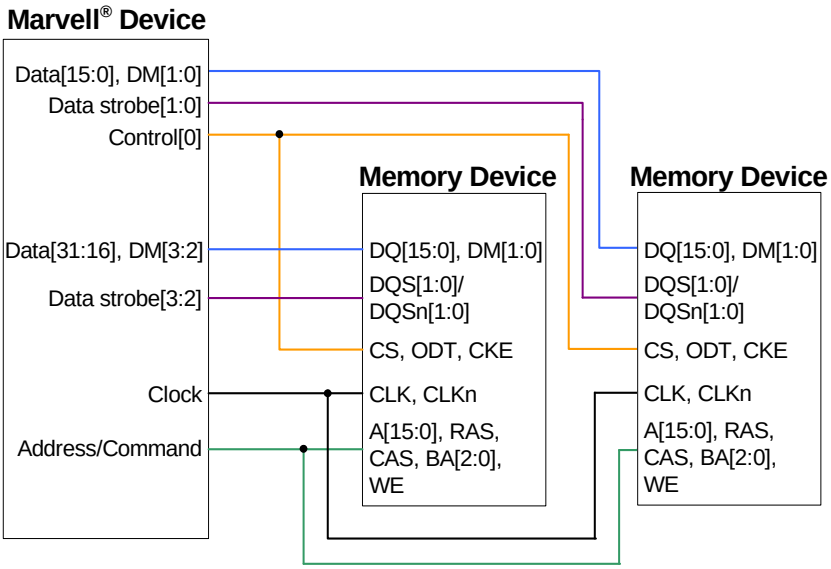

3.1.3 Connectivity with Two Chip Selects 4 x 16-bit Wide Memory

Figure 17 depicts the connection for four 16-bit wide on-board memory devices.

Figure 17: 4 x 16-bit Wide Memory Devices Connected to the Controller

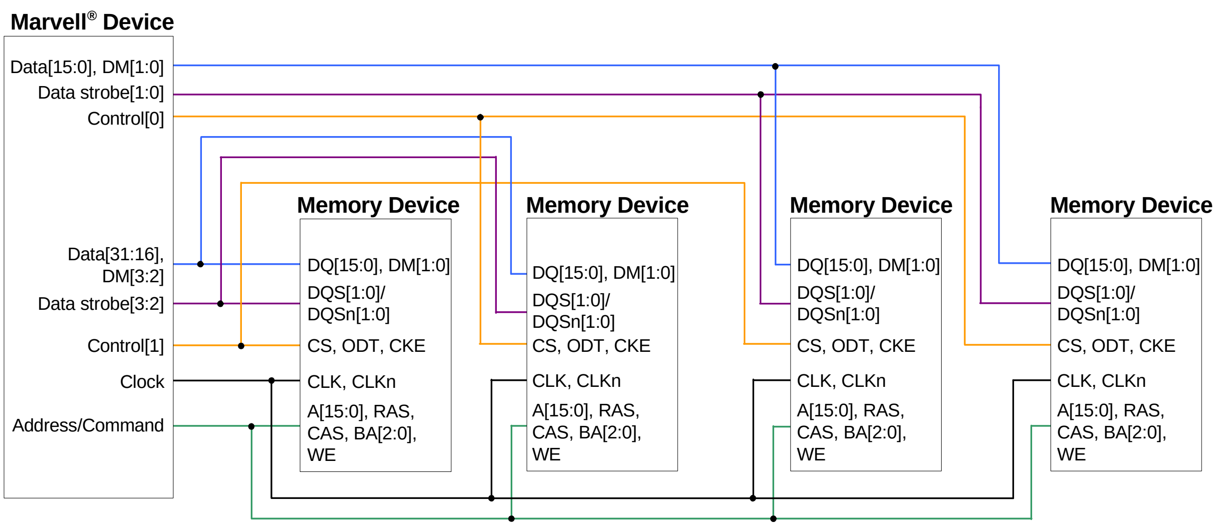

3.1.4 Connectivity with One Chip Select 4 x 8-bit Wide Memory.

Figure 18 depicts the connection for four 8-bit wide on-board memory devices.

Figure 18: 4 x 8-bit Wide Memory Devices Connected to the Controller

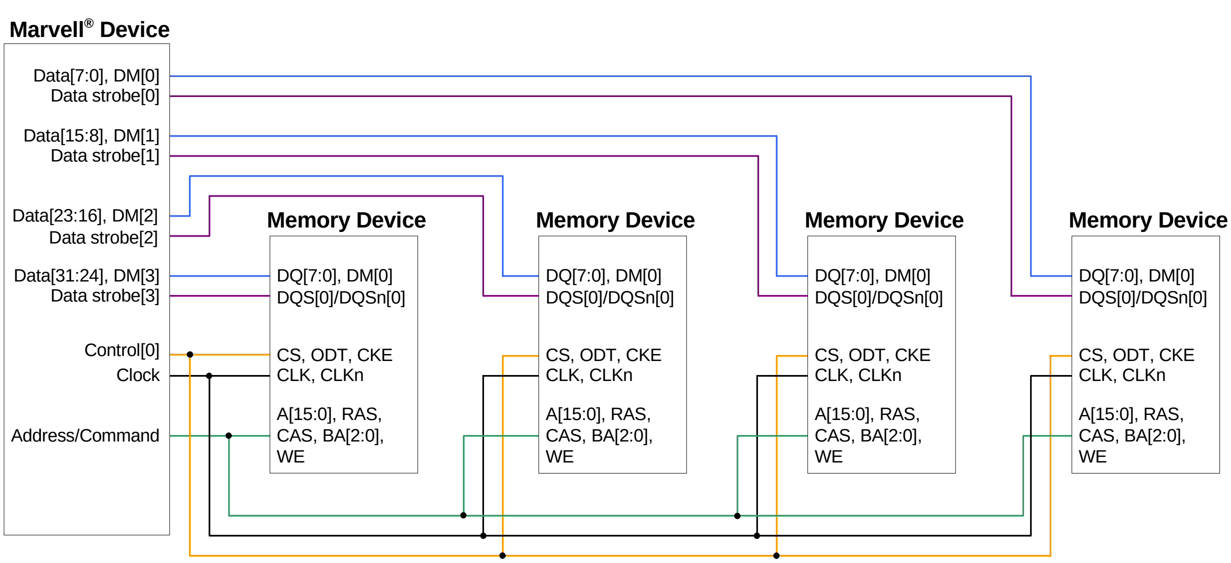

3.2 Interface Signals Layout Guidelines

- For frequencies above 300 MHz:
  - All skew requirements in Table 4 through Table 6 should be calculated with the Marvell® device's package trace length. Package trace length should be taken from the device Hardware Specifications.
  - Keep a continuous reference plane for each signal over its entire path. If the reference plane is a power supply, add a 10 nF X7R decoupling capacitor up to 200 mils from the edges of the trace.
- For frequencies up to 300 MHz, a maximum of two reference polygons is allowed. A decoupling cap is required when changing reference plane.
- If terminations are not used, the voltage swing is 2.4V rather than 1V, as defined by the JEDEC specifications. Get a waiver from DRAM vendor for violating the Vswing (max) and overshoot/undershoot specifications.
- When terminations are not required, Marvell recommends leaving optional terminations for Clock and CSn signals.
- The data signals inside the following groups can be swapped: Data[7:0], Data[15:8], Data[23:16], and Data[31:24].
- The traces of the same group should be routed with the same number of vias, or else the via delay should be taken into account.
- Data and strobe signals should be routed on same layer with the same number of vias. If these recommendations cannot be applied, only route the strobe signals on a separate layer and account for the via timing delay.
- Data and strobe signals should be routed referenced to GND.
- A swap can be performed between any two of the data groups: Byte[0], Byte[1], Byte[2], and Byte[3].
- When the number of ODT/CKE control signals, in the device, are not equal to the number of chip select (CSn) signals, treat these specific control signals as Address/Command signals.
- If the device has multiple output clocks, they may be used to simplify the board layout.
- The DDR CKE signals must have a 4.7 kΩ pulldown resistor in designs, where VTT termination on CKE is not required.

3.2.1 Daisy Chain Topology—One Chip Select 2 x 16-bit Wide Memory

Daisy chain topology is the preferred configuration; it is mandatory for the following configurations:

- Frequency is above 300 MHz
- To enable a one device assembly on a two device board

Figure 19 depicts the connection of two 16-bit wide memory devices with address and controls in a daisy chain topology.

Figure 19: 2 x 16-bit Wide Memory Topology—Address and Control in Daisy Chain

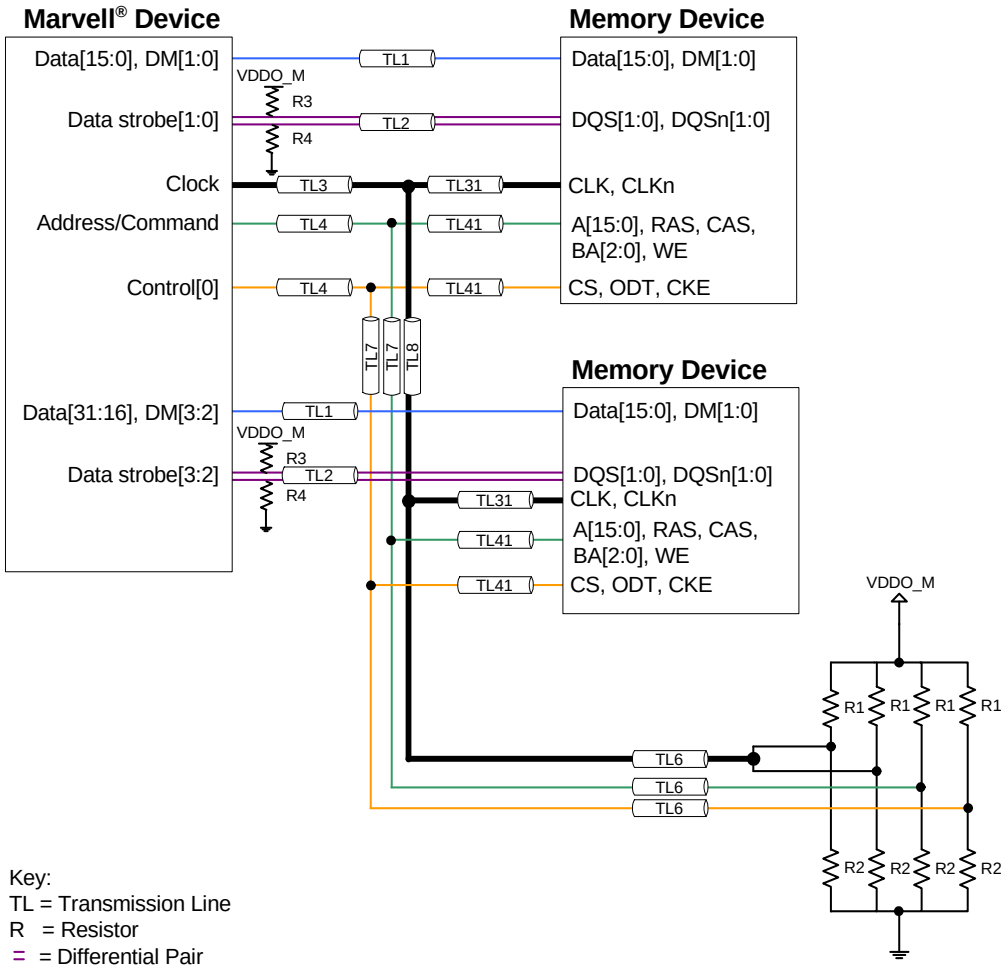

Table 4 lists the routing constraints when using a daisy chain topology.

Table 4: Routing Constraints when Using Daisy Chain Topology

Parameter	Min	Max	Tolerance / Skew	Units	Notes
TL[x] internal-group separation (dimensions)	2			Signal width	1, 2, 3
TL[x] internal-group separation (crosstalk)		-20		dB	1, 2, 3
TL[x] to TL[y] cross-group separation (dimensions)	3			Signal width	1, 2, 3
TL[x] to TL[y] cross-group separation (crosstalk)		-26		dB	1, 2, 3
TL1, TL2 length		4		inch	
DQ to DQS skew TL2–TL1	-50	50		mil	
TL2 in pair skew, TL3+TL31, TL3+TL8+TL31 in pair skew (differential pairs)	-10	10		mil	
DQS0 (first chain) to Clock skew: TL2 – (TL3+TL31)	-1.5	1.5		inch	
DQS1 (second chain) to Clock skew: TL21 – (TL8+TL3+TL31)	-1.5	1.5		inch	
Address/Command/Control to Clock skew (first device): (TL4+TL41) – (TL3+TL31)	-1000	1000		mil	
Address/Command/Control to Clock skew (second device): (TL7+TL4+TL41) – (TL8+TL3+TL31)	-1000	1000		mil	
TL31 length		Min (0.2*TL3, 0.5)		inch	
TL41 length		Min (0.2*TL4, 0.5)		inch	
TL6 length		2		inch	4, 5
R1, R2	100		1%	ohm	4, 5, 6, 7, 8
R3, R4	1			kohm	9
TL[x] impedance (all traces)	50		10%	ohm	10
TL2, TL3, TL8 differential impedance	100		10%	ohm	10, 11
TL2, TL3, TL8 single-ended impedance		60	10%	ohm	10

NOTES

1. If the design cannot meet the above requirements for length and separation, simulation can be performed to ensure proper isolation between signals.
2. The separation is defined from edge-to-edge of two traces.

3. Separation is required between any signals that are not in the same group.
4. Termination is required when the Marvell device is a QFP package in the following cases:
   - The address toggles in 1T mode.
   - If address toggles in 2T mode, termination is needed only on Clock signals.
5. If no termination is required, do not route TL6.
6. The trace length, connecting between the termination resistors R1 and R2, should be less than 100 mils.
7. R1 and R2 may be replaced by a single 50-ohm resistor connected to a VTT power source. In this case, the VTT power should be filtered, to ensure a quiet power source. The VTT power source should have a nominal voltage of VREF±40mV.
8. Place a bypass capacitor between VDDIO and GND, and close to the Thevenin terminations. The ratio of capacitors to terminations is 1:3.
9. The connection from DQSn to R3 (pull up) and R4 (pull down) must be less than 500 mils from the device.
10. The tolerance for the matched impedance value is the manufacturing tolerance. The design tolerance should be kept to a minimum.
11. The target impedance priority is to match the differential impedance of these signal traces and adjust the single-ended impedance accordingly.

3.2.2 Address and Control in Y Topology—One Chip Select 2 x 16-bit Wide Memory

This topology is allowed for frequencies below 300 MHz. The preferable configuration is the Daisy chain that is described in Section 3.2.1, Daisy Chain Topology—One Chip Select 2 x 16-bit Wide Memory, on page 38.

Figure 20 depicts the connection of two 16-bit wide memory devices with address and controls in a Y topology.

Figure 20: 2 x 16-bit Wide Memory Topology—Address and Control in Y Topology

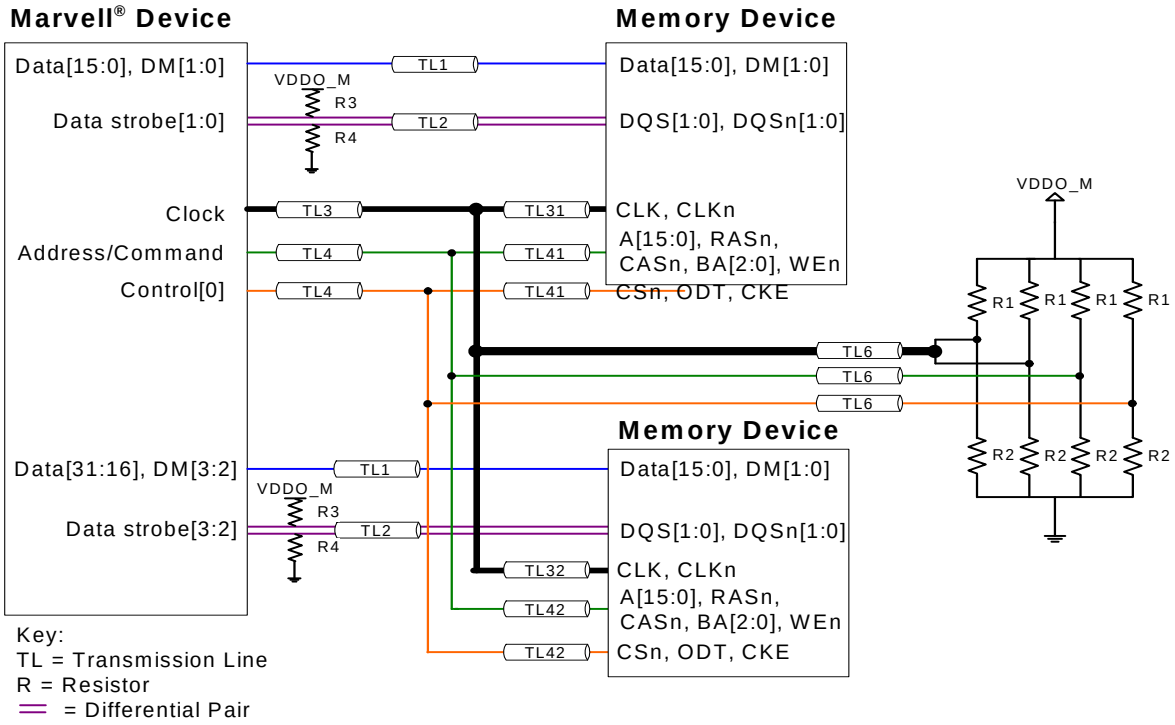

Table 5 lists the routing constraints when using two 16-bit wide memory devices with address and control in a Y topology.

Table 5: Routing Constraints when Using 2 x 16-bit Wide Memory Topology—Address and Control in Y Topology

Parameter	Min	Max	Tolerance / Skew	Units	Notes
TL[x] internal-group separation (dimensions)	2			Signal width	1, 2
TL[x] internal-group separation (crosstalk)		-20		dB	1, 2
TL[x] to TL[y] cross-group separation (dimensions)	3			Signal width	1, 2, 3
TL[x] to TL[y] cross-group separation (crosstalk)		-26		dB	1, 2, 3
TL1, TL2 length		4		inch	
DQ to DQS skew TL2 – TL1	-50	50		mil	
TL2 in pair skew, TL3+TL31, TL3 +TL32 in pair skew (Differential pairs)	-10	10		mil	
TL31 – TL32, TL41 – TL42 (skew between the Y branches)	-0.5	0.5		inch	
DQS0 (first chain) to Clock skew: TL2 – (TL3+TL31)	-1.5	1.5		inch	
DQS1 (second chain) to Clock skew: TL2 – (TL3+TL32)	-1.5	1.5		inch	
Address/Command/Control to Clock skew (first device): (TL4+TL41) – (TL3+TL31)	-1000	1000		mil	
Address/Command/Control to Clock skew (second device): (TL4+TL41) – (TL3+TL32)	-1000	1000		mil	
TL31, TL32 length		Min (0.2*TL3, 0.5)		inch	
TL41, TL42 length		Min (0.2*TL4, 0.5)		inch	
TL6 length		0.5		inch	5, 6
R1, R2	100		1%	ohm	5, 6, 7, 8, 9
R3, R4	1			kohm	10
TL[x] impedance (all traces)	50		10%	ohm	11
TL2 / TL3 differential impedance	100		10%	ohm	11, 12
TL2 / TL3 Single-ended impedance		60	10%	ohm	11

NOTES

1. If the design cannot meet the above requirements for length and separation, simulation can be performed to ensure proper isolation between signals.
2. The separation is defined from edge-to-edge of two traces.
3. Separation is required between any signals that are not in the same group.
4. TL2, TL3, TL31, and TL32 are pairs of differential signals.
5. Termination is required when the Marvell device is a QFP package in the following cases:
   - The address toggles in 1T mode.
   - If address toggles in 2T mode, termination is needed only on Clock signals.
6. If no termination is required, do not route TL6.
7. The trace length, connecting between the termination resistors R1 and R2, should be less than 100 mils.
8. R1 and R2 may be replaced by a single 50-ohm resistor connected to a VTT power source. In this case, the VTT power should be filtered, to ensure a quiet power source. The VTT power source should have a nominal voltage of VREF±40mV.
9. Place a bypass capacitor between VDDIO and GND, and close to the Thevenin terminations. The ratio of capacitors to terminations is 1:3.
10. The connection from DQSn to R3 (pull up) and R4 (pull down) must be less than 500 mils from the device.
11. The tolerance for the matched impedance value is the manufacturing tolerance. The design tolerance should be kept to a minimum.
12. The target impedance priority is to match the differential impedance of these signal traces and adjust the single-ended impedance accordingly.

3.2.3 Routing Topologies—Two Chip Selects 4 x 16-bit Wide Memory

Figure 21 and Figure 22 depict the topology for a 4 x 16-bit wide memory device.

Figure 21: 4 x 16-bit Wide Memory Topology for Clock/Address/Command/Control Implemented as Two Daisy Chain Branches

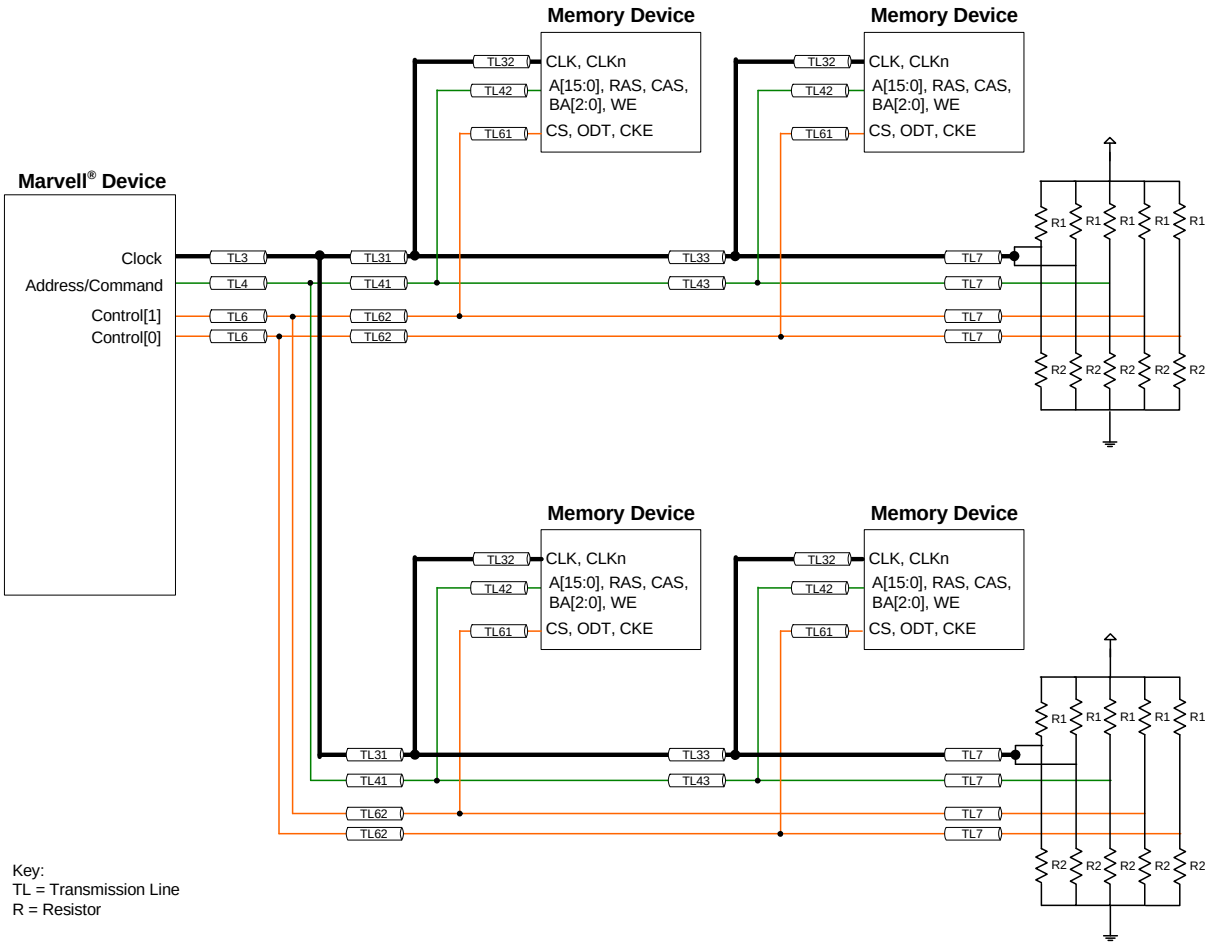

Figure 22: 4 x 16-bit Wide Memory Topology for Data

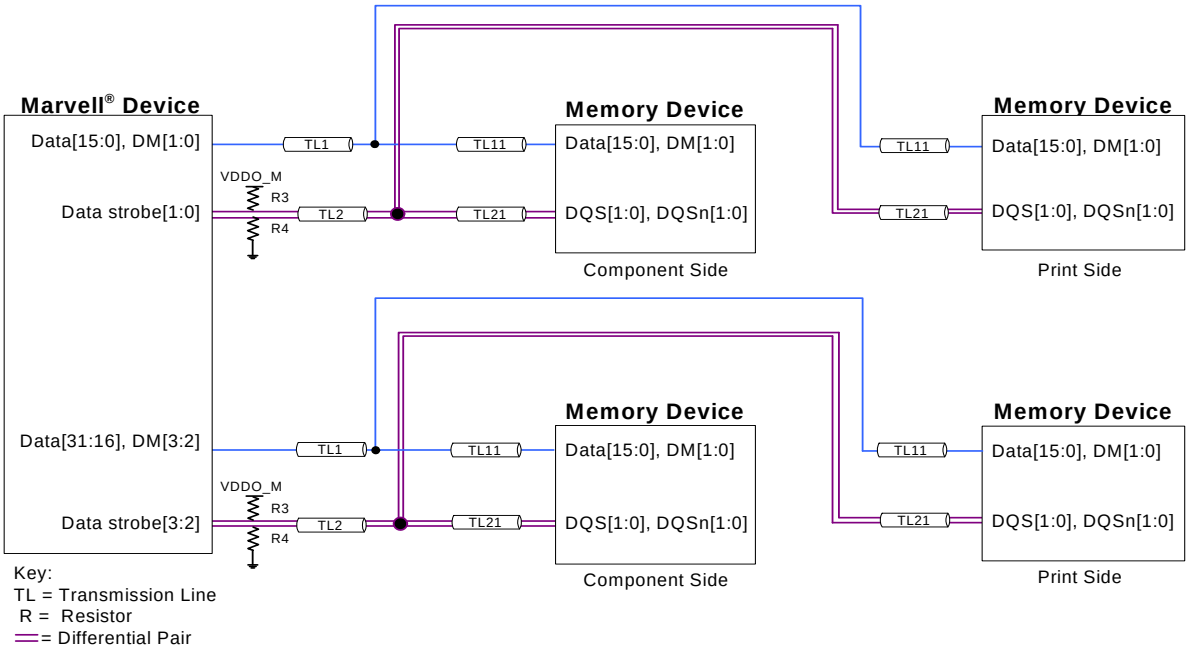

Table 6 lists the routing constraints when using four 16-bit wide memory topology.

Table 6: Routing Constraints Using 4 x 16-bit Wide Memory Topology

Parameter	Min	Max	Tolerance / Skew	Units	Notes
TL[x] internal-group separation (dimensions)	2			Signal width	1, 2
TL[x] internal-group separation (crosstalk)		-20		dB	1, 2
TL[x] to TL[y] cross-group separation (dimensions)	3			Signal width	1, 2
TL[x] to TL[y] cross-group separation (crosstalk)		-26		dB	1, 2
TL1+TL11, TL2+TL21 length		2.0		inch	
TL1+TL12, TL2+TL22 length		3.5		inch	
TL11, TL21 Length		0.5		inch	
DQ to DQS skew (TL2+TL21) – (TL1+TL11) (TL2+TL22) – (TL1+TL12)	-50	50		mil	
In pair skew (differential pairs): TL2+TL21 TL2+TL22 TL3+TL31+TL32 TL3+TL31+TL33+TL32	-10	10		mil	

Table 6: Routing Constraints Using 4 x 16-bit Wide Memory Topology (Continued)

Parameter	Min	Max	Tolerance / Skew	Units	Notes
DQS to Clock skew: (TL2+TL21) – (TL3+TL31+TL32) (TL2+TL22) – (TL3+TL31+TL33+TL32)	-1000	1000		mil	3
Address/Command to Clock skew: (TL4+TL41+TL42) – (TL3+TL31+TL32) (TL4+TL41+L43+TL42) – (TL3+TL31+ TL33+TL32)	-750	750		mil	3
Control to Clock skew: (TL6+TL61) – (TL3+TL31+TL32) (TL6+TL62+TL61) – (TL3+TL31+TL33+TL32)	-750	750		mil	3
TL32 length		0.5		inch	
TL42 length		0.5		inch	
TL7 length, TL61		0.5		inch	4, 5
R1, R2	100		1%	ohm	4, 5, 6, 7, 8
R3, R4	1			kohm	9
TL[x] impedance (all traces)	50		10%	ohm	10
TL2, TL3, TL8 differential impedance	100		10%	ohm	10, 11
TL2, TL3, TL8 single-ended impedance		60	10%	ohm	10

NOTES

1. If the design cannot meet the above requirements for length and separation, simulation can be performed to ensure proper isolation between signals.
2. The separation is defined from edge-to-edge of two traces.
3. The length of the trace to the termination of the clock should not be a part of the calculation.
4. Termination is required when Address toggles in 1T mode.
   Termination on Clock signals is always required.
5. If no termination is required, do not route TL7.
6. R1 and R2 may be replaced by a single 50-ohm resistor connected to a VTT power source. In this case, the VTT power should be filtered, to ensure a quiet power source. The VTT power source should have a nominal voltage of VREF±40mV.
7. The trace length, connecting between the termination resistors R1 and R2, should be less than 100 mil.
8. Place a bypass capacitor between VDDIO and GND, and close to the Thevenin terminations. The ratio of capacitors to terminations is 1:3.
9. The connection from DQSn to R3 (pull up) and R4 (pull down) must be less than 500 mils from the device.
10. The tolerance for the matched impedance value is the manufacturing tolerance. The design tolerance should be kept to a minimum.
11. The target impedance priority is to match the differential impedance of these signal traces and adjust the single-ended impedance accordingly.

3.2.4 Routing Topologies—One Chip Select 4 x 8-bit Wide Memory

Figure 23 and Figure 24 depict the topology for a 4 x 8-bit wide memory device.

Figure 23: 4 x 8-bit Wide Memory Topology for Clock/Address/Command/Control Implemented as Two Daisy Chain Branches

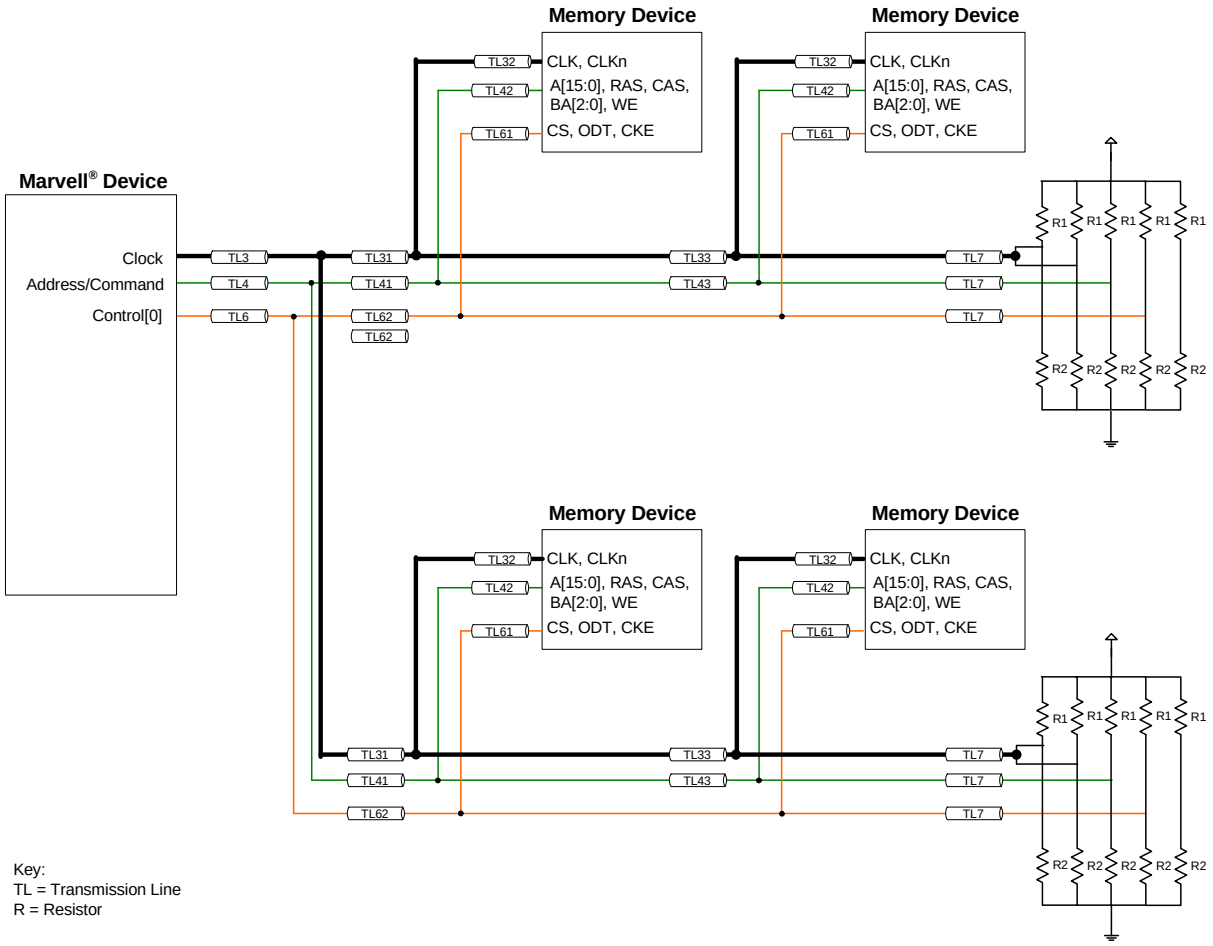

Figure 24: 4 x 8-bit Wide Memory Topology for Data

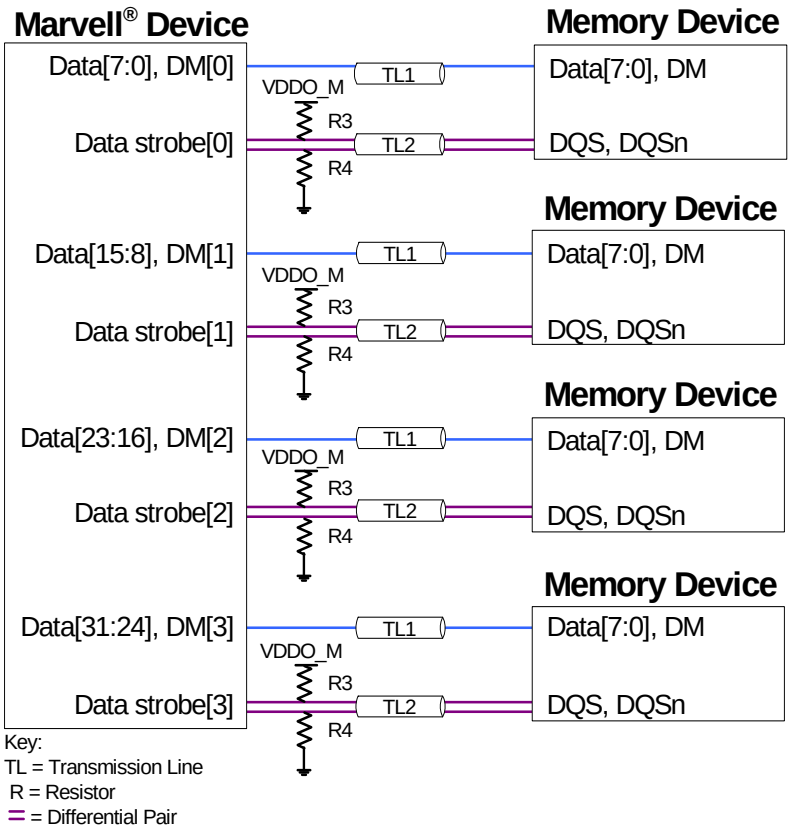

Table 7 lists the routing constraints when using four 8-bit wide memory topology.

Table 7: Routing Constraints Using 4 x 8-bit Wide Memory Topology

Parameter	Min	Max	Tolerance / Skew	Units	Notes
TL[x] internal-group separation (dimensions)	2			Signal width	1, 2
TL[x] internal-group separation (crosstalk)		-20		dB	1, 2
TL[x] to TL[y] cross-group separation (dimensions)	3			Signal width	1, 2
TL[x] to TL[y] cross-group separation (crosstalk)		-26		dB	1, 2
DQ to DQS skew TL2–TL1	-50	50		mil	
In pair skew (differential pairs): TL2 TL3+TL31+TL32 TL3+TL31+TL33+TL32	-10	10		mil	

Table 7: Routing Constraints Using 4 x 8-bit Wide Memory Topology (Continued)

Parameter	Min	Max	Tolerance / Skew	Units	Notes
DQS to Clock skew: (TL2) – (TL3+TL31+TL32) (TL2) – (TL3+TL31+TL33+TL32)	-1000	1000		mil	3
Address/Command to Clock skew: (TL4+TL41+TL42) – (TL3+TL31+TL32) (TL4+TL41+TL43+TL42) – (TL3+ TL31+TL33+TL32)	-750	750		mil	3
Control to Clock skew: (TL6+TL61+TL62) – (TL3+TL31+TL32) (TL6+TL61+TL63+TL62) – (TL3+ TL31+TL33+QTL32)	-750	750		mil	3
TL32 length		0.5		inch	
TL42 length		0.5		inch	
TL7 length, TL61		0.5		inch	4, 5
R1, R2	100		1%	ohm	4, 5, 6, 7, 8
R3, R4	1			kohm	9
TL[x] impedance (all traces)	50		10%	ohm	10
TL2, TL3, TL8 differential impedance	100		10%	ohm	10, 11
TL2, TL3, TL8 single-ended impedance		60	10%	ohm	10

NOTES

1. If the design cannot meet the above requirements for length and separation, simulation can be performed to ensure proper isolation between signals.
2. The separation is defined from edge-to-edge of two traces.
3. The length of the trace to the termination of the clock should not be a part of the calculation.
4. Termination is required when Address toggles in 1T mode.
   Termination on Clock signals is always required.
5. If no termination is required, do not route TL7.
6. R1 and R2 may be replaced by a single 50-ohm resistor connected to a VTT power source. In this case, the VTT power should be filtered, to ensure a quiet power source. The VTT power source should have a nominal voltage of VREF±40mV.
7. The trace length, connecting between the termination resistors R1 and R2, should be less than 100 mil.
8. Place a bypass capacitor between VDDIO and GND, and close to the Thevenin terminations. The ratio of capacitors to terminations is 1:3.
9. The connection from DQSn to R3 (pull up) and R4 (pull down) must be less than 500 mils from the device.
10. The tolerance for the matched impedance value is the manufacturing tolerance. The design tolerance should be kept to a minimum.
11. The target impedance priority is to match the differential impedance of these signal traces and adjust the single-ended impedance accordingly.

3.3 Power Signals

For Vref filtering and scheme, refer to Section 16, Generic Power Board Guidelines, on page 114.

3.4 Calibration Signals

Connect the calibration resistor according to the requirements described in the device Hardware Specifications.

4 32-bit SDRAM DDR3 Interface

The design guidelines presented in this section are relevant for on-board SDRAM devices using a single memory bank (one chip select) and a dual memory bank (two chip selects). Unless stated otherwise all topologies are restricted to a maximal operating frequency of the lower value of:

- 500 MHz
  or
- The frequency stated in the specific device Hardware Specifications

The configurations described below are:

- Two memory devices, each with 16-bit-wide on-board routing, using one chip select signal
- Four memory devices, each with 8-bit-wide on-board routing, using one chip select signal
- Four memory devices, each with 16-bit-wide on-board routing, using two chip select signals
- Eight memory devices, each with 8-bit-wide on-board routing, using two chip select signals

4.1 Interface Connectivity

Note The interface connectivity and routing guidelines apply to an SDRAM interface with a maximum frequency of 500 MHz.

4.1.1 Signal Groups

Table 8 lists the signal groups and indicates which signals are differential pairs.

Table 8: Signal Groups

Group Name	Sub-Group Name	Marvell® Device Signal	Comments
Byte[0]	Data[7:0], DM[0]	DQ[7:0], DM[0]	
	Data Strobe[0]	DQS[0]/DQSn[0]	Differential pair
Byte[1]	Data[15:8], DM[1]	DQ[15:8], DM[1]	
	Data Strobe[1]	DQS[1]/DQSn[1]	Differential pair
Byte[2]	Data[23:16], DM[2]	DQ[23:16], DM[2]	
	Data Strobe[2]	DQS[2]/DQSn[2]	Differential pair
Byte[3]	Data[31:24], DM[3]	DQ[31:24], DM[3]	
	Data Strobe[3]	DQS[3]/DQSn[3]	Differential pair

Table 8: Signal Groups (Continued)

Group Name	Sub-Group Name	Marvell® Device Signal	Comments
Address and Control	Clock	CLK_OUT[1:0]/CLK_OUTn[1:0]	Differential pair
	Address/Command	BA[2:0], A[15:0], WEn, RASn, CASn	The Address bus may have fewer than 16 signals in some devices.
	Control[0]	CSn[0], CKE[0], ODT[0]	
	Control[1]	CSn[1], CKE[1], ODT[1]	
	Reset	RESETn	

4.1.2 Connectivity with One Chip Select 2 x 16-bit Wide Memory

Figure 25 depicts the connection for two 16-bit wide on-board memory devices.

Figure 25: 2 x 16-bit Wide Memory Devices Connected to the Controller

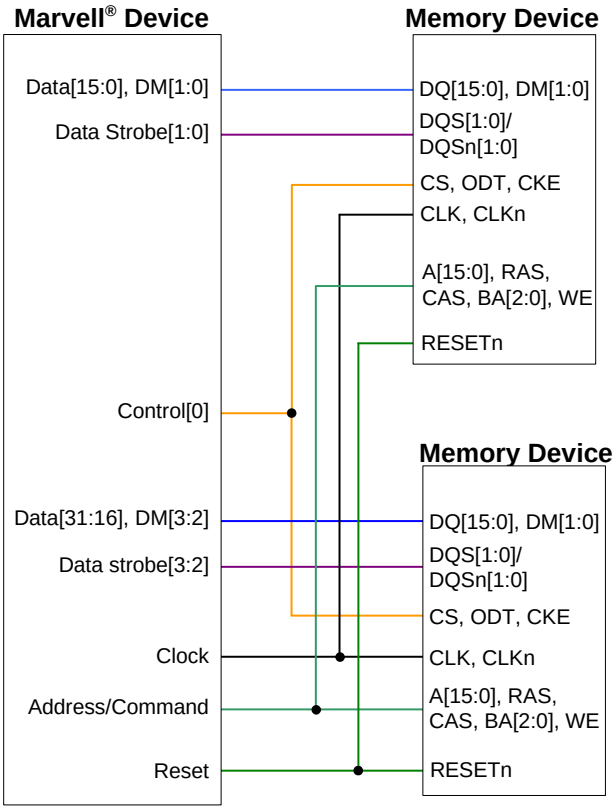

4.1.3 Connectivity with Two Chip Selects 4 x 16-bit Wide Memory

Figure 26 depicts the connection for four 16-bit wide on-board memory devices.

Figure 26: 4 x 16-bit Wide Memory Devices with Two Chip Selects Connected to the Controller

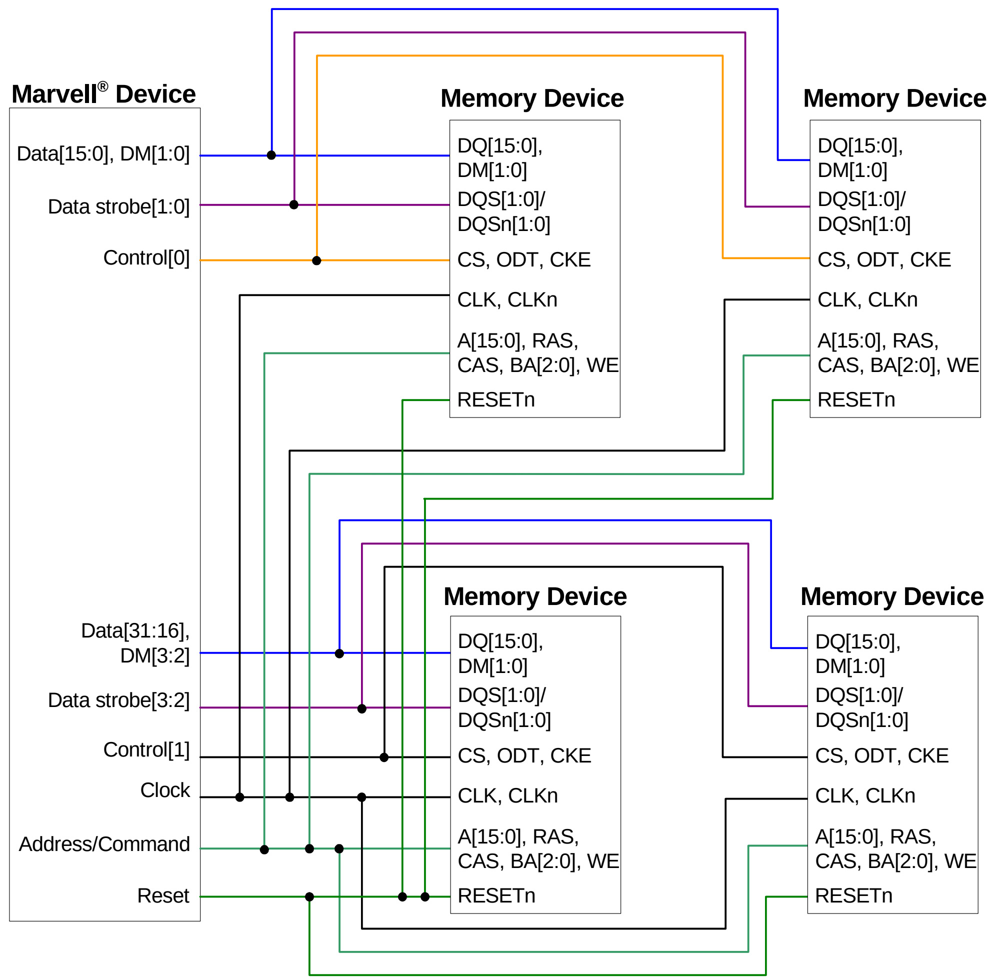

4.1.4 Connectivity with One Chip Select 4 x 8-bit Wide Memory.

Figure 27 depicts the connection for four 8-bit wide on-board memory devices.

Figure 27: 4 x 8-bit Wide Memory Devices Connected to the Controller

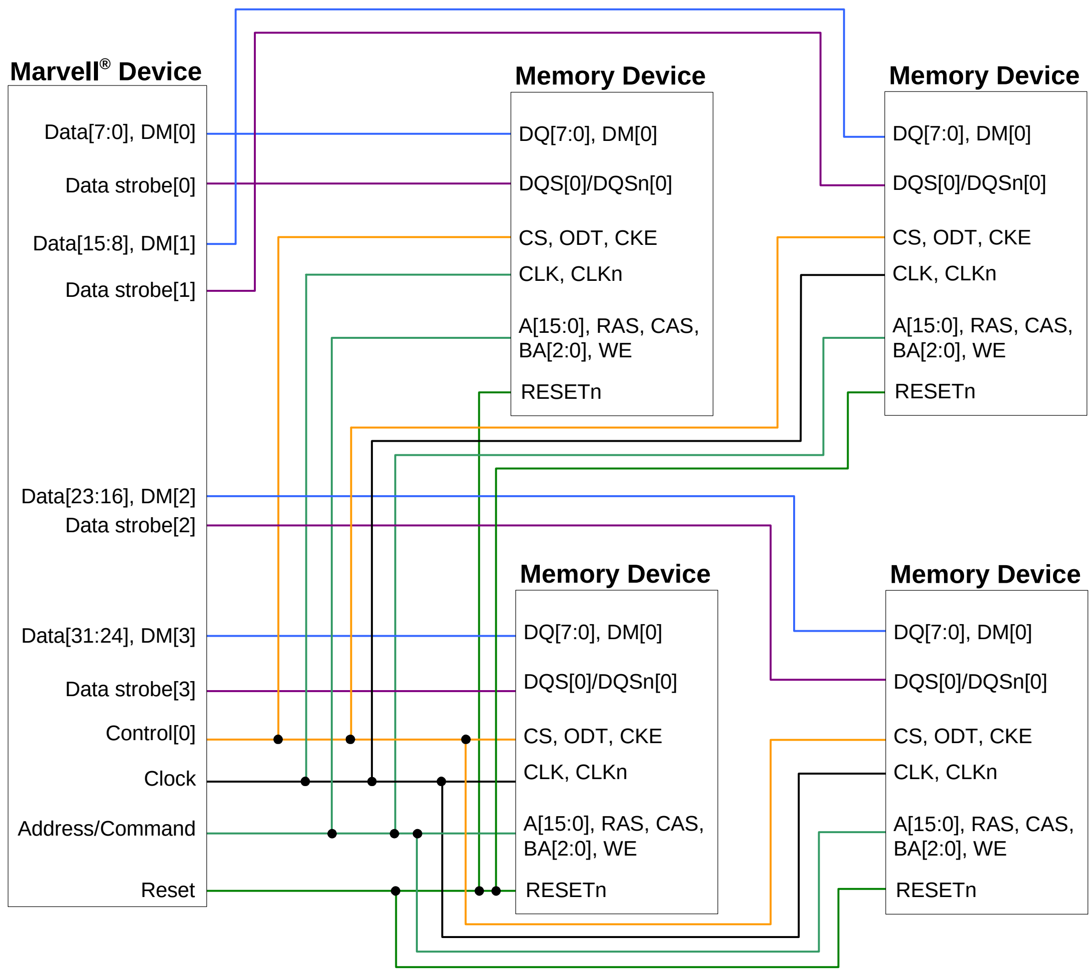

4.1.5 Connectivity with Two Chip Select 8 x 8-bit Wide Memory

Figure 28 depicts the connection for four 8-bit wide on-board memory devices.

Figure 28: 8 x 8-bit Wide Memory Devices Connected to the Controller

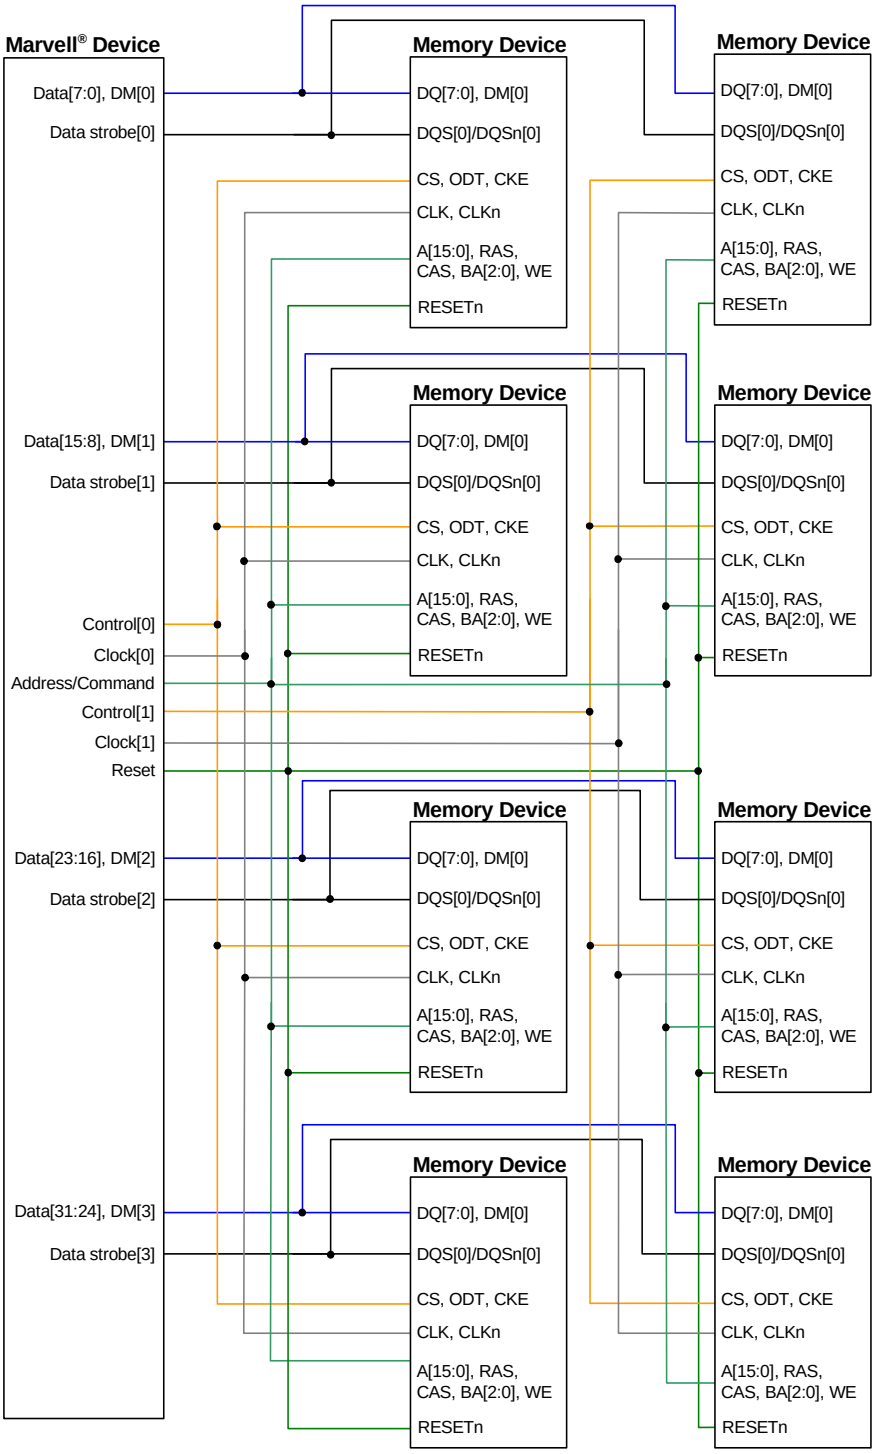

4.2 Interface Signals Layout Guidelines

- All skew requirements in Table 9 through Table 10 should be calculated with the Marvell® device's package trace length. Package trace length should be taken from the Trace Length Appendix to this design guide or the device Hardware Specifications.
- Keep a continuous reference plane for each signal over its entire path. If the reference plane is a power plane, add a 10 nF X7R decoupling capacitor up to 200 mils from the edges of the trace.
- The data signals inside the following groups can be swapped: Data[7:0], Data[15:8], Data[23:16], and Data[31:24].
- The traces of the same group should be routed with the same number of vias, or else the via delay should be taken into account.
- Data and strobe signals should be routed on same layer with the same number of vias. If these recommendations cannot be applied, only route the strobe signals on a separate layer and account for the via timing delay.
- A swap can be performed between any two of the data groups: Byte[0], Byte[1], Byte[2], and Byte[3].
- When the number of ODT/CKE control signals, in the device, are not equal to the number of chip select (CSn) signals, treat these specific control pins as Address/Command signals.
- If terminations are not used, overshoot/undershoot specifications may be violated. Get a waiver from DRAM vendor for violating overshoot/undershoot specifications.
- When terminations are not required, Marvell recommends leaving optional terminations for Clock and CSn signals.
- The RESETn is an asynchronous signal. Therefore, this signal has no AC timing limitation. The only guideline for this signal is to keep it away from all of the other signals to avoid crosstalk.
- If the device has multiple output clocks, they may be used to simplify the board layout.
- When signals are routed with vias, follow the guidelines in Section 2.9, Via Structures, on page 30. Pay particular attention to:
  - Minimizing via stubs
  - Placing a return path for ground vias
- Marvell® recommends having a 4.7 kΩ pulldown resistor on the DDR CKE signals, where VTT termination is not required.

4.2.1 Daisy Chain Topology—One Chip Select 2 x 16-bit Wide Memory

Figure 29 depicts the connection of two 16-bit wide memory devices with address and controls in a daisy chain topology.

Figure 29: 2 x 16-bit Wide Memory Topology—Address and Control in Daisy Chain

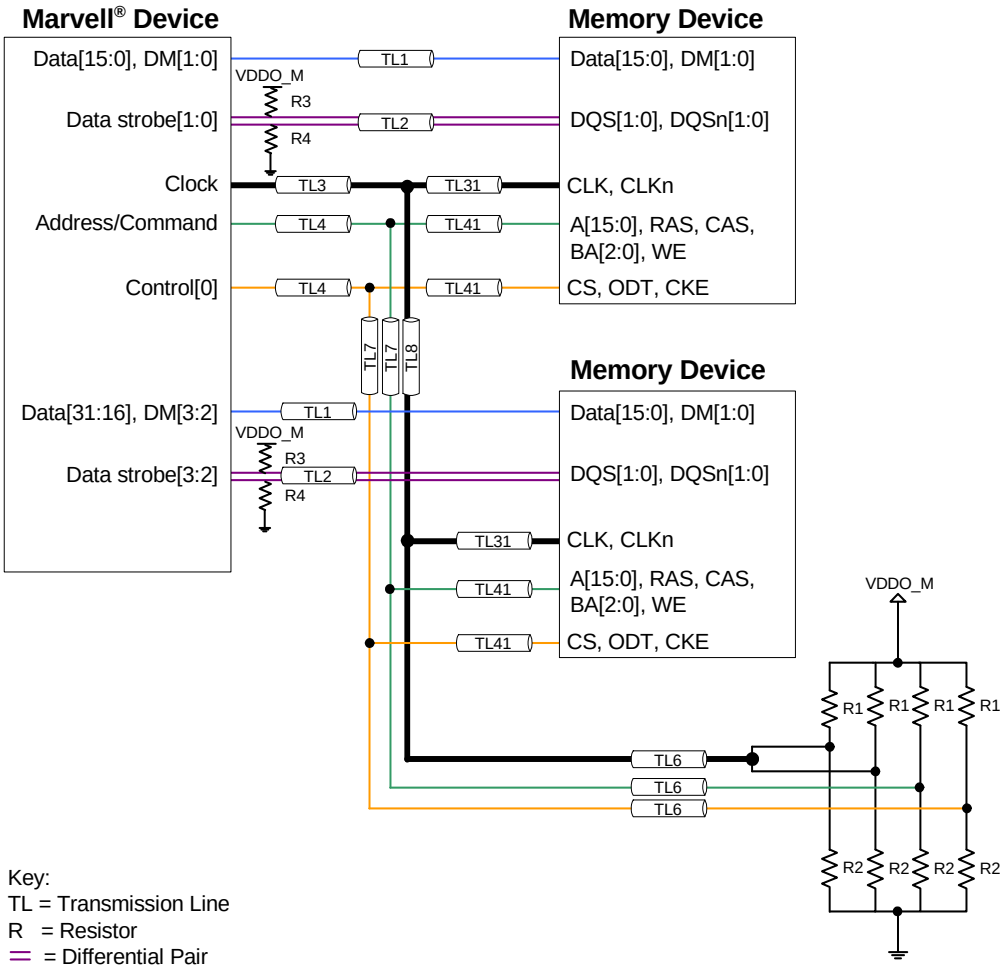

Table 9 lists the routing constraints when using a daisy chain topology.

Table 9: Routing Constraints when Using Daisy Chain Topology

Parameter	Min	Max	Tolerance / Skew	Units	Notes
TL[x] internal-group separation (dimensions)	2			Signal width	1, 2, 3
TL[x] internal-group separation (crosstalk)		-20		dB	1, 2, 3

Table 9: Routing Constraints when Using Daisy Chain Topology (Continued)

Parameter	Min	Max	Tolerance / Skew	Units	Notes
TL[x] to TL[y] cross-group separation (dimensions)	3			Signal width	1, 2, 3
TL[x] to TL[y] cross-group separation (crosstalk)		-26		dB	1, 2, 3
TL1, TL2 length		4		inch	
DQ to DQS skew TL2–TL1	-50	50		mil	
TL2 in pair skew, TL3+TL31, TL3+TL8+TL31, TL3+TL8+TL6 in pair skew (differential pairs)	-10	10		mil	5
DQS0/1 (first chain) to Clock skew: TL2 – (TL3 + TL31)	-8	-1		inch	4
DQS2/3 (second chain) to Clock skew: TL21 – (TL8 + TL3+TL31)	-8	-1		inch	4
Address/Command/Control to Clock skew (first device): (TL4+TL41) – (TL3+TL31)	-500	500		mil	
Address/Command/Control to Clock skew (second device): (TL7+TL4+TL41) – (TL8+TL3+TL31)	-500	500		mil	
TL31 length		Min (0.2*TL3, 0.3)		inch	
TL41 length		Min (0.2*TL4, 0.3)		inch	
TL6 length		1		inch	6, 7
R1, R2	100		1%	ohm	6, 7, 8, 9, 10
R3, R4	1			kohm	11
TL[x] impedance (all traces)	50		10%	ohm	12
TL2, TL3, TL8 differential impedance	100		10%	ohm	12, 13
TL2, TL3, TL8 single-ended impedance		60	10%	ohm	12

NOTES

1. If the design cannot meet the above requirements for length and separation, simulation can be performed to ensure proper isolation between signals.
2. The separation is defined from edge-to-edge of two traces.
3. Separation is required between any signals that are not in the same group.
4. These values assume that Write-Leveling is enabled. If Write-Leveling is disabled, the limitation is +/- 1.5 inches.
5. TL2, TL21, TL3, TL31, and TL8 are pairs of differential signals.
6. Termination is required on Clock signals only.
7. If termination is not required, do not route TL6.
8. The trace length, connecting between the termination resistors R1 and R2, should be less than 100 mils.

9. R1 and R2 may be replaced by a single 50-ohm resistor connected to a VTT power source. In this case, the VTT power should be filtered, to ensure a quiet power source. The VTT power source should have a nominal voltage of Vref±20mV.
10. Place a bypass capacitor between VDDIO and GND, and close to the Thevenin terminations. The ratio of capacitors to terminations is 1:3.
11. The connection from DQSn to R3 (pull up) and R4 (pull down) must be less than 500 mils from the device.
12. The tolerance for the matched impedance value is the manufacturing tolerance. The design tolerance should be kept to a minimum.
13. The target impedance priority is to match the differential impedance of these signal traces and adjust the single-ended impedance accordingly.

4.2.2 Routing Topologies—Two Chip Selects 4 x 16-bit Wide Memory

Figure 30 and Figure 31 depict the topology for a 4 x 16-bit wide memory device.

Figure 30: 4 x 16-bit Wide Memory Topology for Clock/Address/Command/Control Implemented as Daisy Chain

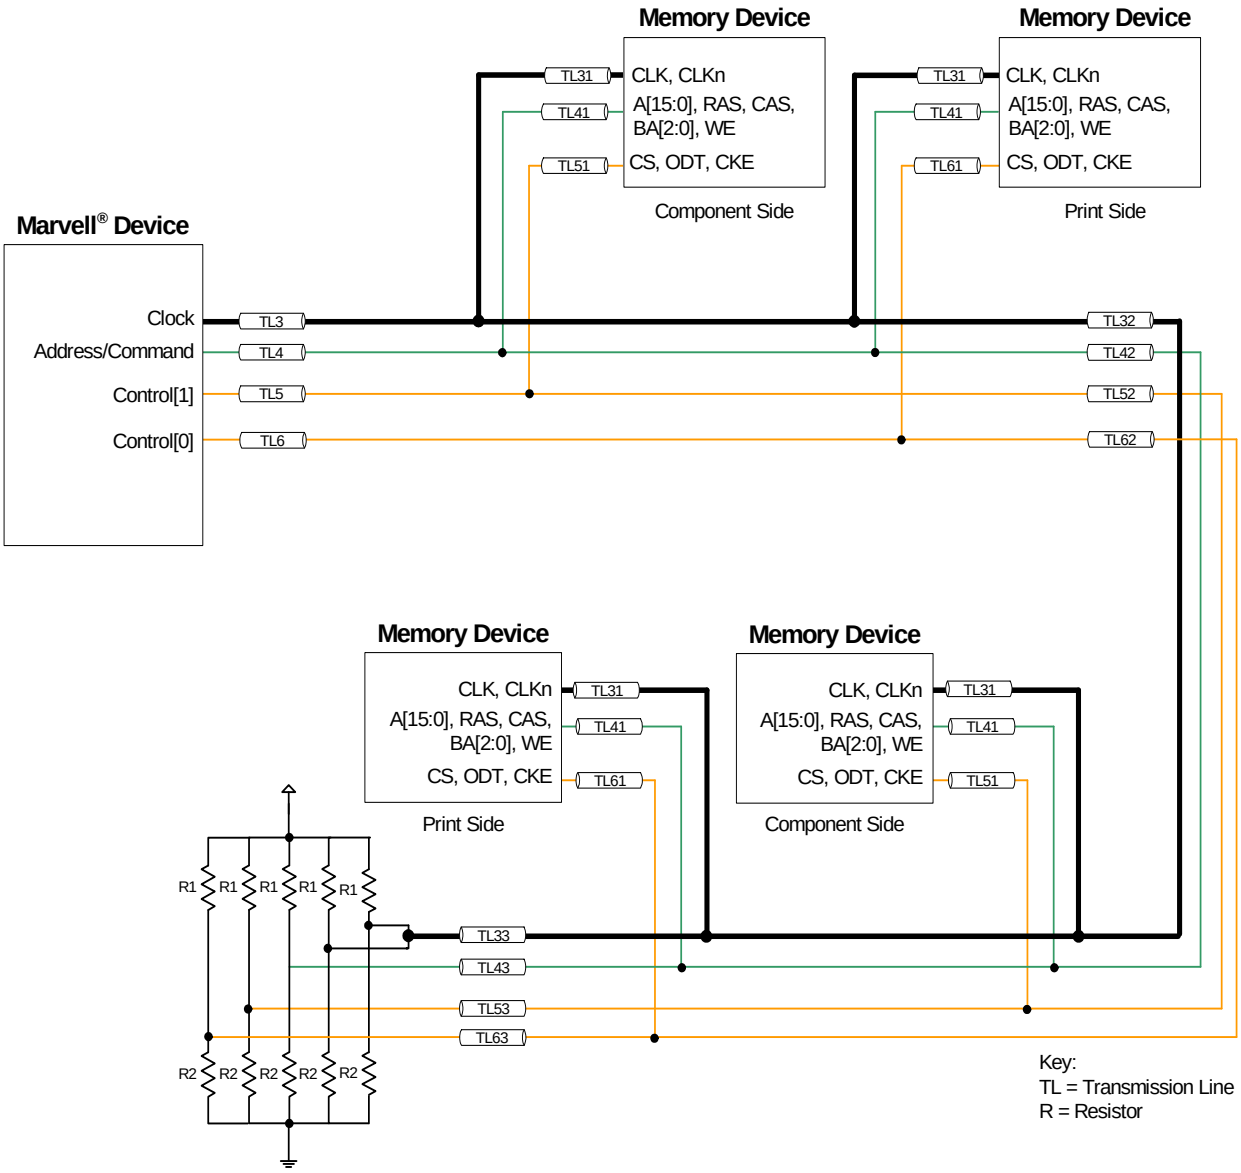

Note Multiple clocks can be used to ease the routing. If multiple output clocks are available in the specific Marvell® device, be sure that the routing does not violate the clock's timing constraints.

Figure 31: 4 x 16-bit Wide Memory Topology for Data

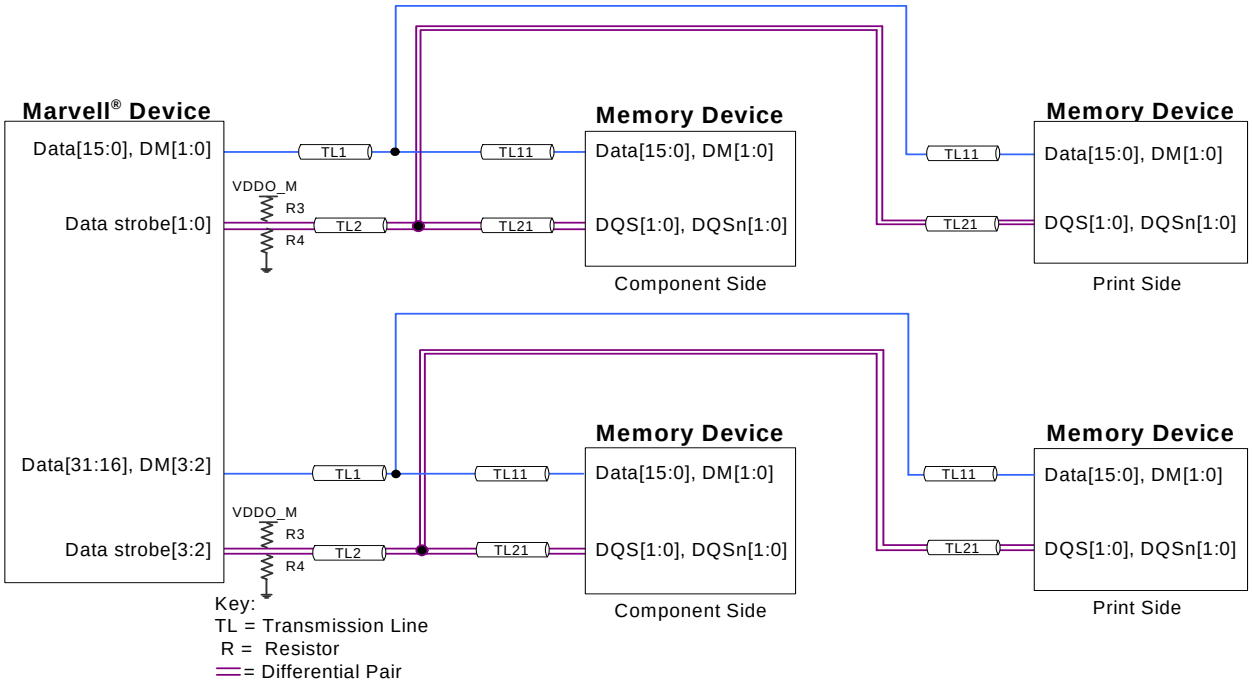

Table 10 lists the routing constraints when using four 16-bit wide memory topology.

Table 10: Routing Constraints Using 4 x 16-bit Wide Memory Topology

Parameter	Min	Max	Tolerance / Skew	Units	Notes
TL[x] internal-group separation (dimensions)	2			Signal width	1, 2
TL[x] internal-group separation (crosstalk)		-20		dB	1, 2
TL[x] to TL[y] cross-group separation (dimensions)	3			Signal width	1, 2
TL[x] to TL[y] cross-group separation (crosstalk)		-26		dB	1, 2
TL1+TL11, TL2+TL21 length		4		inch	
TL11 Length		Min(0.2* TL1,0.2)		inch	
TL21 Length		Min(0.2* TL2,0.2)		inch	
DQ to DQS skew (TL2+TL21) – (TL1+TL11)	-50	50		mil	

Table 10: Routing Constraints Using 4 x 16-bit Wide Memory Topology (Continued)

Parameter	Min	Max	Tolerance / Skew	Units	Notes
In pair skew (differential pairs): TL2+TL21 TL3+TL31 TL3+TL31+TL32 TL3+TL33+TL32	-10	10		mil	4
DQS to Clock skew: (TL+TL21) – (TL3+TL31) (TL2+TL21) – (TL3+TL31+TL32)	-8	-1		inch	3, 5
Address/Command to Clock skew: (TL4+TL41) – (TL3+TL31) (TL4+TL41+TL42) – (TL3+TL31+TL32)	-500	500		mil	5
Control to Clock skew: (TL5+TL51) – (TL3+TL31) (TL5+TL51+TL52) – (TL3+TL31+TL32) (TL6+TL61) – (TL3+TL31) (TL6+TL61+TL62) – (TL3+TL31+TL32)	-500	500		mil	5
Trace length to the termination: TL33, TL43, TL53, TL63 length		0.5		inch	6, 7
R1, R2	100		1%	ohm	6, 7, 8, 9, 10
R3, R4	1			kohm	11
TL[x] impedance (all traces)	50		10%	ohm	12
TL2, TL3, TL31, TL32, TL33 differential impedance	100		10%	ohm	12, 13
TL2, TL3, TL31, TL32, TL33 single-ended impedance		60	10%	ohm	12

NOTES

1. If the design cannot meet the above requirements for length and separation, simulation can be performed to ensure proper isolation between signals.
2. The separation is defined from edge-to-edge of two traces.
3. These values assume that Write-Leveling is enabled. If Write-Leveling is disabled, the limitation is +/- 1.5 inches.
4. TL2, TL21, TL22, TL3, TL31, TL32 and TL33 are pairs of differential signals.
5. The length of the trace to the termination of the clock should not be a part of the calculation.
6. Termination is required when Address toggles in 1T mode.
   Termination on Clock signals is always required.
7. If termination is not required, do not route TL43, TL53, TL63.
8. R1 and R2 may be replaced by a single 50-ohm resistor connected to a VTT power source. In this case, the VTT power should be filtered, to ensure a quiet power source. The VTT power source should have a nominal voltage of Vref±20mV.
9. The trace length, connecting between the termination resistors R1 and R2, should be less than 100 mils.
10. Place a bypass capacitor between VDDIO and GND, and close to the Thevenin terminations. The ratio of capacitors to terminations is 1:3.
11. The connection from DQSn to R3 (pull up) and R4 (pull down) must be less than 500 mil from the device.

12. The tolerance for the matched impedance value is the manufacturing tolerance. The design tolerance should be kept to a minimum.
13. The target impedance priority is to match the differential impedance of these signal traces and adjust the single-ended impedance accordingly.

4.2.3 Routing Topologies—One Chip Select 4 x 8-bit Wide Memory

Figure 32 and Figure 33 depict the topology for a 4 x 8-bit wide memory device.

Figure 32: 4 x 8-bit Wide Memory Topology for Clock/Address/Command/Control Implemented as Two Daisy Chain Branches

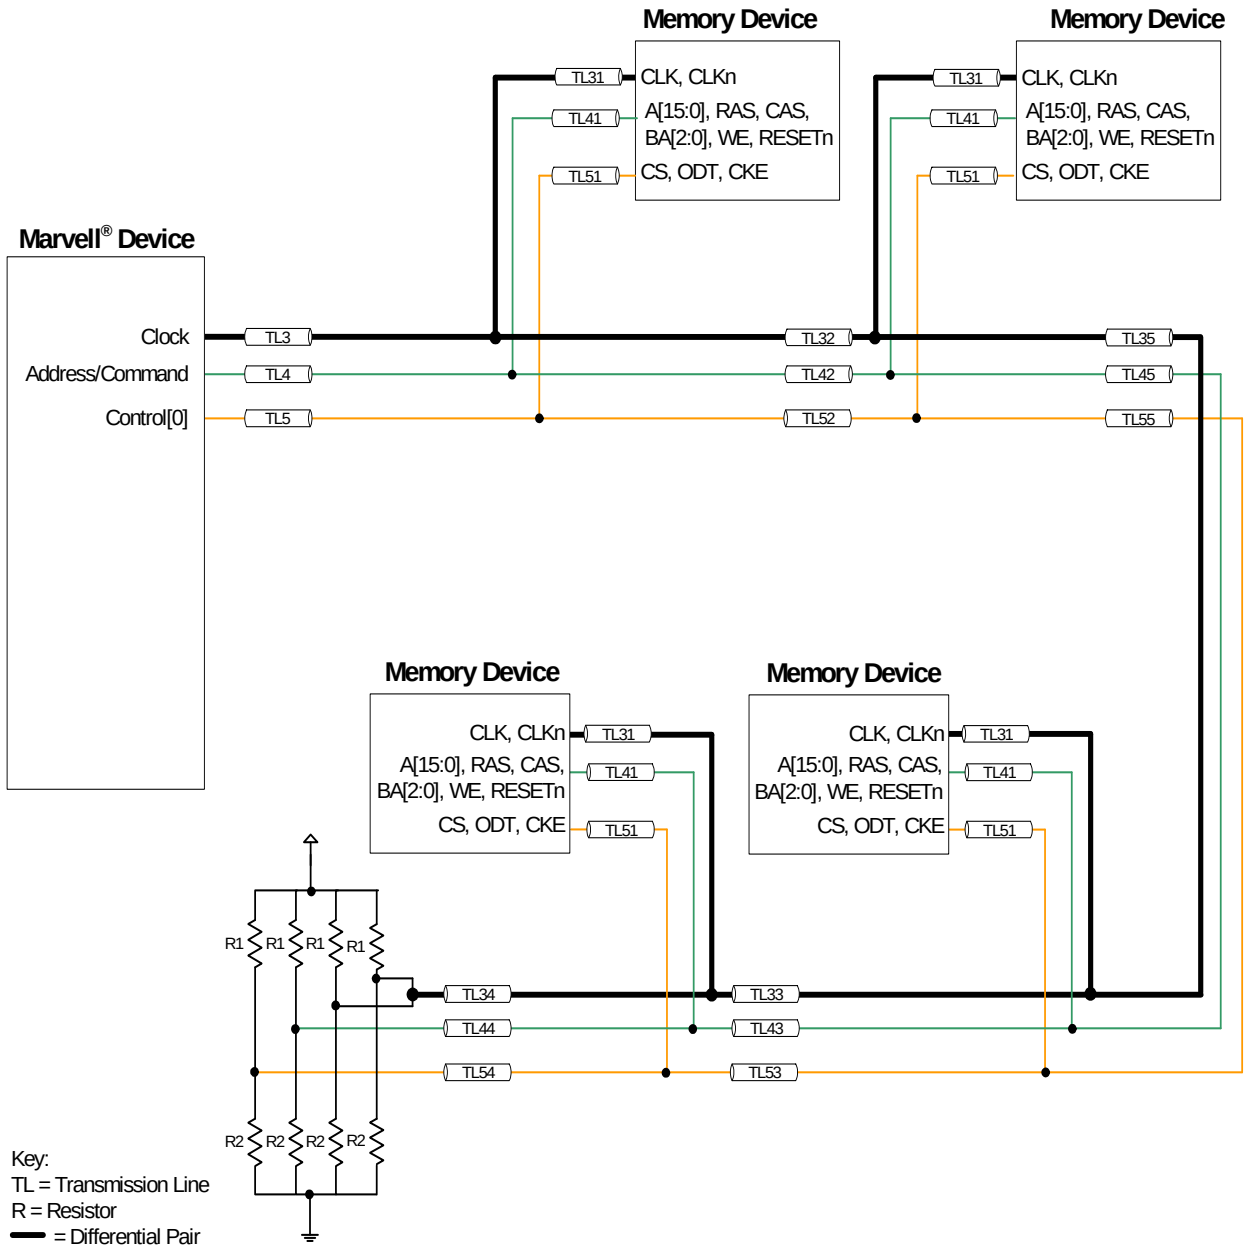

Figure 33: 4 x 8-bit Wide Memory Topology for Data

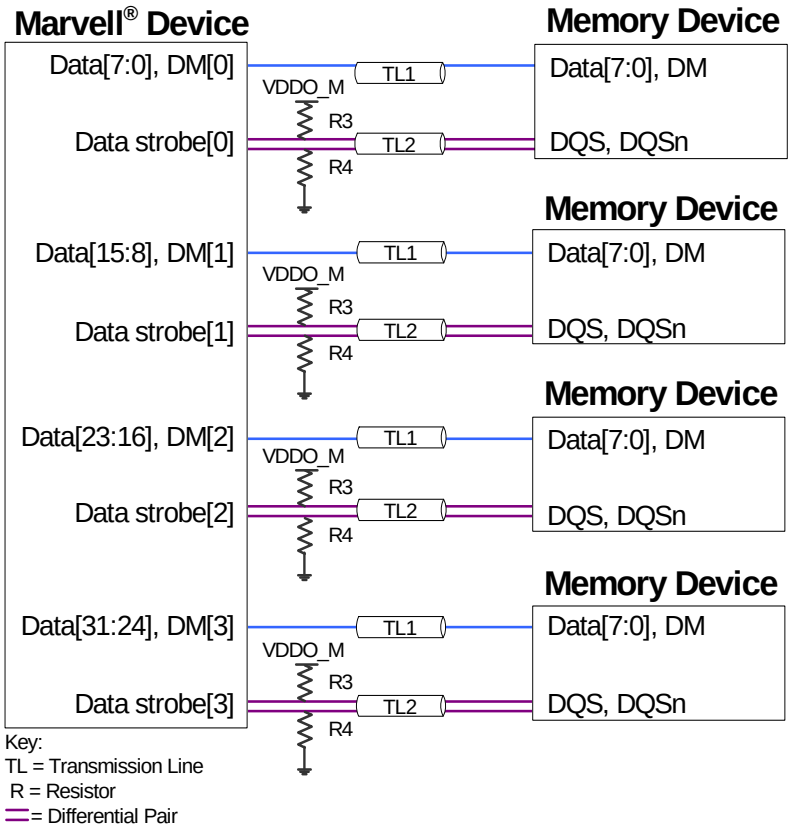

Table 11 lists the routing constraints when using four 8-bit wide memory topology.

Table 11: Routing Constraints Using 4 x 8-bit Wide Memory Topology

Parameter	Min	Max	Tolerance / Skew	Units	Notes
TL[x] internal-group separation (dimensions)	2			Signal width	1, 2
TL[x] internal-group separation (crosstalk)		-20		dB	1, 2
TL[x] to TL[y] cross-group separation (dimensions)	3			Signal width	1, 2
TL[x] to TL[y] cross-group separation (crosstalk)		-26		dB	1, 2
DQ to DQS skew TL2–TL1	-50	50		mil	

Table 11: Routing Constraints Using 4 x 8-bit Wide Memory Topology (Continued)

Parameter	Min	Max	Tolerance / Skew	Units	Notes
In pair skew (differential pairs): TL2 TL3+TL31 TL3+TL32+TL31 TL3+TL32+TL35+TL31 TL3+TL32+TL35+TL33+TL31 TL3+TL32+TL35+TL33+TL34	-10	10		mil	4
DQS to Clock skew (TL2) - (TL3+TL31) (TL2) - (TL3+TL32+TL31) (TL2) - (TL3+TL32+TL35+TL31) (TL2) - (TL3+TL32+TL35+TL33+TL31)	-8	-1		inch	5
Address/Command to Clock skew (TL4+TL41) - (TL3+TL31) (TL4+TL42+TL41) - (TL3+TL32+TL31) (TL4+TL42+TL45+TL41) - (TL3+TL32+TL35+TL31) (TL4+TL42+TL45+TL43+TL41) - (TL3+TL32+TL35+TL33+TL31)	-500	500		mil	5
Control to Clock skew (TL5+TL51) - (TL3+TL31) (TL5+TL52+TL51) - (TL3+TL32+TL31) (TL5+TL52+TL55+TL51) - (TL3+TL32+TL35+TL31) (TL5+TL52+TL55+TL53+TL51) - (TL3+TL32+TL35+TL33+TL31)	-500	500		mil	5
TL31, TL41, TL51 length		0.2		inch	
TL34, TL44, TL54 length		0.5		inch	6, 7
R1, R2	100		1%	ohm	6, 7, 8, 9, 10
R3, R4	1			kohm	10
TL[x] impedance (all traces)	50		10%	ohm	12
TL2, TL3, TL31, TL32, TL33, TL34, TL36 differential impedance	100		10%	ohm	12
TL2, TL3, TL31, TL32, TL33, TL34, TL36 single-ended impedance		60	10%	ohm	12

NOTES

1. If the design cannot meet the above requirements for length and separation, simulation can be performed to ensure proper isolation between signals.
2. The separation is defined from edge-to-edge of two traces.
3. TL2, TL3, TL31, TL32, TL33, TL34 and TL36 are pairs of differential signals.
4. The length of the trace to the termination of the clock should not be a part of the calculation.
5. R1 and R2 may be replaced by a single 50-ohm resistor connected to a VTT power source. In this case, the VTT power should be filtered, to ensure a quiet power source. The VTT power source should have a nominal voltage of Vref±20mV.

6. Termination is required when Address toggles in 1T mode.
   Termination on Clock signals is always required.
7. If termination is not required, do not route TL44, TL54.
8. The trace length, connecting between the termination resistors R1 and R2, should be less than 100 mils.
9. The tolerance for the matched impedance value is the manufacturing tolerance. The design tolerance should be kept to a minimum.
10. The connection from DQSn to R3 (pull up) and R4 (pull down) must be less than 500 mil from the device.
11. Place a bypass capacitor between VDDIO and GND, and close to the Thevenin terminations. The ratio of capacitors to terminations is 1:3.
12. The target impedance priority is to match the differential impedance of these signal traces and adjust the single-ended impedance accordingly.

4.2.4 Routing Topologies—Two Chip Selects 8 x 8-bit Wide Memory

Figure 34 and Figure 35 depict the topology for a 8 x 8-bit wide memory device.

Figure 34: 8 x 8-bit Wide Memory Topology for Data

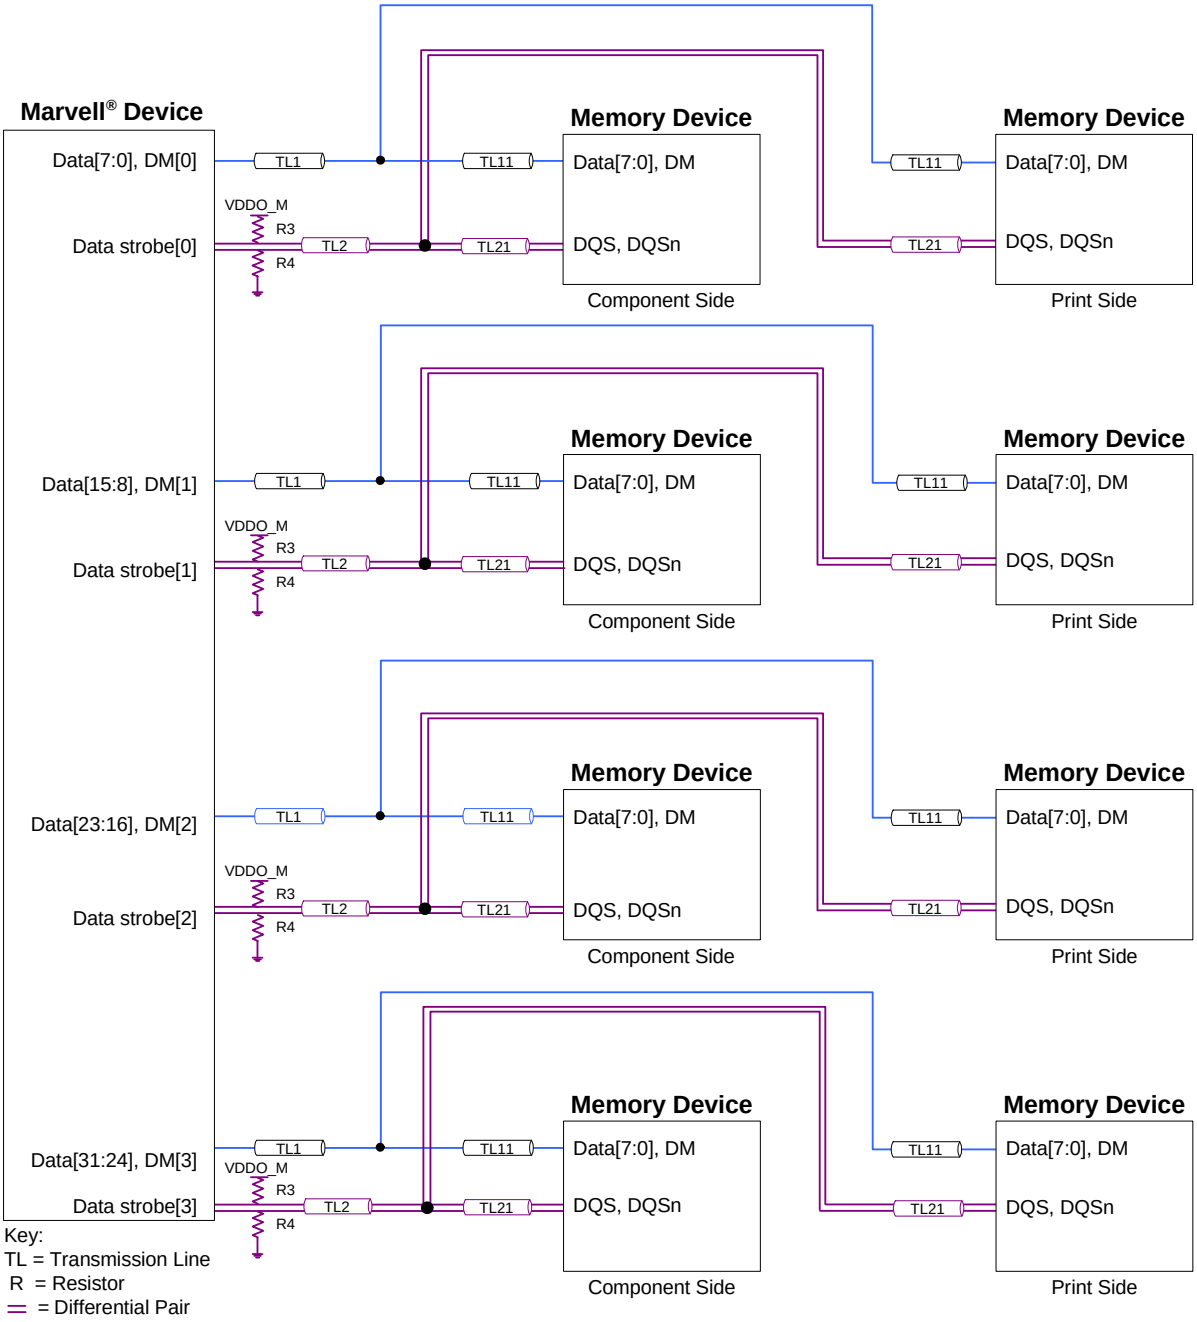

Figure 35: 8 x 8-bit Wide Memory Topology for Clock/Address/Command/Control Implemented as Daisy Chain

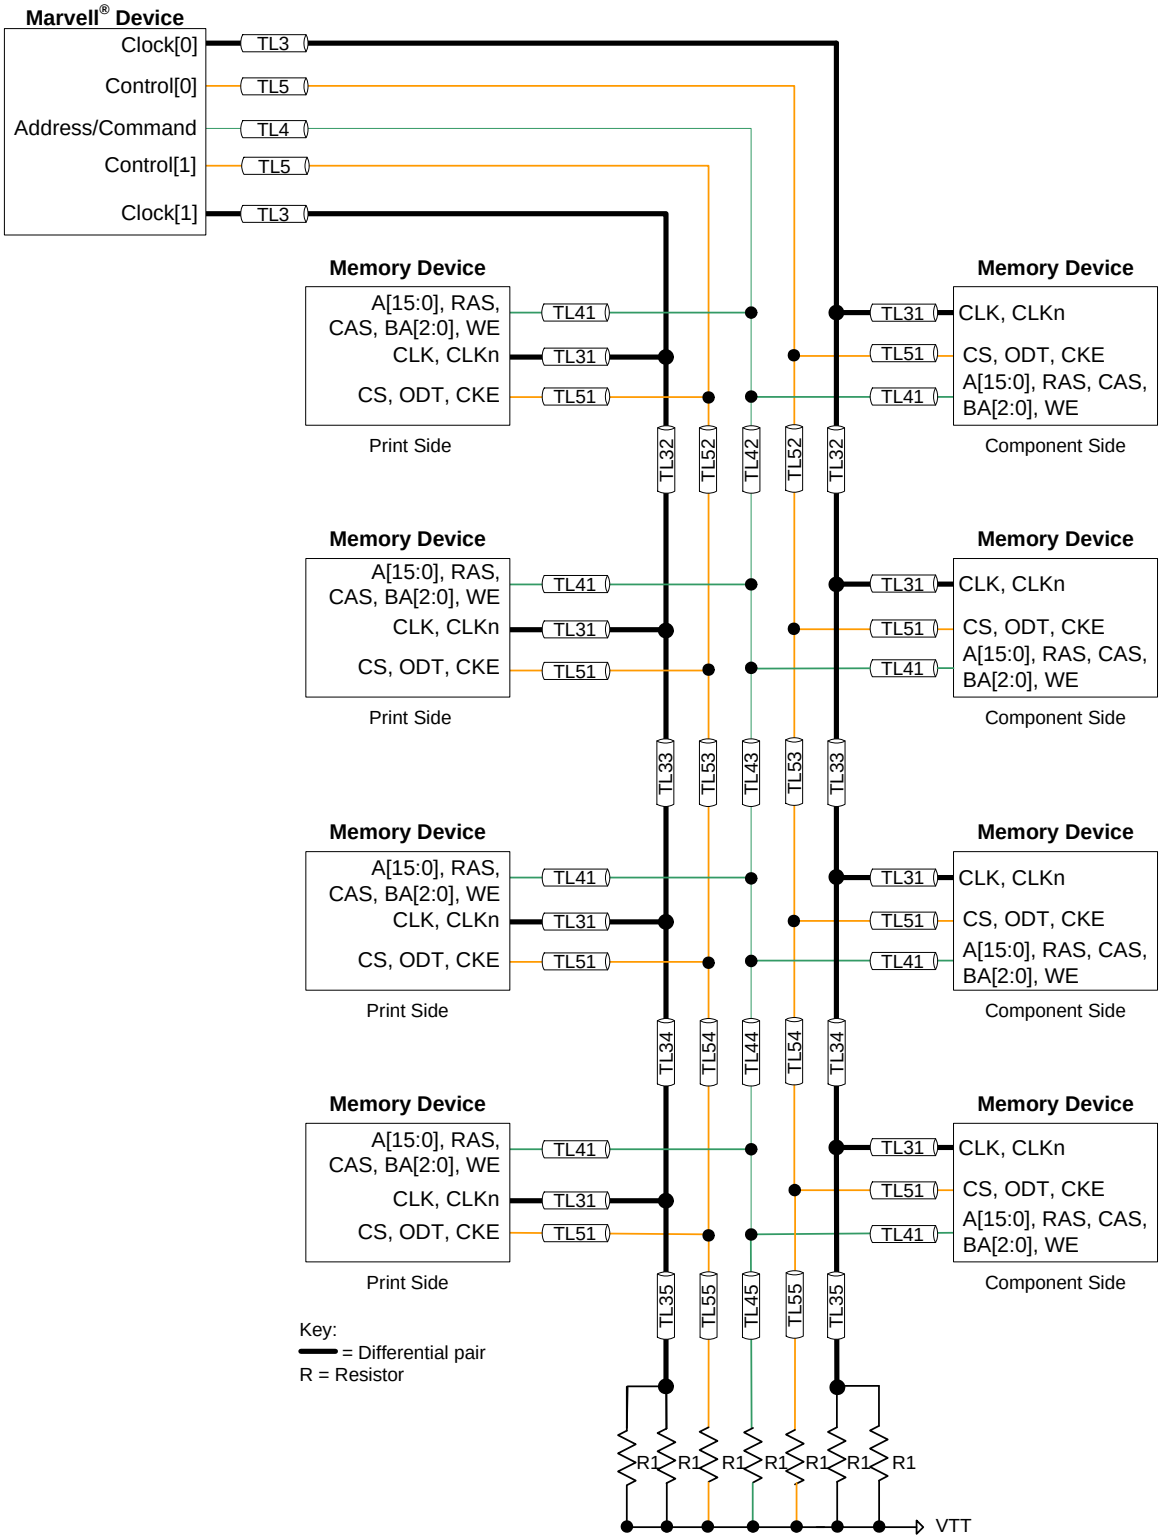

Table 10 lists the routing constraints when using four 8-bit wide memory topology.

Table 12: Routing Constraints Using 8 x 8-bit Wide Memory Topology

Parameter	Min	Max	Tolerance / Skew	Units	Notes
TL[x] internal-group separation (dimensions)	2			Signal width	1, 2
TL[x] internal-group separation (crosstalk)		-20		dB	1, 2
TL[x] to TL[y] cross-group separation (dimensions)	3			Signal width	1, 2, 3
TL[x] to TL[y] cross-group separation (crosstalk)		-26		dB	1, 2
Data total trace length: TL1+TL11, TL2+TL21 length		4		inch	
TL11, TL21, TL31, TL41, TL51 Length		0.2		inch	
DQ to DQS skew (TL2+TL21) – (TL1+TL11)	-50	50		mil	
In pair skew (differential pairs): TL2+TL21 TL3+TL31 TL3+TL31+TL32 TL3+TL32+TL33+TL31 TL3+TL32+TL33+TL34+TL31 TL3+TL33+TL32	-10	10		mil	
DQS to Clock skew: (TL+TL21) – (TL3+TL31) (TL2+TL21) – (TL3+TL31+TL32)	-8	-1		inch	4, 5
Address/Command to Clock skew (TL4+TL41) - (TL3+TL31) (TL4+TL42+TL41) – (TL3+TL32+TL31) (TL4+TL42+TL43+TL41) – (TL3+TL32+TL33+TL31) (TL4+TL42+TL43+TL44+TL41) - (TL3+TL32+TL33+TL34+TL31)	-500	500		mil	5
Control to Clock skew (TL5+TL51) - (TL3+TL31) (TL5+TL52+TL51) – (TL3+TL32+TL31) (TL5+TL52+TL53+TL51) – (TL3+TL32+TL33+TL31) (TL5+TL52+TL53+TL54+TL51) - (TL3+TL32+TL33+TL34+TL31)	-500	500		mil	5
TL35, TL45, TL55 length		0.5		inch	
R1	50		1%	ohm	6, 7
R3, R4	1			kohm	8
TL[x] impedance (all traces)	50		10%	ohm	9
TL2, TL3, TL3x differential impedance	100		10%	ohm	9, 10
TL2, TL3, TL3x single-ended impedance		60	10%	ohm	9

NOTES

1. If the design cannot meet the above requirements for length and separation, simulation can be performed to ensure proper isolation between signals.
2. The separation is defined from edge-to-edge of two traces.
3. For the signal group definitions see Section 4.1.1, Signal Groups, on page 50.
4. These values assume that Write-Leveling is enabled. If Write-Leveling is disabled, the limitation is +/- 1.5 inches.
5. The length of the trace to the termination of the clock should not be a part of the calculation.
6. the VTT power should be filtered, to ensure a quiet power source. The VTT power source should have a nominal voltage of Vref±20mV.
7. R1 may be replaced by two 100-ohm resistors connected as Thevenin termination, then the trace length, connecting between the Thevenin termination resistors to VDD_IO, should be less than 100 mils.
8. The connection from DQSn to R3 (pull up) and R4 (pull down) must be less than 500 mil from the device.
9. The tolerance for the matched impedance value is the manufacturing tolerance. The design tolerance should be kept to a minimum.
10. The target impedance priority is to match the differential impedance of these signal traces and adjust the single-ended impedance accordingly.

4.3 Power Signals

For Vref filtering and scheme, refer to Section 16, Generic Power Board Guidelines, on page 114.

4.4 Special Software Settings

4.4.1 One Chip Select

ODT matrix See Table 13, ODT/RTT Setting Matrix for a Two-Chip-Select Topology, on page 69.

SDRAM configuration Configure to ZQ/6.

4.4.2 Two Chip Selects

ODT matrix See Table 13.

SDRAM configuration Configure to ZQ/7.

Table 13: ODT/RTT Setting Matrix for a Two-Chip-Select Topology

Command	Controller ODT	Memory0 ODT	Memory1 ODT
Write to Memory0	ODT off	60 ohm	ODT off
Write to Memory1	ODT off	ODT off	60 ohm
Read from Memory0	75 ohm	ODT off	ODT off
Read from Memory1	75 ohm	ODT off	ODT off

5 PCI Express Interface 1.0/1.1 (PCIe)

This section provides the connectivity, layout, and power guidelines for the PCI Express interface (PCIe) pins operating in PCIe Version 1.0 or 1.1.
The PCI Express is a serial differential low-voltage, point-to-point interface. The interface consists of one differential pair for transmit, one for receive, and one for reference clock.

5.1 Connectivity

Table 14: Signal Groups

Group Name	PCI Express	Comments
Output Data differential pair	PEX_Tx_P/N[n..0]	The value of n is determined by the specific device interface width.
Input Data differential pair	PEX_Rx_P/N[n..0]	The value of n is determined by the specific device interface width.
Input Clock	PEX_I_CLK_P/N	

Table 15: Specific Device

Group Name	PCI Express	Comments
Output Clock	PEX_CLK_P/N	Optional clock output Can use the same pins as Input Clock. Refer to the Hardware Specifications.

Figure 36 and Figure 37 provide two connectivity options, depending on whether the reference clock is external or supplied by the Marvell® device. Also refer to Section 5.4.2, PCI Express Reference Clock Supplied by the Marvell® Device, on page 77.

Note

- The pin naming conventions may vary for different devices, including PHYs. Verify the connectivity against the naming convention used by the peer device.
- For output clock routing guidelines, refer to Section 5.4.2, PCI Express Reference Clock Supplied by the Marvell® Device, on page 77.

Figure 36: PCI Express Interface Connectivity

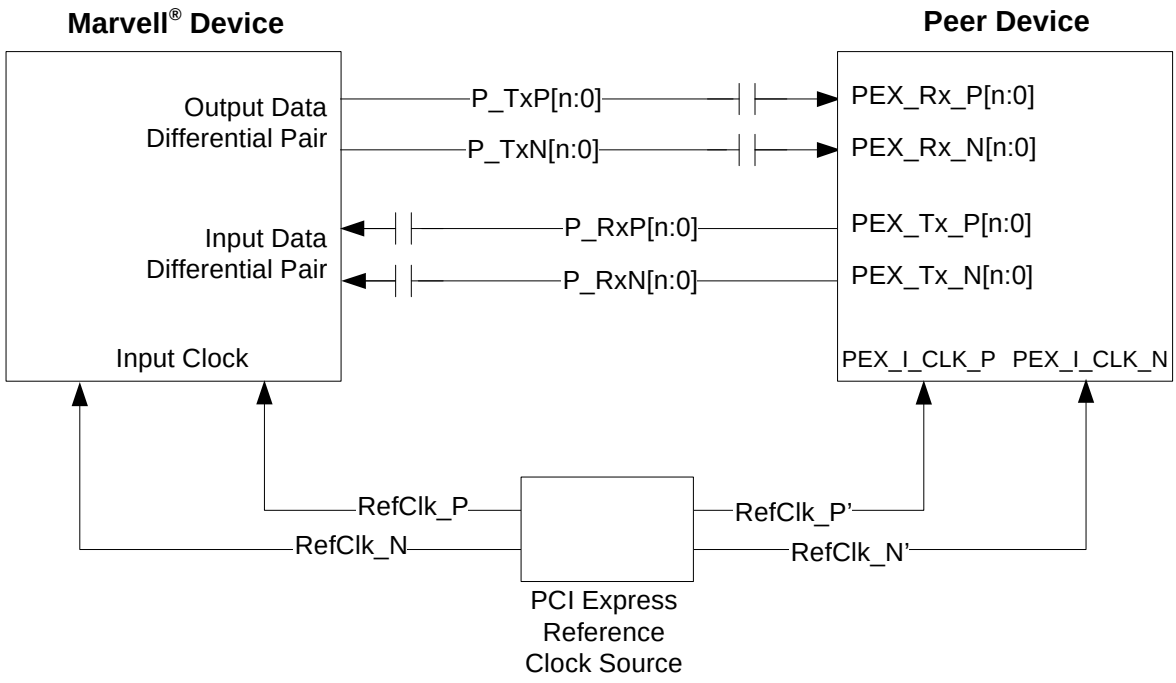

Figure 37: PCI Express Interface Connectivity and Reference Clock Supplied by Marvell® Device

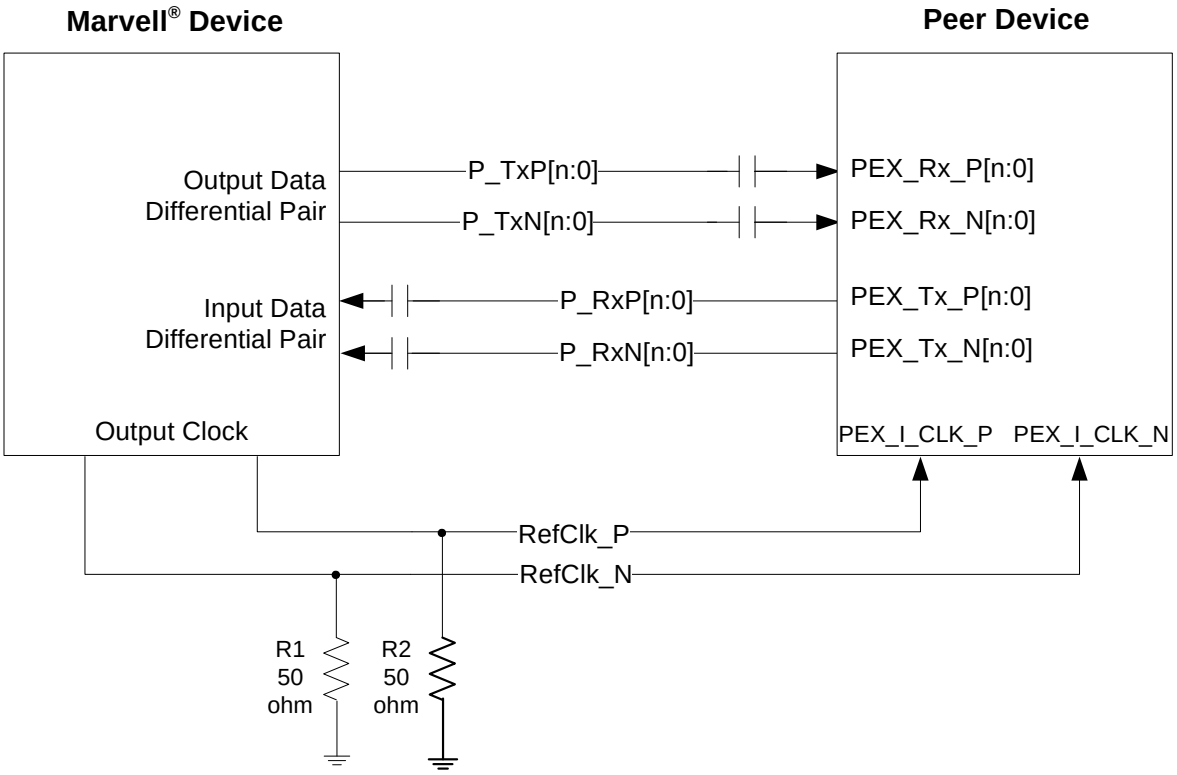

5.2 Interface Signals Layout Guidelines

The location of the AC coupling capacitors varies, as described below, depending on whether the routing is via a connector or on-board routing is used:

When routing via a connector: On the transmitter side, place the capacitors as close as possible to the connector. There is no need to place capacitors on the receiver side.

When using on-board routing: When both ends of the differential traces are located on the same PCB (i.e. no connector along the way), place the capacitor near the receiver device pins.

General Routing Notes

Marvell® recommends following these general routing guidelines:

- It is essential to minimize the effect of capacitive discontinuities on the PCIe differential traces. When routing via a through hole connector, go all the way through the connector via hole and leave a stub as short as possible. Refer to the differential traces routing guidelines in Section 2, Generic Guidelines for SERDES Interfaces, on page 16 for further information.
- Refer to the PCIe CEM specifications, for further information about allowed jitter budgeting.
- The total capacitance detected by the receiver detection circuit must not exceed 3 nF to ground.

5.2.1 Routing Topologies with a Connector and/or an Add-in Card

Figure 38: Topology with an Existing Connector and/or an Add-in Card

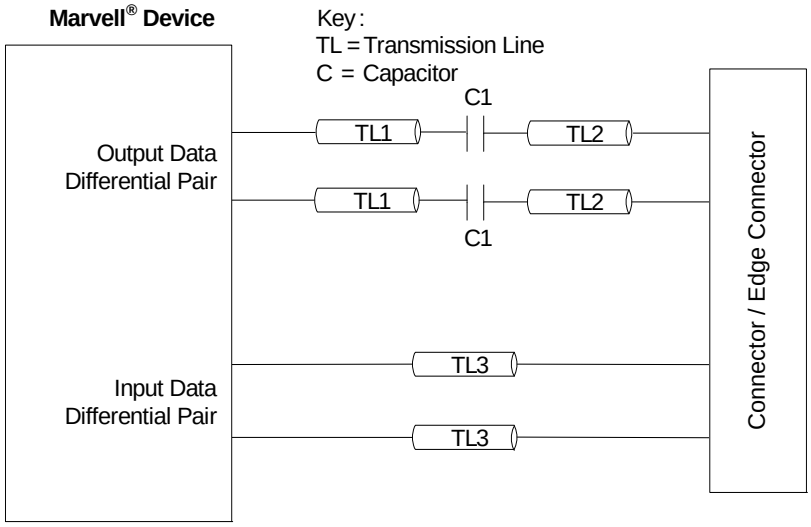

Table 16 and Table 17 provide the values for two topology options, depending on whether an existing connector or an add-in card is used, respectively.

Table 16: System Board with Existing Connector

Parameter	Min	Max	Units	Notes
TL1+TL2 length		11	inch	1, 2, 3, 4
TL1+TL2 loss @ 1.25 GHz		3.6	dB	9
TL2 length		0.4	inch	
TL3 length		11	inch	3, 4
TL1x to TL1y or TL3x to TL3y same direction lane separation (dimensions)	4.5		signal width	5, 6, 7
TL1[x] to TL1[y] separation (far-end crosstalk)		-26	dB	
TL1[x] or TL3[x] to TL[x] lane separation from other traces (dimensions)	4		signal width	8
TL1[x] or TL3[x] to TL[x] lane separation from other traces (crosstalk)		-40	dB	8
TL3 loss @ 1.25 GHz		3.6	dB	1, 2
(TL1n+TL2n) – (TL1m+TL2m) lane-to-lane skew		7	inch	3, 4, 9
(TL3n-TL3m) length		7	inch	3, 4, 9
TL1+TL2, TL3 inner pair skew		10	mil	10
TL1, TL2, and TL3 single-ended impedance	50	60	Ω	1
C1	100	200	nF	11

Notes:

1. Target impedance should be 100Ω differential; any mismatch should be taken as a part of the loss.
2. Derived from the loss value by calculation with 0.33 dB/inch.
3. Discontinuities, such as capacitive type, can cause significant insertion loss degradation, causing the assumption, taken as loss per inch, to be smaller than actual.
4. The maximal allowed trace length can vary according to the material and geometrical characteristics.
5. x and y represent lane numbers, or different, same-direction PCIe SERDES differential lines.
6. The separation is calculated assuming the same type of SERDES and typical board layer stack-up. Make certain to verify that the crosstalk limitation is met in each specific system.
7. The separation value is calculated for a single-ended impedance of 55-ohm, a differential impedance of 100-ohm, u-strip trace length coupled along the entire end-to-end path, and a dielectric loss coefficient of 0.027. If the loss coefficient is lower, a higher separation may be required.
8. The allowed crosstalk is highly dependant of the aggressor's waveform (rise time and swing). The recommended separation value is for an aggressor of the same type of interface.
9. m and n are lane numbers in cases of a multi-lane interface only (m, n may be 0, 1, 2, 3, etc.)
10. In cases of BGA packages, some package length routing skew can be present. This skew can help compensate for pin position within the pinout. Refer to the device's hardware specification for specific package trace length information.
11. It is essential to use high-speed, low-ESL, low-ESR capacitors.

Table 17: Add-in Card

Parameter	Min	Max	Units	Notes
TL1+TL2 length		7.5	inch	1, 2, 3, 4, 5
TL1+TL2 loss @ 1.25 GHz		2.5	dB	10
TL2 length		0.5	inch	
TL3 length		7.5	inch	4, 5
TL1x to TL1y or TL3x to TL3y same direction lane separation (dimensions)	4.5		signal width	6, 7, 8
TL1[x] to TL1[y] separation (far-end crosstalk)		-26	dB	
TL1[x] or TL3[x] to TL[x] lane separation from other traces (dimensions)	4		signal width	9
TL1[x] or TL3[x] to TL[x] lane separation from other traces (crosstalk)		-40	dB	9
TL3 loss		2.5	dB	1, 2
(TL1n+TL2n) – (TL1m+TL2m) lane-to-lane skew		2	inch	4, 5, 11
(TL3n – TL3m) lane-to-lane skew		2	inch	4, 5, 11
TL2+TL1, TL3 inner pair skew		5	mil	12
TL1, TL2, and T3 single-ended impedance	50	60	Ω	1
C1	100	200	nF	13

Notes:

1. Target impedance should be 100Ω differential; any mismatch should be taken as a part of the loss.
2. Derived from the loss value by calculation with 0.33 dB/inch.
3. Refer to the Hardware Specifications for the maximal allowed ISI. For the ISI term definition, refer to Section 2.2, Inter-Symbol Interference (ISI), on page 19.
4. Discontinuities, such as capacitive type, can cause significant insertion loss degradation, causing the assumption, taken as loss per inch, to be smaller than actual.
5. The maximal allowed trace length can vary according to the material and geometrical characteristics.
6. x and y represent lane numbers, or different, same-direction PCIe SERDES differential lines.
7. The separation is calculated assuming the same type of SERDES and typical board layer stack-up. Make certain to verify that the crosstalk limitation is met in each specific system.
8. The separation value is calculated for a single-ended impedance of 55-ohm, a differential impedance of 100-ohm, u-strip trace length coupled along the entire end-to-end path, and a dielectric loss coefficient of 0.027. If the loss coefficient is lower, a higher separation may be required.
9. The allowed crosstalk is highly dependant of the aggressor's waveform (rise time and swing). The recommended separation value should account for all common aggressor types.
10. Loss budget is meant for trace loss and should include far-end crosstalk.
11. m and n are lane numbers in cases of multi-lane interface only (m, n may be 0, 1, 2, 3, etc.)
12. In cases of BGA packages, some package length routing skew can be present. This skew can help compensate for pin position within the pinout. Refer to the device's hardware specification for specific package trace length information.
13. It is essential to use high-speed, low-ESL, low-ESR capacitors.

5.2.2 Routing Topologies with a Same-Board Connection

Figure 39: Topology with a Same-Board Connection

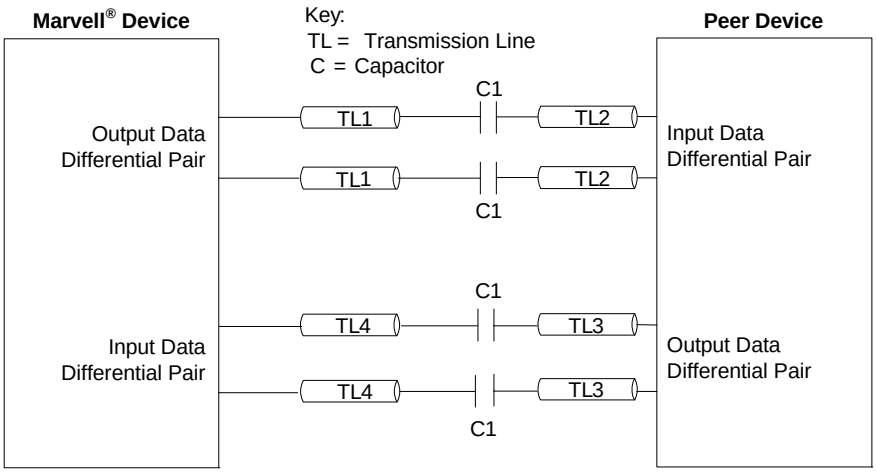

Table 18: Same-Board Connection

Parameter	Min	Max	Units	Notes
TL1+TL2 length		20	inch	1, 2, 3, 3, 4, 5, 6
TL1+TL2 loss @ 1.25 GHz		6.6	dB	5, 7
TL2 length		0.5	inch	
(TL1n+TL2n) - (TL1m+TL2m) lane-to-lane skew (TL3n+TL4n) - (TL3m+TL4m) lane-to-lane skew		5	inch	3, 4, 8
TL1x to TL1y or TL3x to TL3y same direction lane separation (dimensions)	4.5		signal width	9, 10, 11
TL1[x] to TL1[y] separation (far-end crosstalk below loss)		-26	dB	
TL1[x] or TL3[x] to TL[x] lane separation from other traces (dimensions)	4		signal width	12
TL1[x] or TL3[x] to TL[x] lane separation from other traces (crosstalk)		-40	dB	12
TL1+TL2, TL3+TL4 inner pair skew		15	mil	
TL1, TL2, TL3, and TL4 single-ended impedance	50	60	Ω	1
TL3+TL4 length		20	inch	1, 2, 3, 4
TL3+TL4 loss @ 1.25 GHz		6.6	dB	7
TL4 length		0.5	inch	

Table 18: Same-Board Connection (Continued)

Parameter	Min	Max	Units	Notes
C1	100	200	nF	6, 13

Notes:

1. Target impedance should be 100Ω differential; any mismatch should be taken as a part of the loss.
2. Derived from the loss value by calculation with 0.33 dB/inch.
3. Discontinuities, such as capacitive type, can cause significant insertion loss degradation, causing the assumption, taken as loss per inch, to be smaller than actual.
4. The maximal allowed trace length can vary according to the material and geometrical characteristics.
5. The maximal allowed length or loss is assumed for the same Marvell® device on both ends. If using a different peer device, refer to its design guidelines for TL1+ TL2 allowed loss.
6. The interconnect total capacitance to ground as seen by the Receiver Detection circuit must not exceed 3nF, including capacitance added by the attached test instrumentation. This capacitance is separate and distinct from the AC Coupling capacitance value (taken from the *PCI Express Base 1.1 Specification*).
7. Loss budget is meant for trace loss and should include far-end crosstalk.
8. m and n are lane numbers in cases of multi-lane interface only (m, n may be 0, 1, 2, 3, etc.).
9. x and y represent lane numbers, or different, same-direction PCIe SERDES differential lines.
10. The separation is calculated assuming the same type of SERDES and typical board layer stack-up. Make certain to verify that the crosstalk limitation is met in each specific system.
11. The separation value is calculated for a single-ended impedance of 55-ohm, a differential impedance of 100-ohm, u-strip trace length coupled along the entire end-to-end path, and a dielectric loss coefficient of 0.027. If the loss coefficient is lower, a higher separation may be required.
12. The allowed crosstalk is highly dependant of the aggressor's waveform (rise time and swing). The recommended separation value should account for all common aggressor types.
13. It is essential to use high-speed, low-ESL, low-ESR capacitors.

5.3 Power Considerations

The power pins are analog type. See Section 16, Generic Power Board Guidelines, on page 114.

5.3.1 Reference Clock Considerations

Refer to the reference clock distribution device design guidelines for specific layout information.

5.4 Specific Signals

5.4.1 PCI Express Routing Guidelines for the SERDES_ISET Signal

Follow these guidelines for SERDES_ISET signal placement:

- Place the ISET resistor up to 400 mil from the SERDES_ISET pin.
- For resistor parameters, refer to the Hardware Specifications.

5.4.2 PCI Express Reference Clock Supplied by the Marvell® Device

Connect the PCI Express clock out signals of the Marvell® device using a point-to-point connection to the peer device PCI Express input clock pins. Refer to Figure 37, PCI Express Interface Connectivity and Reference Clock Supplied by Marvell® Device, on page 72.

Route the clock out signals using traces with 100 ohm differential impedance.

Supply 50 ohm termination resistors to GND located up to 0.7 inches from the Marvell® device that is driving the clock.

6 Reduced Gigabit Media Independent Interface (RGMII)

The RGMII interface is a DDR source synchronous bus operating at 125 MHz.

The original version of the RGMII standard required the Tx and Rx clocks to be delayed by 1.5–2 ns referenced to the data. To meet this timing requirement, there are different timing schemes that can utilize an internal delay or external PCB delay. These possible timing schemes are detailed below in Table 20 through Table 23.

For possible Tx/Rx internal delays, refer to device Hardware Specifications.

6.1 Interface Connectivity

Table 19: Signal Groups

Group Name	Signal
Output Bus	TXD[3:0] TX_CTL
Output Clock	TX_CLK
Input Bus	RXD[3:0] RX_DV
Input Clock	RX_CLK

6.2 Connectivity

Figure 40 shows the RGMII port connection to a PHY and Figure 41 shows the RGMII port connection to another RGMII MAC.

Note The signal naming convention may differ between devices. Verify the referenced signal names in the device Hardware Specifications.

Figure 40: RGMII Port Connection to PHY

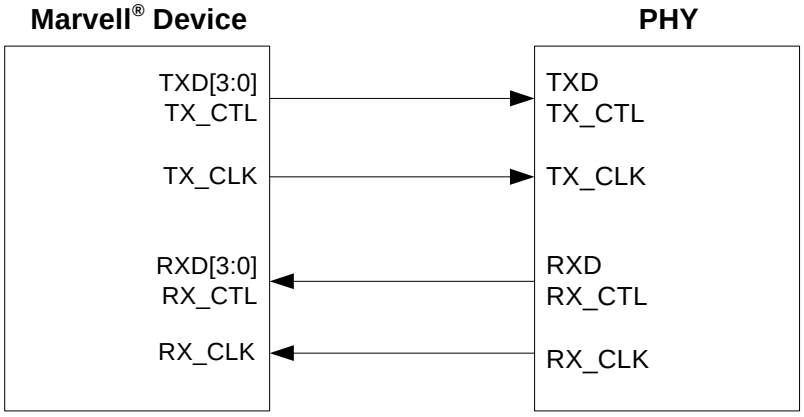

Figure 41: RGMII Port Connection to Another RGMII MAC

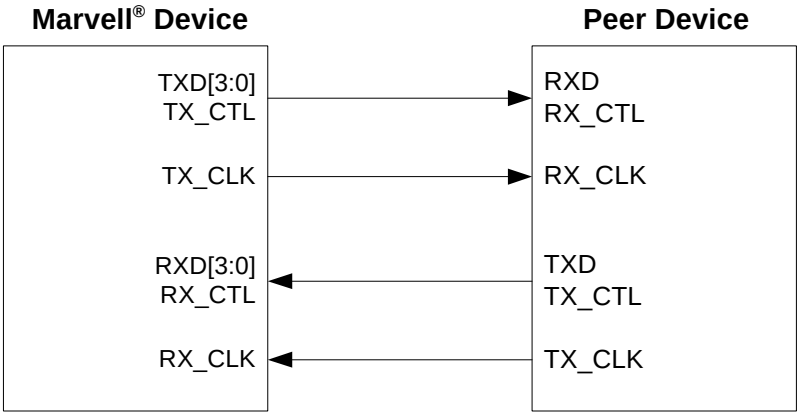

6.3 Interface Signals Layout Guidelines

Table 20 through Table 23 depict the possible timing modes for RGMII interface.

Choose the routing constraints according to the implementation.

The constraints provided are for the Gigabit Ethernet RGMII interface.

6.3.1 Routing Topologies

All signals are connected point-to-point.

Figure 42: RGMII Routing Topology

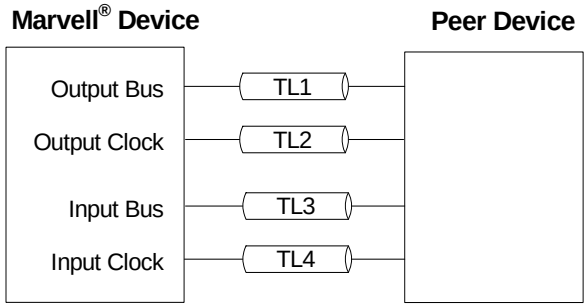

6.3.2 Tx Path Routing Constraints

Table 20: No Internal Delay on Transmitting or on Receiving Peer Side—Tx Path

Parameter	Min	Max	Tolerance / Skew	Units	Notes
TL1 length		6	0.200	inch	1, 2
TL2 length	Restricted by TL1 and the data to the clock relative delay.				1
TL2–TL1 (skew)	10	11		inch	3
TL2–TL1 (skew)	-0.1	0.1		inch	4
TL1 internal separation (dimensions)	2			Signal width	
TL1 internal separation (crosstalk)		-14		dB	
TL1–TL2 cross group separation (dimensions)	4			Signal width	
TL1–TL2 cross group separation (crosstalk)		-26		dB	

Notes:

1. Target impedance 45–55Ω.
2. Use the tolerance/skew value to align the signal of the entire Input/Output bus (data and control).
3. Verify a delay of 1.5–2 ns of the clock relative to the data. Adjust according to the material propagation delay.
4. When using an active element to supply the required external delay, the element device is used according to the delay in note 3 instead of a longer trace.

Table 21: Internal Delay on Transmitting or on Receiving Peer Side—Tx Path

Parameter	Min	Max	Tolerance / skew	Units	Notes
TL1 length		6	0.200	inch	1, 2
TL2 length		6		inch	1
TL1 – TL2 (skew)	-0.1	0.1		inch	
TL1 internal separation (dimensions)	2			Signal width	
TL1 internal separation (crosstalk)		-14		dB	
TL1 – TL2 cross group separation (dimensions)	4			Signal width	
TL1 – TL2 cross group separation (crosstalk)		-26		dB	

Notes:

1. Target impedance 45–55Ω.
2. Use the tolerance/skew value to align the signal of the entire Input/Output bus (data and control).

6.3.3 Rx Path Routing Constraints

Table 22: No Internal Delay on Receiving or on Transmitting Peer Side—Rx Path

Parameter	Min	Max	Tolerance / Skew	Units	Notes
TL3 length		6	0.200	inch	1, 2, 3
TL4 length	Restricted by TL3 and the data to clock relative delay				2
TL4 – TL3 (skew)	10	11		inch	4
TL4 – TL3 (skew)	-0.1	0.1		inch	5
TL3 internal separation (dimensions)	2			Signal width	
TL3 internal separation (crosstalk)		-14		dB	
TL3 – TL4 cross group separation (dimensions)	4			Signal width	
TL3 – TL4 cross group separation (crosstalk)		-26		dB	

Notes:

1. Maximum length is only a recommendation. For actual length, refer to Peer/PHY specifications.
2. Target impedance 45–55Ω.
3. Use the tolerance/skew value to align the signal of the entire Input/Output bus (data and control).

4. Verify a delay of 1.5–2 ns of the clock relative to the data. Adjust according to the material propagation delay.
5. In case of an external delay, the element device is used according to the delay in note #4 instead of a longer trace.

Table 23: Internal Delay on Receiving or on Transmitting Peer Side—Rx Path

Parameter	Min	Max	Tolerance / skew	Units	Notes
TL3 length		6	0.200	inch	1, 2, 3
TL4 length		6		inch	1, 2
TL3 - TL4 (skew)	-0.1	0.1		inch	
TL3 internal separation (dimensions)	2			Signal width	
TL3 internal separation (crosstalk)		-14		dB	
TL3 – TL4 cross group separation (dimensions)	4			Signal width	
TL3 – TL4 cross group separation (crosstalk)		-26		dB	

Note:

1. Maximum length is only a recommendation. For actual length refer to Peer/PHY specifications.
2. Target impedance 45–55Ω.
3. Use the tolerance/skew value to align the signal of the entire Input/Output bus (data and control).

6.3.4 Interface Calibration

Where possible, all output drivers should be configured to 30Ω.

7 Serial Management Interface (SMI)

The Serial Management Interface (SMI) interface complies with IEEE 802.3 Clause 22. However, some Marvell® devices support higher operating frequencies than defined in IEEE 802.3 Clause 22.

There are two modes of SMI: master and slave. The difference between the two modes is that the device, operating in master mode, drives the clock signal.

The SMI comprises a data signal (MDIO) and a clock signal (MDC).

7.1 Interface Connectivity

- MDC and MDIO are connected between the master and all relevant slaves.
- There must be only one master SMI per connectivity.
- MDIO requires a pull-up resistor on board.
- If the SMI voltage is different between the master and the slaves, a voltage level-shifter is required on board.

7.2 Interface Signals Layout Guidelines

7.2.1 General Tips and Guidelines

Minimize the load on the SMI interface.

The maximum amount of slave devices is frequency-dependant and should be determined according to simulations. The simulations should verify timing and clock waveform.

A common rule of thumb is: any MDC with frequency above 2.5 MHz should be connected to a maximum of four slave devices.

When the master device and the slave devices work at two different voltage levels, a level shift is required. Some voltage level shifters have a considerable delay, skew between several lanes, or commit only to a maximal amount of delay. When using level shifting, make certain not to violate the AC timing specifications.

7.2.2 MDIO Tips and Guidelines

The pull-up resistor should be lower than any internal pull-up/down. The external pull-up value on the MDIO signal is usually ~2.0 kΩ.

7.2.3 MDC Tips and Guidelines

Connecting MDC to several slave devices can cause signal integrity issues on the rising edge of the clock.

The user must run topology simulations to verify rising edge monotonic behavior to avoid any double-clocking concerns.

Figure 43 shows examples of waveforms representing good and bad clock transitions.

Figure 43: Clock Transition Examples

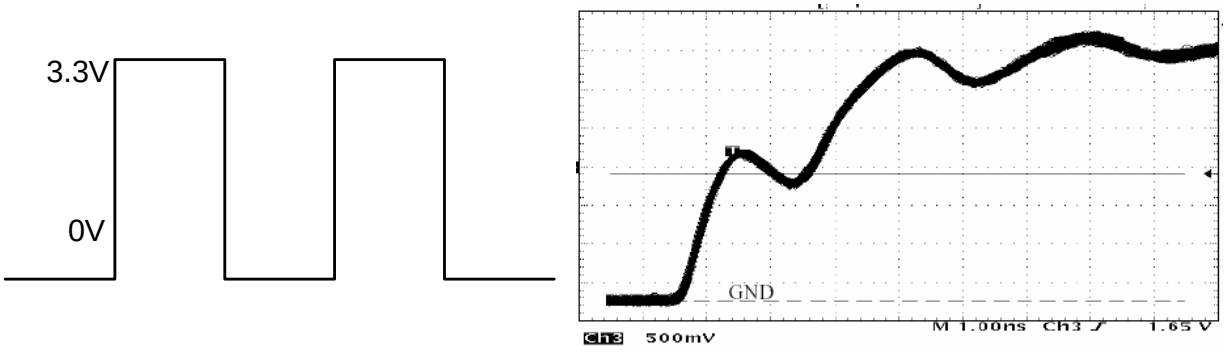

Good Clock Transition

Bad Clock Transition

Tips to avoid this phenomenon are:

- Use an on-board clock buffer.
- Use serial resistor to avoid reflections.
- Route the MDC in daisy chain topology and put the Thevenin termination at the end:
  - Make sure that the voltage levels are still acceptable with Thevenin terminations.
  - The stub's length to each slave device should be minimized.
  - Consider placing a resistor after the first slave device.
  - The Thevenin terminations should be placed after the last slave device.

Figure 44: Master SMI Connectivity Example

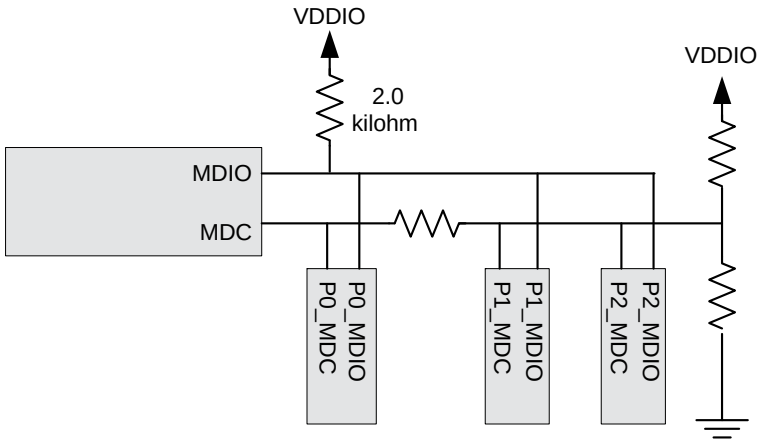

8 Serial ATA (SATA) Interface

This section provides the connectivity and layout guidelines for the Serial ATA interface pins.

The Serial ATA interface is a serial differential low-voltage, point-to-point interface. The interface consists of one differential pair for transmit and one for receive.

8.1 Connectivity

Table 24: Signal Groups

Group Name	Signal
Output Data Differential Pair	SATA_TX_P/N
Input Data Differential Pair	SATA_RX_P/N

Figure 45: SATA Interface Connectivity

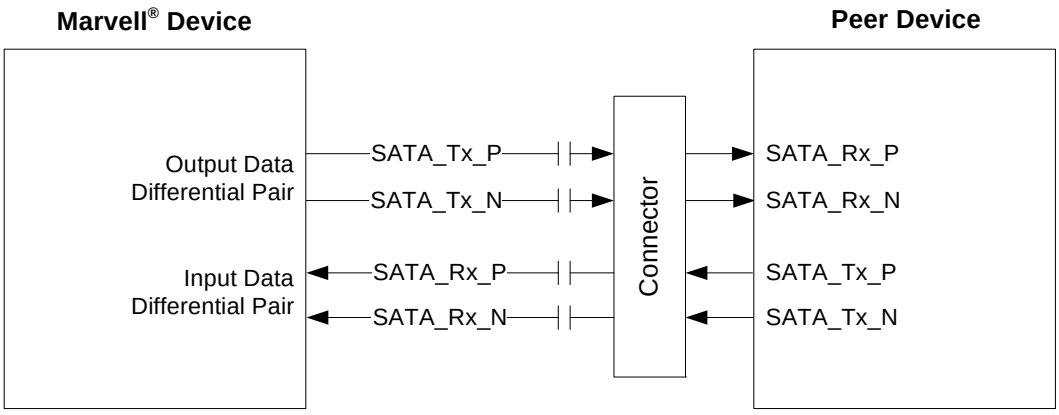

8.2 Interface Signals Layout Guidelines

It is essential to ensure return path continuity, minimize capacitive discontinuities, and to take into consideration the via structure effect. For further information, refer to Section 2, Generic Guidelines for SERDES Interfaces, on page 16.

8.2.1 Routing Topologies for Connection to a Standard SATA II 1-Meter Cable

When routing using a standard SATA II 1-meter cable refer to the connection diagram provided in Figure 46 and the routing constraints listed in Table 25.

Figure 46: Topology for Connection to a Standard SATA II 1-Meter Cable

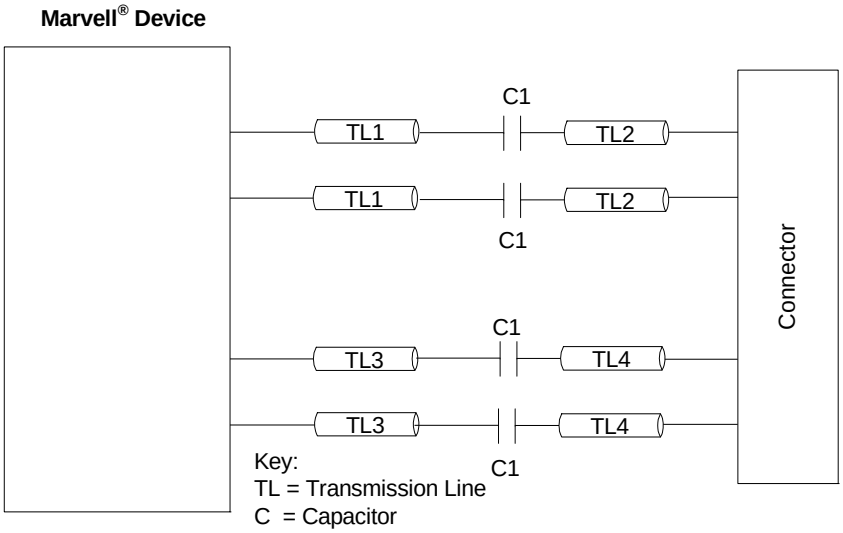

Table 25: Connection to a Standard SATA II 1-Meter Cable

Parameter	Min	Max	Skew	Units	Notes
TL1+TL2 length		4		inch	1, 2, 3, 4, 5, 6
TL1+TL2 loss @ 1.5 GHz		1.33		dB	3, 4, 7
TL1+TL2 loss @ 3 GHz		2.1		dB	3, 4, 7
TL2, TL4 length		0.4		inch	
TL3+TL4 length		4		inch	5, 6
TL1x to TL1y or TL3x to TL3y same direction lane separation (dimensions)	4.5			signal width	8, 9
TL1[x] to TL1[y] separation (far-end crosstalk below loss)		-26		dB	
TL1[x] or TL3[x] to TL[x] lane separation from other traces (dimensions)	5.5			signal width	10
TL1[x] or TL3[x] to TL[x] lane separation from other traces (crosstalk)		-40		dB	10
TL3+TL4 loss @ 1.5 GHz		1.33		dB	3, 4, 7
TL3+TL4 loss @ 3 GHz		2.1		dB	3, 4, 7

Table 25: Connection to a Standard SATA II 1-Meter Cable (Continued)

Parameter	Min	Max	Skew	Units	Notes
TL2+TL1, TL3+TL4 inner pair skew		10		mil	11
Tn differential impedance	100		±10%	Ω	12, 13
Tn single-ended impedance	50	55		Ω	12
C1	10			nF	14, 15

Notes:

1. Target impedance should be 100Ω differential; any mismatch should be included as a part of the loss.
2. Derived from the Loss value by calculation with 0.33 dB/inch.
3. Loss budget is meant for trace loss and should include far-end crosstalk.
4. Loss budget is designed to allow interoperability with an attached standard 1m SATA II cable. For special SATA II cable length or characteristics, high-speed simulations are required.
5. Discontinuities, such as capacitive type, can cause significant insertion loss degradation, causing the assumption, taken as loss per inch, to be smaller than actual.
6. The maximal allowed trace length can vary according to the material and geometrical characteristics.
7. The interconnect loss should have a smooth progression through frequency with no notch-like behavior up to 3 GHz (Fbaud).
8. The separation is calculated assuming the same type of SERDES and typical board layer stack-up. Make certain to verify that the crosstalk limitation is met in each specific system.
9. The separation value is calculated for a single-ended impedance of 55-ohm, a differential impedance of 100-ohm, u-strip trace length coupled along the entire end-to-end path, and a dielectric loss coefficient of 0.027. If the loss coefficient is lower, a higher separation may be required.
10. The allowed crosstalk is highly dependant of the aggressor's waveform (rise time and swing). The recommended separation value should account for all common aggressor types.
11. For BGA packages, some package length routing skew may exist and help compensate for pin position within the pinout.
12. The target impedance priority is to match the differential impedance of these signal traces and adjust the single-ended impedance accordingly.
13. The tolerance is manufacturing tolerance.
14. The capacitor is mandatory for SATA II and recommended for SATA 1. The maximum capacitor value is 12 nF.
15. A low ESL high-speed capacitor.

8.2.2 End-to-End Connection Parameters for a Well-Defined Connection Cable or Board

This topology reflects a generic situation in cases where total loss is well defined and fixed, for example, a very short cable connecting to a long backplane.

When routing using an end-to-end connection refer to the connection diagram provided in Figure 47 and the routing constraints listed in Table 26.

Figure 47: End-to-End Connection Parameters for a Well-Defined Connection Cable or Board

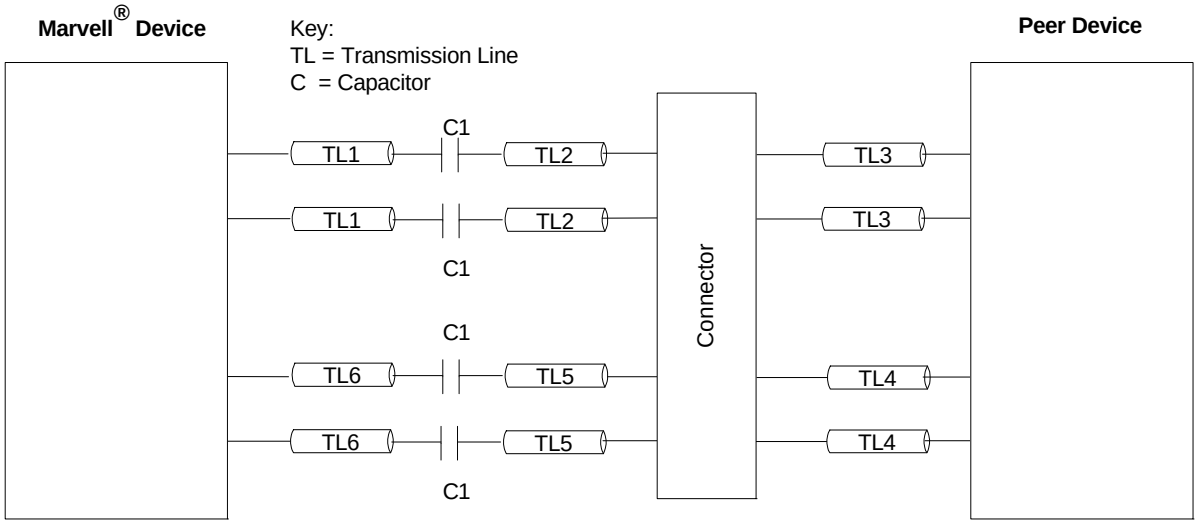

Table 26: End-to-End Connection Parameters for a Well-Defined Connection Cable or Board

Parameter	Min	Max	Skew	Units	Notes
TL1+TL2+TL3 length		18	±0.025	inch	1, 2, 3, 4, 5, 6, 7, 8
TL2, TL6 length		0.4		inch	9
TL1x to TL1y or TL3x to TL3y same direction lane separation (dimensions)	4.5			signal width	10, 11
TL1[x] to TL1[y] separation (far-end crosstalk below loss)		-26		dB	
TL1[x] or TL3[x] to TL[x] lane separation from other traces (dimensions)	5.5			signal width	12
TL1[x] or TL3[x] to TL[x] lane separation from other traces (crosstalk)		-40		dB	12
TL1+TL2+TL3 loss @ 0.3 GHz		2.8		dB	
TL1+TL2+TL3 loss @ 0.75 GHz		5.4		dB	
TL1+TL2+TL3 loss @ 1.5 GHz		7		dB	3, 13, 14

Table 26: End-to-End Connection Parameters for a Well-Defined Connection Cable or Board (Continued)

Parameter	Min	Max	Skew	Units	Notes
TL1+TL2+TL3 loss @ 3.0 GHz		10.7		dB	3, 13, 14
Maximum interconnect ISI		4.5		dB	14
TL4+TL5+TL6 length		18	±0.025	inch	1, 2, 3, 4, 5, 6, 7, 8
TL4+TL5+TL6 loss @ 0.3 GHz		2.8		dB	
TL4+TL5+TL6 loss @ 0.75 GHz		5.4		dB	
TL4+TL5+TL6 loss @ 1.5 GHz		7		dB	3, 13, 14
TL4+TL5+TL6 loss @ 3.0 GHz		10.7		dB	3, 13, 14
TL4+TL5 +TL6, TL1+TL2 +TL3 inner pair skew		25		mil	5, 15
Tn Differential impedance	100		±10%	Ω	16
Tn single-ended impedance	50	55		Ω	16
C1	10			nF	17, 18

Notes:

1. Target impedance should be 100Ω differential; any mismatch should be included as a part of the loss.
2. Derived from the Loss value by calculation with 0.33 dB/inch @ Fbaud/2.
3. Loss budget is meant for trace loss and should include far-end crosstalk.
4. Loss budget is designed to allow interoperability with an attached standard 1m SATA II cable.
5. When using a 1-meter standard cable and TL3 and TL4 are board traces, refer to Section 8.2.1, Routing Topologies for Connection to a Standard SATA II 1-Meter Cable, on page 86.
6. Connector loss of 1 dB @ Fbaud/2 was taken into consideration.
7. Discontinuities, such as capacitive type, can cause significant insertion loss degradation, causing the assumption, taken as loss per inch, to be smaller than actual.
8. The maximal allowed trace length can vary according to the material and geometrical characteristics.
9. The AC coupling capacitor should be placed as close as possible to the receiver or as close as possible to the connector. Therefore, TL2 and TL6 may represent the trace to the receiver or the trace to the connector. If no connector exists, place as close to the receiver as possible, If a connector exists, Marvell® recommends placing the cap close to the connector.
10. The separation is calculated assuming the same type of SERDES and typical board layer stack-up. Make certain to verify that the crosstalk limitation is met in each specific system.
11. The separation value is calculated for a single-ended impedance of 55-ohm, a differential impedance of 100-ohm, u-strip trace length coupled along the entire end-to-end path, and a dielectric loss coefficient of 0.027. If the loss coefficient is lower, a higher separation may be required.
12. The allowed crosstalk is highly dependant of the aggressor's waveform (rise time and swing). The recommended separation value should account for all common aggressor types.
13. The interconnect loss should have a smooth progression through frequency with no notch-like behavior up to 3 GHz (Fbaud).
14. Interconnect ISI is defined as:
    - (TL1+TL2 loss @ 1.5 GHz) - (TL1+TL2 loss @ 0.3 GHz)
      or
    - (TL3+TL4 loss @ 1.5 GHz) - (TL3+TL4 loss @ 0.3 GHz)
15. For BGA packages, some package length routing skew may exist and help compensate for pin position within the pinout.

16. The target impedance priority is to match the differential impedance of these signal traces and adjust the single-ended impedance accordingly.
17. The capacitor is mandatory for SATA II and recommended for SATA 1. The maximum capacitor value is 12 nF.
18. A low ESL high-speed capacitor.

8.3 Clock Considerations

The SATA reference clock is usually jitter sensitive. To minimize clock-source, power-noise oriented jitter, Marvell recommends filtering the clock source's power supply. For power supply filtering recommendations, refer to Section 16, Generic Power Board Guidelines, on page 114.

8.4 Power Considerations

The power pins are analog type. For power supply filtering recommendations, refer to Section 16, Generic Power Board Guidelines, on page 114.

8.5 Specific Signals

8.5.1 SATA Routing Guidelines for the SERDES_ISET Signal

- Place the ISET resistor up to 400 mil from the SERDES_ISET pin.
- For resistor parameters, refer to the device Hardware Specifications.

9 Universal Serial Bus (USB) Interface

The USB port contains a dual role embedded USB 2.0 PHY supporting Host and Peripheral (device) modes.

As a USB Host, the USB interface supports operating in Low-Speed, Full-Speed, and High-Speed modes. As a USB Device, USB interface supports operating in Full-Speed and High-Speed modes.

The USB signals on the board consist of one bidirectional differential pair.

9.1 Interface Connectivity

9.1.1 Signal Groups

Table 27 shows the USB interface pins, by connectivity group.

Table 27: USB Interface Pin Connectivity Groups

Group Name	Signal	Comments
Data	USBx_DP USBx_DM	x = 0, 1, 2… according to the number of USB interface ports integrated in the device.
Control	Vbus	The Vbus has two main features: • Providing a power supply from a USB host to a USB device. • In device mode, sensing the existence of a USB connection. The Vbus can be multiplexed on one of the GPIO/MPP pins.

9.1.2 Data Connectivity

To protect high speed data lines from Electrical Static Discharge (ESD), Electrical Fast Transient (EFT), and lightning, Marvell® recommends adding an additional ESD/EMI (Electromagnetic Interference) protection circuit (see Figure 48).

Figure 48: USB 2.0 Interface Connectivity to a Connector—Including Optional Common Mode Choke and Protection Diodes Circuitry

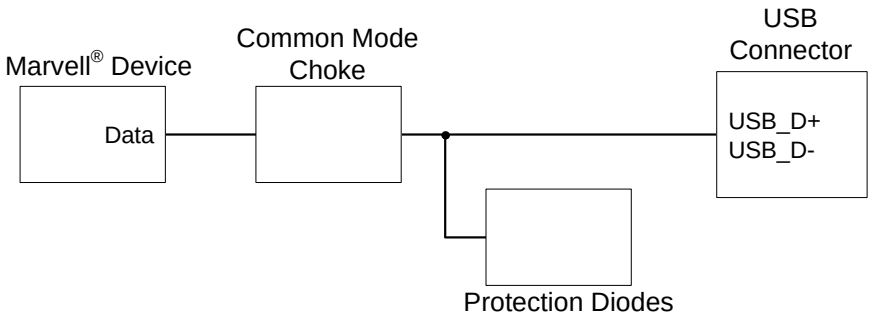

9.1.3 Control Connectivity

When in Host mode, the 5V power is supplied by the board. The board design must protect the 5V power supply from a possible short by using a current limiter as shown in Figure 49.

Figure 49: Vbus Connectivity when Configured as a Host

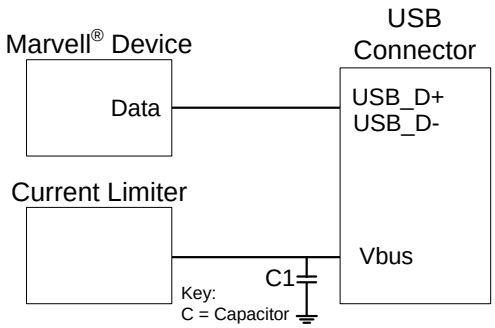

NOTE: The minimal value for C1 should be no less than 120 µF of low-ESR capacitance per hub, as specified in the *Universal Serial Bus Specification*, Revision 2.0.

When the USB port is configured to function as a device, the Vbus can indicate the presence of other devices. To support this sensing functionality, a voltage divider is required, as shown in Figure 50.

Figure 50: Vbus Connectivity when Configured as a Device

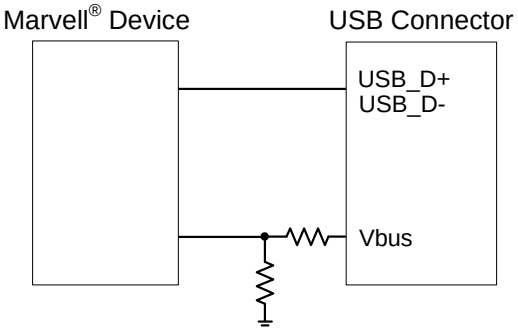

9.2 Interface Signals Layout Guidelines

It is essential to ensure return path continuity, minimize capacitive discontinuities, and take into consideration the via structure effect. For further information, refer to Section 2, Generic Guidelines for SERDES Interfaces, on page 16.

9.2.1 Connectivity to a Connector Including Optional Circuitry when Configured as a Device

Figure 51: Topology—Including Common Mode Choke and Protection Diodes Optional Circuitry—when Configured as a Device

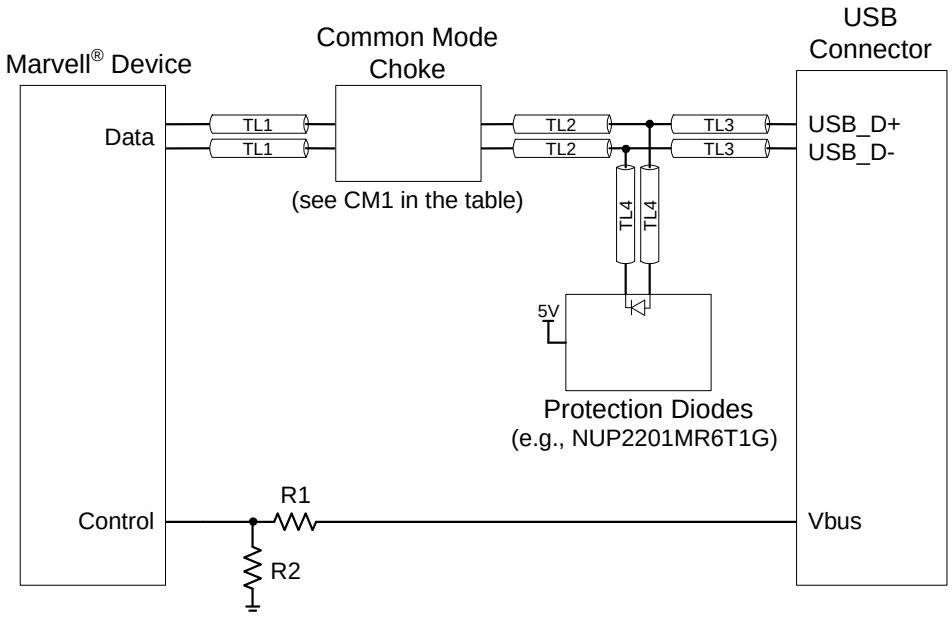

Table 28: Routing Constraints when Configured as a Device

Parameter	Min	Typ	Max	Tolerance / Skew	Units	Notes
TL1+TL2+TL3 length			5		inch	1, 2, 3
TL2+TL3 length			0.6		inch	
TL3+TL4 length			0.4		inch	
TL4 length			0.1		inch	
TL1+TL2+TL3 inner pair skew			150		mil	4
TL[x] differential impedance		90		±10%	ohm	5, 6
TL[x] common mode impedance		30		±20%	ohm	5, 7
R1: R2		1:2			-	
R1	10				kΩ	
CM1 Common mode impedance @ 100 MHz		45			ohm	8, 9
CM1 Common mode impedance @ 500 MHz		70			ohm	8, 9
CM1 DC resistance			0.15		ohm	8

NOTES

1. The target impedance should be 90 ohms differential; any mismatch should be taken as a part of the loss.
2. In the USB routing, it is essential to avoid capacitive discontinuities and via through-hole long stubs. Via stubs can occur on vias or on the USB connector pins. For further information refer to Section 2, Generic Guidelines for SERDES Interfaces, on page 16.
3. Discontinuities, such as capacitive type, can cause significant insertion loss degradation, causing the assumption, taken as loss per inch, to be smaller than actual.
4. In BGA packages, some package length routing skew may exist and help compensate for the ball position within the package.
5. The target impedance priority is to match the differential impedance of these signal traces and adjust the single-ended impedance and the common mode impedance, accordingly.
6. The tolerance is manufacturing tolerance.
7. Common mode impedance is the impedance measured when stimulating a pair of traces with identical signals.
8. Examples of a common mode (CM) filter (choke) part numbers are: Coilcraft 0603USB-251ML and Coilcraft 0805USB-421ML.
9. Common mode chokes can distort signal quality. The signal distortion increases as common mode impedance increase. Different common mode impedance may be needed to pass EMI testing. The suggested common mode choke parts were proven by laboratory testing to retain an acceptable signal waveform.

9.2.2 Connectivity to a Connector Including Optional Circuitry when Configured as Host

Figure 52: Topology—Including Optional Common Mode Choke and Protection Diodes Circuitry—when Configured as a Host

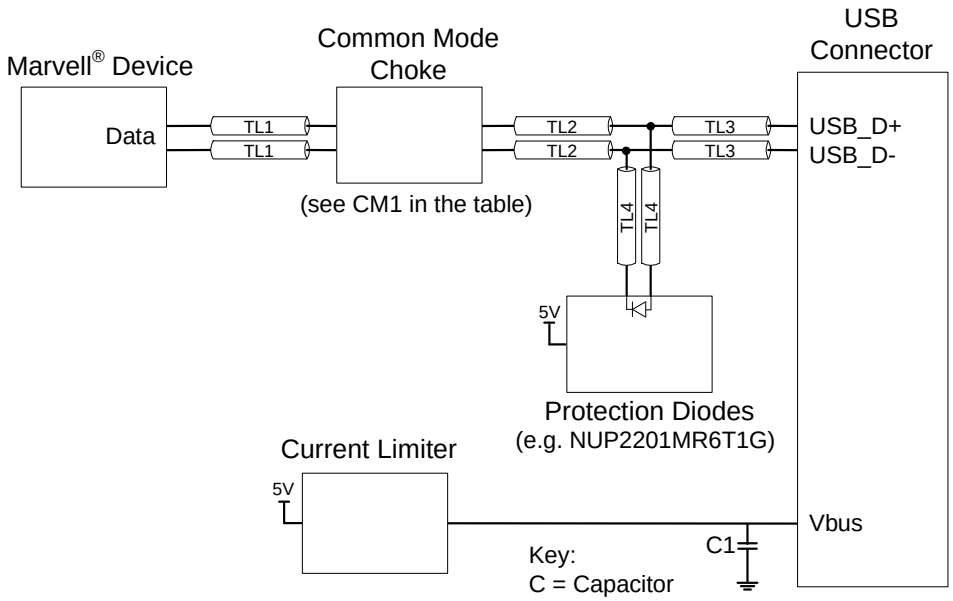

NOTE: The minimal value for C1 should be no less than 120 µF of low-ESR capacitance per hub, as specified in the *Universal Serial Bus Specification*, Revision 2.0.

Table 29: Routing Constraints when Configured as a Host

Parameter	Min	Typ	Max	Tolerance / Skew	Units	Notes
TL1+TL2+TL3 length			5		inch	1, 2, 3
TL2+TL3 length			0.5		inch	
TL3+TL4 length			0.3		inch	
TL4 length			0.1		inch	
TL1+TL2+TL3 inner pair skew			150		mil	4
TL[x] differential impedance		90		±10%	ohm	5, 6
TL[x] common mode impedance		30		±20%	ohm	5
CM1 Common mode impedance @ 100 MHz		45			ohm	7, 8
CM1 Common mode impedance @ 500 MHz		70			ohm	7, 8
CM1 DC resistance			0.15		ohm	7

NOTES

1. The target impedance should be 90 ohms differential; any mismatch should be taken as a part of the loss.
2. In the USB routing, it is essential to avoid capacitive discontinuities and via through-hole long stubs. Via stubs can occur on vias or on the USB connector pins. For further information refer to Section 2, Generic Guidelines for SERDES Interfaces, on page 16.
3. Discontinuities, such as capacitive type, can cause significant insertion loss degradation, causing the assumption, taken as loss per inch, to be smaller than actual.
4. In BGA packages, some package length routing skew may exist and help compensate for the ball position within the package.
5. The target impedance priority is to match the differential impedance of these signal traces and adjust the single-ended impedance and the common mode impedance, accordingly.
6. The tolerance is manufacturing tolerance.
7. Examples of a common mode (CM) filter (choke) part numbers are: Coilcraft 0603USB-251ML and Coilcraft 0805USB-421ML.
8. Common mode chokes can distort signal quality. The signal distortion increases as common mode impedance increase. Different common mode impedance may be needed to pass EMI testing. The suggested common mode choke parts were proven by laboratory testing to retain an acceptable signal waveform.

9.3 Power Considerations

The power pins are analog type. For power supply filtering recommendations, refer to Section 16, Generic Power Board Guidelines, on page 114.

9.4 Specific Signals

9.4.1 USB Routing Guidelines for the SERDES_ISET Signal

Follow these guidelines for SERDES_ISET signal placement:

- Place the ISET resistor up to 400 mil from the SERDES_ISET pin.
- For resistor parameters, refer to the Hardware Specifications.

10 Sony/Philips Digital Interface (S/PDIF)

The S/PDIF interface (IEC-958) is a “consumer” version of the Audio Engineering Society (AES) and the European Broadcasting Union (EBU) AES/EBU interface.

The electrical interface for S/PDIF signals can be either a 75 ohm coaxial cable or an optical fiber (usually called TOSLINK).

10.1 Interface Connectivity

Table 30: Signal Groups

Group Name	Signal
Output Data	S/PDIF out (SPDIFO)

10.2 Connectivity

Figure 53: Fiber Optical Connection

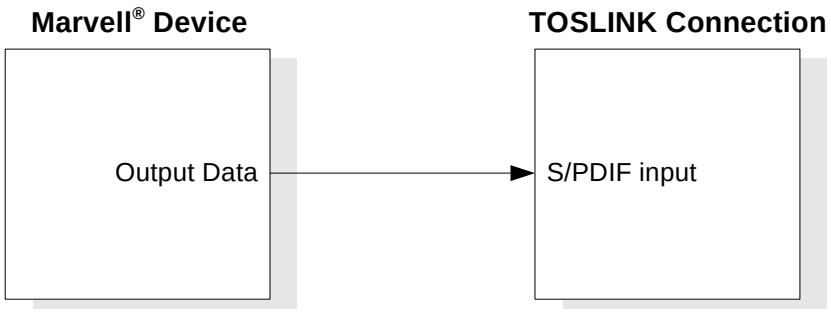

Figure 54: Coaxial Cable Connection

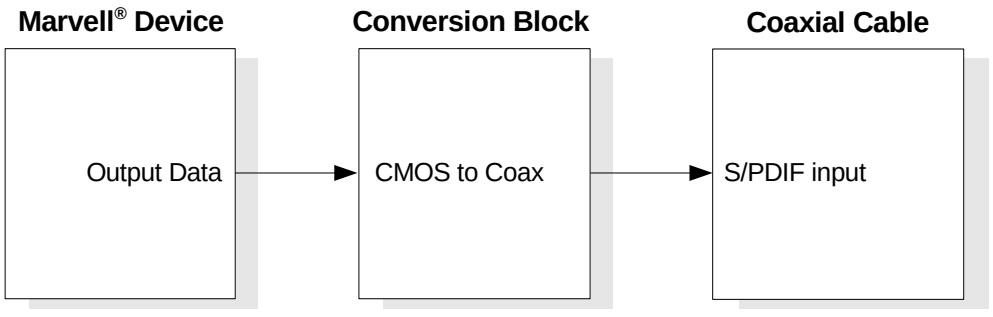

10.3 Interface Signals Layout Guidelines

10.3.1 Routing Topologies

- Coaxial S/PDIF connections, using good 75 ohm coaxial cable, typically work at distances of up to 10–15 meters.
- The output impedance is 75 ohm.
- Ordinary coaxial cable designed for video applications can be used.
- The minimal input level of the S/PDIF interface is 200 mV that allows some cable loss.
- A common 75 ohm cable can be used, eliminating the need for special quality cable (a good video accessory cable works also as a good S/PDIF cable).

The circuit in Figure 55 is presented as a building block that has at one end an S/PDIF standard signal and at the other end a CMOS level signal.

10.3.2 S/PDIF Output Buffer Circuits

Figure 55 is an output circuit that takes in a CMOS level S/PDIF signal and outputs a standard S/PDIF coax cable output signal.

Figure 55: CMOS-to-S/PDIF COAX Conversion Circuit

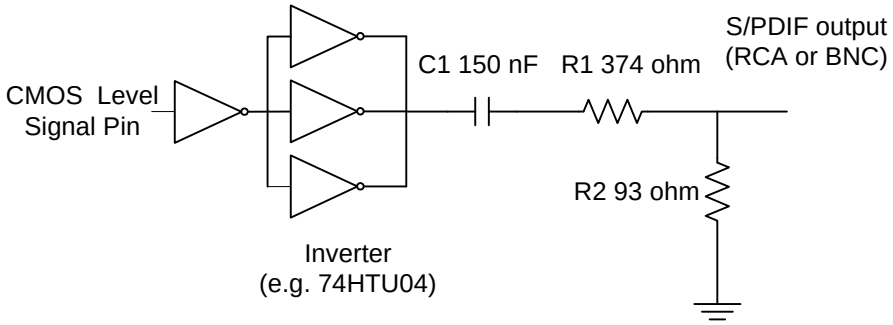

This circuit does not provide any isolation on the output. The CMOS level signal from the inverter (74HCT04 buffer) is attenuated to the levels specified in the S/PDIF optical interface by using the voltage divider consisting of 374 ohm and 93 ohm resistors.

11 Liquid Crystal Display (LCD) Interface Guidelines

The LCD interface is a parallel SDR source synchronous interface that can be connected to several display options:

- LCD panel
- HDMI or DVI transmitter
- LCD-to-LVDS converter
- LCD-to-VGA converter
- LCD-to-SVideo (component or composite) converter
- Other Marvell® devices, such as the DE2700 family of adaptive digital video format converters

11.1 Interface Connectivity

LCD data bus and clock are connected between the transmitter and all relevant receivers.

There must be only one transmitter LCD per connectivity.

According to the device's color depth, some applications may not use all of the data bus connectivity pins.

Table 31: Signal Groups

Group Name	LCD Express	Comments
Data bus	LCD_D[23:0]	
Control bus	LCD_E LCD_HSYNC LCD_VSYNC	
Clock	LCD_CLK	

Figure 56 shows the LCD port point-to-point connection to an external LCD device.

Figure 56: LCD Port Connectivity for Point-to-Point Connection

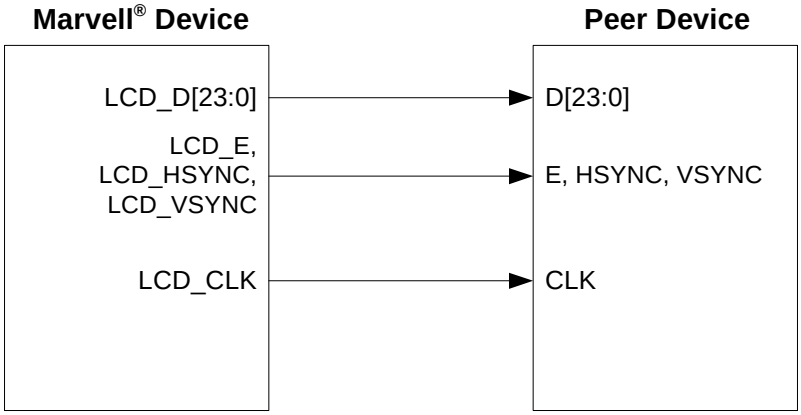

11.2 Interface Signals Layout Guidelines

11.2.1 General Tips and Guidelines

Marvell recommends following these general layout guidelines:

- Minimize the load on the LCD interface. If necessary, use buffers to minimize the load.
- The maximum number of receiver devices is frequency-dependant and should be determined according to simulations. The simulations should verify the timing and clock waveform.
- If an LCD interface is used with a frequency greater than 100 MHz, it should be connected to a maximum of two slave devices.
- When the transmitter device and the receiver devices work at two different voltage levels, a level shift is required. Some voltage level shifters have a considerable delay, skew between several lanes, or commit only to a maximal amount of delay. When using level shifting, do not violate the AC timing specifications provided in the Marvell® device Hardware Specifications, and in the target device specification.

11.2.2 LCD Clock Tips and Guidelines

Marvell® recommends following these clock layout guidelines:

- Connecting a clock to several receiver devices may cause signal integrity issues on the rising edge of the clock.
- The user must run topology simulations to verify rising-edge monotonic behavior to avoid any double-clocking concerns.

Tips to avoid double-clocking:

- Consider adding an on-board clock zero delay PLL buffer, if there is a high load.
- Use a serial resistor to avoid reflections.
- Route the clock using a daisy chain topology and put the Thevenin terminations at the end:
  - Make sure that the voltage levels are still acceptable with the Thevenin terminations.
  - The stub's length to each slave device should be minimized.

- Consider placing a resistor after the first slave device.
- The Thevenin terminations should be placed after the last slave device.

11.2.3 Routing Topologies

All signals are connected point-to-point.

Figure 57: LCD Point-to-Point Routing Topology

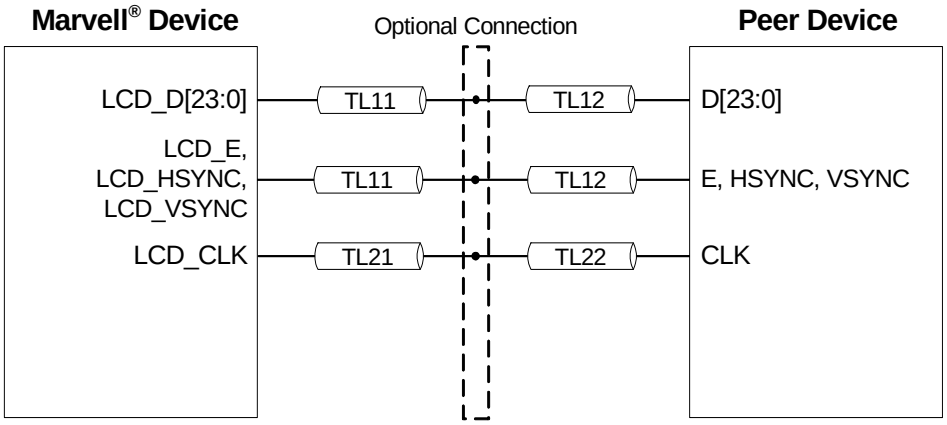

Table 32 lists the routing constraints when using a daisy chain topology.

Table 32: Routing Constraints when Using Daisy Chain Topology

Parameter	Min	Max	Tolerance / Skew	Units	Notes
TL[x] internal-group separation (dimensions)	2			Signal width	1, 2
TL[x] internal-group separation (crosstalk)		-20		dB	1, 2
TL1[x] to TL2[y] cross-group separation (dimensions)	3			Signal width	1, 2, 3
TL1[x] to TL2[y] cross-group separation (crosstalk)		-26		dB	1, 2, 3
TL11, TL21 length		5		inches	4
TL12, TL22 length		12		inches	5
		7.5		inches	6
Data Bus and Control Bus to Clock (TL11+TL12) – (TL21+TL22)		0.75		inches	
TL[x] impedance (all traces)	50		10%	Ohm	7

NOTES

1. If the design cannot meet the above requirements for length and separation, simulation can be performed to ensure proper isolation between signals.
2. The separation is defined from edge-to-edge of two traces.

3. Separation is required between any signals that are not in the same group.
4. For longer trace lengths, a zero-delay clock PLL buffer is required.
5. For a flat cable with no more than 120 ps/inch delay time.
6. For FR4 board connectivity.
7. The tolerance for the matched impedance value is the manufacturing tolerance. The design tolerance should be kept to a minimum.

12 Video Grid Array (VGA) Interface

The design guidelines presented in this section are relevant for Video Grid Array (VGA) connection from a Marvell® device to a VGA standard connector (HVRGB).

12.1 Interface Connectivity

Table 33 lists the VGA interface signals by connectivity group.

Table 33: Signal Groups

Group Name	Signal	Notes
Data	VGA_R, VGA_G, VGA_B	Analog pins
Control	VGA_Hsync, VGA_Vsync	Digital pins
Analog Ground	VGA_RGND, VGA_GGND, VGA_BGND	Return path for analog pins
Digital Ground	VGA_SGND, VGA_GND	Return path for digital pins

Figure 58 depicts VGA interface connectivity.

Figure 58: VGA Interface Connectivity

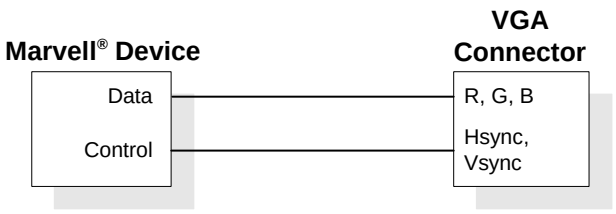

12.2 Interface Signals Layout Guidelines

Figure 59 depicts the VGA connector topology and Table 34 lists the routing constrains for this connection.

Figure 59: VGA Connector Topology

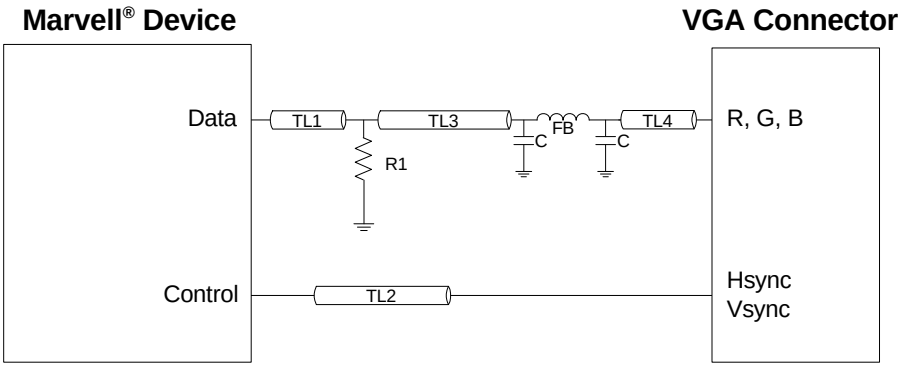

Table 34: Routing Constraints for VGA Connection

Parameter	Min	Typ	Max	Skew/ Tolerance	Units	Notes
TL[x] internal-group separation (dimensions)	2				Signal Width	1
TL[x] internal-group separation (crosstalk)		-20			dB	1
TL[x] to TL[y] cross-group separation (dimensions)	3				Signal Width	1
TL[x] to TL[y] cross-group separation (crosstalk)		-26			dB	1
TL2, TL3 length			10	±0.5	inch	2
TL4 length			500		mil	
Control to Data skew TL2 – (TL1+TL3+TL4)			1		inch	2
TL1 length			0.5		inch	
TL1, TL3, TL4 single-ended impedance	50	75	80	±10%	ohm	2
TL2 single-ended impedance		50		±10%	ohm	2
R1		75		±5%	ohm	3
C and FB – RGB signal filters						4

Notes:

1. If the design cannot meet the above requirements for length and separation, simulation can be performed to ensure proper isolation between signals.

2. The tolerance for the matched impedance value is the manufacturing tolerance. The design tolerance should be kept to a minimum.
3. Connecting R1 should not add a stub longer than 100 mils.
4. The RGB signal filters components' values can vary according to the specific system requirements. an example of such a filter can be: C = 3.3 pF and FB = 75 ohm @ 100 MHz

12.3 Analog Signal Reference Plane and Return Path Considerations

The following analog signal reference plane and return path factors should be considered when designing the system:

- Form a reference plane using a continuous ground plane.
- Use the reference plane as the sole reference plane and return path of the analog signals (data group).
- Connect the reference ground plane to the VGA analog ground pins of the Marvell device package.
- Connect the distant edge of the return path plane to the D-Sub connector ground pins or the analog ground pins of the peer device.
- Form ground shielding around the RGB signals. Between the shielding and the RGB traces, keep the spacing of at least: 5 x the height above the reference plane. Stitch the shielding to the reference plane with vias, while keeping the distance between the vias lower than 0.5".

Note

- Connect the analog ground to connector pins 6, 7, and 8, the red, green, and blue grounds, respectively (see Figure 60).
- Connect the digital ground to connector pins 5 and 10, the Hsync and Vsync grounds, respectively.

Figure 60: Example of D-Sub Connector

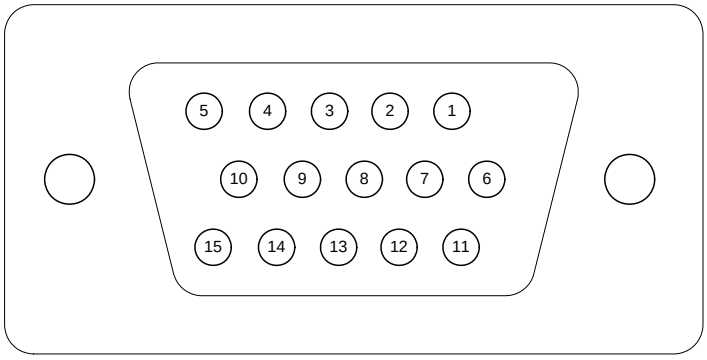

12.4 Power Considerations

The power pins are analog type. For power supply filtering recommendations, refer to Section 16, Generic Power Board Guidelines, on page 114.

12.5 Specific Signals

12.5.1 VGA Routing Guidelines for the VGA_REXT Signal

- Place the resistor up to 400 mils from the VGA_REXT pin.
- For resistor parameters, refer to the Hardware Specifications.
- Separation between VGA_REXT and all other signals is at least three times the width of the signal.

13 NAND Flash Interface Guidelines

The design guidelines presented in this section are relevant for on-board NAND Flash devices, using a one-device load or a two-device load. When two 8-bit NAND flash devices are connected to a single Marvell 16-bit NAND flash interface, this is referred to as a NAND flash gang configuration.

13.1 Interface Connectivity

13.1.1 Signal Groups

Table 35 lists the signal groups.

Table 35: NAND Flash Signal Groups

Group Name	Nand Flash Interface	Comments
Input/Output[0]	NF_IO[7:0]	
Input/Output[1]	NF_IO[15:8]	Used in Gang mode or with 16-bit flash devices.
Command	NF_CLE, NF_ALE, NF_REn, NF_WEn	
Control0	NF_CEn[0], NF_CEn[2], NF_RDY[0]	
Control1	NF_CEn[1], NF_CEn[3], NF_RDY[1]	Some devices may not have more than one control bus.

13.1.2 Connectivity with Two 8-Bit Flash Devices in Gang Mode

Figure 61 depicts the connection for two 8-bit flash devices, connected in Gang mode.

Figure 61: Connectivity with Two 8-Bit Flash Devices Connected in Gang Mode

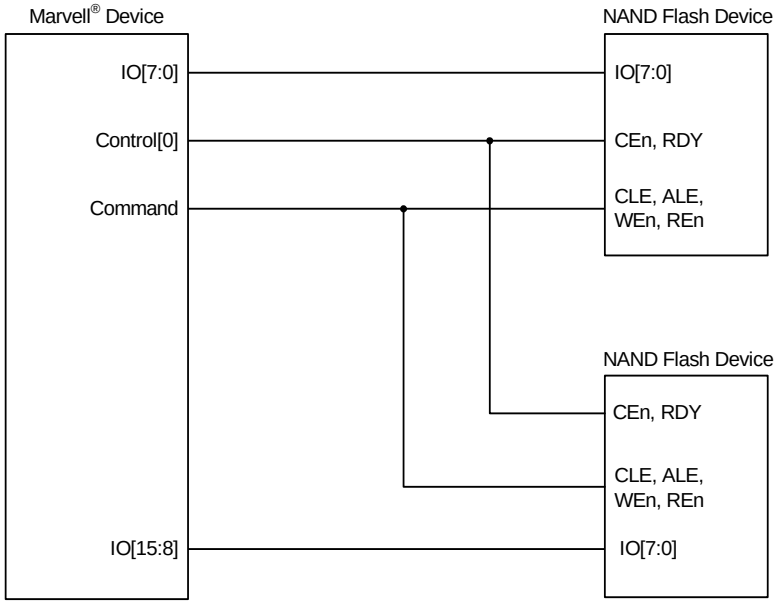

13.1.3 Connectivity with Four 8-Bit Devices in Gang Mode

Figure 62 depicts four 8-bit wide on-board NAND Flash devices connected in Gang mode.

Figure 62: Connectivity with Four 8-Bit Devices in Gang Mode

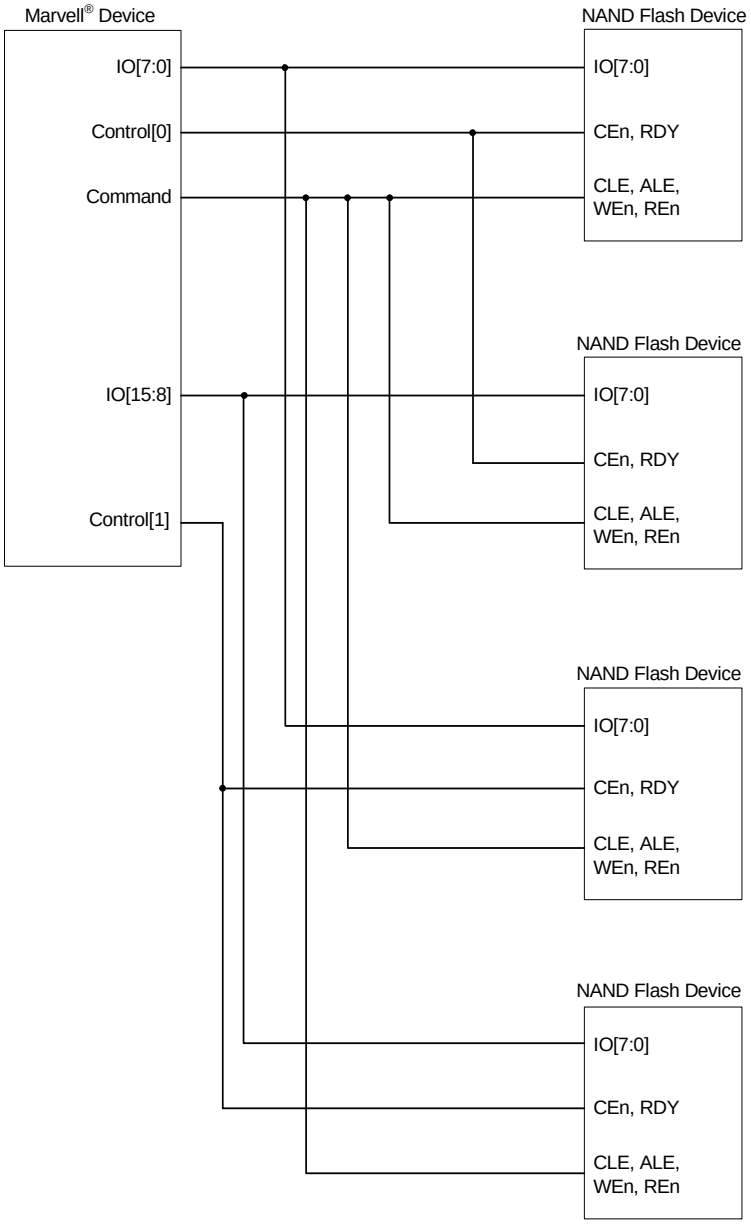

13.2 Interface Signals Layout Guidelines

A swap can be performed between IO[0] and IO[1].

13.2.1 Layout Guidelines for Two 8-Bit Flash Devices in Gang Mode

Figure 63 depicts the signal layout for a two 8-bit flash devices in Gang mode.

Figure 63: Layout Guidelines for Two 8-Bit Flash Devices in Gang

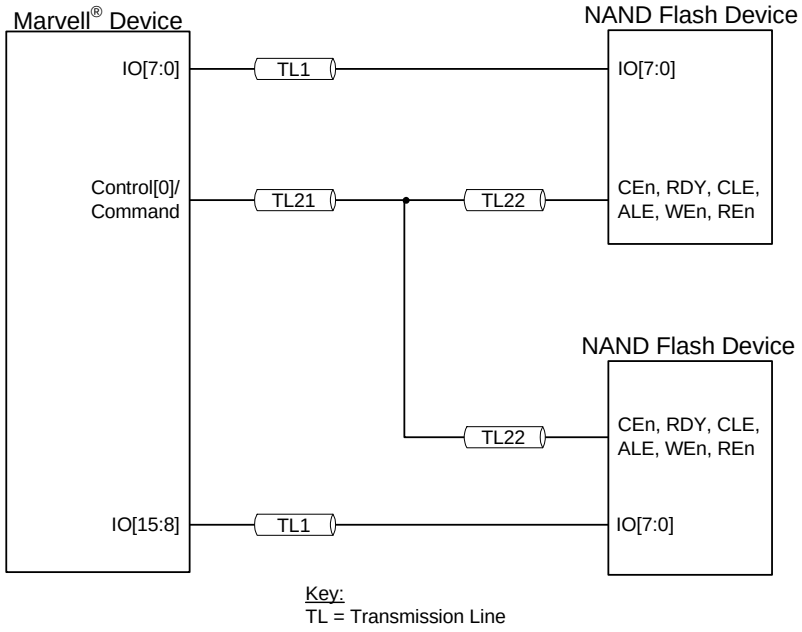

Table 36 lists the routing constraints for a two 8-bit flash devices in Gang mode.

Table 36: Routing Constraints Using Two Flash Devices in Gang Mode Topology

Parameter	Min	Max	Tolerance / Skew	Units	Notes
TL[x] internal-group separation (dimensions)	2			Signal width	1, 2
TL[x] internal-group separation (crosstalk)		-20		dB	1, 2
TL[x] to TL[y] cross-group separation (dimensions)	3			Signal width	1, 2, 3
TL[x] to TL[y] cross-group separation (crosstalk)		-26		dB	1, 2, 3
TL[22] length		Min (0.2* TL[21], 1)		inch	
TL[x] impedance (all traces)		50	10%	ohm	4

Table 36: Routing Constraints Using Two Flash Devices in Gang Mode Topology

Parameter	Min	Max	Tolerance / Skew	Units	Notes
IO to NF_REn/NF_WEn TL[1] – NF_WEn TL[1] – NF_REn	-2	+2		inch	
Command to NF_REn/NF_WEn TL[21+TLL22] – NF_WEn TL[21]+TL[22] – NF_REn	-2	+2		inch	

NOTES

1. If the design cannot meet the above requirements for length and separation, simulation can be performed to ensure proper isolation between signals.
2. The separation is defined from edge-to-edge of two traces.
3. Separation is required between any signals that are not in the same group.
4. The tolerance for the matched impedance value is the manufacturing tolerance. The design tolerance should be kept to a minimum.

13.3 Special Software Settings

Refer to the functional specification for register settings to get maximum performance.

14 JTAG Connection Information

To work with external JTAG hardware, note the following considerations during the board design phase.

- Figure 64 shows the recommended JTAG connection. Use this configuration on any board that has a JTAG interface. Use this configuration whether the JTAG interface is used or not used.
- When the JTAG is used, there is an external JTAG controller connected to the JTAG connector. All of the JTAG signals are driven by that JTAG controller.

Figure 64: JTAG Connection

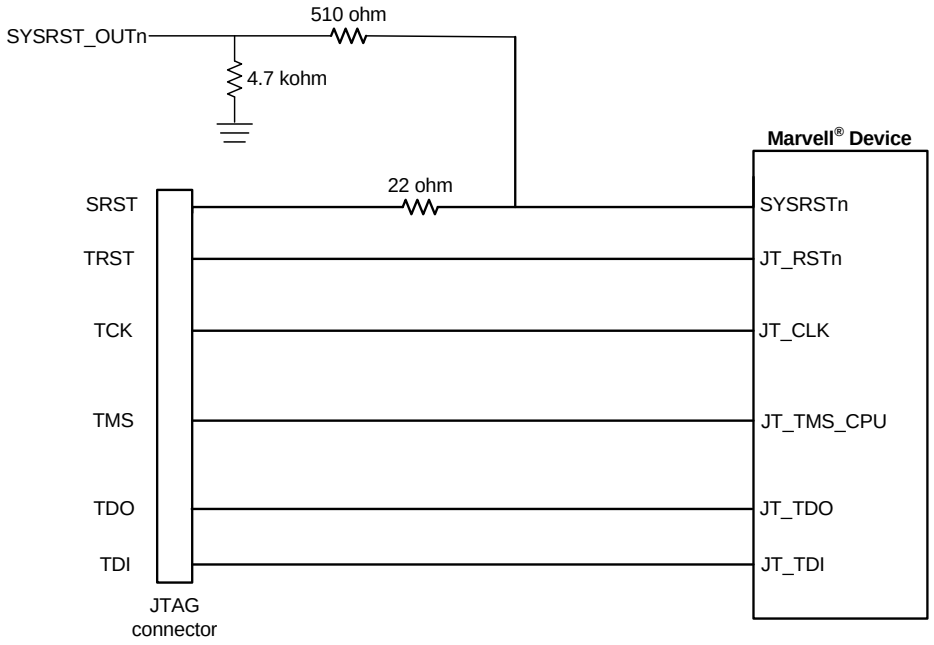

15 Unused Interface Guidelines

The devices are versatile devices with interfaces that support different types of user applications. Most applications and designs do not use all of the device interfaces. If an interface is left unused, some of its signals cannot be left unconnected and floating.

The “Unused Interface” section in the Hardware Specification (see Related Documentation on page 14) provides very specific instructions about what needs to be done with the signals of any unused interface.

Caution It is necessary to follow the instructions Marvell® provides for unused interfaces. If these guidelines are not followed, it might lead to unpredictable results, and may even cause damage to the device.

16 Generic Power Board Guidelines

Each device requires adequate power filtering, decoupling, and bypassing. The power delivery system can include:

- Analog power filtering
- Core/CPU power decoupling
- I/O power bypassing

This section refers to both Core and CPU power. The CPU power guidelines are only relevant for device that contain an embedded CPU.

16.1 Generic Power Network Guidelines

16.1.1 Loop Inductance

Loop inductance is the main physical characteristic that usually limits power supply network capability to supply current from a specific frequency and above. It is essential to design a power supply network with low connectivity loop inductance.

Loop inductance is composed of the following:

- The effective loop inductance of the vias connecting the device pins to the capacitors.
- The equivalent serial inductance (ESL) of the decoupling capacitors.
- Planes and polygons loop inductance

Usually the main contributors to high-loop inductance are the power and ground vias, and therefore, the main loop inductance consideration is usually to provide a low-inductance connection to the closest (and usually lowest value) decoupling capacitors.

16.1.2 Implementation Notes

- Supply as many power and ground vias as possible between the device power pins and the decoupling capacitors. Refer to the board reference design for an implementation example.
- Couple power vias to ground vias. Place them as close to each other as design and manufacturing limitations permit. See Figure 65.
- Place both vias of a bypass capacitors close to the signal reference change location.
- Increase the distance between each of the ground vias from other ground vias and each of the power vias from other power vias. Usually a distance equal to the thickness of the board is adequate.
- Minimize the routing length between the vias and the capacitors' pads as shown in Figure 66.
- Use wide traces to connect the capacitors' pads to the power and ground vias.
- Use an adequate amount of decoupling capacitors according to the specific device design guidelines for each power supply network to lower the total ESL.

Figure 65: Coupling of Vias

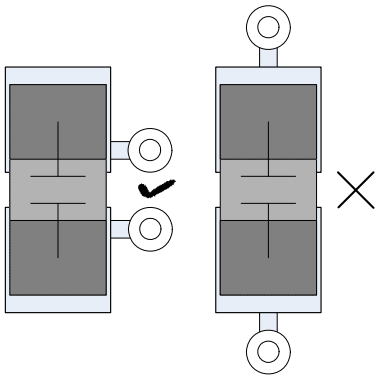

Figure 66: Routing Length Between the Capacitor and the Vias

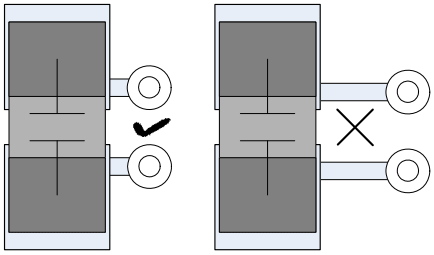

16.2 DC Voltage Drop

The conductivity of copper is a finite number, and therefore it causes DC voltage drop. It is essential to provide high DC current conductivity between the power supply and the device pins. Refer to the Hardware Specifications to determine the average current of each power rail and determine the polygon width accordingly.

16.3 Polygon Shape Considerations

The polygon underneath a BGA can suffer from many holes due to many via anti-pads causing it to become Swiss-cheese like. Other layout restrictions can also cause a polygon to have narrow sections. In these situations, surface current density should be estimated, and it should be verified that no areas create fuses (There are no areas of high current density).

Note Simulation and extraction tools may be needed to perform this task.

16.4 Analog Power Filtering

Analog power supply is usually used to nourish internal PLL circuitry and/or analog circuitry that is sensitive to power noise. The current consumption from the analog power device pins is usually characterized by very low content in the high frequency. Nevertheless, due to the power-noise sensitive internal circuitry, an adequate external power noise filter should be provided.

16.4.1 Implementation Notes

Note

Specific interface circuitry power filtering may be required. Refer to the specific device filtering/power section for further information.

Generic board design notes—preceding the filter

Design the board's power delivery network to supply the nominal voltage with up to ±5%, but not more than 100 mV peak-to-peak noise/ripple as sampled in time domain (Broadband noise that incorporates all frequencies ranging from near DC to 10 GHz).

Filter implementation notes

Design a power filter with the following characteristics:

- DC voltage drop of up to 30 mV—refer to the Hardware Specifications for current consumption values.
- A low pass voltage transfer function with DC transfers as close as possible to 0 dB, and at least -18 dB at frequencies from 0.5 MHz and up. The transfer function is presented in Figure 67.
- Verify that the filtering scheme is resonance free. Examine the impedance of the filtering scheme as seen from the device balls, and verify low impedance over the frequency range.

Figure 67: Required Analog Power Filter Voltage Transfer Function

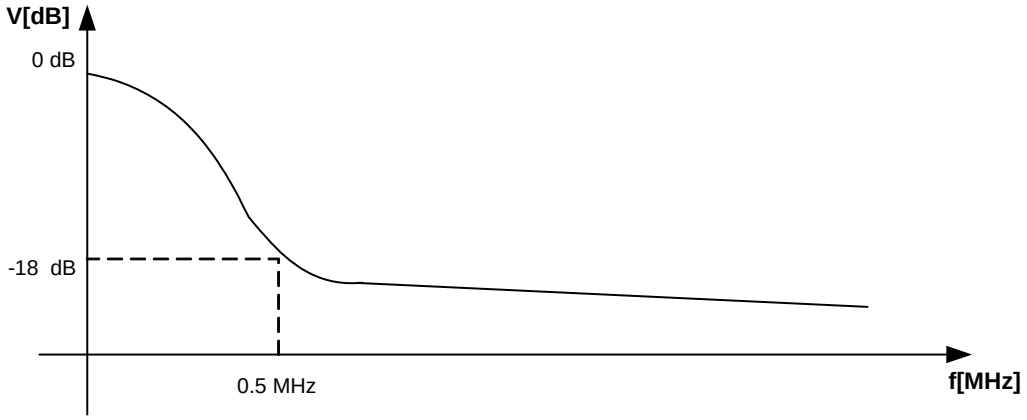

Decoupling implementation notes

Provide decoupling capacitors with:

- A low-inductance routing connection to the device analog power supply pins.
- A usual value of 100 nF, placed as close as possible to the device analog power pins.

The usual number of decoupling capacitors is one for each two device analog power pins. An exception to this rule is when individual pins are geometrically distant from each other. In that case, provide a decoupling capacitor for each power pin.

16.4.2 Analog Power Filter Example

Figure 68 depicts one possible example of an analog power filter that complies with the above required voltage transfer function. Table 37 defines the symbols that appear in Figure 68.

Figure 68: Analog Power Filter Example

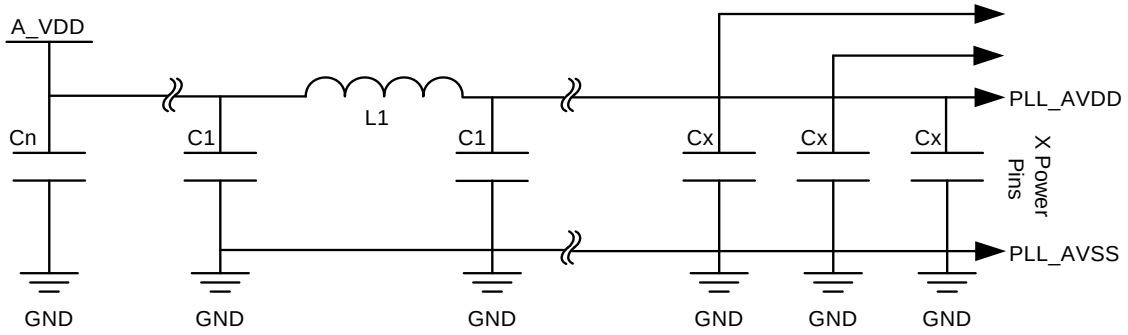

Key
L = Ferrite bead
C = Capacitor

Table 37: Definition of Analog Power Filter Symbols

Symbol	Value	Amount	Notes
A_VDD	Nominal power supply voltage		1
L1	600 ohm @ 100 MHz ferrite bead	1	2, 3
C1	1 µF capacitor	2	4
Cn	22 µF capacitor	1	5
Cx	100 nF capacitor	One for every two power pins	6, 7, 8

NOTES:
1. A_VDD should have up to ±5%, but not more than 100 mV peak-to-peak voltage ripple/noise.
2. If the DC impedance of the selected ferrite bead is too high, use a lower impedance ferrite bead @ 100 MHz and increase the C1 capacitance value accordingly. The lowest recommended impedance is 60 ohms.
3. It is essential to verify the DC resistance of the ferrite versus the estimated current and the desired DC voltage drop.
4. C1 indicates the Pi filter capacitors. Place them close to the ferrite.
5. Cn is a bulk capacitor placed between the power supply output and the filter. The minimum number is one capacitor per 1.5 square inches, and it is related to the polygon geometrical shape.
6. Cx indicates decoupling capacitors. Place them near the device pins.
7. Supply a low loop-inductance connection of these capacitors to the device power pins.
8. The amount is applicable only if physical location permits both pins sharing the capacitor to have low loop-inductance connections to the capacitor.

16.5 Core/CPU Power Decoupling

Core/CPU power pins may be characterized by high current surges with wide frequency bandwidth. It is therefore essential to supply the Core/CPU power pins with a low impedance power supply.

16.5.1 Implementation Notes

- Supply a low loop inductance connection according to Section 16.1.1, Loop Inductance, on page 114. Make certain to consider the power and ground via placement. Provide high coupling between GND and power vias and minimize the coupling of GND/power vias to vias of the same type.
- Follow the specific device decoupling recommendations (in the specific device guidelines) in terms of:
  - Number of decoupling capacitors.
  - Location of the decoupling capacitors.
  - Connectivity of the decoupling capacitors to the device pins.
  - The range of values of the decoupling capacitors.
- In cases where the design is very different from the recommended configuration, it is most important to perform frequency domain extraction / modeling of the power supply network and to verify that the power impedance versus the frequency does not shows high impedance resonance at the desired bandwidth.

16.6 I/O Power Bypassing

I/O power pins may be characterized by high current surges during the transition time. These high currents are usually intended to supply return current continuity. The amount of charge needed for the I/O buses is usually low. The main concerns are:

- Allowing return path currents to have a low-inductance closing path.
- Supplying the chip with a stable power to allow minimal in-chip power noise to the I/O and the in-chip terminations.
- Accounting for the physical location of the return path current when placing I/O power bypassing capacitors.

16.6.1 Implementation Notes

- Supply a low-loop inductance connection according to Section 16.1.1, Loop Inductance, on page 114. When the return current path is interrupted due to reference to a non-continuous polygon shape, or due to a signal changing the routing layer, it is important to place a bypass capacitor with a low-inductance connection as close as possible to the return current path disturbance (see Figure 69).
- During the capacitor-placement phase, take the arrangement of the device pinout into consideration. Note the location of the power and ground pins, both at the device periphery and in the middle of the device, and place the bypass capacitors to allow for a low-inductance connection to these device pins (see Figure 70).
- Take into account the return path reference supplied to the signals on the board. Supply a low-inductance capacitor connection between the return path reference power and the ground pin of the package on the device periphery (see Figure 70).

Figure 69: Return Path Discontinuity—Bypass Capacitor

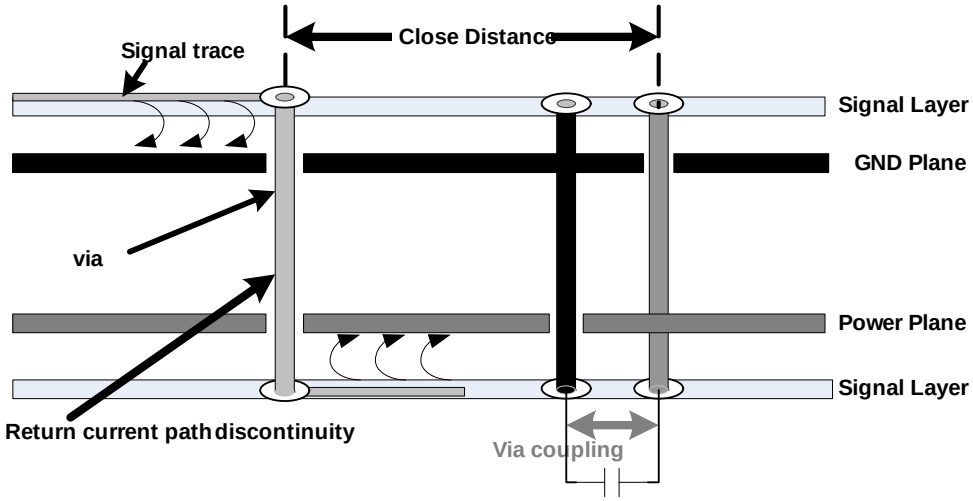

Figure 70: Placement of a Bypass Capacitor in Relation to Device Ground Pins

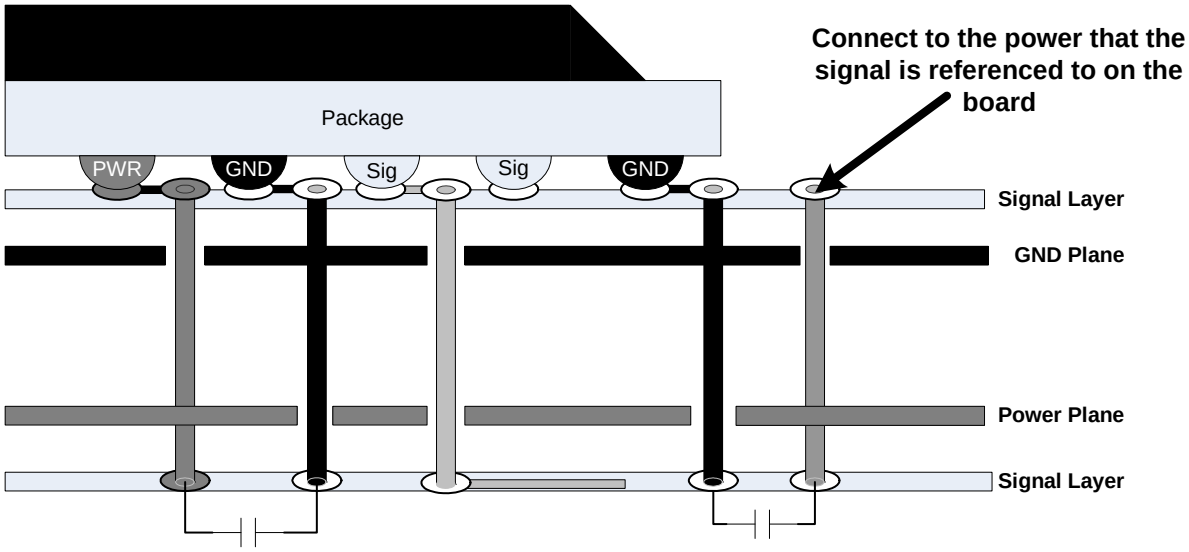

16.7 Bulk Capacitors

It is essential to provide adequate capacitance, according to geometrical considerations for each of the power polygons.

Bulk capacitance should provide needed charge as well as suppress polygon resonance.

- As a general rule place one bulk capacitor every 2 square inches (the Coffee Cup rule).

Make sure that the bulk capacitance does not introduce an inrush current issue.

16.8 Power Signals—Vref

The Vref is used as a reference voltage to the input cell. It defines the threshold point at which the signal changes its logic state from 1 to 0 and vice versa. R1 and R2 do not have specific, mandatory values. Choose the resistors value to be within the range of a few hundred ohms to avoid high current flowing from VDDIO to GND through R1 and R2.

Note

Marvell® recommends supplying a dedicated voltage divider for the Marvell device Vref pin.

The Vref typical value is VDDIO/2.

16.8.1 Routing and Placement Guidelines

Placement of the Vref voltage divider and decoupling capacitors is ideally recommended to be as close as possible to the device pins. However, in situations where there is a trade-off between the Vref divider and the placement of other decoupling/bypass capacitors, it is possible to place R1, R2, and C1 outside the device area. Place C2 close to the device pin, and connect C2 and the device pin with a short trace. Make certain to avoid crosstalk aggressors on the trace.

Figure 71 depicts Vref generation and filtering. Both R1 and R2 should have the same resistance value, between 100 ohm and 600 ohm (recommended value 500 ohm). R1 and R2 should have 1% accuracy.

Figure 71: Vref Generation and Filtering

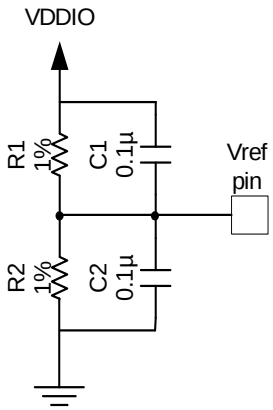

17 Power Management Unit (PMU) Board Design Guidelines

The 88AP510 integrates an advance power management capability that allows low Thermal Design Power (TDP) and long battery life.

Using the power management capabilities allows the user to extend the battery life of the system while sufficient power remains available to drive power-hungry applications, such as video and web browsing.

17.1 Power Features

This device incorporates the following power features:

Dynamic Voltage Scaling (DVS)	The device integrates an SDI interface that can control the voltage level of Marvell® regulators such as the 88PH845. At any time, the PMU can control the voltage level of the CPU. Reduction of operating voltage to the minimum required reduces the power consumption significantly.
Dynamic Frequency Scaling (DFS)	The CPU dynamic power depends on the CPU frequency. The device can control the CPU frequency on the fly. Combining the DFS and the DVS (DVFS) results in a significant reduction in CPU power consumption.
Idle Mode	
Wait for Interrupt (WFI)	In most use case, the CPU is idle most of the time. In addition to the DVFS capability, the CPU clocks can be gated by a hardware mechanism, and then consume only static power.
CPU power down (Deep Idle and eBook modes)	The device can eliminate even the CPU static power consumption by shutting OFF its power and then waking up following an interrupt.
CPU and Core power down (Standby mode)	The SDRAM is in Self-Refresh mode. The VDDO_M is powered down. A board that supports Standby mode should separate the power planes of the Marvell and the SDRAM devices.
Power gating of idle units	The power of the Video Decoding Unit (VMeta™) and Graphics Processing Unit (GPU) can be turn OFF by an internal power switch, controlled by the PMU.
Clock gating of idle units	The peripheral unit clocks can be gated to save dynamic power.

17.2 Power Modes

The 88AP510 power modes are listed below, in order of their recovery time back to Run mode, with Idle mode being the fastest and Hibernate mode being the slowest.

- Run mode
- Idle mode
- eBook mode

- Deep Idle mode
- Standby mode
- Hibernate mode

Note For additional information about the power management modes, refer to the device Hardware Specifications and the device Functional Specifications.

Table 38 lists the properties of each of the 88AP510 power modes.

Table 38: 88AP510 Power Modes

Power domain	Run	Idle	eBook	Deep Idle	Standby (Suspend to RAM)	Hibernate (Suspend to Disk)
CPU	ON DVFS	ON - WFI	OFF	OFF	OFF	OFF
CPU PLL	ON	ON	ON	ON	OFF	OFF
VMeta™, GPU (Part of the SoC core domain)	Optionally OFF	Optionally OFF	Optionally OFF	Optionally OFF	OFF	OFF
LCD	ON	ON	LCD Refresh	OFF	OFF	OFF
DDR	ON	ON	ON	Self-refresh	Self-refresh	OFF
SoC CORE	ON	ON	ON	ON	OFF	OFF
PMU	ON	ON	ON	ON	ON	OFF
RTC	ON	ON	ON	ON	ON	ON (Battery)
Enter By	N/A	OS Timer	OS Timer	OS Timer	User	User
Exit By	N/A	SoC Interrupt	SoC Interrupt	SoC Interrupt	External Wake-ups/ RTC/ Battery alarm	Power-up
Boot From	N/A	N/A	DDR	DDR	DDR— System Boot	Flash/Disk

The above states are implemented using the Power Management Unit (PMU) signals that are multiplexed over MPP[15:0] as shown in Table 39.

Table 39: Power Management Signals

I/O	Type	Multiplexed on MPP	Description
PMU_CPU_PWRDWN	Output	MPP[15:0]	CPU_PWRDWN[1]: CPU external voltage supplier powering down
PMU_CORE_PWR_GOOD	Input	MPP[7:0]	CORE_PWR_GOOD[1]:VDD_CORE external voltage supplier power good indication.
PMU_SDI	Output	MPP[15:0]	SDI[1]: Serial interface for programming the Marvell® PMIC. **NOTE:** The relevant power rail for this signal depends on which MPP pin is used. It can be multiplexed on MPP[15:8] or MPP[7:0].
MRn	Input	N/A	Manual Reset Input.Used to reset all units (including PMU and POR) to their initial state.SYSRST_OUTn is asserted low as long as the MRn signal is asserted low, and for additional 20 ms after MRn de-assertion. **NOTE:** Must be pulled up through a resistor.
PMU_STBY_PWRDWN	Output	MPP[7:0]	STBY_PWRDWN[1]: Standby Power Down for disabling external regulators.
PMU_EXT_WU[2:0]	Input	MPP[7:0]	EXT_WU[1]: External low-power modes wake-up pins.
PMU_BLINK	Output	MPP[7:0]	BLINK[1]: Blink option with configurable duty cycle for LED control.
PMU_GPO	Output	MPP[7:0]	Drive_0 or Drive_1[1]: General Purpose Output
PMU_BAT_FAULT	Input	MPP[7:0]	BAT_FAULT[1]: Battery Fault external indication.
DDR_TERM	Output	MPP[71:0]	Control DDR external termination power
CKE_MASK	Output	MPP[7:0]	CKE masking control
M_RESET_MASK	Output	MPP[7:0]	Memory reset for DDR3

NOTES:
1. Name of signal control option in PM Signal Selectors Control 0 and PM Signal Selectors Control 1 registers.
2. MPP[7:0] are functional during standby.
3. MPP[15:8] are used for the Core/SoC, as described in device Hardware Specifications.
4. DDR_TERM is controlled by CPU as a GPO. It is not controlled by the PMU.

17.3 System Power Topology Implementation

The power topology of a system that implements the 88AP510 can be divided into four types of power domains (see Figure 72). Each power domain is active in different power modes:

Hibernate power: Power domains that are always ON will be used for logic, such as if the Embedded Controller (EC) is used, or other logic that should power ON the system from Hibernate mode.
During schematic drawing it is recommended to mark those power domains with the suffix *_hiber* for better understanding the system power topology.

Standby power: Power domains that will be power ON during Standby mode, such as: the PMU digital and IO powers, DRAM device power, and SDRAM VRef power supply. During the schematic drawing, it is recommended to mark those power domains with the suffix *_stby* for better understanding the system power topology.

Run power: Power domains that will be powered ON during the below modes:

- Run mode
- Idle mode
- Deep Idle mode
- eBook mode

Figure 72: Power Modes

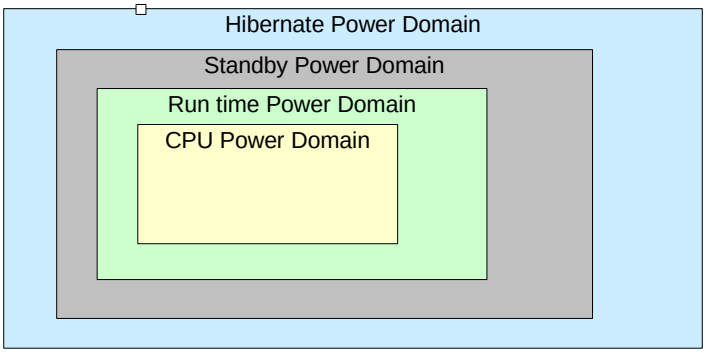

The peripheral and CPU power can be controlled separately (see Figure 73):

Peripheral power domains: It is possible to turn OFF the power supply of unused devices in the system. Controlling the power supply of those devices is accomplished using the GPIO ports of the 88AP510 or the EC.

For better understanding of the system power topology during the schematic drawing, it is recommended to mark the run time power domains by adding the postfix “_run” on the power domain name.

CPU power: The CPU has a separate power domain. This allows shutting down the CPU power during Deep Idle and eBook modes and changing its level during Run mode (DVS). The delay of getting in and out of those power saving modes is critical and should be as short as possible. The 88AP510 cannot exit Deep Idle/eBook until the CPU power is up and stable. To minimize this delay, it is recommended to use a MOS FET as a switch on the CPU power supply. The power-up delay of this topology will make the CPU power stability delay to be by far faster than using the power regulator enable and feeding the CPU power directly from the power regulator.

Figure 73: Good and Bad Power Examples

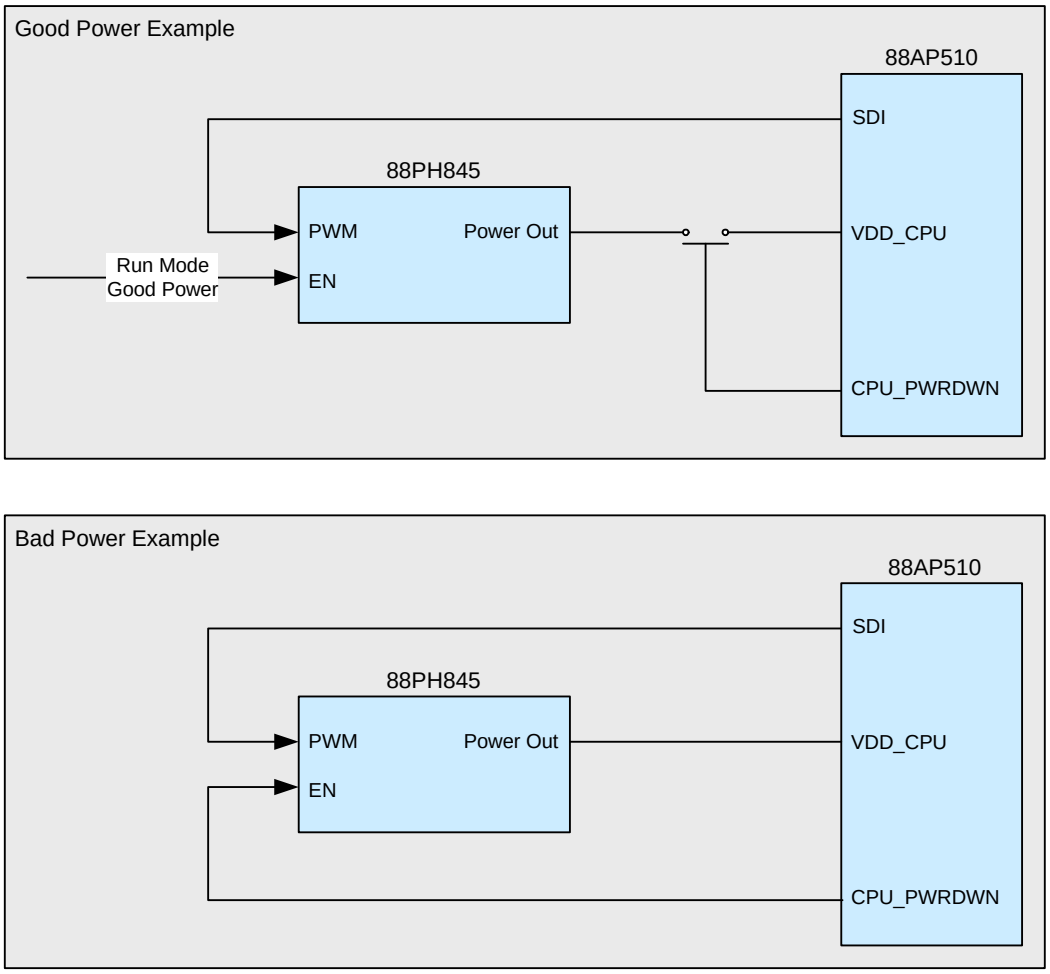

Note

- To avoid a possible IR drop in the CPU and DDR power, Marvell® recommends using FET transistors with resistance between the source and the drain of the transistor (Rds) lower than 20 mΩ.
- When using the 88AP510 with a CPU frequency of 1 GHz, the CPU power supply recommended operating condition is stricter. To meet those specifications, Marvell® recommends using the Marvell® 88PH845-T184 regulator.

17.4 Run Time Power Management

In a typical smart book system, the power consumption will be spread unevenly between the different power consumers. Each device will consume its share of the power and the total power consumption will be the sum of all devices plus the power regulators efficiency.

Power management can be split into a few levels. The highest layer is managed by the application layer during run time. Any peripheral device such as the LCD backlight, audio decoder, or camera may go to power down or power saving mode when not used. The methods of powering down those devices may differ from one device to another. However, the common power down method will be by sending special power down commands, asserting an external signal, or by shutting down its power.

The 88AP510 incorporates 70 MPP pins that may function as GPIO pins or as dedicated pins. The GPIO pins can be used to control each of those peripheral devices.

In most of the GPIO pins are used as dedicated pins, the designer might find it hard to locate a spare GPIO and may consider adding an embedded controller to the system.

- The advantages of using an embedded controller:
  - Frees a GPIO pin for power management and other uses such as LED control
  - OFF loads tasks such as battery charging
  - Programmable devices may add flexibility to the board
- The disadvantages of using an embedded controller:
  - Adding hardware (Bill of Materials—BOM) to the board
  - Additional firmware that requires maintenance

Note The general power saving concept should be: If a device is not used, turn it OFF.

17.4.1 CPU DVFS (Dynamic Voltage and Frequency Scaling)

The CPU is the heart of the system. The dynamic voltage frequency scaling (DVFS) allows on-the-fly frequency adjustment according to the performance requirements of the system. By reducing the CPU frequency, the user will save a significant portion of the dynamic power. The DFS does not require any board level consideration.

Once the CPU frequency was reduced, the more significant power saving will come by reducing the CPU power (VDD_CPU).

Note

The dynamic power is a combination of few factors.

As a general rule use the formula $P = f \times C \times V^2$

- P = Total dynamic power
- f = Frequency
- C = Gate and metal capacity
- V = Operating voltage

From the equation above, it can be seen that when reducing the frequency from 800 MHz to 400 MHz and the reducing the voltage from 1.1V to 1.0V, the user will save ~60% of the CPU power.

To implement the DVS, The CPU power regulator should support dynamic voltage scaling through the SDI protocol.

It is recommended to use the Marvell® 88PH845 as the CPU power supply.

The SDI can be multiplexed on any of the pins MPP[15:0]. Connect the selected MPP pin to the PWM pin of the 88PH845 as shown in Figure 74. The PMU controls the SDI pin and changes the CPU power according to the current frequency.

Figure 74: CPU Dynamic Voltage and Frequency Scaling Connection

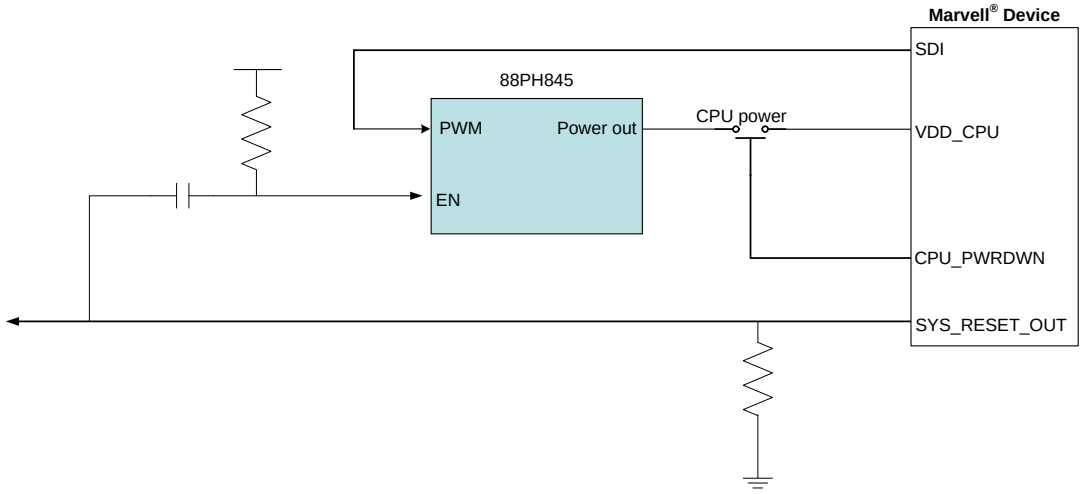

If the system is reset, while the CPU is running at a slow speed and using low voltage, it is important to raise the CPU power before trying to boot.

It is recommended to connect the EN pin of the power regulator to the system reset, through a capacitor, then the regulator will sense the reset and return to the default power level (see Figure 74).

When using DVFS, take into account the system considerations

- The CPU power must be increased first, before scaling up the CPU frequency.
- If the power is not stable, the CPU cannot operate at a faster frequency.
- When running process that requires heavy CPU resources on bursts (such as playing HD movie), the performance might be affected by a delay.

17.4.2 DDR Size and Type

The DDR is a major power consumer in a typical system. On the other hand, DDR provides an efficient way to boost performance. There are a few trends that have become quite clear in recent years:

- The performance requirements have grown significantly
- The DDR prices per megabyte have been greatly reduced
- The DDR frequency and capacity keeps increasing

However, it is clear that the faster the DDR frequency becomes and the larger the capacity of the DDR, the greater the DDR power consumption will be. There are a few recommended topologies, nevertheless, the major difference in the power consumption is determined by a single factor: The DRAM operating power.

- DDR2 is available at two major voltage levels: 1.8V and 1.5V
- DDR3 is available at 1.5V or 1.35V.

The power saving is exponential. For example:

- If two systems use the same DDR frequency and DDR size, it is clear the users' experience will be the same.
- If one system uses DDR2 at 1.8V and the other uses DDR3 at 1.35V. The DDR power consumption of the first system will be more than 75% higher than the second system.

17.4.3 DDR Termination Control

When working with DDR in frequencies higher than 300 MHz, termination on the control and address signals is mandatory. The termination reduces the reflections and improves the signal integrity. However, it does not come for free. As long as the termination is active, power is constantly be consumed.

For example, a standard DDR2 (1.8V), eight devices, requires two sets of terminations for a total of 62 Thevenin VTT termination points. A quick calculation shows that the termination alone constantly consumes ~1W.

The ways to implement DDR termination are by using a:

- **Thevenin termination:** (see Figure 75)

Figure 75: Implement DDR Termination Using Thevenin Termination

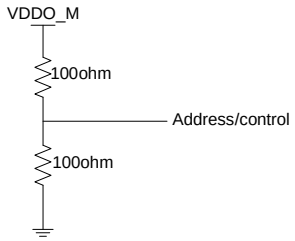

- **VTT plane:** (see Figure 76)

Figure 76: Implement DDR Termination Using a VTT Plane

When the termination is not needed it is recommended to power down the termination power supply. By doing so, the user saves up to 1W. This power saving enables longer battery life and lower heat dissipation.

The cases where termination is not required are:

- DDR frequency is lower than 300 MHz— both if this is because the configuration use DDR DFS or the DDR frequency was initially set to a lower value.
- When the 88AP510 is in Deep Idle or Standby mode, which sets the DDR to Self Refresh mode.

Powering OFF the termination can be controlled by any of the 72 GPIOs in the 88AP510. The CPU turns the DDR_TERM signal OFF and ON when necessary. The DDR_TERM signal controls the DDR termination power supply.

Figure 77 provides an example of DDR_TERM signal connectivity.

17.4.4 VDDO_M Power Control

The VDD_CORE must be powered up within 20 ms after VDDO_M is powered up, or any time earlier.

During standby, the VDD_CORE is powered OFF. This means that VDDO_M and the controller VREF must be turned OFF. The DRAM devices must be kept alive in Self-Refresh mode.

The usual method is to split the SDRAM power between the 88AP510 and the SDRAM devices

using a FET switch, while controlling the switch using the STBY_PWRDWN signal (see Figure 77).

Figure 77: DDR Termination Control Connection

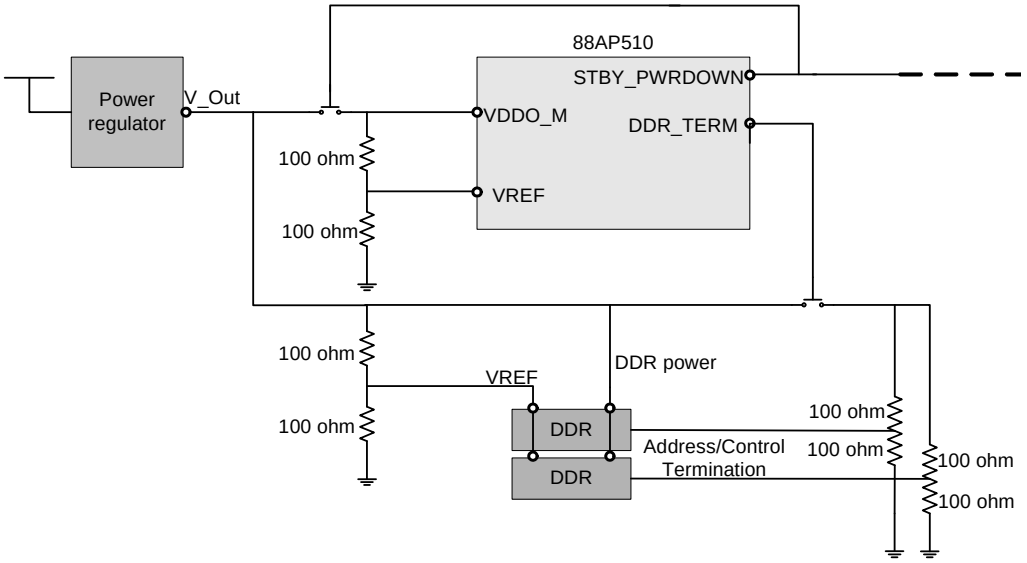

17.4.5 Special Consideration for DDR Reset and CKE During Standby

Unlike the other control and address signals, the M_CKE[1:0] and M_RESETn signals must be pulled low/high during standby. This maintains the DDR in Self-Refresh mode.

Since the 88AP510 memory controller is off during standby, the DDR clock enable and reset signals are driven from the MPP PMU signals. These signals are configured on any of the MPP[7:0] signals:

MPPx Used as the CKE masking control, connect it according to the scheme in Figure 78.

The BootROM drives the signal low after entering the Self-Refresh mode, but before entering the Standby mode. Change the signal to an input after a wake event and before exiting the Self-Refresh mode.

MPPy Used as M_RESETn. Connect this signal directly to the DDR reset signal, and add a pull-up resistor to the DDR standby voltage (see Figure 78).

At power-up, The BootROM drives the signal low for 300 us and then change it to an input. For DDR3 DRAM, the pull-up keeps the signal high during standby.

Figure 78:CKE and Reset Implementation for Standby Support

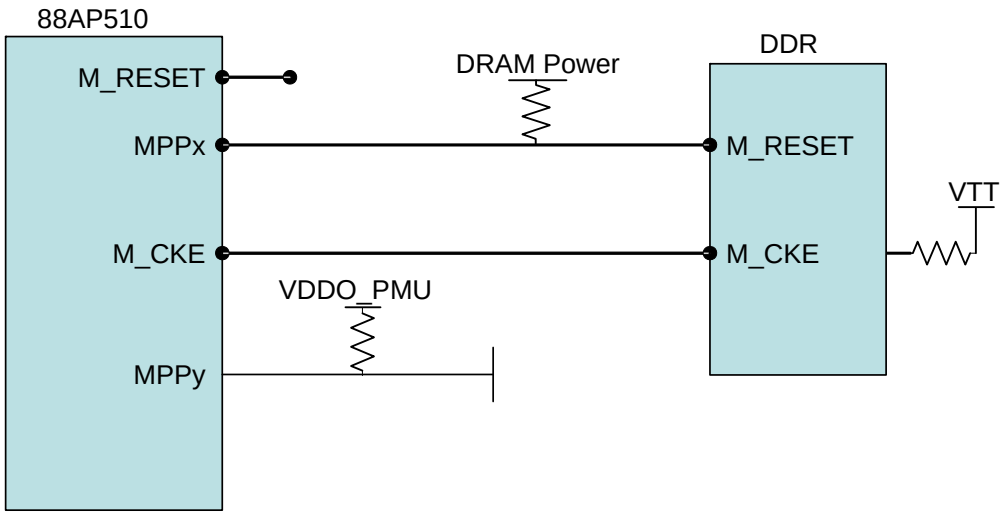

17.4.6 IO State During Standby Mode

When the Core power is turned OFF and the IO power of any interface is ON, the pin will turn to High-Z state. In most cases, this should not effect the peripheral devices. However, in rare situations, if the related device is powered ON, floating signals might cause undesired results.

For example, if the NAND flash is still powered ON during standby, and the control signals are floating, it might lead to data corruption. It is the responsibility of the design to pull the relevant signals to the desired value when they are floating.

The above case will probably be relevant in the power up stage, when the IO power is up and the Core power is still down. It is expected that most of the system peripheral devices will be powered OFF during standby. However, it is also expected that some logic will be powered ON during standby while the 88AP510 IO power and Core power are powered OFF.

In a typical system during standby, it is expected that the EC will be kept powered ON, while the rest of the system is powered OFF.

When the Core power is OFF and the interface IO power is OFF as well, the ESD diode might break and short the connected signal to GND. This might lead to undesired results. For example, if the design connected the 88AP510 the EC and the battery charger over the I^2C interface. During standby, the 88AP510 Core and IO power will be turned OFF, the ESD diode that protects the I^2C interface will break down, and the bus will not be function. The EC will not be able to communicate with the battery charger, without the 88AP510 running.

There are a few ways to avoid such a situation. One option is to design the system so that the 88AP510 does not share a bus that should remain functional during standby. Another option is, that during standby, isolate the 88AP510 from other devices that should remain functional (see Figure 79).

Figure 79: IO State During Standby Mode Design Options

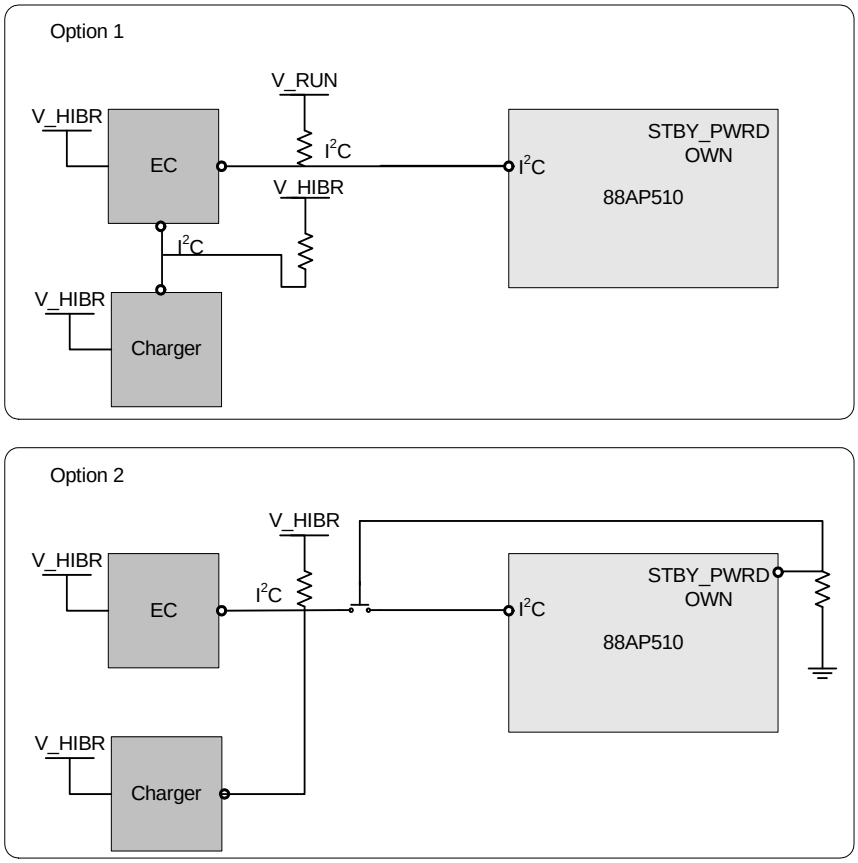

17.5 Power Saving Methods

This section describes system power saving methods for the 88AP510 device. It also describes software procedures for setting the relevant interfaces and system components to a power saving state. For each interface, the following information includes:

- The required device resources.
- The default power-on values.
- The procedure for
  - Entering the Power Saving State.
  - Exiting the Power Saving State.
  - Where applicable, the complete shutdown of an unused interface.

17.5.1 PCI Express Power Saving

The PCI Express Specification defines a software driven power management mechanism, and an autonomous hardware based Active State Power Management (ASPM) mechanism. Refer to the *PCI Express Base Specification.1.1, Copyright © 2002-2005 PCI-SIG*.

The two mechanisms can be activated together, or separately.

Software driven power management defines the D0, D1, D2, and D3 Device states. The power savings increase as the Device state transitions from D0 to D3.

The PCI Express-PM defines the following link power management states:

- L0 (normal operation state)
- L0s (standby state)
  NOTE: L0 support is not applicable to PCI-PM compatible power management.
- L1, L2/L3 Ready, L2, L3, and LDn link states

PCI Express-PM also defines the ASPM, which requires mandatory L0s support and optional L1 support. Refer to the *PCI Express Base Specification.1.1, Copyright © 2002-2005 PCI-SIG, Sec.* 5.2.

Table 40 lists the available PCI Express power saving resources.

Table 40: PCI Express Power Saving Resources

Register	Offset	Field	Bits	Power-on Value
IO Power-Down Control	0xD0058	<PEX0PwrDwn> <PEX1PwrDwn>	[16] [17]	SAR
CPU Control and Status	0x20104	<PCIDs0> <PCIDs1>	[0] [3]	0x0 0x0
Power Management Control and Status	0x40044 0x80044	<PMEStat> <No_Soft_Reset> <PMState>[1]	[15] [3] [1:0]	0x0 0x0 0x0
PCI Express Link Control Status	0x40070 0x80070	<AspmCnt>[2] <EnClkPwrMng>	[1:0] [8]	0x0 0x0
PCI Express Status	0x41A04 0x81A04	<DLDown>	[0]	0x0
PCI Express Root Complex PME	0x41A14 0x81A14	<PME Pending> <PME Status> <PME Requester ID>	[17] [16] [15:0]	0x0 0x0 0x0
PCI Express Power Management	0x41A18	<Ref_clk_off_en> <Send Turn Off Msg> <Send Turn Off Ack Msg> <L1_aspm_ack (RC)> <L1_aspm_en (EP)>	[16] [5] [4] [1] [0]	0x0 0x0 0x0 0x1 0x0
PCI Express Interrupt Cause	0x41900	RcvTurnOff RcvPmPme DstateChange	[29] [28] [11]	0x0 0x0 0x0

1. The <PMState> field encodes the following values: 0 = D0, 1 = D1, 2 = D2, 3 = D3.
2. The <AspmCnt> field encodes the following values:
   0 = Disabled: Disabled.
   1 = L0Support: L0s entry supported.
   2 = Reserved
   3 = L0L1Support: L0s and L1 entry supported.

PCI Express link states are not directly visible to the legacy bus driver software. The link states are derived from the power management state of the components residing on those links. The software may only enable or disable the ASPM mechanism in the setup phase. After this phase, if enabled,

the ASPM acts autonomously, and there is a full handshake between the peers of the PCI Express link.

The PCI Express specification requires support for all PCI-PM device power management states. All functions must support the D0 and D3 states (including $D3_{hot}$ and $D3_{cold}$). The D1 and D2 states are optional.

The D0 state is divided into two distinct sub-states:

- $D0_{uninitialized}$
- $D0_{active}$

When initially applying power to a PCI Express component, it defaults to the $D0_{uninitialized}$ state. The PCI Express Hierarchy Enumeration process enumerates and configures components that are set to this state. Following the completion of the enumeration and configuration process, the function enters the $D0_{active}$ state; the fully operational state for a PCI Express function.

On entering the power saving state, the Root Complex side of the link is always the initiator. The Endpoint side is always the executor (see Figure 80).

Figure 80: PCI Express Root Complex to Endpoint Relationship

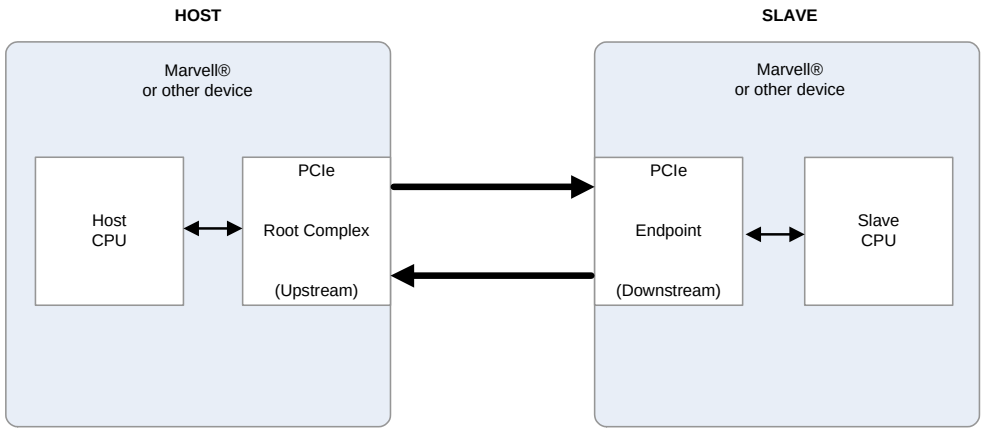

The 88AP510 device contains a CPU. The CPU is a Host CPU when it controls the Root Complex, and a Slave CPU when it controls the Endpoint.

When examining the PCI Express interface inside the 88AP510 device configured as a Root Complex, it is not important if the Endpoint on the other side is a Marvell device or another vendor's device. It is also not important if there is a CPU on the other side or not.

Conversely, examining the PCI Express interface inside the 88AP510 device configured as an Endpoint, it is not important if the Root Complex on the other side is a Marvell device or another vendor's device. It is also not important if there is a CPU on the other side or not.

The PCI Express interface provides the following power management support mechanisms:

- Software power management states D0, D1, D2, $D3_{hot}$, and $D3_{cold}$.
- Active State Power Management L0s and L1.

The 88AP510 devices also supports PMEn message generation. When a wake-up event occurs, the local CPU triggers a PMEn message by writing the value 0x1 to the PCI Power Management Control and Status register <PMEStat> field (bit [15]). In response, the system host sets the 88AP510 device back to the D0 state, and clears the PMEn interrupt by writing the value 0x1 to the <PMEStat> field.

17.5.1.1 D-State Software Power Management Options

The 88AP510 supports the PCI Express software power management options D1, D2, and D3, as described in the PCI Express Specifications.

As an Endpoint

The device implements the Power Management Configuration registers, including the PME message.

Also, the device supports the turnoff process. Upon receiving a turnoff message from the Root Complex, a maskable interrupt is set. The device CPU acknowledges the turnoff request by setting <Send Turn Off Ack Msg> field (bit [4]) in the PCI Express Power Management register.
NOTE: If the slave CPU does not set this bit, the Host continues after 10 ms, as if the slave CPU has approved it.

As a Root Complex

the PCI Express Interface sets the device to a turnoff state by sending a turnoff message. To send the turnoff message, set the <Send Turn Off Msg> (bit [5]) in the PCI Express Power Management Register. When the turnoff process completes, a maskable interrupt is set in the main cause register.

the PCI Express Interface responds to a Power Management Event (PME) generated by the device. When a PME message is received, a maskable interrupt is set. The PME data is saved in the PCI Express Root Complex PME register.

17.5.1.2 Active State Power Management

The 88AP510 PCI Express Interface supports all ASPM options defined by the PCI Express specifications as a Root Complex and Endpoint. It allows the hardware to lower the link power state.

L0 Power Saving

The device supports PCI Express L0 Receive and Transmit power saving states, as defined in the PCI Express specification. If there is no traffic on the link, the PCI Express PHY initiates power down on its Transmit side.

Similarly, if an external PCI Express device initiates an L0 Power Down, the PCI Express on the device also initiates an L0 Power Down on the Receive side.

L1 ASPM Endpoint Mode

If this feature is enabled, the device starts the L1 ASPM event when the conditions for this event are met, as defined in the PCI Express Specification, and when the <L1_aspm_en> field (bit [0]) in the PCI Express Power Management Register is set.

L1 ASPM Root Complex Mode

As a root complex, the device acknowledges an L1 ASPM event when the <L1_aspm_ack> field (bit [1]) in the PCI Express Power Management Register is set. If this field is not set, an L1 ASPM non-acknowledge (NAK) response is sent upon the L1 ASPM request.

17.5.1.3 Activating PCI Power Management

The 88AP510 supports PCI Power Management under the following use cases:

- There is no device connected to the PCI Express interface.
- The PCI Express is used as Root Complex (controlling external PCI Express Endpoints).
- The PCI Express is used as an Endpoint.

Note

- The 88AP510 Includes two PCIe units (PCIe 0 and PCIe 1).
- PCIe 0 can be configured only as Root Complex port.
- PCIe 1 Can be configured as a Root Complex or Endpoint port.

If it is required to activate the ASPM, these preliminary steps must be performed:

1. Set the <L1_aspm_ack> field of the PCI Express Power Management register to 0x1.
   The Root Complex allows acknowledging an L1 ASPM request from an Endpoint.
2. Set the <AspmCnt> field of the PCI Express Link Control Status register to 0x3.
   This setting means that the L0s and L1 entries are enabled, Tx-L0s and L1-ASPM entries are supported.

17.5.1.3.1 No Device Connected

Entering the Power Saving State

When there is no device connected to the 88AP510 PCI Express interface, the hardware polls the <DLDown> field (bit[0]) of the PCI Express Status register link for a fixed time interval (usually 1 second), and shuts down the link.

It is also possible to shut down the link and reduce the 88AP510 power consumption by disabling the PCI Express clock. Set the <PEX0PwrDwn> or <PEX1PwrDwn> fields to 0x1 in the IO Power-Down Control register.

Exiting the Power Saving State

Exiting the power saving state actually means connecting a device while the host is powered on, i.e. Hot Plug. There is no support for Hot Plug in the PCI Express Specification.

The solution is completely dependent on user implementation.

17.5.1.3.2 The Device as a PCI Express Root Complex

Entering the Power Saving State

The initiative to enter the power saving state must come from the Host software.

To enter the power saving state, the Host software performs the following steps (as recommended by the PCI Express specification):

1. Direct all functions of a downstream component to the $D3_{hot}$ state. Implementation of this step is customer-specific.
   NOTE: The PCI Express Specification also presents an alternate sequence to remove power without first placing all devices into the $D3_{hot}$ state. This sequence begins with step (2).
2. The Host Software sets the <Send Turn Off Msg> field (bit [5]) of the PCI Express Power Management register to 0x1. This setting causes the Root Complex to broadcast the PME_Turn_Off message in preparation for removing the main power source.
   NOTE: The Downstream component is expected to respond with PME_TO_Ack.
3. Upon receiving the PME_TO_Ack, a PME event is signaled to the Host Software.
   NOTE: The Host Hardware sets the <RcvTurnOff> field (bit [29]) of the PCI Express Interrupt Cause register, signaling that the downstream component is ready for TurnOff.

4. The Host Software can turn down the Endpoint Power on the slave board, if the board is equipped to support that action.

Note In a D3 power state, the external Endpoint is only allowed to respond to configuration cycles. The software must not access the external device with memory or I/O transactions. This kind of access results in a master abort condition.

Exiting the Power Saving State

Exiting the power saving state to a fully operative L0 state is triggered by traffic appearing on either side of the link.

Initiated by the External PCI Express

If the external PCI Express Endpoint requires a wake-up (such as a NIC device that received a wake-up packet) the following events occur:

1. The Endpoint sends a PMEn message. As a result, the PCI Express Interrupt Cause register <RcvPmPme> field (bit [28]) is set, and a CPU interrupt is asserted, if not masked.
2. The operating system responds with the following actions:
   - Sets the Endpoint back to D0 for normal operation.
     Switching from $D3_{hot}$ state to D0 state is accomplished by writing to the <PMState> field (bits [1:0]) in the PCI Express Power Management Control and Status register.
   - Clears the device's PMEn interrupt by writing the value 0x1 to the device PCI Express Power Management Control and Status register <PMEStat> field (bit [15]).
3. A signal is sent to the Endpoint software driver to resume normal operation

Note In the D3 state, the Endpoint software looses its context. This implies that switching back from the D3 state is always performed in the $D0_{uninitialized}$ state.

Initiated by the Internal PCI Express

Any access to the PCI Express Endpoint, such as the CPU writing to the Endpoint's <PMState> field (bits [1:0]), results in a link resuming its full power state.

Note The D-state must also be switched back to the D0 state, otherwise the link remains in a low power state.

17.5.1.3.3 The Device as a PCI Express Endpoint

As an Endpoint, the external system host directs the 88AP510 to enter the power saving state.

Entering the Power Saving State

To set the 88AP510 Endpoint device into one of the power saving states, the external System Host must follow these steps.

1. Send a request for the 88AP510 software driver to complete all of its tasks.
2. Set the PCI Express Power Management Control and Status register <PMState> field (bits [1:0]) to the required power state on the Endpoint side.

As a result, the 88AP510 PCI Express interface initiates a link power down process, which eventually results in a link power down.

Due to the <PMState> field change, the 88AP510 sets the PCI Express Interrupt Cause register <DStateChange> field (bit [11]), and asserts a CPU interrupt, if not masked. The local CPU, attached to the 88AP510, can use this interrupt to set the board components into a power saving state.

Note

While in a D3 power state, the 88AP510 resets the PCI Express configuration registers to their default values, as required by the PCI Express specification.

Power Down

To power down the 88AP510 Endpoint device, the external System Host must follow these steps:

1. The operating system requests that the Endpoint software driver complete all of its tasks, and sets the device to a Power Down state.
2. The operating system sets the Endpoint PCI Power Management Control and Status register <PMState> bits [1:0] to 11, the D3 power saving state.
3. Read the D-state register from the Endpoint device, and verify that the D3 is written in the register.

At this point, the link is down in a D3Hot state.

Additional Power Saving

The Host system can achieve a deeper power saving state by performing the following actions:

1. Disable the PCI Express Ports as follows:
   - Send Turn Off Msg.
   - Clear the <PCIDs0> field (bit [0]) to 0 in the CPU Control and Status register.
   - Clear the <PCIDs1> field (bit [3]) to 0 in the CPU Control and Status register
2. Wait 20 ms.
3. Turn off the power to the Endpoint device.

This procedure depends on the system board implementation. For example, a board designer can use an MPP output to disable/enable the dedicated Endpoint power supplies, as displayed in Figure 81, “Endpoint Power Supply Control.

The Status is lost after the above actions are performed.

Re-enumeration is required at wake-up.

Figure 81: Endpoint Power Supply Control

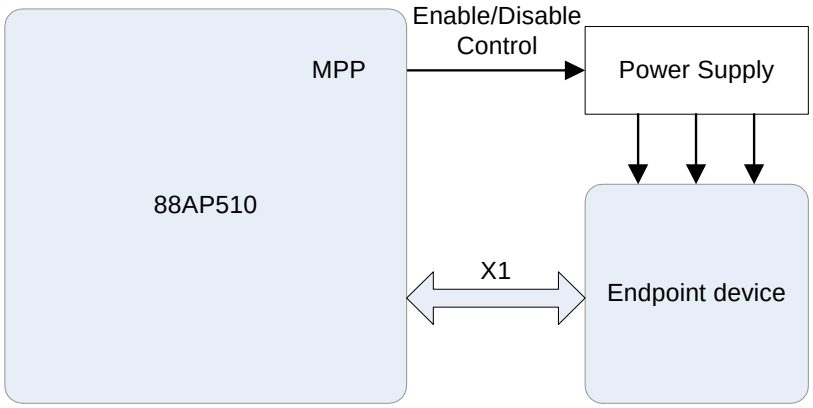

Clock Request Procedure

The 88AP510 PCI Express Interface provides a mechanism that shuts down the reference clock to the downstream device via CLKREQ# signaling. Shutting down the reference clock may increase power savings.

This mechanism is called CLKREQ# Support, as defined in the PCI Express specification.

To enable CLKREQ# Support, the following handshake must be completed:

1. The Host is responsible for setting the <EnClkPwrMng> field (bit [8]) in the PCI Express Link Control Status register of the Endpoint (Host accesses Endpoint's register over the link) to 0x1.
2. The Host sets the MPP control registers to enable the CLKREQ# pin.
3. By default, the CLKREQ# pin is asserted (driven low).
4. To turn off the reference clock, the Slave CPU is responsible for setting the <Ref_clk_off_en> (bit [16]) in the PCI Express Power Management Register to 0x1.

Note

If the ASPM mechanism is enabled, with no more pending transactions in the unit, and the link is in the L1 state, the CLKREQ# pin is de-asserted (driven high). This instructs the external clock driver to turn off the reference clock. The penalty for using this feature is a delay in the exit from L1 due to RefClock recovery time.

17.5.1.4 Exiting the Power Saving State

To exit from the power saving state following the Power Down procedure on page 137, and to restore the device to normal operations, follow these steps:

1. Set the PCI Express Power Management Control and Status register <PMState> field (bits [1:0]) back to 0x0 ($D0_{uninitialized}$ state).
2. Signal the 88AP510 software driver to resume normal operation.

As a result of a <PMState> field change, the 88AP510 sets the PCI Express Interrupt Cause register <DStateChange> field (bit [11]), and asserts a CPU interrupt, if not masked. The device's embedded CPU uses this interrupt to return the board components to normal operation.

To power up the Endpoint device following the Additional Power Saving procedure on page 137, follow these steps:

1. Turn on power to the device.

The period required to power up is dependant on the reset sequence.

2. Wait for the <DLDown> field (bit [0]) in the PCI Express Status Register to be cleared to 0x0. This setting means that the DL link is active.

The Endpoint device is ready for enumeration and software initialization.

Note When designing the software, take into account the fact that the Endpoint device initialization is software-dependent.

17.5.1.5 Complete Shutdown of an Unused PCI Express Interface

1. Disable training of the PCI Express interface by clearing the <PCIDs0> field (bit [0]) and the <PCIDs1> field (bit [3]) in the CPU Control and Status Register to 0x0.
2. Clear the <PEX0PwrDwn> or <PEX1PwrDwn> fields for the relevant port to 0x1 (power down) in the IO Power-Down Control Register.

17.5.2 USB Power Saving

The device supports powering-down of any attached USB device, as specified in the EHCI specification.

Upon connection of an external device to a USB port in an inactive state (no external device connected), the PORTSC1 register <ConnectStatusChange> field (bit [1]) is set, and an interrupt is asserted, if not masked.

At this point, the software wakes up the bus.

Upon a disconnect event, the software turns off the USB port.

Table 41 lists the USB power-saving resources that are available.

Table 41: USB Power-Saving Resources

Register	Offset	Field	Bits	Power-on Value
Clock Gating Control	0xD0038	USB0EnClock USB1EnClock	[0] [1]	0x1 0x1
USB 2.0 Power Control	0x5n400 n=port or port 1	PWRCTL_WAKEUP SUSPENDM PuPII Pu	[4] [2] [1] [0]	0x0 0x1 0x1 0x1

17.5.2.1 Complete Shutdown of an Unused USB Interface

If a a USB port is not used, the entire USB Controller and PHY of that port can be shut down.

To shut down the PHY, clear the USB 2.0 Power Control register <PuPll> field, bit[1] and <Pu> field, bit [0] to 0x0.

To shut down the Port controller, clear the <USB0EnClock> field of the Clock Gating Control Register bit [0] for port 0 or the <USB1EnClock> field of the Clock Gating Control Register bit [1] for port 1.

17.5.3 Ethernet Power Saving

It is possible to reduce the device power consumption by disabling the Ethernet port and its clock.

It is also recommended refer to the *Green Ethernet with Marvell® Alaska® PHYs and Integrated Switches* application note.

Table 42 lists the Ethernet power-saving resources that are available.

Table 42: Ethernet Power-Saving Resources

Register	Offset	Field	Bits	Power-on Value
IO Power Control	0xD0058	Giga IO Power Down	[2]	0x1
Clock Gating Control	0xD0038	GEnClock	[2]	0x1
Transmit Queue Command	0x72448	DISQ ENQ	[15:8] [7:0]	0x0 0x0
Receive Queue Command	0x72680	DISQ ENQ	[15:8] [7:0]	0x0 0x0
Ethernet Port Status	0x72444	TxFIFOEmp TxInProg LinkUp	[10] [7] [1]	0x0 0x0 0x0
Port Serial Control	0x7243C	ForceLinkFail ForceLinkPass PortEN	[10] [1] [0]	0x0 0x0 0x0

17.5.3.1 Entering the Power Saving State

To disable the Ethernet port, complete these steps:

1. Set the active queues in the Transmit Queue Command (TQC) and Receive Queue Command (RQC) registers <DISQ> fields. There is one bit per queue. This action stops all active Transmit and Receive queues.
2. Verify that all the <ENQ> bits in the Transmit Queue Command (TQC) and Receive Queue Command (RQC) registers are set to 0. If so, then all of the queues have stopped.
3. Read the Ethernet Port Status (PS) register to confirm that the <TxFIFOEmp> field (bit [10]) and the <TxInProg> field (bit [7]) have a value of 0x1. This value means that the Transmit FIFO is empty, and that all Transmit activity has stopped.
4. Clear the <ForceLinkFail> field (bit [10]) in the Port Serial Control (PSC) register to force the link down.
5. Clear the <PortEN> field (bit [0]) in the Port Serial Control (PSC) register to disable the port.

To disable the Ethernet port clock, clear the <GEnClock> field (bit [2])in the Clock Gating Control Register.

To disable the Ethernet I/O power, clear the <Giga IO Power Down> field (bit [2]) in the IO Power Control Register.

17.5.3.2 Exiting the Power Saving State

To exit the power saving state, and enable the port and the clock, complete the following steps:

1. Set the <Giga IO Power Down> field (bit[2]) in the IO Power Control register to enable the Ethernet IO power.
2. Set the <GEnClock> field in the Clock Gating Control register to enable the Ethernet port clock.
3. Set the <PortEN> field (bit [0]) in the Port Serial Control (PSC) register to enable the port.
4. Set the <ForceLinkPass> field (bit [1]) in the Port Serial Control (PSC) register to cancel the force linkdown mode, and allow the linkup mode.

5. Verify that the <LinkUp> field bit [1] in the Ethernet Port Status (PS) register has a value of 0x1, indicating that the link is up.
6. Set the active queues in the Transmit Queue Command (TQC) and Receive Queue Command (RQC) registers <ENQ> bits [7:0]. There is one bit per queue. This setting enables all active Transmit and Receive queues.

17.5.4 SATA Power Saving

When using the SATA interface, the power consumption can be reduced by disabling the SATA PHY. If the SATA interface is not used, it is still recommended to disable the interface to achieve minimum power consumption.

To disable the PHY, shut down the link, and disable the port PHY.

Table 43 lists the SATA power-saving resources that are available.

Table 43: SATA Power-Saving Resources

Register	Offset	Field	Bits	Power-on Value
Clock Gating Control	0xD0038	SunitEnClock	[3]	0x1
SControl	0xA2308	SPM IPM	[15:12] [11:8]	0x0 0x0
PHY Mode 2	0xA2330	PU_IVREF PU_PLL FORCE_PU_RX FORCE_PU_TX	[3] [2] [1] [0]	0x1 0x1 0x1 0x1
Serial-ATA Interface Configuration	0xA2050	PhyShutdown	[9]	0x0

17.5.4.1 Entering the Power Saving State

To shut down the link, complete the following steps:

1. Verify that value of the <IPM> field (bits [11:8]) in the SControl register represents the enabled Interface Power Management states that can be invoked via the Serial ATA Interface power management capabilities.
2. Set the <SPM> field (bits [15:12]) in the SControl register to 0x2. This setting initiates the SLUMBER power management state.

Note

It is also possible to set the <SPM> field to 0x1. This initiates a partial power management mode that is a faster, but less efficient, state.

To disable the port PHY after the <SPM> field is set, complete the following steps:

1. Clear the following fields in the PHY Mode 2 register to 0:
   - <PU_IVREF> (bit [3])
   - <PU_PLL> (bit [2])
   - <FORCE_PU_RX> (bit [1])
   - <FORCE_PU_TX> (bit [0])
2. Set the <PhyShutdown> field (bit [9]) in the Serial-ATA Interface Configuration register to 0x1. This setting places the PHY in Shutdown mode.

17.5.4.2 Exiting the Power Saving State

To reactivate the SATA port, complete the following steps:

1. Clear the <PhyShutdown> field in the Serial-ATA Interface Configuration register to 0x0. This setting places the PHY in a Functional mode.
2. Set the following fields in the PHY Mode 2 register to 1:
   - <PU_IVREF> (bit [3])
   - <PU_PLL> (bit [2])
   - <FORCE_PU_RX> (bit [1])
   - <FORCE_PU_TX> (bit [0])
3. Set <SPM> field (bits [15:12]) in the SControl register to 0x3. This setting places the PHY in Active mode.

17.5.4.3 Complete Shutdown of an Unused Interface

If a SATA interface is not used, the entire SATA port (PHY and Controller) can be shut down by following this procedure in sequence:

1. Clear the following fields in the PHY Mode 2 register to 0:
   - <PU_IVREF> (bit [3])
   - <PU_PLL> (bit [2])
   - <FORCE_PU_RX> (bit [1])
   - <FORCE_PU_TX> (bit [0])
2. Set the <PhyShutdown> field in the Serial-ATA Interface Configuration register to 0x1. This sets the PHY Analog section in Shutdown mode.
3. Clear the Clock Gating Control register <SunitEnClock> field (bit [3]). This shuts down the SATA PHY digital section.

17.5.5 Clock Gating for Unused Interfaces

The device also has the option to stop the clocks of any unused interfaces by setting the relevant bits in the Clock Gating Control register (Offset: 0xD0038) that are listed in Table 44.

To enable the power saving state, clear the relevant interface's bit to 0.

To disable the power saving state, set the relevant interface's bit to 1.

Table 44: Clock Gating and I/O Power Down Fields

Interface	Field	Bits
SDIO 0	SD0EnClock	[8]
SDIO 1	SD1EnClock	[9]
NAND Flash	NFEnClock	[10]
Camera	CamEnClock	[11]
Audio 0	AD0EnClock	[12]
Audio 1	AD1EnClock	[13]
CESA	TuEnClock	[15]
PDMA	PdmaEnClock	[22]
XOR 0	XE0EnClock	[23]
XOR 1	XE1EnClock	[24]

18 Additional Recommendations

This section provides additional design recommendations.

18.1 Clock Topology

This section presents all of the required clocks for the devices.

Figure 82 shows all of the required clocks for the device.

Figure 82: Device Clock Topology

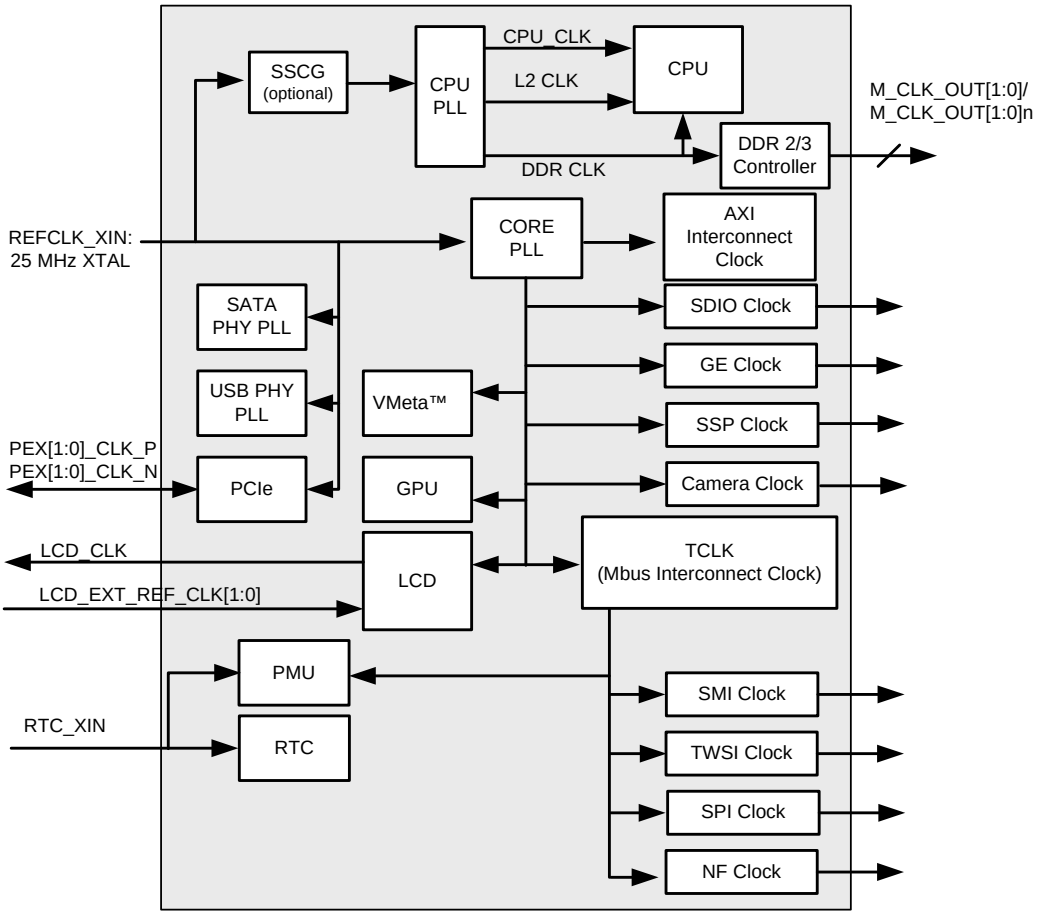

Note If a device interface is not being used, refer to the Unused Interface Strapping section in the *Hardware Specifications* (see Related Documentation on page 14) for the unused interface signal strapping recommendations.

The device integrates an internal clock generator that significantly reduces the use of external clock buffers and high frequencies oscillators.

All of the clocks need to be stable and quiet, and respect the requirements for jitter, skew rate, duty cycle, and other parameters as defined in the *Hardware Specification* (see Related Documentation on page 14).

18.2 System Clock Guidelines for Crystal Implementation

The device can use a crystal as an input to the core reference clock. This section provides design guidelines for crystal implementation. A crystal is also required for operation of the Real-Time Clock (RTC) in the 88AP510.

18.2.1 Crystal Oscillator Circuit for REF_CLK_XIN

Figure 83 provides a schematic example as a guideline for the XTAL_IN / XTAL_OUT (REFCLK_XIN / REFCLK_XOUT) connections.

Figure 83: Example of XTAL_IN/XTAL_OUT Connections

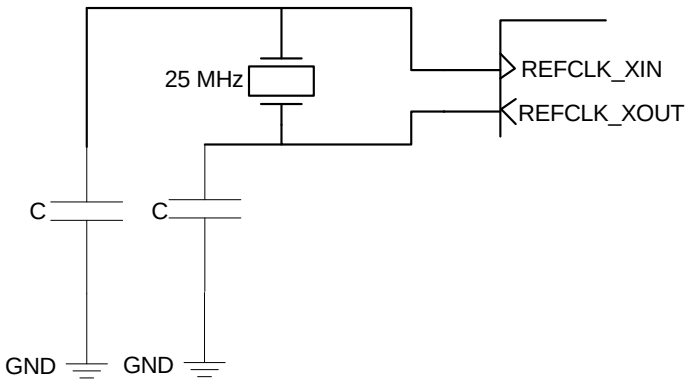

NOTE: C = $2C_L$ - C_{pin} where C_L is the load capacitance of the crystal.

Note The RTC design was optimized for a standard CL = 12.5 pF crystal. No passive components are provided internally. Connect the crystal and the passive network as recommended by the crystal manufacturer.

18.2.2 PCB Layout Guidelines

Since the Marvell® device crystal input has very high impedance, the leads of the crystal can behave like an antenna. The crystal leads are able to couple high frequency signals from the rest of the

system to the input comparator of the clock source and generate clock noise. It is important to follow the following PCB layout guidelines to minimize the clock noise.

- Place the crystal as close as possible to the XTAL1 and XTAL2 pins to minimize the trace length. Short traces help to reduce the parasitic inductance and stray capacitance, and decrease the amount of noise coupling.
- Make the width of the trace to XTAL1 and XTAL2 to be approximately 12 mil. Thicker traces increase the stray capacitance as well as the likelihood of capacitive noise coupling. Thin traces can cause extra parasitic inductance and resistance.
- Avoid routing any other signal trace directly underneath the crystal or the XTAL1 and XTAL2 traces, especially not running parallel with the clock traces. Try to keep all other signal traces at least 200 mil away from the XTAL1 and XATL2 traces.
- Do not use the crystal output to drive any other device. The crystal oscillator is designed to drive the crystal in a direct feedback loop. Any extra branching or loading may effect the oscillator's ability to function properly.
- It is highly recommended to have a ground ring (a signal trace connected to ground plane) on the signal layer surrounding the area formed by the XTAL1 pin, XATL2 pin, and the footprint of the crystal (see Figure 84).

Figure 84: Ground Ring Surrounding the Crystal Oscillator Circuit

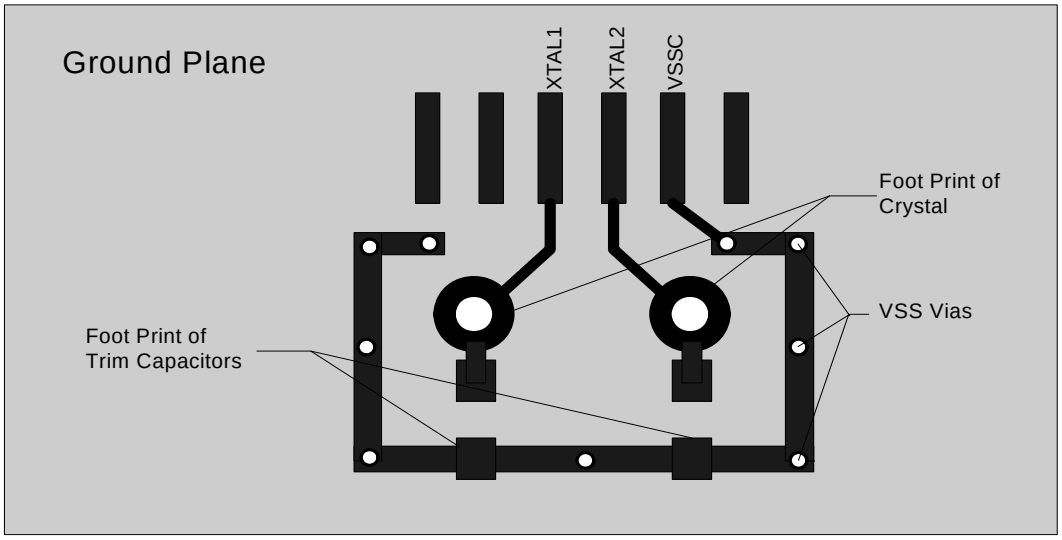

18.3 System Reset Sequence

This section describes the required system reset timing and sequence from the point at which the power supplies and clocks are stabilized and activated.

The following information does not contradict the information provided in the hardware specifications, and should be treated as a visual clarification of the information written in the hardware specifications. Figure 85 shows the required reset de-assertion sequence with the minimum timing requirements.

Figure 85: Reset De-assertion Sequence and Timing

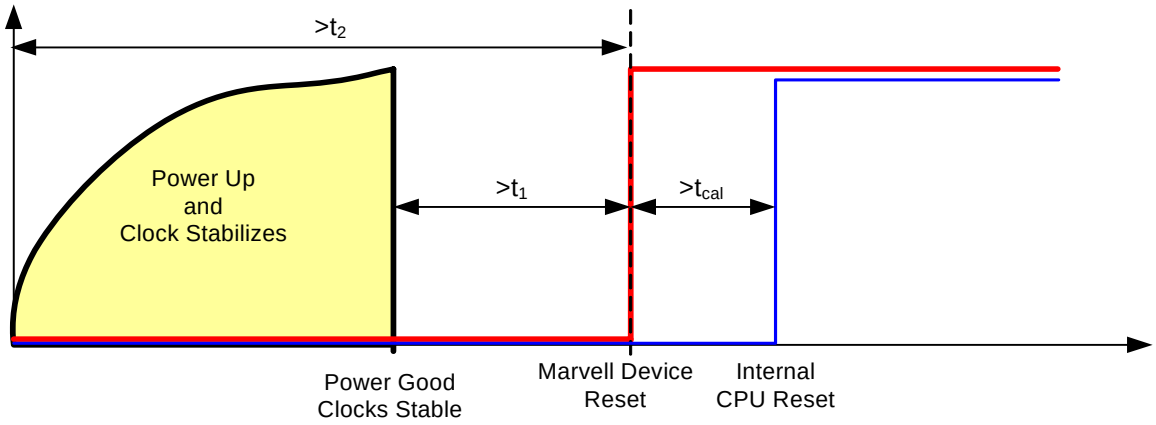

- t_1: The period from good power and a stable clock until device SYSRSTn reset de-assertion should be at least 300 us.
- t_2: The period from power up until device SYSRSTn reset de-assertion should be at least 100 ms.
- t_{CAL}: The internal data path initialization, and the for the output drivers impedance auto-calibration for the SDRAM and GbE interfaces. Requires at least 512 TCLK cycles, or 3.072 us if the TCLK frequency equals 166 MHz.

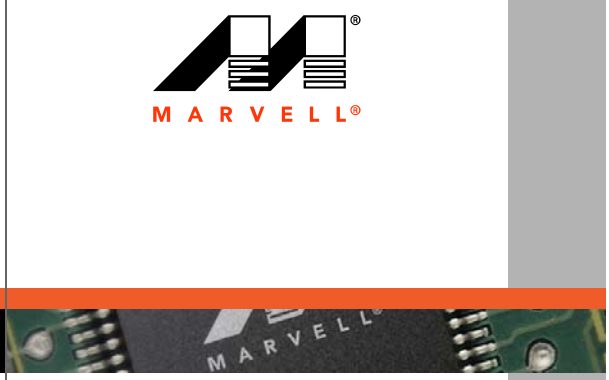

88AP510

Hardware Design Guide

Appendix

Doc. No. MV-S301669-U0, Rev. E
July 5, 2011, Preliminary

A Hardware Abstraction Layer Power Management Procedures Appendix

This application details the Hardware Abstraction Layer (HAL) procedures that are required to support the power management features in the 88AP510 device.

The 88AP510 integrates an advance power management unit (PMU) with advanced capabilities that allows for a longer battery life. These features improve the power versus performance ratio to ensure energy is available for driving power-hungry applications, such as video or web browsing.

The PMU uses a “Micro-Sequencer” logic that executes “Micro-Instructions”. All Micro-Instructions are loaded into the “Instruction space” at boot time and on recovery from Standby mode. The Instruction space has 512 bytes, and includes several routines. Each routine performs a specific flow (Standby entry, Standby exit, Deep-idle entry, etc.).The entry point of each routine, offset from the instruction space start, is configured in a dedicated register.

The PMU also includes 2 KB of SRAM that is used by the HAL procedures if the DRAM inaccessible.

Figure 86: 88AP510 Power Domains and PMU Control Signals

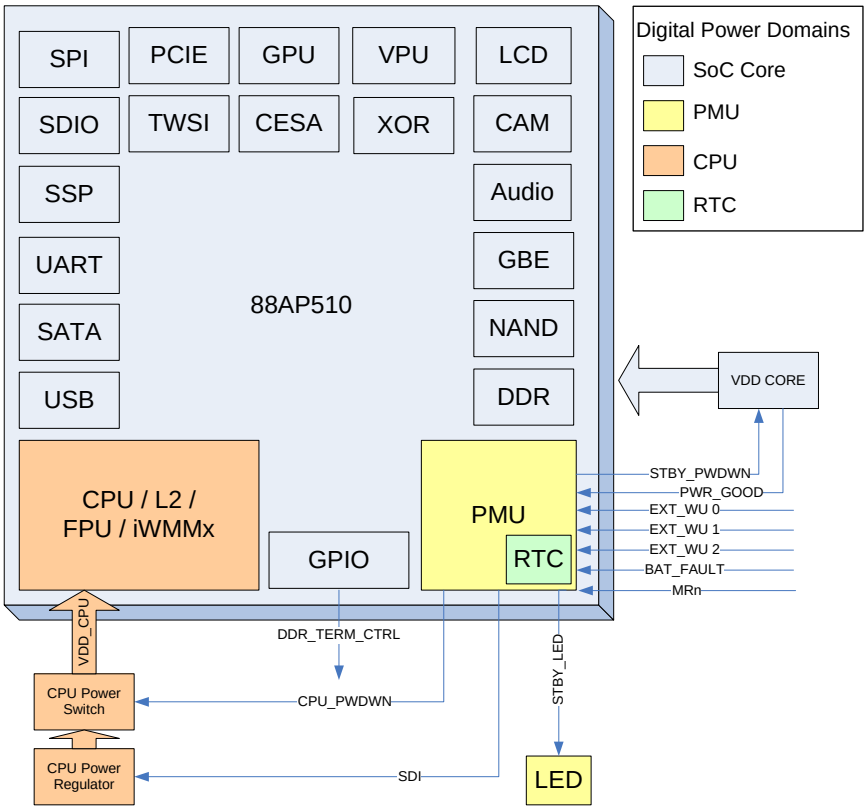

Note

- For details regarding the Power Management Unit (PMU) internal operations, see the *88AP510 Functional Specification.*
- For the device's power consumption values, see the *88AP510 Hardware Specification.*
- The term "device" is used in this document to refer to the 88AP510.

A.1 Software Model

Figure A.1 shows the software part that this document describes and its interfaces to the higher software layers.

Figure 87: 88AP510 Hardware Abstraction Layer Power Management Block

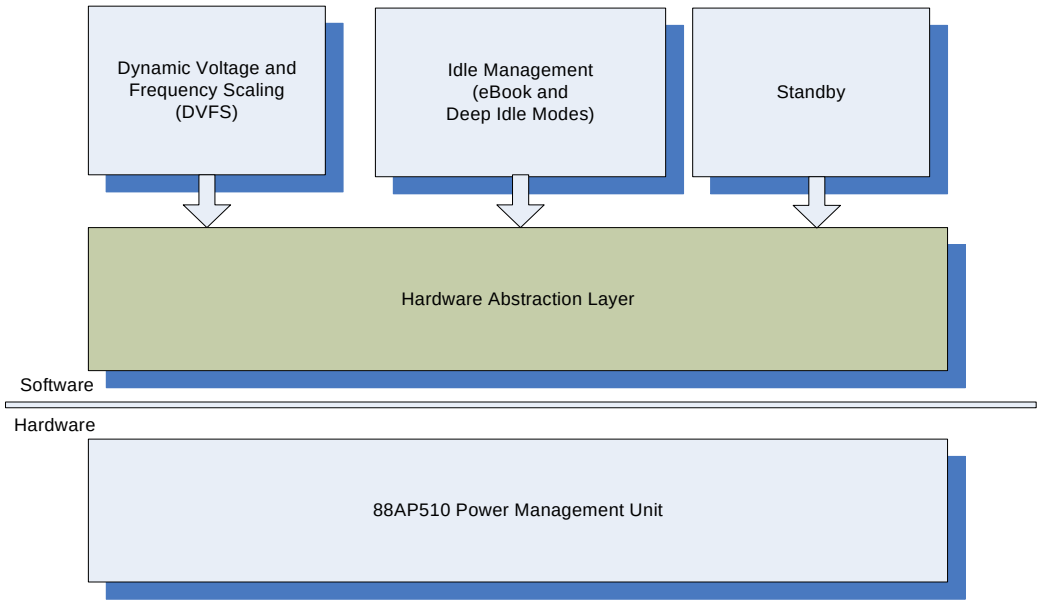

The Power Management (PM) procedures are a set of routines that implement the flow needed to operate each of the 88AP510 PM features. These procedures include:

- Configuring the PMU instruction space.
- Copying the necessary CPU code to the PMU SRAM (to be executed while the CPU has no access to the DDR).
- Preparing the necessary configuration to be passed to the hardware or to the device's BootROM.
- Completing the necessary register configurations.

These routines are implemented in the Hardware Abstraction Layer (HAL) provided with the device's BSP. As part of the explanation, the descriptions refer directly to the routine names, file names, and PMU registers offsets.

A.2 Initialization

This flow refers to the specific implementation on the 88AP510 development board.

A.2.1 Power Management Unit Initialization

The PMU initialization process is completed once at boot time by running the mvPmuInit(...) routine. The following are the initialization steps:

1. Configure the Standby Power Delay to indicate the number of RTC clock cycles (32 Khz) that the PMU must wait after enabling the power from Standby mode to resume. This period required for the power to stabilize, and for the core PLL to lock. The trigger to this counter is the PWR_GOOD signal. On the 88AP510 development board, the Core Power switch is used as the external power good indication.

   ```
   StbyPwrDel 0xD8018[31:0] = 0x140 // Wait 10ms after powering up.
   ```

2. Upon a battery fault event, configure the Battery Fault handling to control the wake-up.

   ```
   BattFaultMngDis 0xD801C[0] = 0x0 // Assertion of the Battery Fault Pin wakes up the system from Standby.

   BattFaultStbyExitEn 0xD801C[1] = 0x1 // The system can wake-up from Standby regardless of the status of the Battery Fault Pin. Otherwise, the system will not wake-up from Standby until deasserting the Battery Fault Event.
   ```

3. Configure the polarity of the CPU_PWDWN and STBY_PWDWN power down pins. This configuration depends on the polarity of the external regulator supplying the Core and CPU digital power (VDD_CORE and VDD_CPU). In the configuration below, the CORE and CPU are powered down through asserting the relevant pin LOW.

   ```
   CPUPwrDnEn 0xD8010[1] = 0x0 // CPU Power Down active LOW. Deep-Idle active.

   cPUPwrDwnDis 0xD8010[4] = 0x1 // CPU Power Up when HI. Deep-Idle not active.

   StbyPwrDnEn 0xD8010[7] = 0x0 // CORE Power Down active LOW. Standby active.

   StbyPwrDnDis 0xD8010[5] = 0x1 // CORE Power Up when HI. Standby not active.
   ```

4. Enable PMU control on CPU reset to allow the PMU to control the CPU reset.

   ```
   PMUCPURstEn 0xD025C [1] = 0x1 // PMU controls the CPU reset.
   ```

5. Configure the PMU signal selection to determine the functionality of each MPP pin used by the PMU. This depends on the board connectivity. The following is an example that uses MPP3 as an external Standby
   wake-up event:

   ```
   0xD802C[15:12] = 0xB // EXT0 Wake-up.
   ```

6. Configure the PMU to use the Power Good indication, if an external MPP is used for this purpose.

   ```
   PwrGoodPINEn 0xD8010[0] = 0x1 // External Power Good indication is used.
   ```

7. Configure the Standby LED blinking period to set the blinking speed and duty cycle of the Standby LED.

   ```
   PmuBlinkTotalPeriod 0xD8034[31:0] = 0x10000 // 2 seconds (64K RTC 32Khz clock cycles).

   PmuBlinkOnePeriod 0xD8038 [31:0] = 0xE000 // 1.75 seconds (The LED is active low, thus the duty cycle is 1/8)
   ```

8. Configure the Dynamic Voltage Scaling (DVS) delay to set the time the PMU waits for the voltage to stabilize before resuming execution.

   ```
   DVSDel 0xD0018[31:0] 0x0 // Do not wait after changing the CPU voltage.
   ```

9. Load the PMU instruction space and PMU routine pointers.

   ```
   0xD0500 - 0xD06FC // 512 bytes representing the entire instruction space used by the PMU.

   0xD0100 - 0xD0120 // offset of the different PMU routines within the instruction space.

   A list of all instruction space changes is found in the uCfix[] array in the mvPmu.c

   The value of all PMU routine offsets can be found in the file mvPmuRegs.h
   ```

10. Read the initial PLL:L2 ratio value. This value is needed in the frequency scaling routines, and it changes depending on the reset sampling.
11. Relocate all necessary routines to the SRAM, so the routines are executed from that location. This is done by calling the mvPmuSRAMInit(...) routine located in the mvSRAM.c file. Section A.2.2 describes the SRAM initialization.

The following settings are system level configurations, and not part of mvPmuInit(...) routine. These configuration should be completed after the previous configurations in the mvPmuInit(..).

1. Configure which MPP pins are used by the PMU based on the board connectivity. The following is an example that uses MPP3 as a an external STANDBY wake-up event.

   ```
   0xD0210[3] = 0x1 // MPP3 is used as a PMU pin and not a regular MPP.
   ```

2. Configure the allowed PMU wake-up sources and their polarity. The following uses both RTC and external pin #0.

   ```
   RTC_WUD 0xD8008[0] = 0x1 // wake-up on RTC events
   EXT0_WUD 0xD8008[2:1] = 0x2 // wake-up on the falling edge of EXT0.
   ```

A.2.2 SRAM Initialization

The SRAM initialization process is performed once at boot time. It is part of the PMU initialization, as the last step in the mvPmuInit(...) routine. The purpose of the SRAM initialization is to allocate space for all configurations parameters and routines. Table 45 lists all of the information contained in the SRAM.

Table 45: Information Located in the SRAM

Offset	Size (Byte)	Content	Details
0x0	0x4	DDR Termination Power Shutdown Control	Enable(1) /Disable(0) VTT shutdown during Deep-Idle.
0x4	0x4	DDR Termination Power Shutdown GPIO Mask	Mask of GPIO used to disable VTT.
0x8	0x18	Reserved	Alignment to 32B cache line.
0x20	0x120	DDR Configuration Parameters	BootROM Resume descriptor #1 Thirty-six (36) DDR configuration couples (Address, Value) padded with 0x0 to be aligned to cache line. This descriptor is used in resume from Standby mode. The allocated space represents the maximum size used in all possible DDR configurations (DDR2/DDR3 and frequency).
0x140	0xC	DDR Initialization Completed	BootROM Resume descriptor #2 One (Address, Mask, Value) to be used by the BootROM in Standby resume for polling the DDR initialization completion.
0xx14C	0x14	Reserved	Alignment to 32B cache line.
0x160	0xC0	SYS DFS Routine	Low level routine to perform system (CPU/DDR/L2) frequency scaling.
0x220	0x120	Deep-Idle Entry Routine	Entry code for both Deep-Idle and eBook modes.
0x340	0x160	Deep-Idle Exit Routine	BootROM entry point in resume from Deep-Idle and eBook modes.
0x4A0	0xC0	Standby Entry Routine	Entry code for Standby mode.

Table 45: Information Located in the SRAM (Continued)

Offset	Size (Byte)	Content	Details
0x560	0xC0	Standby Exit Routine	BootROM entry point in resume from Standby mode.
0x620	0x60	CPU DFS Routine	Code that performs CPU frequency scaling mode.
0x680	0x180	Not Used	Empty space left in SRAM.

A.3 Hardware Abstraction Layer Flows

This section describes each of the PM software flows, referring to the specific implementation on the 88AP510 development board.

A.3.1 Standby PM Flow

The Standby PM flow is divided into several parts that starts from the high level system flow level, continues through the HAL and PMU steps, and into the exit sequence.

Figure 88 shows the Standby mode PM flow.

Figure 88: Standby PM Flow

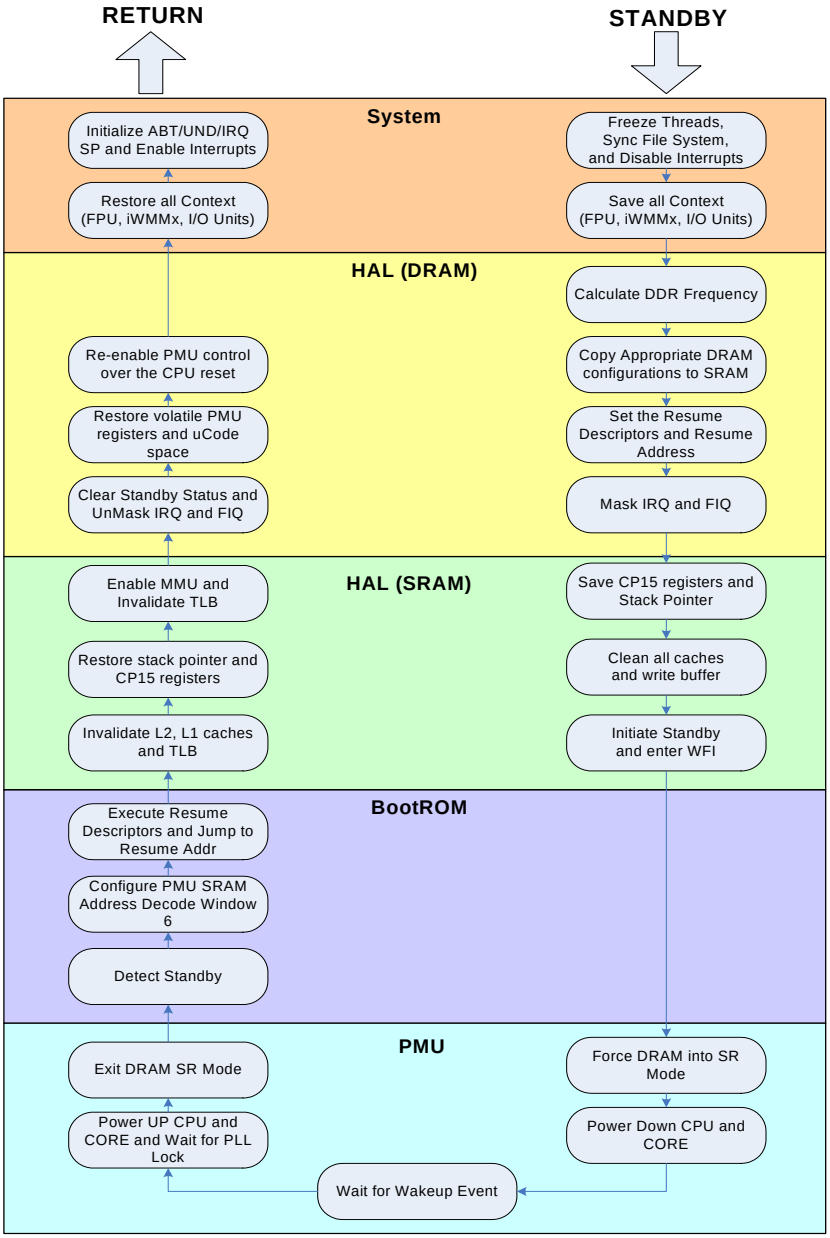

A.3.1.1 System Entry

1. The Standby flow is triggered by a user request.
2. All threads are frozen except for the thread performing the Standby flow.
3. All temporary buffers are flushed to the storage (FileSystem sync).
4. All I/O devices are set in an IDLE state and their context is saved in the DRAM.
5. Disable IRQ and FIQ at the CPSR level.
6. Save iWMMX context.
7. Save FPU context.
8. Save context of the entire CPU environment including:

- Address decoding windows
- CPU configuration
- Upstream bridge configuration
- Timer configurations
- Interrupt configurations.

A.3.1.2 HAL Entry

The PMU Low Level flow description (through calling mvPmuStandby() found in mvPmu.c) is as follows:

1. Read the CPU PLL Control 0 register and calculate the PLL frequency.

   ```
   apll_cfg_ndiv_stts 0xD003C[8:0] // N = apll_cfg_ndiv_stts + 1
   apll_cfg_mdiv_stts 0xD003C[17:9] // M = apll_cfg_ndiv_stts + 1
   apll_cfg_kdiv_stts 0xD003C[20:18] // K = apll_cfg_kdiv_stts
   PLL Frequency = ((25 * N) / (M * K))
   ```

2. Read the CPU CLock Divider Control 0 register to get the PLL:DDR ratio.

   ```
   dpratio 0xD0044[29:24] // PLL:DDR ratio
   DDR frequency = PLL Frequency / PLL:DDR ratio
   ```

3. Set the BootROM resume address. This is the physical address of the Standby mode exit routine in the SRAM.

   ```
   Resume Address 0xD8104[31:0] =0xF000C560 // Standby resume address. This is the physical address of the Standby resume routine in the PMU SRAM.
   ```

4. Set the BootROM Resume Control register to instruct the BootROM how to resume.

   ```
   Descriptor Count 0xD8100[2:0] = 0x2 // Two BootROM resume descriptors (listed in the SRAM tabled).
   Target ID 0xD8100[7:4] = 0xD // This is the PMU SRAM Target Interface ID.
   Target Attribute 0xD8100[15:8] = 0x0 // This is the PMU SRAM Target Interface Attribute. This should always be 0x0.
   Target Base 0xD8100[31:16] = 0xF000 // This is the base address of the PMU SRAM address decoding window.
   The previous 3 parameters (ID, Attribute and Base) are used by the BootROM firmware to setup the address decoding window to the SRAM before resuming.
   ```

5. Initialize all of the space allocated for the DDR parameters with 0x0.
6. Copy into the appropriate SRAM location all DDR initialization parameters (Address/Value couples) that are suitable for the detected DDR frequency and type (DDR2 or DDR3). These configurations, and their count, can differ due to the DDR type and frequency used. The space allocated in the SRAM should cover the worst case (i.e. maximum count) instance.

7. Setup BootROM Resume Descriptor #1. This is the descriptor used to initialize the DRAM.

   ```
   Desc#1: Type 0xD8108[3:0] = 0x1 // Descriptor pointing to Address/Value couples.
   Desc#1: Configurations count 0xD8108[14:4] = 0x24 // 36 Address/Value couples are used.
   Desc#1: Delay 0xD8108[31:15] = 0x0 // This is the delay introduced after each Address = Value setting.
   Desc#1: Address 0xD810C [31:16] = 0xF000 // Base address of PMU SRAM window.
   Desc#1: SRAM Offset 0xD810C [15:12] = 0xC // This is the offset of the SRAM and should not be changed.
   Desc#1: offset 0xD810c[11:0] = 0x20 // Offset of DRAM initialization couples in the SRAM (listed in the PMU SRAM table).
   The previous 3 parameters make up the physical address 0xF000C020.
   ```

8. Setup BootROM Resume Descriptor #2. This is the descriptor used to poll the DRAM initialization completion bit.

   ```
   Desc#2: Type 0xD8108[3:0] = 0x7 // Descriptor pointing to the polling information represented in the form of Address/Mask/Value.
   Desc#2: Configurations count 0xD8108[14:4] = 0x1 // 1 register is polled only.
   Desc#2: Delay 0xD8108[31:15] = 0x0 // This is the delay introduced after each Address = Value setting.
   Desc#2: Address 0xD810C [31:16] = 0xF000 // Base address of PMU SRAM window.
   Desc#2: SRAM Offset 0xD810C [15:12] = 0xC // This is the offset of the SRAM and should not be changed.
   Desc#2: offset 0xD810c[11:0] = 0x140 // Offset of the polling information in the PMU SRAM (listed in the PMU SRAM table).
   The previous 3 parameters make up the physical address 0xF000C140.
   ```

9. Mask out IRQ and FIQ from being asserted towards the CPU.

   ```
   MaskIRQ 0xD8000[24] = 0x1 // IRQ signal is masked out at the PMU level.
   MaskFIQ 0xD8000[27] = 0x1 // FIQ signal is masked out at the PMU level.
   ```

10. Set the CPU power delay to minimum.

    ```
    CPUPwrDel 0xD8014[31:0] = 0x1 // Wait 1 RTC (32Khz) clock cycle after powering up the CPU.
    ```

11. Jump to the SRAM and continue the procedure. The address to jump to is the virtual address of the Standby mode entry routine in the SRAM (virtual address of 0xF000C4A0).
12. Save all general purpose registers and LR (r0–r12 and LR) on the stack.
13. Save all CP15 registers on the stack.This includes:

    ```
    CR: Control Register // mrc    p15, 0, r0, c1, c0, 0
    DACR: Domain Access Control Register // mrc   p15, 0, r1, c3, c0, 0
    FCSE: Fast Context Switch Register // mrc   p15, 0, r2, c13, c0, 0
    TTBR: Translation Table Base Register // mrc   p15, 0, r3, c2, c0, 0
    Context ID Register // mrc   p15, 0, r4, c13, c0, 1
    Extra Features Register // mrc   p15, 1, r5, c15, c1, 0
    Stack Pointer
    ```

14. Save the Physical address of Stack Pointer on the SRAM (within the resume code).

15. Flush all caches:

```
DCache clean & invalidate // mcr   p15, 0, r0, c7, c14, 0
flush BTAC/BTB // mcr   p15, 0, r0, c7, c5, 6
I-BTB cache invalidate // mcr   p15, 0, r0, c7, c5, 0
Clean L2 cache // mcr   p15, 1, r0, c7, c11, 0
Drain write buffer // mcr   p15, 0, r0, c7, c10, 4
Drain memory barrier // mcr   p15, 0, r0, c7, c10, 5
```

16. Initiate the Standby flow.

```
StbyPwrEn 0xD8000[4] = 0x1 // Enable Standby on next entry to WFI.
DDRSelfRefEn 0xD8000[5] = 0x1 // Force DRAM into SR mode on next entry to WFI.
```

17. Enter the WFI mode. This triggers the PMU to start the Standby flow.

```
Enter WFI // mcr   p15, 0, r0, c7, c0, 0
```

A.3.1.3 PMU Entry and Exit

The following flow is performed by the PMU to:

1. Power down the CPU and CORE domains.
2. Wait for one of the configured wake-up events.
3. Power up the CORE and CPU power domains.
4. Wait for the CPU to exit reset.

A.3.1.4 BootROM Exit

The following flow is performed by the CPU upon exiting reset.

1. The BootROM checks if the CPU is resuming from Standby mode. Then, the BootROM continues with the resume flow, otherwise a normal boot is performed.
2. Configure the address decoding window to the PMU SRAM.

```
Window 6 win_en 0x20060[0] = 0x0 // Close the window before updating it.
Window 6 Base 0x20064[31:16] = 0xF000 // Taken from Target Base 0xD8100[31:16].
Window 6 Target 0x20060[7:4] = 0xD // Taken from Target ID 0xD8100[7:4].
Window 6 Attribute 0x20060[15:8] = 0x0 // Taken from Target Attribute
0xD8100[15:8].
Window 6 Size 0x20060[31:61] = 0x0 // 64KB of space
```

3. Read the number of descriptors to execute.

```
Descriptor Count 0xD8100[2:0] = 0x2 // 2 descriptors to be executed.
```

4. Execute the first descriptor by writing all 36 Address value couples, starting from the location indicated in the PMU Resume Descriptor Address 0 Register.

```
pointer 0xD810C[31:0] // pointer to first Address/Value couple of the 36.
```

5. Execute the second descriptor by polling the indicated address and comparing it to the masked value. The location of the first (and only) Address/Mask/Value set is indicated in the PMU Resume Descriptor Address 1 register.

```
pointer 0xD81014[31:0] // pointer to Address/Mask/Value
```

6. Jump to the resume address indicated in the SRAM.

```
Resume Address 0xD8104[31:0] // resume physical address in SRAM (0xF000C560)
```

A.3.1.5 HAL Exit

1. Invalidate all caches.

   ```
   Invalidate both caches and BTC // mcr   p15, 0, r0, c5, c7, 0
   TLB invalidate // mcr   p15, 0, r0, c8, c7, 0
   L2 cache invalidate // mcr   p15, 1, r0, c7, c7, 0
   Prefetch flush (ISB) // mcr   p15, 0, r0, c7, c5, 4
   ```

2. Read the physical address of the Stack Pointer (stored within the code in the SRAM).
3. Remove all registers saved on the stack.
4. Restore the saved CP15 registers

   ```
   Extra Features Register // mcr   p15, 1, r5, c15, c1, 0
   Context ID Register // mcr   p15, 0, r4, c13, c0, 1
   TTBR: Translation Table Base Register // mcr   p15, 0, r3, c2, c0, 0
   FCSE: Fast Context Switch Register // mcr   p15, 0, r2, c13, c0, 0
   DACR: Domain Access Control Register // mcr   p15, 0, r1, c3, c0, 0
   ```

5. Change the PMU SRAM entry in the page table (indicated in the TTBR) so that VA=PA. The original mapping is saved in one of the unused registers.
6. Change the page table base address entry in the page table (indicated in the TTBR) so VA=PA. The original mapping is saved in an unused register.
7. Align the code to a cache line.
8. Enable MMU by writing the CP15 CR value.

   ```
   CR: Control Register // mcr    p15, 0, r0, c1, c0, 0
   ```

9. Wait for MMU to be enabled.

   ```
   Prefetch Flush (ISB) // mcr   p15, 0, r0, c7, c5, 4
   ```

10. Jump to the Virtual domain. This is virtual address of the next instruction.
11. Restore the original page table mapping for the entry representing PMU SRAM.
12. Restore the original page table mapping for the entry representing page table base.
13. Invalidate TLB.

    ```
    TLB Invalidate // mcr   p15, 0, r0, c8, c7, 0
    ```

14. Restore the original Stack Pointer.
15. Restore all general purpose registers and go to LR (r0-r12), back to the DRAM.
16. Clear the Standby bit in the PMU Status register. Otherwise, the BootROM starts a resume sequence on the next reset.

    ```
    StbyPwrSt 0xD8004[4] = 0x0
    ```

17. Clear the Standby Enable bit in the PMU Control register.

    ```
    StdbyPwrEn 0xD8000[4] = 0x0
    ```

18. Unmask the IRQ and FIQ gating at the PMU level.

    ```
    MaskIRQ 0xD8000[24] = 0x0 // ungate IRQ signal.
    MaskFIQ 0xD8000[27] = 0x1 // ungate FIQ signal.
    ```

19. Since the PMU instruction space is volatile (instruction space is powered down in Standby mode), the same instruction space configurations done at the initialization must be executed again.

    ```
    0xD0500 - 0xD06FC // 512 bytes representing the entire instruction space used by the PMU.
    0xD0100 - 0xD0120 // offset of the different PMU routines within the instruction space.
    A list of all instruction space changes is found in the uCfix[] array in the mvPmu.c
    The value of all PMU routine offsets can be found in the file mvPmuRegs.h
    ```

20. Re-enable the PMU control on CPU reset. Setting this bit allows the PMU to control the CPU reset. This is the same step at the boot time initialization.

```
PMUCPURstEn 0xD025C [1] = 0x1 // PMU controls the CPU reset.
```

A.3.1.6 System Exit

1. Re-initialize the SP for all mode (IRQ, FIQ, ABT, UND). These are the same values used at initialization.
2. Restore context of the CPU environment including:
   - Address decoding windows
   - CPU configuration
   - Upstream bridge configuration
   - Timer configurations
   - Interrupt configurations.
3. Restore FPU context.
4. Restore iWMMX context.
5. Enable IRQ and FIQ at the CPSR level.
6. Restore the original context of all I/O devices saved in the DRAM.

A.3.2 Deep-Idle and eBook Modes PM Flow

The Deep-Idle and eBook modes are very similar. The only difference relates to the DRAM state in the specific mode. In eBook mode, the DRAM is working normally, and any can access the DRAM (i.e., the LCD controller refreshing its frame buffer). In Deep-Idle mode, the DRAM is set in its deepest power saving state (Self Refresh Mode - SR).

The Deep-Idle and eBook flows are divided into several parts that start from the system high level, continue through the HAL and PMU steps, and returns to the exit sequence. Some steps (related to entering/exiting DDR SR mode) are skipped in eBook mode. Figure 89 shows the Deep-Idle and eBook PM flow.

Figure 89: Deep-Idle and eBook Flow

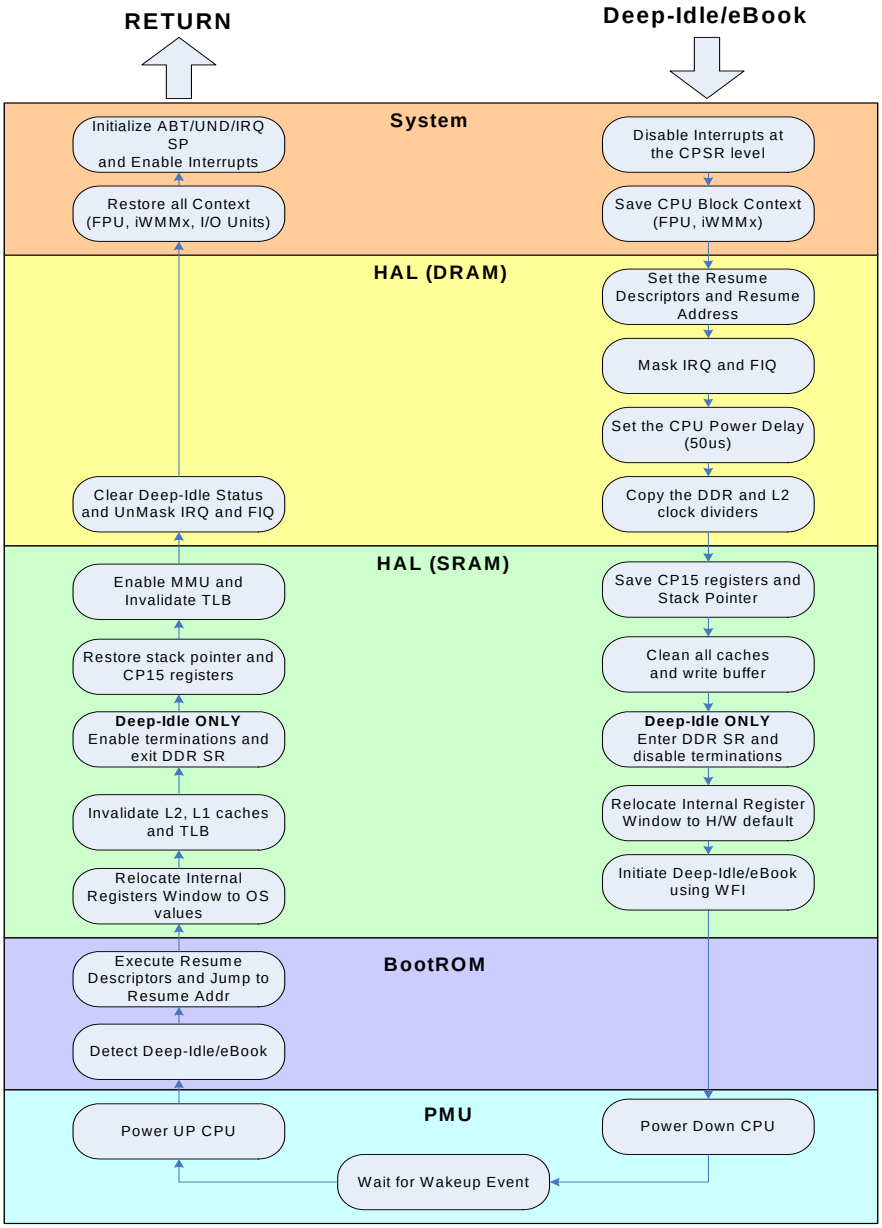

18.3.1 System Entry

Entry to Deep-Idle/eBook mode is triggered automatically by the system whenever the system is idle and have no active threads for execution. This is usually done by replacing the idle loop in the operating system based on certain criteria. The following is an example of one criteria for entering each of the Deep-Idle and eBook modes.

- The system is idle in the last X milliseconds, and no schedule work for the coming Y milliseconds and the LCD is active ==> use eBook mode instead of the idle loop on the next idle occurrence.
- The system is idle in the last X milliseconds, and no schedule work for the coming Y milliseconds and the LCD is not active (blanked due to user idleness) ==> use Deep-Idle mode instead of the idle loop on the next idle occurrence.

The following are the steps that should be performed at the system level:

1. Disable interrupts at CPSR level.
2. Save FPU, iWMMx context.

18.3.2 HAL Entry

The following are the steps performed by PMU Low Level HAL. This is triggered through calling the mvPmuDeepIdle() routine (found in mvPmu.c). This routine takes a single parameter that differentiates between Deep-Idle and eBook.

1. Set the BootROM resume address. This is the “Physical” address of the Deep-Idle exit routine in the SRAM.

   ```
   Resume Address 0xD8104[31:0] =0xF000C340 // Refer to the SRAM allocation table
   ```

2. Set the BootROM Resume Control Register. This instructs the BootROM how to perform the resume.

   ```
   Descriptor Count 0xD8100[2:0] = 0x0 // No resume descriptors needed in this mode.
   ```

3. Mask out IRQ and FIQ from being asserted towards the CPU.

   ```
   MaskIRQ 0xD8000[24] = 0x1 // IRQ signal is masked out at the PMU level.
   MaskFIQ 0xD8000[27] = 0x1 // FIQ signal is masked out at the PMU level.
   DDRSelfRefEn 0xD8000[5] = 0x0 // Make sure that the DDR will not be set in SR
   ```

4. Set the CPU power delay in TCLK cycles. This parameter depends on how much time does the CPU power need to raise up and stabilize.

   ```
   CPUPwrDel 0xD8014[31:0] = 0x2080 // Wait 50 micro second after powering up the CPU.
   ```

5. Decide the DDR and L2 ratios (PLL:DDR and PLL:L2) that will be used after exiting Deep-Idle. In normal operation, these parameters will not be changed and thus should be copied from the “PMU Clock Divider Control 0 Register”

   ```
   DDRLRatio 0xD0000[8:3] = 0xD044[29:24] // current PLL:DDR ratio
   CpuL2CR 0xD0000[14:9] = 0xD0044[21:16] // current PLL:L2 ratio
   ```

6. Jump to the SRAM and continue execution from there. The address to jump to is the virtual address of the Deep-Idle entry routine in the SRAM (virtual address of 0xF000C220).
7. Save all general purpose and link registers (r0-r12 and LR) on the stack.
8. Save all CP15 registers on the stack.This includes:

   ```
   CR: Control Register // mrc    p15, 0, r0, c1, c0, 0
   DACR: Domain Access Control Register // mrc   p15, 0, r1, c3, c0, 0
   FCSE: Fast Context Switch Register // mrc   p15, 0, r2, c13, c0, 0
   TTBR: Translation Table Base Register // mrc   p15, 0, r3, c2, c0, 0
   Context ID Register // mrc   p15, 0, r4, c13, c0, 1
   Extra Features Register // mrc   p15, 1, r5, c15, c1, 0
   Stack Pointer
   ```

9. Save the Physical address of Stack Pointer on the SRAM (within the resume code).
10. Flush all caches.

    ```
    DCache clean & invalidate // mcr   p15, 0, r0, c7, c14, 0
    flush BTAC/BTB // mcr   p15, 0, r0, c7, c5, 6
    I-BTB cache invalidate // mcr   p15, 0, r0, c7, c5, 0
    Clean L2 cache // mcr   p15, 1, r0, c7, c11, 0
    Drain write buffer // mcr   p15, 0, r0, c7, c10, 4
    ```

11. If the Deep-idle flag is enabled (Deep-Idle selected instead of eBook), set the DDR in SR mode. This is done through setting the SR request bit and polling the acknowledge bit until it is set. This will guarantee that the DDR is in SR mode before proceeding. **This step is skipped in eBook mode.**

    ```
    McSleepReq 0xD8000[18] = 0x1 // enter SR mode
    Verify that McSleepReqAck is set 0xD8000[19]?= 0x1 // Poll until this bit is set.
    ```

12. Disable the power on the onboard DDR terminations. The purpose of this step is to save the power dissipated on DDR termination since the DRAM interface is not used during Deep-Idle (DDR in SR mode). This is done through setting the level of the GPIO controlling the terminations power to LOW. The address of the GPIO register and the mask of the bit (that is the GPIO pin) controlling the terminations power is written at offset 0x0 and 0x4 in the SRAM during initialization. **This step is also skipped in eBook mode.**

    ```
    reg_addr = SRAM[0x0] // Read 32bit address written in SRAM offset 0x0
    gpio_mask = SRAM[0x4] // Read 32bit mask written in SRAM offset 0x4
    reg_val = [reg_addr] // read value of gpio out register
    reg_val &= ~gpio_mask // clear the bit according to the mask
    [reg_addr] = reg_val // write the modified value
    ```

13. Initiate the Deep-Idle/eBook flow.

    ```
    DeepIdlePwrEn 0xD8000[3] = 0x1 // Enable Deep-Idle/eBook on next entry to WFI
    ```

14. Relocate the internal registers address decoding window to its hardware default. This is necessary since the BootROM will access the internal registers space according to the hardware default configurations after resuming.

    ```
    Internal Registers Base Address 0x20080[31:0] = 0xD0000000 // h/w default
    ```

15. Enter WFI mode. This triggers the PMU to start it Deep-Idle/eBook flow.

    ```
    Enter WFI // mcr   p15, 0, r0, c7, c0, 0
    ```

18.3.3 PMU Entry and Exit

The following steps are performed by the PMU:

1. PMU powers down the CPU domain.
2. PMU waits for on of the configured wake-up events.
3. PMU powers up the CPU power domain.
4. PMU releases the CPU reset.

18.3.4 BootROM Exit

The following flow is performed by the CPU upon exiting reset.

1. The BootROM checks if the CPU is resuming from Deep-Idle/eBook otherwise normal boot is performed.
2. Read the number of descriptors to execute.

   ```
   Descriptor Count 0xD8100[2:0] = 0x0 // no descriptors to be executed.
   ```

3. Jump to resume address indicated in the

   ```
   Resume Address 0xD8104[31:0] // resume physical address in SRAM (0xF000C340)
   ```

18.3.5 HAL Exit

1. Reconfigure the internal registers address decoding window to its original value used before entering Deep-Idle/eBook mode. This is necessary since the rest of the resume process accesses the internal registers space according to the original configuration.

   ```
   Internal Registers Base Address 0x20080[31:0] = 0xF1000000 // depends on system configuration
   ```

2. Invalidate all caches.

   ```
   Invalidate both caches and BTC // mcr   p15, 0, r0, c5, c7, 0
   TLB invalidate // mcr   p15, 0, r0, c8, c7, 0
   L2 cache invalidate // mcr   p15, 1, r0, c7, c7, 0
   Prefetch flush (ISB) // mcr   p15, 0, r0, c7, c5, 4
   ```

3. Enable the power on the onboard DDR terminations before accessing the DDR. This is done through setting the level of the GPIO controlling the terminations power to HI. The address of the GPIO register and the mask of the bit (that is the GPIO pin) controlling the terminations power is written at offset 0x0 and 0x4 in the SRAM during initialization. **This step is skipped in eBook mode.**

   ```
   reg_addr = SRAM[0x0] // Read 32bit address written in SRAM offset 0x0
   gpio_mask = SRAM[0x4] // Read 32bit mask written in SRAM offset 0x4
   reg_val = [reg_addr] // read value of gpio out register
   reg_val |= gpio_mask // set the bit according to the mask
   [reg_addr] = reg_val // write the modified value
   ```

4. Exit DDR SR mode. This is done through clearing the SR request bit and polling the acknowledge bit until is cleared. This will guarantee that the DDR is not in SR mode before proceeding. **This step is skipped in eBook mode.**

   ```
   McSleepReq 0xD8000[18] = 0x0 // exit SR mode
   Verify that McSleepReqAck is clear 0xD8000[19]?= 0x0 // Poll until this bit is cleared.
   ```

5. Read the physical address of the Stack Pointer (stored within the code in the SRAM).
6. Pop out all register saved on the stack.
7. Restore saved CP15 registers

   ```
   Extra Features Register // mcr   p15, 1, r5, c15, c1, 0
   Context ID Register // mcr   p15, 0, r4, c13, c0, 1
   TTBR: Translation Table Base Register // mcr   p15, 0, r3, c2, c0, 0
   FCSE: Fast Context Switch Register // mcr   p15, 0, r2, c13, c0, 0
   DACR: Domain Access Control Register // mcr   p15, 0, r1, c3, c0, 0
   ```

8. Re synchronize the DDR clocks.

   ```
   Phy_sync_en 0x800240[31] = 0x1 // synchronize dclk2x and dclk
   ```

9. Change the PMU SRAM entry in the page table (indicated in the TTBR) so that VA=PA. The original mapping is saved in one of the unused registers.
10. Change the page table base address entry in the page table (indicated in the TTBR) so that VA=PA. The original mapping is saved in one of the unused registers.
11. Align the code to a cache line.
12. Enable MMU through writing the saved value of the CR in CP15.

    ```
    CR: Control Register // mcr    p15, 0, r0, c1, c0, 0
    ```

13. Wait for MMU to be enabled.

    ```
    Prefetch Flush (ISB) // mcr   p15, 0, r0, c7, c5, 4
    ```

14. Jump to the Virtual domain. This is virtual address of the next instruction.
15. Restore the original page table mapping for the entry representing PMU SRAM.
16. Restore the original page table mapping for the entry representing page table base.

17. Invalidate TLB.

    ```
    TLB Invalidate // mcr   p15, 0, r0, c8, c7, 0
    ```

18. Restore the original Stack Pointer.
19. Restore all general purpose registers and jump to LR (r0-r12). This step jumps back to the DRAM.
20. Clear the Standby bit in the PMU Status Register other wise the BootROM will perform resume on the next reset.

    ```
    DeepIdlePwrSt 0xD8004[3] = 0x0
    ```

21. Unmask the IRQ and FIQ gating at the PMU level.

    ```
    MaskIRQ 0xD8000[24] = 0x0 // ungate IRQ signal.
    MaskFIQ 0xD8000[27] = 0x1 // ungate FIQ signal.
    ```

18.3.6 System Exit

1. Re-initialize the SP for all CPU modes (IRQ, FIQ, ABT, UND). These are the same values used at initialization.
2. Restore FPU context.
3. Restore iWMMX context.
4. Enable IRQ and FIQ at the CPSR level.

B Package Trace Length Appendix

Note The attached Package Trace Length is PRELIMINARY and SUBJECT TO CHANGE.

he 88AP510 package trace lengths are provided as an Excel file attachment.

To open the attached Excel package trace length file, double-click the pin icons below:

88AP510 Package Trace Length

Note File attachments are only supported by Adobe Reader 6.0 and above.

To download the latest version of free Adobe Reader go to http://www.adobe.com.

C LCD Clock Tree Architecture Appendix

This appendix provides information about the 88AP510 LCD clock tree architecture and its features.

C.1 LCD and Display Controller Block

The 88AP510 incorporates two LCD controllers. Both LCD controllers drive their outputs to the LCD and Display Controller (LCD/DCON) unit. This unit drives the digital parallel, 24-bit RGB interface (Digital Video Output or DVO) and the Digital to Analog Converter (DAC) analog RGB interface (VGA). The analog RGB is generated by the Digital to Analog Converter (DAC).

Figure 90: 88AP510 LCD and Display Controller Block and Interfaces

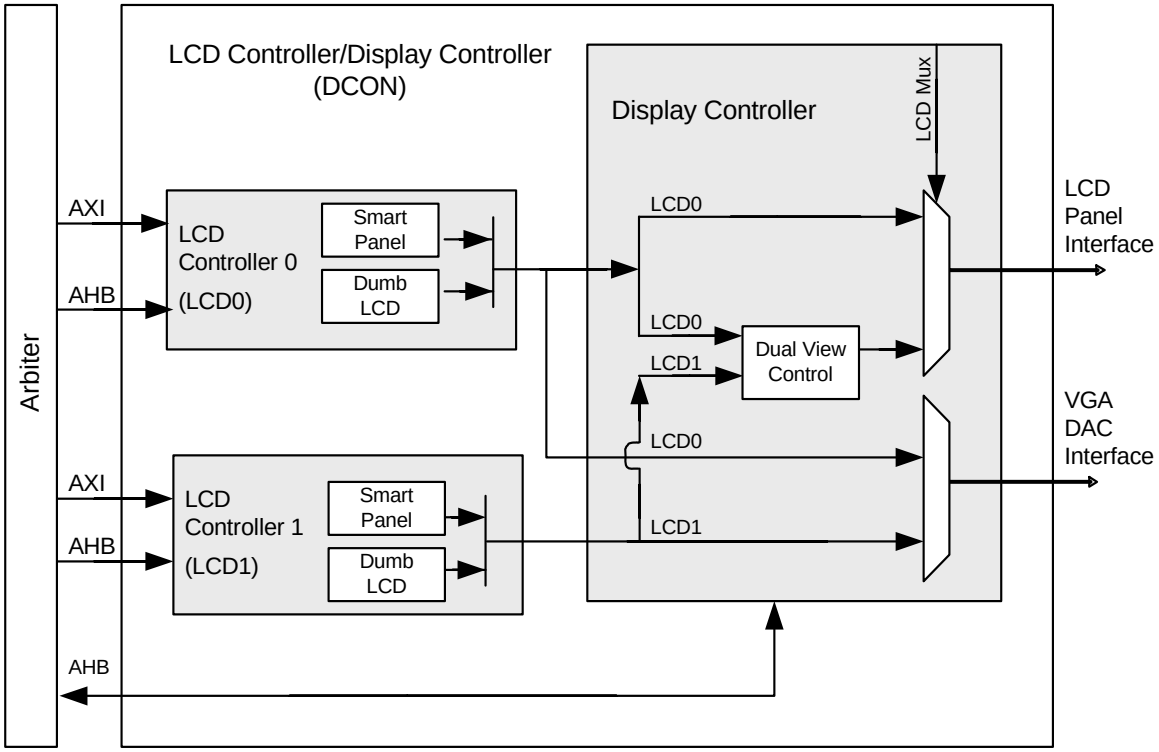

C.2 Maximum Pixel Clock Frequencies

The DAC internal pixel clock can run up to 200 MHz. This speed allows for resolutions up to 1920 x 1080 pixels @ 60 Hz (that is, 1080p).

Generally, the digital RGB interface pixel clock can run up to 100 MHz when using the VDDO_LCD with a 3.3V power supply voltage. This configuration supports resolutions up to 720p, or 1024 X 768 pixels @ 85 Hz.

Under some condition it is possible to run the pixel clock up to 150 MHz using 3.3V, while allowing for the following constraints:

- The maximum trace length is five inches.
- There can only be one device on the LCD bus that has a maximum pad capacity of 3 pF.
- A serial termination of 20 ohm must be placed on the pixel clock.

If the VDDO_LCD is using a 1.8V power supply voltage, the LCD pixel clock can run up to 160 MHz. This allows for a 1080p display.

C.3 Clock Tree

Both LCD displays are fed from a single reference PLL, running at 2 GHz, located in the SoC core.

To generate the correct pixel clock according to the connected display resolution and refresh rate, there are two stages of the clock divider.

The first stage divides the 2 GHz PLL using a 2:N divider.

The output of the first stage is then driven to both LCD controllers.

Each LCD controller can divide the clock input using a 1:N clock divider.

Figure 91 shows the stages of the clock divider.

Figure 91: LCD Clock Divider Stages

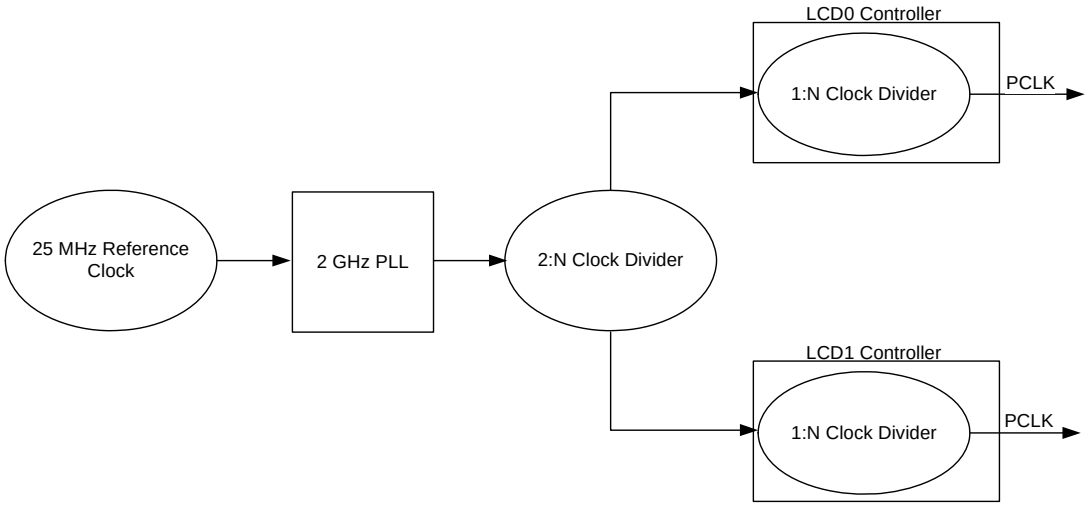

C.4 Single Display with Internal PLL

When using a single display, the 2 GHz PLL can generate most of the standard resolutions, as defined by the VESA standard Display Monitor Timing (DMT) or Generalized Timing Formula (GTF)

The VESA specification defines 0.5% as the maximum allowed clock error. However, empirically most displays can work with up to 1% maximum jitter.

Table 46 lists the VESA GTF resolutions supported by the 88AP510 using the internal PLL.

Table 46: VESA GTF Computability Summary Using 200 MHz Internal PLL

Number of Pixels	Resolution\Refresh (Hz)					
	60	70/72	75	85	100	120
640 x 480 pixels	Y	Y	Y	Y	Y	Y
800 x 600 pixels	Y	Y	Y	Y	R	R
1024 x 768 pixels	R	Y	Y	R	R	R
1280 x 1024 pixels	R	Y	Y	Y	ES	ES
1600 x 1024 pixels	N	R	ES	ES	ES	N

- N = Not supported
- Y = The timing meets the VESA specification (within 0.5%)
- R = The timing is within 1% of the VESA specification
- ES = The timing exceeds the device specification (160 MHz)

Table 47 For VESA CEA-861E

Table 47: CEA-861E Compatibility Summary (HDTV)

Number of Pixels	Resolution\ Refresh 60 (Hz)
1280 x 720	74.25
1920 x 1080	148.5

C.5 Dual Display with Internal PLL as Reference

If two displays are used in parallel, it is more difficult to accurately set the pixel clock.

The first 2:N division must be as high as possible, to meet as many frequencies as possible. The maximum frequency generated by the first clock divider is 400 MHz.

For example, if the LCD is driving 1280 x 1024 pixels @ 60 Hz display, the required pixel clock for such a resolution is 108 MHz. This pixel clock can be generated by dividing the 2 GHz PLL by a 2:37 ratio. The result is 108.1 MHz. There is no other option to meet this pixel clock using the two stage division of clocks.

The second 1:N LCD divider will have to be configured to a 1:1 ratio.

If another display is added to the other port, the pixel clock of the second display is generated from 108.1 MHz divided by 1:N. A configuration like this cannot meet most of the display resolutions.

C.6 Enabling Two Display Support (Applicable to the 88AP510-A0)

The 88AP510-A0 device revision has two pins that replace pins in the 88AP510-X0 device revision:

- LCD_EXT_CLK[0] replaces VGA_E
- LCD_EXT_CLK[1] replaces VGA_CLK

With LCD_EXT_CLK[0] and LCD_EXT_CLK[1], it is possible to configure the device so these pins are the source clock of each one of the LCD controllers, see Figure 92.

Marvell recommends using only LCD_EXT_CLK[0] for one display, while using the internal PLL for the secondary display. Connecting LCD_EXT_CLK[1] to GND reduces the noise and jitter on the external pixel clock.

If an external fixed clock is used for the dumb LCD, the second LCD controller is not be limited by the clock division of the other controller. The second LCD controller can generate any clock frequency, see Single Display with Internal PLL.

Another option is to use external clock generator such as the IDT5V49EE503. There should be no limitation of any pixel clock requirement with such a device, see Figure 93.

Finally, it is possible to use an external LCD controller that incorporates a PLL, and can scale the picture, such as the KG2.

Figure 92: 88AP510-A0 LCD Clock Topology

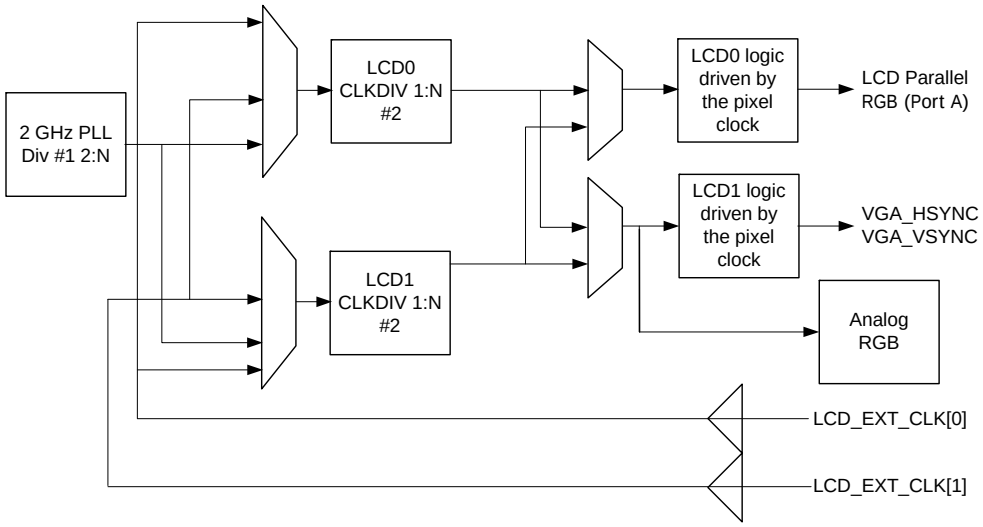

Figure 93: External Clock Generator

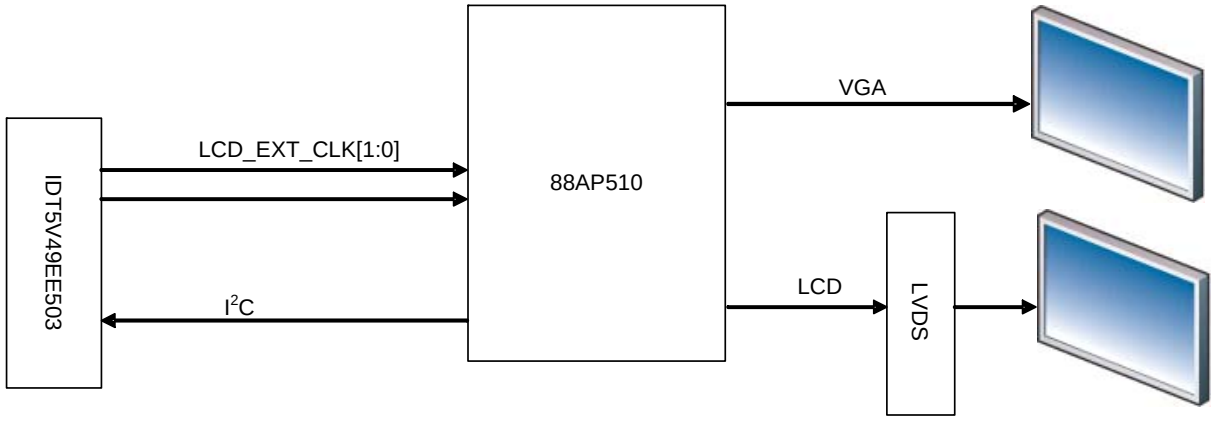

Marvell Semiconductor, Inc.
5488 Marvell Lane
Santa Clara, CA 95054, USA

Tel: 1.408.222.2500
Fax: 1.408.988.8279

www.marvell.com

Marvell. Moving Forward Faster