



Socket AM2 Motherboard Design Guide

Publication # 33165 Revision: 1.11 Issue Date: February 2007
--

© 2005—2007 Advanced Micro Devices, Inc. All rights reserved.

The contents of this document are provided in connection with Advanced Micro Devices, Inc. (“AMD”) products. AMD makes no representations or warranties with respect to the accuracy or completeness of the contents of this publication and reserves the right to make changes to specifications and product descriptions at any time without notice. The information contained herein may be of a preliminary or advance nature and is subject to change without notice. No license, whether express, implied, arising by estoppel or otherwise, to any intellectual property rights is granted by this publication. Except as set forth in AMD’s Standard Terms and Conditions of Sale, AMD assumes no liability whatsoever, and disclaims any express or implied warranty, relating to its products including, but not limited to, the implied warranty of merchantability, fitness for a particular purpose, or infringement of any intellectual property right.

AMD’s products are not designed, intended, authorized or warranted for use as components in systems intended for surgical implant into the body, or in other applications intended to support or sustain life, or in any other application in which the failure of AMD’s product could create a situation where personal injury, death, or severe property or environmental damage may occur. AMD reserves the right to discontinue or make changes to its products at any time without notice.

Trademarks

AMD, the AMD Arrow logo, AMD Athlon, AMD Opteron, and combinations thereof, and AMD-8132, AMD-8151 are trademarks of Advanced Micro Devices, Inc.

HyperTransport is a licensed trademark of the HyperTransport Technology Consortium.

PCIe is a trademark, and PCI Express and PCI-X are registered trademarks of the PCI-Special Interest Group (PCI-SIG).

Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.

Contents

Revision History	15
Chapter 1 System Overview	19
Chapter 2 General Motherboard Layout Guidelines	21
2.1 Impedance	21
2.2 Crosstalk	22
2.3 General Power Supply Guidelines	22
2.3.1 Power Supply Type Classification	22
2.3.2 Plane Splits	23
Chapter 3 Socket AM2 Processor Pinout and Package	25
3.1 Pinout Assignment	25
3.2 Pin List	25
3.3 Package Information	25
Chapter 4 Schematic and Layout Guidelines for Input Clock and HyperTransport™ Link	27
4.1 Processor System Clock Layout Recommendations	27
4.1.1 Processor Clocks	27
4.1.1.1 Termination	27
4.1.1.2 Layout Requirements	28
4.2 HyperTransport Technology Design Guidelines	29
4.2.1 HyperTransport Technology Overview	29
4.2.1.1 Signal Groups	30
4.2.1.2 Frequencies	31
4.2.1.3 Electrical Specifications	31
4.2.1.4 HyperTransport Link Configuration	32
4.2.2 HyperTransport Technology Routing Guidelines	33
4.2.2.1 General Routing Rules	33
4.2.2.2 HyperTransport Technology Motherboard Impedance	33
4.2.2.3 Microstrip Signal Trace Width and Separation	34
4.2.2.4 HyperTransport Technology Signal Length Matching	34
4.2.2.5 Trace Length Mismatch Control	42
4.2.2.6 HyperTransport Trace Referencing	43

4.2.2.7	Changes in HyperTransport Trace Referencing	44
4.2.2.8	Layer Changes	45
4.2.2.9	Miscellaneous Signals	45
4.3	HyperTransport Technology and Processor Breakout	46
4.4	Recommended HyperTransport Technology Logic Analyzer Connection	48
Chapter 5	Socket AM2 Design Guidelines for Two DDR2 SO-DIMMs	49
5.1	DDR2 SDRAM Introduction	49
5.1.1	Signal Names and Descriptions	49
5.1.2	Layer Assignments	52
5.1.3	DDR2 Supply Voltages	53
5.2	Overview of DDR2 SDRAM Interface	54
5.2.1	Interface Overview and Block Diagram	54
5.2.2	Valid Configurations – Two Unbuffered SO-DIMMs	55
5.2.3	Frequency	55
5.3	Unbuffered Two SO-DIMM DDR2 System Layout Guidelines	55
5.3.1	Address/Command Termination and Layout Guidelines	57
5.3.2	Data Group Connection and Termination Guidelines	60
5.3.3	Clock Termination and Layout Guidelines	63
5.4	SO-DIMM Routing with Five-Mil Traces	67
5.5	Basic Idea for SO-DIMM DDR2 Memory Routing	68
5.5.1	Routing Procedure	68
5.5.2	Routing and Length-Matching Rules	68
5.5.2.1	CLK_H and CLK_L Signals	68
5.5.2.2	Data, DQS, DM Signals	69
5.5.2.3	ADD, BANK, RAS, CAS, WE, CS, CKE, ODT	69
Chapter 6	Design Guidelines for Four Unbuffered DDR2 DIMMs	71
6.1	DDR2 SDRAM	71
6.1.1	Signal Names and Descriptions	71
6.1.2	Four Layer Assignments	72
6.1.3	Six Layer Assignments	75
6.1.3.1	Stripline Only Routing	75
6.1.3.2	Stripline and Microstrip Routing	76
6.2	Overview of DDR2 SDRAM Interface	77

6.2.1	Interface Overview and Block Diagram	77
6.2.2	Valid DIMM Configurations – Four Unbuffered DIMMs	79
6.2.3	Frequency	79
6.3	System Configuration Notes	80
6.4	Unbuffered Four-DIMM DDR2 Layout Guidelines	80
6.4.1	Address/Command Layout Guidelines	83
6.4.1.1	Address/Command and Control Termination	83
6.4.1.2	Address/Command Group Layout Guidelines	83
6.4.1.3	CKE Guidelines	85
6.4.1.4	CS and ODT Guidelines	85
6.4.1.5	Address/Command and Control Trace Width and Separation	86
6.4.2	Data Group Guidelines	87
6.4.2.1	Data Group Termination	89
6.4.2.2	Data Bus Trace Width and Separation	89
6.4.3	Clock Layout Guidelines	91
6.4.3.1	DIMM Clock Assignment	92
6.4.3.2	Clock Group Termination	92
6.4.3.3	Clock Group Trace Width and Separation	93
6.5	Basic Idea for Unbuffered DDR2 Memory Routing	93
6.5.1	Routing Procedure	93
6.5.2	Routing Rules	94
6.5.2.1	CLK_H and CLK_L Signals	94
6.5.2.2	Data, DQS, Check Signals	94
6.5.2.3	ADD, BANK, CS, CKE, RAS, CAS, WE, ODT Signals	94
6.5.3	Unbuffered 4-DIMM DDR2 Routing Example	96
Chapter 7	BTX Form Factor Design Guidelines for Four Unbuffered DDR2 DIMMs	103
7.1	DDR2 SDRAM	103
7.1.1	DDR Routing on BTX Form Factor	103
7.2	Unbuffered Four-DIMM DDR2 Layout Guidelines	104
7.2.1	Address/Command Group Layout Guidelines	107
7.2.1.1	CKE Guidelines	109
7.2.1.2	CS and ODT Guidelines	110
7.2.1.3	Address/Command and Control Trace Width and Separation	111
7.2.2	Data Group Guidelines	111

7.2.2.1	Data Bus Trace Width and Separation	114
7.2.3	Clock Layout Guidelines	115
7.2.3.1	Clock Group Trace Width and Separation	117
7.3	Basic Idea for Unbuffered DDR2 Memory Routing	117
7.3.1	Routing Procedure	117
7.3.2	Routing Rules	118
7.3.2.1	CLK_H and CLK_L Signals	118
7.3.2.2	Data, DQS, Check Signals	118
7.3.2.3	ADD, BANK, CS, CKE, RAS, CAS, WE, and ODT Signals	118
7.3.3	Unbuffered 4-DIMM DDR2 Routing Examples on a BTX Board	120
Chapter 8	Design Guidelines for Miscellaneous Signals	127
8.1	CLKIN_H and CLKIN_L Differential Input	130
8.2	Auto-Compensation	130
8.2.1	HyperTransport I/O Compensation	130
8.2.2	Memory I/O Compensation	130
8.3	Hardware Debug Tool	131
8.3.1	Single-Processor HDT	131
8.3.2	TMS	133
8.3.3	TCK	133
8.3.4	TRST_L	133
8.3.5	TDI	133
8.3.6	TDO	133
8.3.7	DBREQ_L	133
8.3.8	DBRDY	133
8.3.9	RESET_L	133
8.4	Voltage Regulator Signals	134
8.4.1	VID [5:0]	134
8.4.2	PSI_L	134
8.4.3	VDD_FB_H and VDD_FB_L	134
8.4.4	VDDIO_FB_H AND VDDIO_FB_L	134
8.4.5	VTT_SENSE	135
8.4.6	CPU_PRESENT_L	135
8.5	Global Signals	135

8.5.1	PWROK	135
8.5.2	RESET_L	136
8.5.3	LDTSTOP_L	136
8.6	Thermal-Related Miscellaneous Signals	136
8.6.1	THERMTRIP_L	136
8.6.2	PROCHOT_L	137
8.6.3	SIC and SID	138
8.6.4	THERMDA and THERMDC	138
8.6.4.1	Temperature Sensor Layout Recommendations	139
8.7	No Connects	139
Chapter 9	Socket AM2 Processor Power Requirements	141
9.1	High-Frequency Motherboard Decoupling	141
9.2	Power Generation and Distribution Guidelines	143
9.2.1	Socket AM2 Processor Power and Ground Overview	144
9.2.2	VDD Power Requirements	144
9.2.3	VDDIO, VTT, and M_VREF Routing and Decoupling	149
9.2.3.1	VDDIO Decoupling	149
9.2.3.2	VTT Power Delivery and Decoupling	150
9.2.3.3	M_VREF Decoupling and Routing Guidelines	154
9.2.3.4	VDDSPD	154
9.2.4	VDDA Power Requirements	155
9.2.5	VLDT Power Requirements	156
Chapter 10	EMI Design Guidelines	159
10.1	I/O Signal Partitioning and Separation	159
10.2	Decoupling, Bypass, Stitching, and Filtering Capacitors	159
10.2.1	Voltage Plane Decoupling Capacitors	159
10.2.2	Voltage Filtering Requirements	160
10.2.3	High-Frequency Bypass and Stitching Capacitors	160
10.2.4	DDR VDDIO to VSS Stitching Capacitor Requirements	161
10.3	Motherboard Grounding	161
10.3.1	Motherboard-to-Chassis Grounding	161
10.3.2	I/O Filtering and Shield Grounding	161
10.3.3	Processor Heatsink Grounding	162

10.4	Clocks and EMI	162
10.4.1	Clock Circuit Placement and Wiring	162
10.4.2	Clock Signal Termination	162
10.4.3	Spread Spectrum Clocking	163
10.4.4	Special Considerations for Clock Generator Modules	163
10.4.5	Disable Unused Clock Outputs	164
Chapter 11	Thermal and Mechanical Motherboard Requirements	165
11.1	Socket AM2 ATX Processor Component Keepout and Height Restrictions	165
11.2	Socket AM2 BTX Processor Component Keepout and Height Restrictions	172

List of Figures

Figure 1.	Block Diagram Example—Socket AM2 Processor-Based System	20
Figure 2.	Trace Crossing Plane Split with Stitching Capacitors as High-Frequency Return Path	24
Figure 3.	Motherboard Clock Termination	28
Figure 4.	CLKIN Board Stackup	29
Figure 5.	HyperTransport™ Technology Block Diagram	31
Figure 6.	HyperTransport Technology Differential Pair	32
Figure 7.	Microstrip-Board Stackup.	34
Figure 8.	Compensation Points	35
Figure 9.	Case 1: True is Longer	36
Figure 10.	Case 2: Complement is Longer.	37
Figure 11.	Case 3: True and Complement are the Same Length	37
Figure 12.	Added Routing Length to the Breakout	41
Figure 13.	VSS Referenced Layout	43
Figure 14.	Changing Reference Layout	44
Figure 15.	Crossing Reference Plane Splits.	45
Figure 16.	16 x 16-Bit Link Topside Breakout for the Socket AM2 Processor.	46
Figure 17.	16 x 16-Bit Link Bottomside Breakout for the Socket AM2 Processor.	47
Figure 18.	Six-Layer Board Stackup Overview	52
Figure 19.	Five-Mil Board Stackup	53
Figure 20.	Block Diagram for Socket AM2 Mobile System with Two Unbuffered SO-DIMM Modules	54
Figure 21.	Undesirable Net Routing	56
Figure 22.	Definition of Lead-in Length for Channel A and B Routing	57
Figure 23.	DDR2 Address/Command Routing	59
Figure 24.	DQS Routing and Termination	61
Figure 25.	Data and DM Routing and Termination	62
Figure 26.	Clock Routing and Termination	64
Figure 27.	Clock Routing and Termination	65
Figure 28.	Board Stackup Overview	73

Figure 29.	Board Stackup Parameters for a Four-Layer Board	74
Figure 30.	Six-Layer Board Stackup Overview—Stripline Only	76
Figure 31.	Six-Layer Board Stackup Overview—All Layers Used.	77
Figure 32.	Unbuffered 4-DIMM Block Diagram.	78
Figure 33.	Unbuffered 2-DIMM Block Diagram.	79
Figure 34.	Unbuffered 4-DIMM General Layout Picture	80
Figure 35.	Undesirable Net Routing	81
Figure 36.	Lead-In Length for Memory Routing	82
Figure 37.	Channel A Command and Address Routing.	84
Figure 38.	Channel B Command and Address Routing.	84
Figure 39.	CKE Routing and Termination.	85
Figure 40.	CS and ODT Routing and Termination	86
Figure 41.	Data Group Routing and Termination	88
Figure 42.	Routing for DQS Pair	89
Figure 43.	Clock Termination and Routing	91
Figure 44.	Clock Pair Assignment	92
Figure 45.	General DDR2 Component Placement and Routing for the Socket AM2 Processor.	96
Figure 46.	Example of Three Nets Routing Between the Two Processor Pads.	97
Figure 47.	Example of a Socket AM2 Processor Breakout—Topside.	98
Figure 48.	Example of a Socket AM2 Processor Breakout—Bottomside	99
Figure 49.	Socket AM2 Processor Four DIMM Routing Example—Topside.	100
Figure 50.	Socket AM2 Processor Four DIMM Routing Example—Bottomside.	101
Figure 51.	Unbuffered 4-DIMM General Layout Picture for BTX Boards.	104
Figure 52.	Undesirable Net Routing	105
Figure 53.	Lead-In Length for Memory Routing	106
Figure 54.	Channel A Command and Address Routing.	107
Figure 55.	Channel B Command and Address Routing.	108
Figure 56.	CKE Routing and Termination.	109
Figure 57.	CS and ODT Routing and Termination	110
Figure 58.	Data Group Routing and Termination	112
Figure 59.	Routing for DQS Pair	113
Figure 60.	Clock Termination and Routing	116
Figure 61.	General DDR2 Component Placement and Routing for a BTX Form Factor	120

Figure 62.	Example of Three Nets Routing Between the Two Processor Pads	121
Figure 63.	Example of a Socket AM2 Processor Breakout on a BTX Board—Topside	122
Figure 64.	Example of a Socket AM2 Processor Breakout on a BTX Board—Bottomside	123
Figure 65.	Socket AM2 Processor Four DIMM Routing Example on a BTX Board—Topside . .	124
Figure 66.	Socket AM2 Processor Four DIMM Routing Example on a BTX Board— Bottomside	125
Figure 67.	HDT Header	132
Figure 68.	Thermtrip Circuit	137
Figure 69.	Example PROCHOT_L Schematic	138
Figure 70.	High-Frequency Decoupling Routing Comparison	142
Figure 71.	General View of the Processor Power Plane Cuts	143
Figure 72.	Example of VDD Power Fill on the Top Layer	145
Figure 73.	VDD Decoupling Placement With Backplate Windows—Bottom Signal Layer . . .	147
Figure 74.	Decoupling Placement Strategy on Bottom Signal Layer—Without Backplate	148
Figure 75.	VDDIO Decoupling Placement Between Processor and DIMMs	149
Figure 76.	VTT Delivery on Single Side	152
Figure 77.	VTT Decoupling Placement Between Processor and DIMMs	153
Figure 78.	M_VREF Decoupling Placement	154
Figure 79.	VDDA Filtering Circuit	155
Figure 80.	PCB Mounting Hole	161
Figure 81.	Top Layer Beneath the Clock Generator Module.	163
Figure 82.	Inner Layer Beneath the Clock Generator Module.	164

List of Tables

Table 1.	Power Plane Type Activity Over Various Operating States	23
Table 2.	Clock Termination Values	28
Table 3.	CLKIN Board Routing Parameters	29
Table 4.	Microstrip-Board Routing Parameters	34
Table 5.	True-to-Complement Skew for HyperTransport™ Signals	38
Table 6.	Differential Pair Matching Rules	42
Table 7.	Signal Group Length-Matching Rules	42
Table 8.	DDR2 Signal Descriptions	51
Table 9.	Termination Values for Unbuffered SO-DIMM Address/Command Signals	58
Table 10.	Parameters for Address/Command and Control Trace Width and Space for a Board With Five-Mil Inner Dielectric	60
Table 11.	DDR2 Data Groups	60
Table 12.	Parameters for MA_DQS and MB_DQS Trace Width and Separation for a Board With Five-Mil Inner Dielectric	62
Table 13.	Parameters for Data and DM Trace Width and Separation for a Board With Five-Mil Inner Dielectric	63
Table 14.	CLK_H and CLK_L Termination	64
Table 15.	Parameters for CLK Trace Width and Separation for a Board With Five-Mil Inner Dielectric	66
Table 16.	DDR2 Signal Descriptions	72
Table 17.	Parameters for Stackup on a Four-Layer Board	75
Table 18.	Termination Values for Unbuffered DIMM Address/Command Signals	83
Table 19.	Parameters for Address/Command and Control Trace Width and Space	87
Table 20.	DDR2 Data Groups	88
Table 21.	Parameters for DQS Trace Width and Space	90
Table 22.	Parameters for Data and Check Trace Width and Space	90
Table 23.	CLK_H and CLK_L Termination	92
Table 24.	Parameters for CLK Trace Width and Space	93
Table 25.	Parameters for Address/Command and Control Trace Width and Space	111
Table 26.	DDR2 Data Groups	112
Table 27.	Parameters for DQS Trace Width and Space	114

Table 28.	Parameters for Data and Check Trace Width and Space	115
Table 29.	Parameters for CLK Trace Width and Space	117
Table 30.	Miscellaneous Signals Quick Reference	128
Table 31.	Compensation Resistor Routing Rules for HyperTransport Technology	130
Table 32.	Compensation Resistor Routing Rules for M_ZP and M_ZN	130
Table 33.	HDT Header Pin-Out	132
Table 34.	VDD Decoupling Recommendations	146
Table 35.	VDDIO Decoupling Recommendations	150
Table 36.	VTT Decoupling Recommendations.	151
Table 37.	M_VREF-to-Processor Decoupling Recommendations	154
Table 38.	VDDA Filter Recommendations.	156
Table 39.	VLDT Decoupling Recommendations	157

Revision History

Date	Revision	Description
February 2007	1.11	<p>Added the following:</p> <ul style="list-style-type: none"> • In Chapter 6, Figure 33 • In Chapter 10, Section 10.4.4 and Figures 81 and 82 <p>Revised the following:</p> <ul style="list-style-type: none"> • In Chapter 2, the Introduction • In Chapter 4, Introduction and Section 4.2.2.2 • In Chapter 5, Figure 23 • In Chapter 8, Section 8.6.4.1 • In Chapter 10, Sections 10.2.1, 10.2.4, 10.3.1, 10.3.3, 10.4.1, and Figure 80
June 2006	1.10	<p>Revised the following:</p> <ul style="list-style-type: none"> • In Chapter 6, Section 6.3 • In Chapter 8, Sections 8.3
April 2006	1.09	<p>Added the following:</p> <ul style="list-style-type: none"> • Sections 6.1.3, 6.1.3.1, 6.1.3.2, Figure 30 and 31 <p>Revised the following:</p> <ul style="list-style-type: none"> • In Chapter 1, the Introduction • In Chapter 4, Section 4.2.1.1, 4.2.2.1, 4.2.2.7, 4.2.2.9, and Figure 15 • In Chapter 5, Sections 5.3.3 and 5.4 • In Chapter 6, the Introduction, Sections 6.1.2, 6.4.1.5, 6.4.2.2, 6.4.3.2, 6.4.3.3 • In Chapter 7, Sections 7.2.1.3, 7.2.2.1, and 7.2.3.1 • In Chapter 8, Table 30, Sections 8.2.2, 8.3, 8.4.3, 8.4.4, 8.6.2, 8.6.3, and 8.6.4.1 • In Chapter 9, the Introduction and Section 9.2.5
March 2006	1.08	<p>Revised the following:</p> <ul style="list-style-type: none"> • In Chapter 5, Section 5.1.2, 5.2.2, 5.2.3, Tables 10, 12, 13, 15, and Section 5.5.2.2 • In Chapter 6, Section 6.1.2, 6.2.2, 6.2.3, Tables 19, 21, 22, 24, Sections 6.5.2.1 and 6.5.2.2 • In Chapter 7, Tables 25, 27—29, Sections 7.3.2.1 and 7.3.2.2 • In Chapter 8, Sections 8.2.1, 8.6.2 and 8.6.3 • In Chapter 9, Table 37, Sections 9.2.3.4, 9.2.4, Figure 78, and Section 9.25 • In Chapter 11, Section 11.1

Date	Revision	Description
January 2006	1.07	Revised socket name to AMD Socket AM2
November 2005	1.06	Revised the following: <ul style="list-style-type: none"> • In Chapter 11, BTX Keepout Drawings
September 2005	1.05	Revised the following: <ul style="list-style-type: none"> • In Chapter 4, Sections 4.1.1, 4.2.1.2, and 4.2.2.4.1 • In Chapter 5, Table 10, Table 12, Table 13, and Table 15 • In Chapter 6, Table 19, Table 21, and Table 24 • In Chapter 7, Table 25, Table 27, Table 28, and Table 29 • In Chapter 8, Table 30 and Sections 8.3.9, 8.4.1, 8.4.5, 8.6.1, 8.6.2, and 8.6.3 • In Chapter 9, Table 34, Table 35, Table 34, Table 36, Table 39, Sections 9.2.4 and 9.2.3.2 • In Chapter 10, Sections 10.2.3 and 10.2.4 • In Chapter 11, “Socket AM2 BTX Processor Component Keepout and Height Restrictions” in Section 11.2 Added the following: <ul style="list-style-type: none"> • In Chapter 7, Figure 51, Figure 52, Figure 53, Figure 61, Figure 63, Figure 64, Figure 65, and Figure 66
September 2005	1.04	Revised the following: <ul style="list-style-type: none"> • Tables 19, 21, 22 and 24 in Chapter 6 • Tables 25, 27, 28, and 29 in Chapter 7 • AM2 ATX processor component keepout and height restriction images in Chapter 11
July 2005	1.03	Removed change bars

Date	Revision	Description
July 2005	1.02	Added: <ul style="list-style-type: none"> • Chapter 7 “BTX Form Factor Design Guidelines for Four Unbuffered DDR2 DIMMs” • THERMDA and THERMDC to Table 30 • Section 8.6.4 “THERMDA and THERMDC” • Section 11.2 “Socket AM2 BTX Processor Component Keepout and Height Restrictions” Revised: <ul style="list-style-type: none"> • Section 2.3.2 “Plane Splits” • Section 4.1.1.1 “Termination” • Section 4.2.2.1 “General Routing Rules” • Section 4.2.2.2 “HyperTransport™ Technology Motherboard Impedance” • Table 7 • Section 4.2.2.5 “Trace Length Mismatch Control” • Figure 15 • Section 5.1.2 “Layer Assignments” • Section 6.1.2 “Four Layer Assignments” • Section 8.4.5 “VTT_SENSE” • Section 9.2.3.2 “VTT Power Delivery and Decoupling” • Chapter 11.1 “Socket AM2 ATX Processor Component Keepout and Height Restrictions”
May 2005	1.01	Minor corrections made and removed change bars
May 2005	1.00	First NDA release

Chapter 1 System Overview

This chapter provides a description of the socket AM2 processor and a system block diagram example.

The socket AM2 processor is a superb, high-performance product for mobile and desktop systems. With HyperTransport™ technology bus and on-chip memory controller, the processor provides outstanding processing power, significant performance, and feature improvements.

Note: Do not put AMD-8111™, AMD-8131™, AMD-8132™, or AMD-8151™ chipsets on the same HyperTransport link as devices with MMIO BARs (PCI Express® bridges).

Figure 1 on page 20 shows a block diagram of a typical socket AM2 processor-based system with a generic I/O hub and PCI Express® Root Complex.

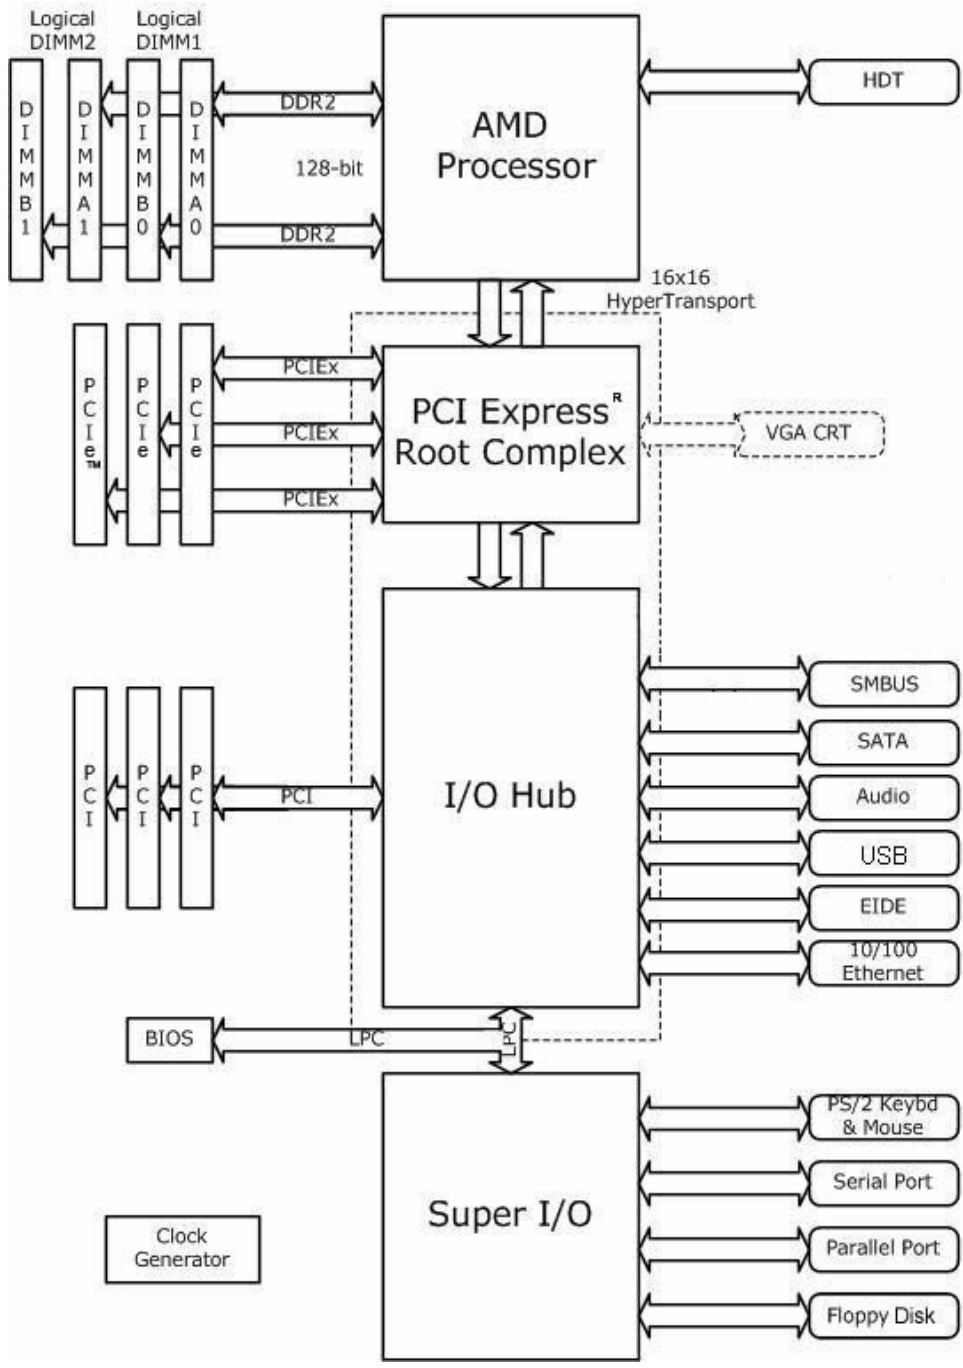


Figure 1. Block Diagram Example—Socket AM2 Processor-Based System

Chapter 2 General Motherboard Layout Guidelines

This chapter contains general guidelines that, when followed, help to ensure good signal quality. These guidelines include recommendations for lowering impedance, connecting power-plane pins to the motherboard, reducing crosstalk, and, when unavoidable, how best to cross plane splits.

A Socket AM2 platform should use routing rules that have been developed based on a motherboard's dielectric thickness, trace width, and trace spacing. The HyperTransport™ routing needs to achieve a $93\text{-}\Omega \pm 10\%$ differential impedance.

The DDR2 nets have different impedance targets, depending on their net class (Data, DQS, Address/Command and CLK), which is explained in Chapters 5 through 8.

2.1 Impedance

In a high-speed signaling environment, signal-trace impedances must be controlled in order to maintain good signal quality across the motherboard. Signal-trace impedance is a function of the following three factors:

- Motherboard stackup
- Dielectric constant of the PCB substrate
- Signal-trace width and thickness

The reliance of signal-trace impedance on these three factors demonstrates the importance of following the recommended stackup and routing rules outlined in this document.

The processor routing guidelines were developed with the aid of signal-integrity simulations. These simulations assumed controlled impedance motherboards. In general, wider signal traces have lower trace impedances.

For recommended stackup, see Sections 5.1.2 on page 52 and 6.1.2 on page 72.

2.2 Crosstalk

The following list of recommendations should be followed to help reduce the crosstalk on the motherboard:

- Do not allow high-speed signals to cross plane splits.
- Reference critical signals to ground planes.
- Do not cut ground planes, unless it is absolutely necessary.
- Reduce the length of signals that are routed parallel.
- Provide analog signals with guard shields or guard rings.
- Keep analog signals away from digital signals.

2.3 General Power Supply Guidelines

The following subsections provide an overview of the power supplies required for the processor. This section also recommends good design practices for connecting the power-supply pins to the motherboard, and it discusses how to handle power-plane splits.

For more information on processor power requirements, see Chapter 9.

2.3.1 Power Supply Type Classification

When designing a system, it helps the designer to have a method of classifying power signals that describes both *when* each particular power rail is active and *how* that signal is sourced. The AMD method uses abbreviations to assign a type to a signal or power plane. For example, a signal labeled VDDIO_RUN should be connected to the VDDIO power supply and exhibit the same behavior as defined under the RUN classification. There are four power types—ALW, DUAL, SUS, and RUN. These are defined below. Table 1 on page 23 summarizes which operating states activate the various power plane types.

ALW—(Always) Any power plane or line with the ALW classification maintains power through power states (S0–S5). The 5-V ALW supply line from the silver box power supply provides power for these lines. The only time power is removed from these signals is when the system is unplugged from the external power source. These signals do not have a battery backup.

DUAL—(Larger Capacity Always) Any power plane with the DUAL classification maintains power through power states (S0–S5). The DUAL classification refers to how these lines are powered. DUAL voltage planes are powered from the RUN outputs of the silver box power supply during the S0 and S1 states and are powered from the 5-V ALW output of the silver box power supply during the S3, S4, and S5 states. There is a DUAL power type, because during the S0 and S1 states the ALW output of

the silver box power supply cannot provide enough current to meet device power consumption requirements.

SUS—(Suspend) Any power plane with the SUS classification maintains power through power states (S0–S3). These signals are typically powered from DUAL supplies.

RUN—(Running) Any power plane with the RUN classification maintains power through power states (S0–S1). These signals are powered by the main outputs from the silver box power supply. The majority of system power is drawn from lines of this classification.

Table 1 summarizes the various operating states.

Table 1. Power Plane Type Activity Over Various Operating States

				Power Plane Type			
Global State	Description	System State	Description	ALW	DUAL	SUS	RUN
G0	Running	S0	Running	ON	ON	ON	ON
G1	Sleeping	S1	Suspend	ON	ON	ON	ON
		S3	Suspend to RAM	ON	ON	ON	OFF
		S4	Suspend to Disk	ON	ON	OFF	OFF
G2	Soft-Off	S5	OFF	ON	ON	OFF	OFF
G3	Mechanical-Off		OFF	OFF	OFF	OFF	OFF

2.3.2 Plane Splits

The high-frequency return path for any signal lies directly beneath the signal on the adjacent layer. Providing a solid plane underneath a signal greatly reduces problems with signal integrity and timing, because the plane provides a direct return path for that signal. If crossing a plane split is unavoidable, 0.01 μ F stitching capacitors should be used. When using stitching capacitors, place them as close as possible to the trace crossing the split, as shown in Figure 2 on page 24.




Figure 2. Trace Crossing Plane Split with Stitching Capacitors as High-Frequency Return Path

Chapter 3 Socket AM2 Processor Pinout and Package

For pinout and package information, see the following references.

3.1 Pinout Assignment

The socket AM2 processor's pinout allows all the signals to break out on two signal layers. The exact pin assignments are documented in the *Socket AM2 Processor Functional Data Sheet*, order# 31117.

3.2 Pin List

The *Socket AM2 Processor Functional Data Sheet*, order# 31117, should be referenced for the exact location of the pin assignments.

3.3 Package Information

The *Socket AM2 Processor Functional Data Sheet*, order# 31117, should be referenced for the packaging information.

Chapter 4 Schematic and Layout Guidelines for Input Clock and HyperTransport™ Link

This chapter contains specific termination, routing, and layout rules for socket AM2 processor clocks and the HyperTransport™ link.

The “processor breakout region” is defined for each pin, as the area on the board from the pin towards the package outline for a distance of up to 500 mils, where traces need to be routed around BGA pads and miscellaneous components.

4.1 Processor System Clock Layout Recommendations

This section describes the processor clock design requirements.

4.1.1 Processor Clocks

The processor bus interfaces—HyperTransport technology link and DDR2 memory—are both source-synchronous. Therefore, the system clock input to the processor is not skew-controlled and is used only to lock the PLL inside the processor. The CPU clock input frequency is 200 MHz and connects to a PLL inside the processor, which generates the core, DDR2, and HyperTransport technology clocks. The processor clock receiver is a differential, VDDIO-tolerant input. The following guidelines are for clock generator devices that follow the *Clock Generator Specification for AMD64 Processors*, order# 24707.

4.1.1.1 Termination

The clock inputs for the processor (CLKIN_H and CLKIN_L) should be driven with a 200-MHz differential clock source. It is necessary to provide DC isolation between the processor clock input and the clock generator. The isolation is accomplished with two AC coupling capacitors, as shown in Figure 3 on page 28.




Figure 3. Motherboard Clock Termination

The clock termination values are shown in Table 2.

Table 2. Clock Termination Values

Instance	Value	Units	Tolerance
R_{Series}	47.5	Ω	$\pm 1\%$
R_{Shunt1}	261	Ω	$\pm 1\%$
C_{Series}	3900	pF	$\pm 10\%$
R_{Shunt2}	169	Ω	$\pm 1\%$

The following rules apply to the layout of the clock termination:

- The trace connecting the clock generator to R_{Series} must not exceed 0.5 inches.
- R_{Shunt1} should be placed less than 100 mils from R_{Series} .
- Place R_{Shunt2} less than 600 mils from the processor.
- The trace connecting C_{Series} to the processor must not exceed 1250 mils.

4.1.1.2 Layout Requirements

Optimally, the layout of the processor clocks should be such that they always reference the ground layer (except for up to 500 mils in the processor breakout area) with no plane split crossings. If it is necessary to route the processor clocks over a plane split, minimize the split width that the signals cross. Cross the signals perpendicularly to the split. High-frequency capacitors (0.01 μ F X7R) should be used to AC-couple the planes across the split crossings (within 20 mils of where the signals cross the split).

CLKIN_H and CLKIN_L have the routing requirement of 100- Ω differential impedance. To achieve the target differential impedances, the layer stackup configuration in Figure 4 on page 29 is recommended for microstrip routing. If desired, the clocks may be routed as striplines with 100- Ω differential impedance target. The length of CLKIN_H and CLKIN_L should be minimized to ensure processor input requirements are met.

The mismatch tolerance for CLKIN_H and CLKIN_L is 35 mils.




Figure 4. CLKIN Board Stackup

Table 3 shows the CLKIN board routing parameters.

Table 3. CLKIN Board Routing Parameters

Type	H ₁ (mils)	Th ₁ (mils)	W _{SIG} (mils)	W _{DP} (mils)	W _{SPACE} ¹ (mils)	E _R
Microstrip	4.0	2.0	4.0	6.0	16.0	4.0 @ 1.0 GHz
Microstrip	4.5	2.0	5.0	6.0	20.0	4.0 @ 1.0 GHz

Note:
1. If the net is serpentine, then the space to itself should be 20 mils or more.

4.2 HyperTransport™ Technology Design Guidelines

This section describes the design requirements for the HyperTransport system link interface to the socket AM2 processor as well as the general requirements for connecting any HyperTransport technology device.

4.2.1 HyperTransport™ Technology Overview

The HyperTransport link is the channel for the processor to connect to the chipsets. The processor has a single 16 χ 16 link for connecting to a single 16 χ 16 tunnel device (e.g., AMD-8151™)

HyperTransport AGP3.0 graphics tunnel). The HyperTransport link on the processor is naturally compensated. For more information on HyperTransport natural compensation, see Section 4.2.2.4.1 on page 35.

The HyperTransport link is a source synchronous, fully differential, point-to-point interface. The source synchronous nature removes the difficulty of length-matching system clocks to all devices. The only timing factors to consider are the individual true-to-complement skew, clock-to-data skew, and group-to-group skew for each point-to-point link. For improved signal integrity, the termination resistors for each differential pair are embedded in the receiving device at the input terminals. This setup removes any termination requirements on the motherboard and improves signal integrity by removing stubs from the termination circuit.

4.2.1.1 Signal Groups

There are four signal groups that compose a HyperTransport link, and there are two copies of these groups required for each link (one transmit group and one receive group). The CAD, CLK, and CTL groups are differential and the miscellaneous group is single-ended. The CAD group may use an 8-bit or 16-bit wide bus, depending on the device connected to the processor.

- CLK—There is one differential clock signal per eight CAD pairs.
- CTL— On each HyperTransport link, there are two control signals for each direction. CTLIN_H[1:0] and CTLIN_L[1:0] are two differential pairs used as input CTL. CTLOUT_H[1:0] and CTLOUT_L[1:0] are two differential pairs used as output CTL. If the processor is connected to a device which does not support two control differential pairs, then CTLIN_H[1] should be pulled up to VLDT through a 51- Ω resistor and CTLIN_L[1] should be pulled low to VSS through a 51- Ω resistor.
- MISC (LDTSTOP_L, PWROK, and RESET_L)—These three signals must be provided to each device in the system.

Figure 5 on page 31 illustrates a generic hookup between a socket AM2 processor and a HyperTransport technology device.

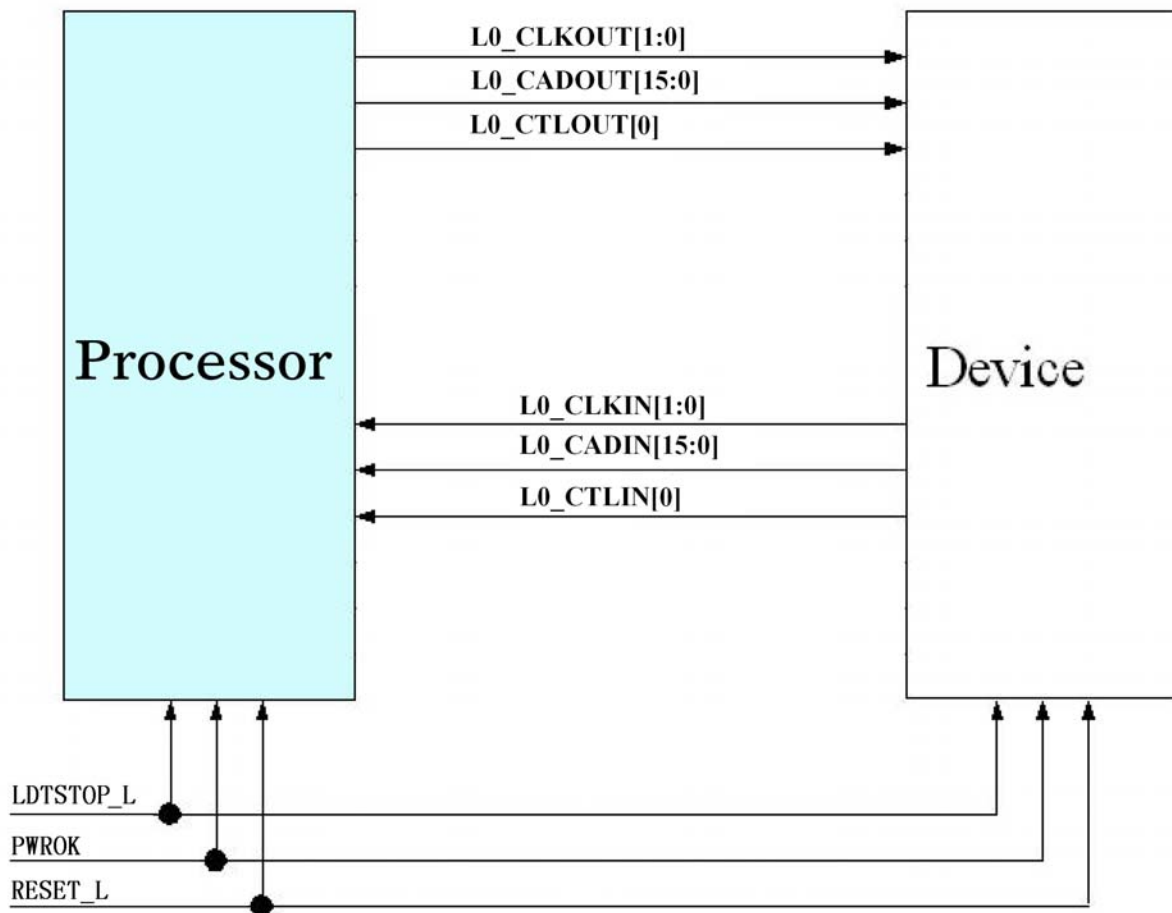


Figure 5. HyperTransport™ Technology Block Diagram

4.2.1.2 Frequencies

This design guide is valid for the frequencies supported in the *AMD NPT Family 0Fh Processor Electrical Data Sheet*, order # 31119, which can be referenced for a listed of supported frequencies.

4.2.1.3 Electrical Specifications

Each HyperTransport signal consists of a differential driver. The differential pair requires a balanced driver setup and a procedure to set the driver impedance for both PMOS and NMOS sides of the driver.




Figure 6. HyperTransport™ Technology Differential Pair

As seen on Figure 6, the HyperTransport technology electrical interface is a low swing differential signal from an impedance-controlled driver (R_{ON}) that drives a differential receiver with an on-die differential termination (R_{TT}) through a differential impedance trace (Z_{OD}). In Figure 6, both R_p and R_n have the same value as R_{ON} . R_{ON} acts as a source termination and R_{TT} acts as the load termination with the following relationships:

- $R_{ON} = Z_{OD} / 2$
- $R_{TT} = Z_{OD}$

All termination for the driver and receiver (R_{TT} and R_{ON}) is provided on-die. No termination is required on the motherboard. The I/O controller strength is actually compensating for variation in process, voltage, and temperature (PVT). The compensation of the receiver R_{TT} and driver impedance R_{ON} on the processor is performed through the HTREF0 and HTREF1 pins.

For detailed information regarding the connection of the HTREF0 and HTREF1 pins, refer to Section 8.2.1 on page 130.

4.2.1.4 HyperTransport™ Link Configuration

Below are system configuration options for a socket AM2 processor HyperTransport link.

- A 16×16 link connecting to either another processor or a HyperTransport device with a 16×16 HyperTransport link. In this configuration, all CAD and CLK nets are point-to-point, and the CTLIN[1] pins are not used. The CTLIN[1] pins should be terminated to logic high (true is pulled high and its complement is pulled low with $51\text{-}\Omega \pm 5\%$ resistors).
- A 16×16 link connecting to another device with an 8×8 HyperTransport link (the 16×16 link is only partially used). Each of the unused CLK and CAD inputs of the link must be pulled to a logic low, and the unused CTLIN[1] differential pair should be pulled to logic high. For each pullup or pulldown use a resistor with value of $51\ \Omega (\pm 5\%)$. For the CAD and CLK nets pull the true low

and pull the complement high. For the CTLIN[1] nets pull the true high and pull the complement low.

4.2.2 HyperTransport™ Technology Routing Guidelines

HyperTransport signaling uses low-voltage swing differential signaling with source synchronous clocking. The transmitted source synchronous clock is shifted by nominally one-half bit-time in order to facilitate data recovery in the receiver. Maintaining good signal quality and low skew requires using the following general approaches to implement the physical interconnect. These rules should be adhered to on all motherboards designed for the socket AM2 processor.

Group electrical “lengths” should be matched. This means all signals within a clock group should be routed nominally alike so that the input timing of the receiving device is the same for all signals.

4.2.2.1 General Routing Rules

The following list shows rules for general signal routing.

- General
 - CAD_L, CAD_H and CTL_L, CTL_H within a clock group should be treated identically. Clock groups are as follows:
 - CLK[0], CAD[7:0], and CTL[0]
 - CLK[1] and CAD[15:8]
 - The effective length of the differential pair is the average length of the pair of traces. The average is [True+Complement]/2.
 - The vias for differential pairs must be placed within 50 mils, center-to-center, of each other.
- For rates of 400 MT/s and 800 MT/s
 - The bus length must be greater than 1.0 inch and less than 24.0 inches.
 - A group layer change (e.g., in the middle of the bus) is discouraged but allowed to happen.
- For rates of 1.2 GT/s, 1.6 GT/s, and 2.0 GT/s and higher
 - The shortest net in the link must be greater than 1 inch, and the longest net in the link must be less than 12.0 inches for 5-mil traces and less than 8.0 inches for 4-mil traces.
 - All CAD/CTL/CLK signals within a clock group must route on the same signal layer.
 - A group layer change (e.g., in the middle of the bus) is not permitted.

4.2.2.2 HyperTransport™ Technology Motherboard Impedance

AMD recommends targeting a differential impedance range of $93 \Omega \pm 10\%$ for all HyperTransport differential pair traces (inclusive of manufacturing tolerances). Operation outside of this range of impedance is possible but appropriate simulations need to be completed to define appropriate routing

guidelines. The following stackup and trace geometries (refer to Table 4 and Figure 7) are an example of how to achieve a 93- Ω differential impedance. If desired, HyperTransport nets may be routed as symmetric or asymmetric stripline. In any case, the target for trace differential impedance should be $93 \Omega \pm 10\%$.

4.2.2.3 Microstrip Signal Trace Width and Separation

Table 4 contains the recommendations for routing the HyperTransport link on the top or bottom layer of the motherboard. H_1 is the height above the plane that the HyperTransport technology trace should be. Th_1 is the trace thickness. W_{SIG} is the trace width. W_{DP} is the differential trace separation. W_{SPAC} is the general trace separation.

In cases where a HyperTransport link is routed through a connector, special care must be taken in selecting the type of connector and in the implementation of the link routing. A 10-mil addition to W_{SPAC} may be needed to achieve a robust system.

Table 4. Microstrip-Board Routing Parameters

Type	H_1 (mils)	Th_1 (mils)	W_{SIG} (mils)	W_{DP} (mils)	W_{SPAC} (mils)	E_R
Microstrip	4.0	2.0	4.0	4.0	16.0 ¹	4.0 @ 1.0 GHz
Microstrip	4.5	2.0	5.0	5.0	15.0 ¹	4.0 @ 1.0 GHz

Note:
1. If the net is serpentine, then the space to itself should be 20 mils.




Figure 7. Microstrip-Board Stackup

4.2.2.4 HyperTransport™ Technology Signal Length Matching

Both timing uncertainty and signal integrity must be considered when routing the HyperTransport bus. Therefore, signals must be matched within a differential pair, and signals must be matched both within a group and across groups. CAD_L, CAD_H and CTL_L, CTL_H within the low clock group should be treated identically. The signal group length-matching rules are based upon propagation rates of 150 ps/in for microstrip and 180 ps/in for stripline. All the length-matching numbers are direct derivatives of the HyperTransport specification for PCB skew assumptions. Please refer to the *HyperTransport I/O Link Specification*, document# HTC20031125-0035-0003 Revision 2.00 from the HyperTransport Technology Consortium for more details.

The signal group length-matching rules are in Table 7 on page 42.

4.2.2.4.1 HyperTransport™ Natural Compensation

All HyperTransport™ signals for 1.27-mm pitch AMD packages are naturally compensated. A naturally compensated package has mismatched substrate routing for all the differential pairs that are oriented radial to the package core. This routing mismatch is approximately equal to the pin (or ball) pitch.

The HyperTransport interface is differential; therefore, the differential pair's true and complement signal routing must be matched as close as possible. Natural compensation allows known routing differences to exist (as in the mismatched substrate routing mentioned above). Instead of each net segment on the board being exactly matched, the known difference is corrected for elsewhere. In the case of natural compensation, both the board and the package cooperate to make the length matching easier.

When routing between two naturally compensated devices, the compensation point is a line parallel to the package edge, mid-way between the two pins (or balls) in the pair. Whenever using natural compensation, *it is critical that the different pieces use the same compensation point at the interfaces.* Figure 8 illustrates this requirement.




Figure 8. Compensation Points

The signal breakout illustrated in Figure 8 can be implemented, if the pair arrangement is radial (horizontal) to the core of the first package and tangential (vertical) to the core in the second package.

The mismatch in the package is equal to 8.5 ps or about 50 mils (170 ps/in) of the PCB route. This 50 mils is equal to the pin pitch of the package. The package skew is always 8.5 ps, regardless of whether the package type is ceramic or organic. Because of processor variation, this mismatch can vary ± 5 ps (3.5 ps to 13.5 ps).

For the PCB, there are two types of skew that exist—intra-pair skew and inter-pair skew. Intra-pair skew is the variation in propagation speed within the halves of a differential pair. Inter-pair skew is the variation in propagation speed between pairs. Intra-pair skew can cause common mode noise due to the timing mismatch of signal edges, that is, between the rising edge of one half of the pair and the falling edge of the other half of the pair. In a perfectly balanced differential system, both halves of the pair are matched perfectly in both time and voltage. Inter-pair skew causes a reduction in timing margin for either setup or hold time between the CLK pair of the signal group and the CAD/CTRL pair. When routing the HyperTransport nets, it is important to take into account the routing in both the package and PCB to ensure the bus always passes the HyperTransport specification for skew. For the motherboard, this means that the PCB traces need to be routed such that any package-trace induced skew is matched or compensated for. The combined package and PCB skews will affect both the inter-pair and intra-pair skews.

There are three different cases that can occur in a naturally compensated package: true is longer, complement is longer, or both traces are the same length. The three cases are illustrated in the following figures: Figure 9 on page 36, Figure 10 on page 37, and Figure 11 on page 37. In these figures, X is the package length of the device.




Figure 9. Case 1: True is Longer




Figure 10. Case 2: Complement is Longer

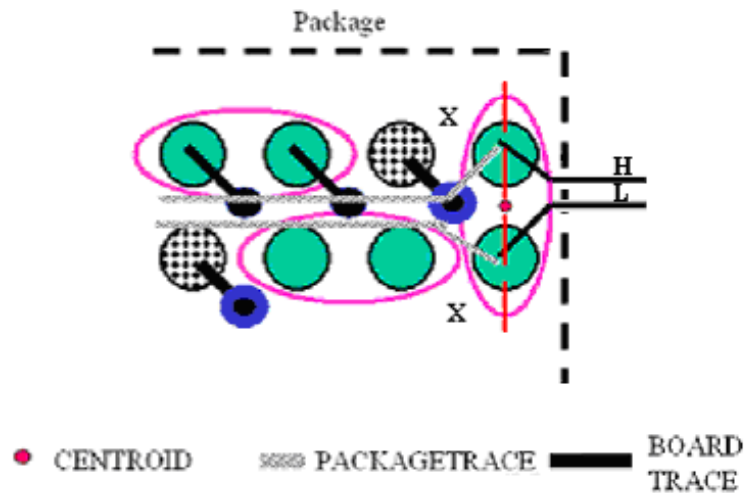


Figure 11. Case 3: True and Complement are the Same Length

Table 5 lists the signal names, the total true-complement skew, and the case to which the signal belongs.

Table 5. True-to-Complement Skew for HyperTransport™ Signals

Naturally Compensated Package Signal Name	Total True-Complement Skew (ps)	Type
Group 0		
CLKIN_H[0]	-8.5	Case 2
CLKIN_L[0]		
CADIN_H[0]	-8.5	Case 2
CADIN_L[0]		
CADIN_H[1]	0.0	Case 3
CADIN_L[1]		
CADIN_H[2]	-8.5	Case 2
CADIN_L[2]		
CADIN_H[3]	0.0	Case 3
CADIN_L[3]		
CADIN_H[4]	0.0	Case 3
CADIN_L[4]		
CADIN_H[5]	-8.5	Case 2
CADIN_L[5]		
CADIN_H[6]	0.0	Case 3
CADIN_L[6]		
CADIN_H[7]	-8.5	Case 2
CADIN_L[7]		
CTLIN_H[0]	0.0	Case 3
CTLIN_L[0]		
Group 1		
CLKIN_H[1]	0.0	Case 3
CLKIN_L[1]		
CADIN_H[8]	0.0	Case 3
CADIN_L[8]		
CADIN_H[9]	8.5	Case 1
CADIN_L[9]		
CADIN_H[10]	0.0	Case 3
CADIN_L[10]		

Table 5. True-to-Complement Skew for HyperTransport™ Signals (Continued)

Naturally Compensated Package Signal Name	Total True-Complement Skew (ps)	Type
CADIN_H[11]	8.5	Case 1
CADIN_L[11]		
CADIN_H[12]	8.5	Case 1
CADIN_L[12]		
CADIN_H[13]	0.0	Case 3
CADIN_L[13]		
CADIN_H[14]	8.5	Case 1
CADIN_L[14]		
CADIN_H[15]	0.0	Case 3
CADIN_L[15]		
CTLIN_H[1]	8.5	Case 1
CTLIN_L[1]		
Group 2		
CLKOUT_H[0]	0.0	Case 3
CLKOUT_L[0]		
CADOUT_H[0]	0.0	Case 3
CADOUT_L[0]		
CADOUT_H[1]	8.5	Case 1
CADOUT_L[1]		
CADOUT_H[2]	0.0	Case 3
CADOUT_L[2]		
CADOUT_H[3]	8.5	Case 1
CADOUT_L[3]		
CADOUT_H[4]	8.5	Case 1
CADOUT_L[4]		
CADOUT_H[5]	0.0	Case 3
CADOUT_L[5]		
CADOUT_H[6]	8.5	Case 1
CADOUT_L[6]		
CADOUT_H[7]	0.0	Case 3
CADOUT_L[7]		
CTLOUT_H[0]	8.5	Case 1
CTLOUT_L[0]		

Table 5. True-to-Complement Skew for HyperTransport™ Signals (Continued)

Naturally Compensated Package Signal Name	Total True-Complement Skew (ps)	Type
Group 3		
CLKOUT_H[1]	-8.5	Case 2
CLKOUT_L[1]		
CADOUT_H[8]	-8.5	Case 2
CADOUT_L[8]		
CADOUT_H[9]	0.0	Case 3
CADOUT_L[9]		
CADOUT_H[10]	-8.5	Case 2
CADOUT_L[10]		
CADOUT_H[11]	0.0	Case 3
CADOUT_L[11]		
CADOUT_H[12]	0.0	Case 3
CADOUT_L[12]		
CADOUT_H[13]	-8.5	Case 2
CADOUT_L[13]		
CADOUT_H[14]	0.0	Case 3
CADOUT_L[14]		
CADOUT_H[15]	-8.5	Case 2
CADOUT_L[15]		
CTLOUT_H[1]	0.0	Case 3
CTLOUT_L[1]		

In addition to minimizing the skew between the true and complement traces, the design needs to minimize the skew between the CAD and CLK signals within a CAD group. For HyperTransport, the skew between differential pairs is calculated by taking the average of the true and complement traces.

Note: The natural compensation only applies to the true-complement skew.

For all the cases in Table 5, the true-complement average remains the same; therefore, the CAD–CLK skew is unaffected by natural compensation.

4.2.2.4.1.1 Routing a Naturally Compensated Package to a Non-Compensated Package

If a package is length-matched (most third-party chipsets and all AMD chipsets that do not have a 1.27-mm pin pitch) and it is being routed to a naturally compensated chip, then the motherboard designer must take into account the package skew between the true and complement traces. The

designer needs to ensure that the flight times from die to die do not violate the HyperTransport specification for true-complement and CAD-CAD skew. Since a naturally compensated device has package routes that are 50 mils longer, extra routing length must be added to make up for the difference the outer ball must travel. Figure 12 illustrates this solution. The extra routing length is equal to 50 mils and is only applied to the signals that fall into either Case 1 or Case 2 (see Figure 9 on page 36, Figure 10 on page 37, and Table 5 on page 38). Extra routing length is not needed for signals that are length matched on the package, as in Case 3 (see Figure 11 on page 37). When adding the extra routing length, the designer must try to maintain the differential routing.

However, maintaining the differential routing is not straightforward, since adjusting one differential pair for true-complement skew can affect the skew between differential pairs. The motherboard designer should use the *HyperTransport™ Technology Link Length Checker*, order# 30867, when performing this type of routing.




Figure 12. Added Routing Length to the Breakout

4.2.2.4.2 True-to-Complement

Table 6 contains the differential pair length matching rules, which define the absolute length difference allowed between true and the complement signal traces on the PCB. The mismatch tolerance is what is recorded in the table.

Table 6. Differential Pair Matching Rules

Processor I/O Connections	Target PCB Trace Mismatch	Mismatch Tolerance	
		Stripline	Microstrip
AMD Athlon 64™ processor chipsets (excluding AMD-8132™ tunnel)	Ball pitch (50 mils for a 1.27-mm package)	28 mils	33 mils
Third-party chipset	Depends on processor ball orientation	See the <i>HyperTransport™ Link Length Checker</i> , order# 30867.	
AMD-8132™ tunnel	Depends on processor ball orientation	See the <i>HyperTransport™ Link Length Checker</i> , order# 30867.	

4.2.2.4.3 CAD-to-CLK and CTL-to-CLK

Table 7 contains CAD-to-CLK and CTL-to-CLK length matching rules for each clock group. CAD-to-CLK length matching defines the length difference allowed between CAD (the average of True and Complement) signals within a clock group and the associated CLK (the average length of True and Complement) signals. The same rules apply to the CTL nets.

Table 7. Signal Group Length-Matching Rules

Case	Frequency	Mismatch Tolerance	
		Stripline	Microstrip
Processor to AMD chipsets (excluding AMD-8132™)	1.6—2.0 GT/s	55 mils	66 mils
Processor to third-party chipset	1.6—2.0 GT/s	See <i>HyperTransport™ Link Length Checker</i> , order# 30867	
Processor to AMD-8132™ tunnel	1.6—2.0 GT/s	See <i>HyperTransport™ Link Length Checker</i> , order# 30867 (Pin-matched option)	

4.2.2.5 Trace Length Mismatch Control

For a detailed list of length-matching requirements, refer to the *HyperTransport™ Technology Link Length Checker*, order# 30867.

4.2.2.6 HyperTransport™ Trace Referencing

HyperTransport is a ground-referenced differential link. Differential signal pairs are weakly coupled and reference the coupled plane more than each other. Ideally this reference plane is VSS, but it can be VSS, VLDT, or VDD. See Figure 13.

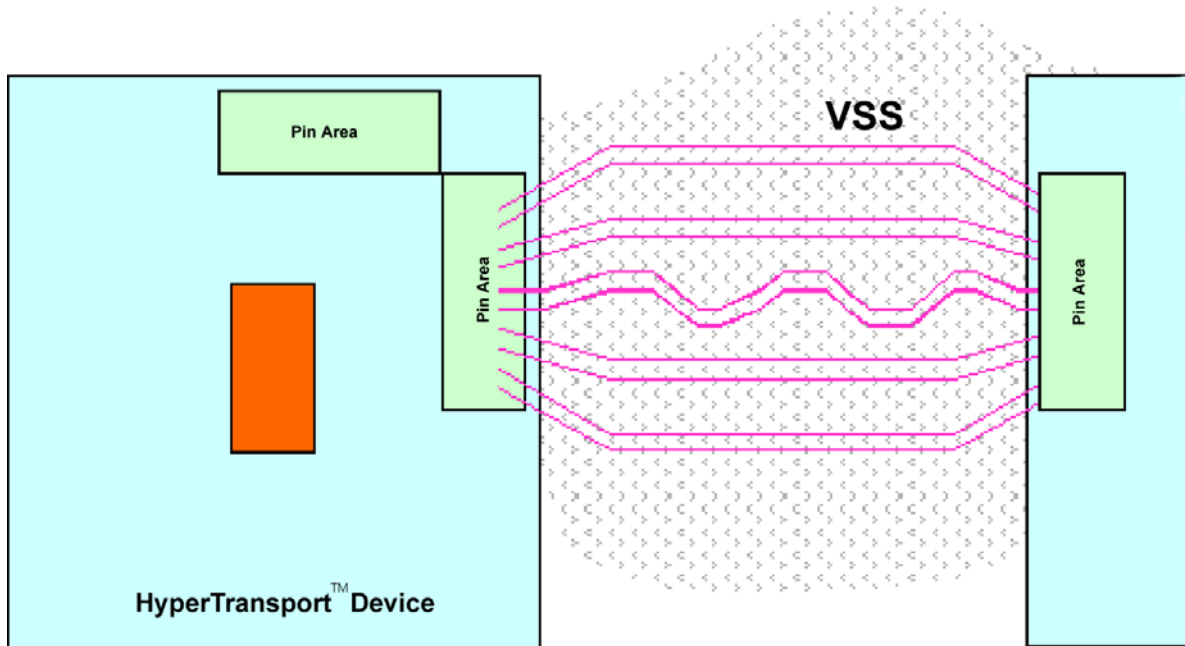


Figure 13. VSS Referenced Layout

4.2.2.7 Changes in HyperTransport™ Trace Referencing

Designers are discouraged from crossing plane splits with HyperTransport signals. If the HyperTransport signals cross a plane split between the source and destination voltages, then the following rules should be followed:

- The plane separation should be 10 mils or less.
- For plane decoupling, the design must use, at a minimum, one 0.01- μ F capacitor for every four differential pairs. Each capacitor should be centered within the four differential pairs so that it is not further than 70 mils from the farthest net in this group of four nets. This design helps to minimize crosstalk in the return path signal. For an example of how to place the stitching capacitor that bridges a split plane, refer to Figure 14.




Figure 14. Changing Reference Layout

Note: The package capacitors are the stitching capacitors for the plane referencing change from motherboard to the package.

For more information on crossing a plane split, see Section 2.3.2 on page 23.

4.2.2.8 Layer Changes

Only links of 800 MT/s or slower are allowed to change layers, with the exception of vias for breakout under the processor.

- If a clock/data group must change layers, place the via pair as close as possible in order to prevent a significant Z_{OD} discontinuity. The via anti-pad plane clearances can overlap in this instance.
- Use one 0.01- μ F 0603 X7R ceramic capacitor to bridge the referenced planes at the crossing boundary for every four differential pairs, as shown in Figure 15. This design helps minimize plane bounce and crosstalk.




Figure 15. Crossing Reference Plane Splits

4.2.2.9 Miscellaneous Signals

All miscellaneous signals should be routed and laid out according to the general layout and routing guidelines of Chapter 8 on page 127. These signals include HTREF0, HTREF1, LDTSTOP_L, RESET_L, and PWROK.

4.3 HyperTransport™ Technology and Processor Breakout

Figure 16 on page 46 and Figure 17 on page 47 are the trace-routing breakouts (topside and bottomside respectively) of a processor HyperTransport link configured as a 16 x 16 link.




Figure 16. 16 x 16-Bit Link Topside Breakout for the Socket AM2 Processor




Figure 17. 16 x 16-Bit Link Bottomside Breakout for the Socket AM2 Processor

4.4 Recommended HyperTransport™ Technology Logic Analyzer Connection

For ease of implementation, AMD recommends placing route-through, probe footprints on the HyperTransport links for logic analysis. Examples of this technology are the Agilent SoftTouch Connectorless or Tektronix Connectorless Compression probe. Please consult your preferred logic analyzer supplier for specific implementation details.

For Agilent Pro Series Soft Touch Connectorless or Tektronix D-Max Probing Technology users, AMD has published a document specifying a recommended pinout, placement, and layout details. Please refer to the *Industry-Standard, Connectorless Probe-Technology Pinout Specification for HyperTransport™ Links*, order# 32693, for details.

For Agilent SoftTouch users, AMD has published a document specifying a recommended pinout, placement, and layout details. Please refer to the *HyperTransport™ Probe8/16 Pinout Specification for SoftTouch*, order# 31133, for details.

It may be possible to leverage the logic analyzer footprint pads for ICT or electrical probing as well.

Chapter 5 Socket AM2 Design Guidelines for Two DDR2 SO-DIMMs

This chapter describes the bus, layout and routing rules, termination resistors, and stackup for unbuffered SO-DIMM designs.

5.1 DDR2 SDRAM Introduction

The socket AM2 processor contains an integrated 128-bit dual-channel DDR2 SDRAM memory controller, which supports up to two unbuffered DDR2 SDRAM SO-DIMMs.

5.1.1 Signal Names and Descriptions

There are four signal groups for the DDR2 bus interface—a data group, an address/command group, a control group, and a clock group. There are two independent channels, labeled A and B.

- Data A—This group consists of the data bits MA_DATA[63:0], the differential strobe signals MA_DQS_H[7:0] and MA_DQS_L[7:0], and the datamask bits MA_DM[7:0].
- Control A—This group consists of chip select signals MA0_CS_L[1:0], MA1_CS_L[1:0], MA_CKE[1:0], and ODT (on-die termination) signals. The ODT pins for control group A are MA0_ODT[0] and MA1_ODT[0]. The processor has connections for check bits MA_CHECK[7:0] that are not used in this configuration.

Notes: The Control A group follows the design guidelines for the address/command A group.

- Address/Command A—This group consists of MA_ADD[15:0], MA_BANK[2:0], MA_RAS_L, MA_CAS_L, and MA_WE_L.
- Clock A—This group consists of MA0_CLK_H[2:1] and MA0_CLK_L[2:1]. Unbuffered SO-DIMMs require two differential clock pairs to each SO-DIMM. MA0_CLK_H[0], MA0_CLK_L[0], MA1_CLK_H[2:0], and MA1_CLK_L[2:0] are not used when connecting the processor to two SO-DIMMs.
- Data B—This group consists of the data bits MB_DATA[63:0], the differential strobe signals MB_DQS_H[7:0] and MB_DQS_L[7:0], and datamask bits MB_DM[7:0].
- Control B—This group consists of chip select signals MB0_CS_L[1:0], MB1_CS_L[1:0], MB_CKE[1:0], and ODT signals. The ODT pins for control group B are MB0_ODT[0] and MB1_ODT[0]. The processor has connections for check bits MB_CHECK[7:0] that are not used in this configuration.

Notes: The Control B group follows the design guidelines for the address/command B group.

- Address/Command B—This group consists of MB_ADD[15:0], MB_BANK[2:0], MB_RAS_L, MB_CAS_L, and MB_WE_L.
- Clock B—This group consists of MB0_CLK_H[2:1] and MB0_CLK_L[2:1]. Unbuffered SO-DIMMs require two differential clock pairs to each SO-DIMM. MB0_CLK_H[0], MB0_CLK_L[0], MB1_CLK_H[2:0], and MB1_CLK_L[2:0] are not used when connecting the processor to two SO-DIMMs.

The four signal groups are treated differently in the motherboard layout. Each group has a particular set of layout and routing guidelines.

The signal pin names for the processor's system controller and the DDR2 SDRAM SO-DIMM are cross-referenced in Table 8 on page 51.

Table 8. DDR2 Signal Descriptions

System Controller Signal Name	DDR2 SO-DIMM Signal Name	Signal Description	Direction ¹
MA_DATA[63:0], MB_DATA[63:0]	DQ[63:0]	Memory Data	BI
MA_DM[7:0], MB_DM[7:0]	DM[7:0]	Memory Datamask	OUT
MA_ADD[15:0], MB_ADD[15:0]	A[15:0]	Memory Address	OUT
MA_BANK[2:0], MB_BANK[2:0]	BA[2:0]	Bank Address	OUT
MA_DQS_H[7:0], MB_DQS_H[7:0]	DQS[7:0]	Data Strobe	BI
MA_DQS_L[7:0], MB_DQS_L[7:0]	DQS[7:0]#	Data Strobe	BI
MA_RAS_L, MB_RAS_L	RAS#	Row Address Select	OUT
MA_CAS_L, MB_CAS_L	CAS#	Column Address Select	OUT
MA_WE_L, MB_WE_L	WE#	Write Enable	OUT
MA_CKE[1:0], MB_CKE[1:0]	CKE[1:0]	Clock Enable	OUT
MA0_ODT[0], MA1_ODT[0], MB0_ODT[0], MB1_ODT[0]	ODT[1:0]	DRAM Termination	OUT
MA0_CS_L[1:0], MB0_CS_L[1:0], MA1_CS_L[1:0], MB1_CS_L[1:0]	S[3:0]#	Chip Select	OUT
MA0_CLK_H[2:1], MB0_CLK_H[2:1]	CK[1:0]	Differential Clock	OUT
MA0_CLK_L[2:1], MB0_CLK_L[2:1]	CK[1:0]#	Differential Clock	OUT
Note:			
1. Direction is defined relative to the processor.			

5.1.2 Layer Assignments

For Socket AM2 mobile system stackups, AMD recommends a minimum of six layers. The DDR2 signal assignments are shown in Figure 18. Data and clock signals are referenced to VSS, and all address/command signals are referenced to VDDIO. Routing net segments too close to plane splits can cause signal integrity issues. For a dielectric height, H, a net segment in the bus channel that is less than 3H from a power-plane split must not exceed 100 mils in length for that spacing. The layout is easier to implement if the plane splits are finalized before routing the signal layers.

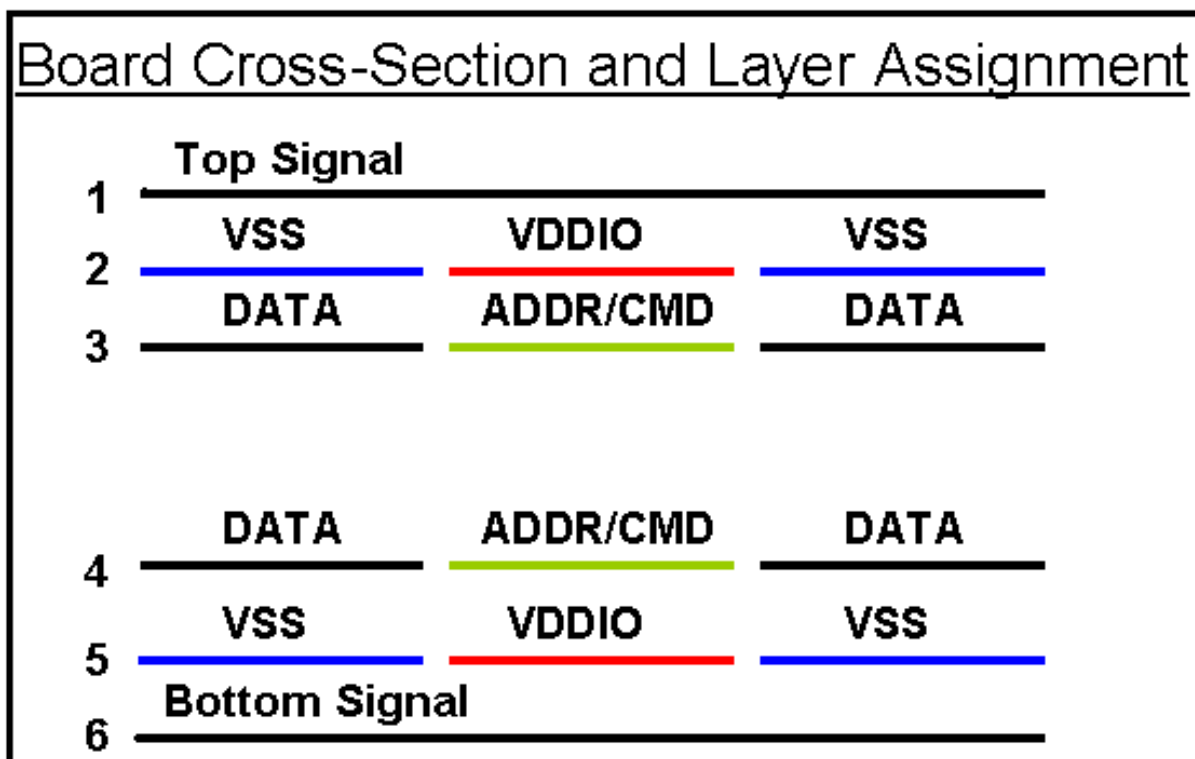


Figure 18. Six-Layer Board Stackup Overview

The DDR2 memory interface may be routed on two internal layers using 4-mil traces (5-mil inner dielectric) or on three internal layers with 5-mil traces (6-mil inner dielectric). One possible stackup option studied in this design guide—for a board with 4-mil dielectric on its outer layers—is illustrated in Figure 19 on page 53. In this chapter, we refer to this stackup as a “5-mil dielectric board” (because it has 5-mil dielectric for its internal layers, which are used for DDR2 routing).

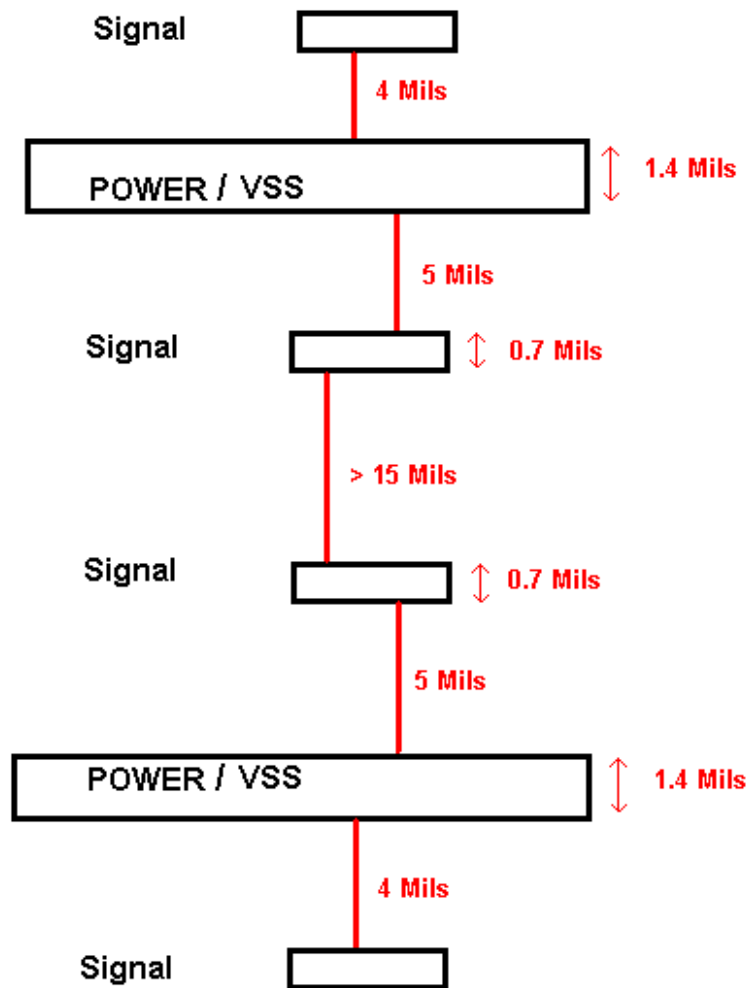


Figure 19. Five-Mil Board Stackup

5.1.3 DDR2 Supply Voltages

The DDR2 controller requires three voltages—VDDIO, M_VREF, and VTT. For more information on the required VDDIO decoupling and decoupling placement, refer to Section 9.2.3.1 on page 149. For more information on the required decoupling for VTT and decoupling placement, refer to Section 9.2.3.2 on page 150. For more information on the required decoupling for M_VREF and decoupling placement, refer to Section 9.2.3.3 on page 154.

5.2 Overview of DDR2 SDRAM Interface

The following sections provide an overview of the processor-memory interface and interconnect, and a description of valid SO-DIMM configurations when operating a system with two DDR2 SO-DIMM modules.

5.2.1 Interface Overview and Block Diagram

Figure 20 shows a block diagram for a mobile system with two unbuffered DDR2 SO-DIMMs.

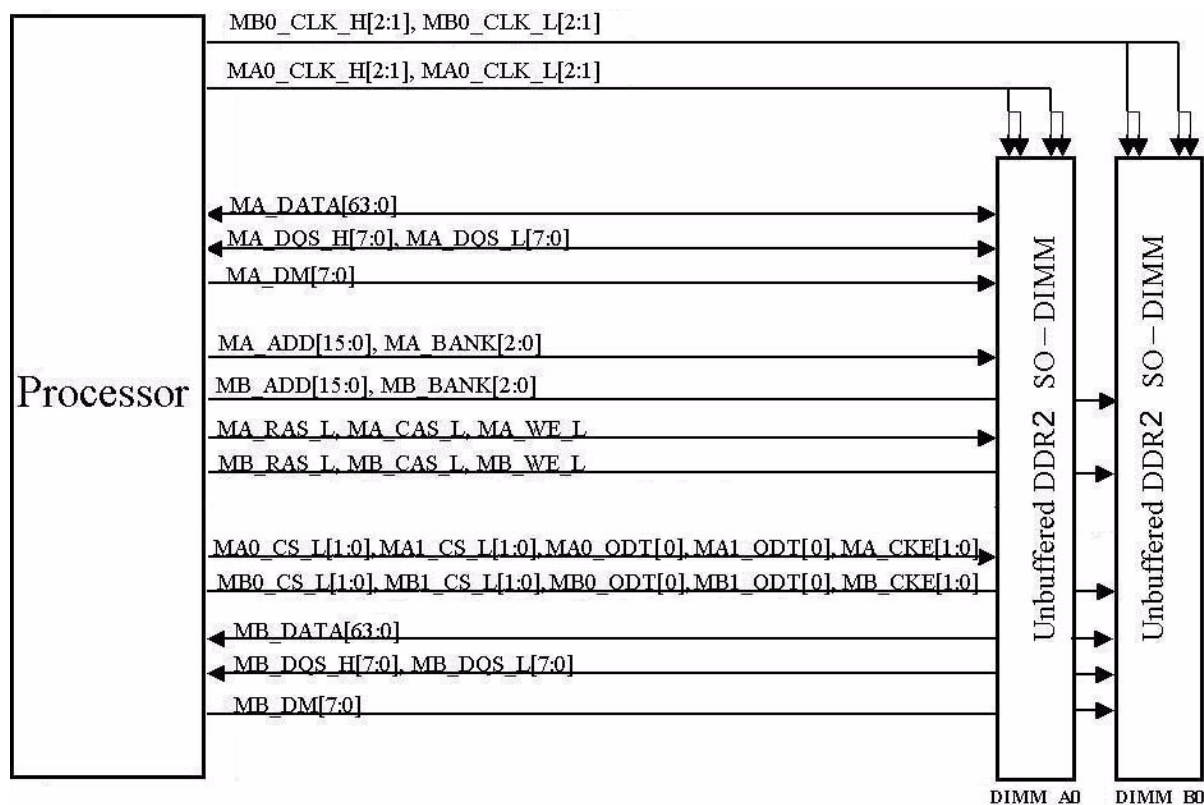


Figure 20. Block Diagram for Socket AM2 Mobile System with Two Unbuffered SO-DIMM Modules

5.2.2 Valid Configurations – Two Unbuffered SO-DIMMs

The socket AM2 processor can support up to two, unbuffered SO-DIMM slots. Refer to the *AMD BIOS and Kernel Developers Guide for AMD NPT Family 0Fh Processors*, order# 32559, for specific memory size configurations.

5.2.3 Frequency

For the most up-to-date supported frequency, refer to the *AMD BIOS and Kernel Developers Guide for AMD NPT Family 0Fh Processors*, order# 32559.

5.3 Unbuffered Two SO-DIMM DDR2 System Layout Guidelines

The following sections describe the recommended termination, placement and routing, as well as the trace width and separation guidelines for the unbuffered mobile DDR2 memory interface design.

General routing rules are as follows:

1. The variable “x” is used to define a target lead-in length of CLK to each SO-DIMM. A range for processor and SO-DIMM placement can be specified, if a variable is used for the lead-in length. All recommended trace lengths are defined relative to “x”.
2. The lead-in length is defined as the distance from the processor pins to the SO-DIMM pins.
3. To determine the lead-in length, route the longest and the shortest data nets and set the lead-in length target to be an average of the two lengths: (longest data net length + shortest data net length)/2.
4. The DDR2 channel B should be routed on one inner signal layer, and the DDR2 channel A should be routed on another inner signal layer.
5. To reduce cross talk from the pins to the nets, designers should avoid adding net length to the nets under the processor as a means to achieve length matching. Additionally, nets should not be routed around processor pins. Refer to Figure 21 on page 56 for an illustration of what *not* to do.

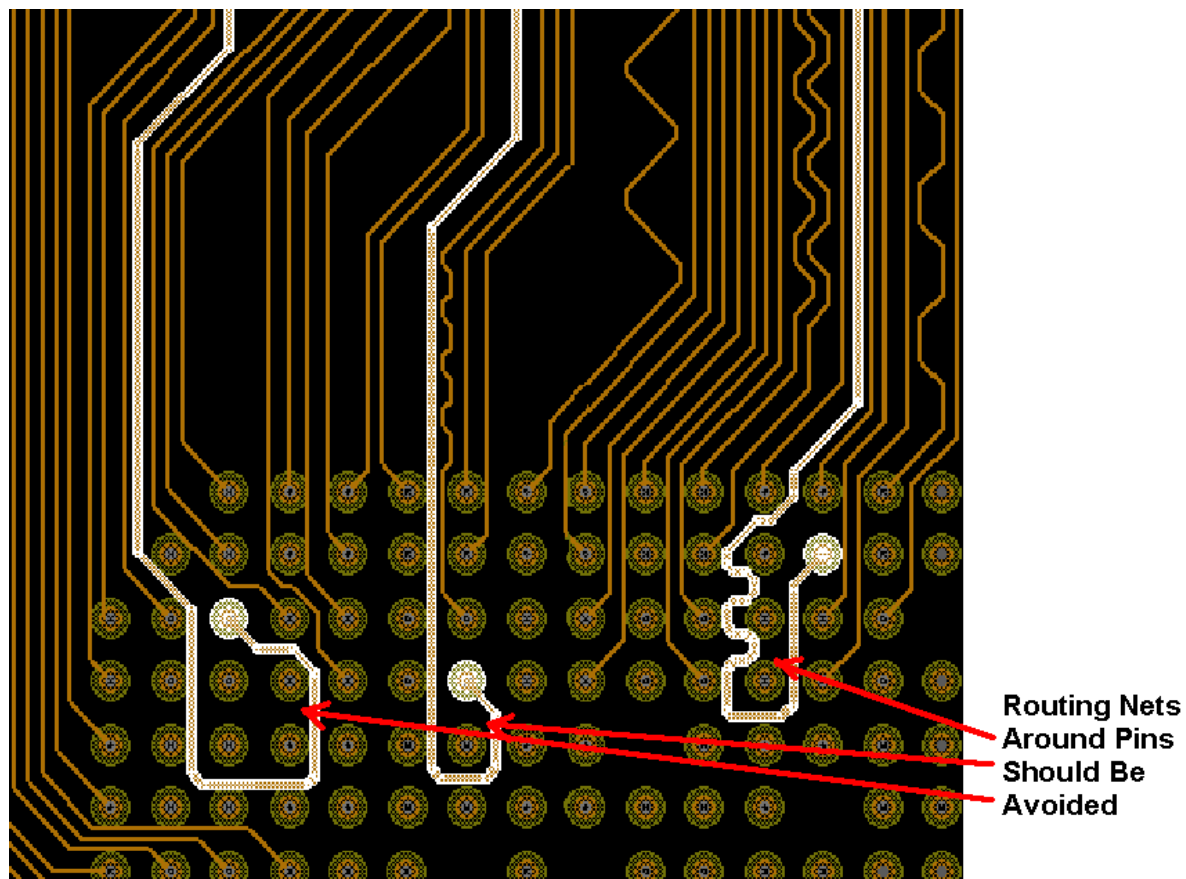


Figure 21. Undesirable Net Routing

Figure 22 on page 57 shows the lead-in length requirements for memory routing (from the processor to the SO-DIMMs).

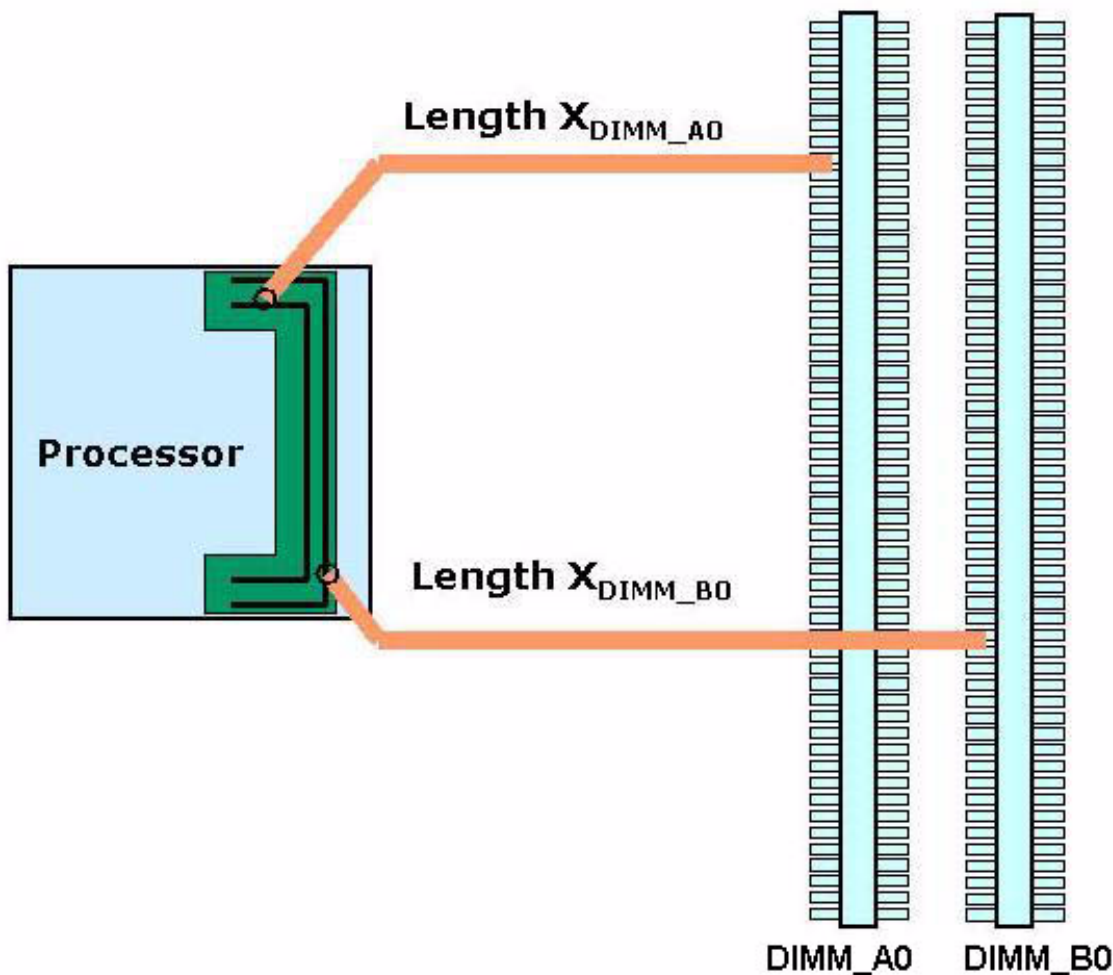


Figure 22. Definition of Lead-in Length for Channel A and B Routing

5.3.1 Address/Command Termination and Layout Guidelines

This section describes termination and routing rules for the following signals:

MA_ADD[15:0] and MB_ADD[15:0]

MA_BANK[2:0] and MB_BANK[2:0]

MA_RAS_L and MB_RAS_L

MA_CAS_L and MB_CAS_L

MA_WE_L and MB_WE_L

MA_CKE[1:0] and MB_CKE[1:0]

MA0_CS_L[1:0], MA1_CS_L[1:0], MB0_CS_L[1:0], and MB1_CS_L[1:0]

MA0_ODT[0], MA1_ODT[0], MB0_ODT[0], and MB1_ODT[0]

The address/command group is terminated with an RTT resistor pulled to VTT.

Recommended values are listed in Table 9.

Table 9. Termination Values for Unbuffered SO-DIMM Address/Command Signals

Parameter	Value	Tolerance	Pull Voltage
RTT	47 Ω	5%	VTT

There are two address/command channels. The A channel should be routed to SO-DIMM_A0, and the B channel to SO-DIMM_B0. Figure 23 on page 59 shows the DDR2 address/command termination for channel A or B.




Figure 23. DDR2 Address/Command Routing

Note: Figure 23 shows equivalent diagrams. Any of the three can be implemented.

Table 10 defines the routing parameters for address/command and control nets. Target impedance for this group of signals is 60 Ω with 10% tolerance, including manufacturing tolerances.

Table 10. Parameters for Address/Command and Control Trace Width and Space for a Board With Five-Mil Inner Dielectric

Definition	Breakout ^{2, 3, 4}	Bus Channel ⁵	SO-DIMM Via Field	Units
Trace Width	4.0	4.0	4.0	mils
Trace Spacing ¹	6.0	8.0	8.0 ⁶	mils
Notes:				
1. If the net is serpentine, then the space to itself should be 12 mils or more.				
2. The length of the nets in the breakout area should be less than 500 mils.				
3. If three nets are routed between two processor pins, then they may be routed as 4-mil traces with 4-mil spacing.				
4. Maximize the space between the nets.				
5. If the lead-in length of a trace is longer than 4 inches, then 4 mils should be added to the spacing to other traces in the bus channel.				
6. The maximum length of parallel nets with less than the designated trace spacing is 100 mils.				

5.3.2 Data Group Connection and Termination Guidelines

This section defines the data groups and describes their connections. There are a total of 16 data groups (8 groups per channel), each consisting of eight data bits, one differential data strobe pair, and one datamask signal. All data groups on each channel should be routed on the same layer.

Table 11 on page 60 shows the DDR2 data groups for a single channel. However, both channels have identical signals. When ganged in 128-bit mode, channel A serves as the lower 64-bits and channel B serves as the upper 64-bits.

Table 11. DDR2 Data Groups

Group	Data	DQS_H/DQS_L	DM
0	[7:0]	0	0
1	[15:8]	1	1
2	[23:16]	2	2
3	[31:24]	3	3
4	[39:32]	4	4
5	[47:40]	5	5
6	[55:48]	6	6
7	[63:56]	7	7

Figure 24 shows DQS routing. Figure 25 on page 62 shows Data and DM routing. As seen in these figures, no termination resistors are needed for the DATA, DQS, or DM nets.




Figure 24. DQS Routing and Termination

Table 12 defines the DQS routing parameters. The differential impedance target for DQS nets is 93 Ω with 10% tolerance, including manufacturing tolerances.

Table 12. Parameters for MA_DQS and MB_DQS Trace Width and Separation for a Board With Five-Mil Inner Dielectric

Definition	Breakout ^{2, 3, 4}	Bus Channel ⁵	SO-DIMM Via Field	Units
Trace Width	4.0	4.0	4.0	mils
DQS Spacing to other nets ¹	6.0	12.0	8.0 ⁶	mils
Differential spacing for DQS nets	4.0	4.0	4.0	mils

Notes:

1. If the net is serpentine, then the space to itself should be 12 mils or more.
2. The length of the nets in the breakout area should be less than 500 mils.
3. If three nets are routed between two processor pins, then they may be routed as 4-mil traces with 4-mil spacing.
4. Maximize the space between DQS and other nets.
5. If the lead-in length of a trace is longer than 4 inches, then 4 mils should be added to the spacing to other traces in the bus channel.
6. The maximum length of parallel nets with less than the designated trace spacing is 100 mils.

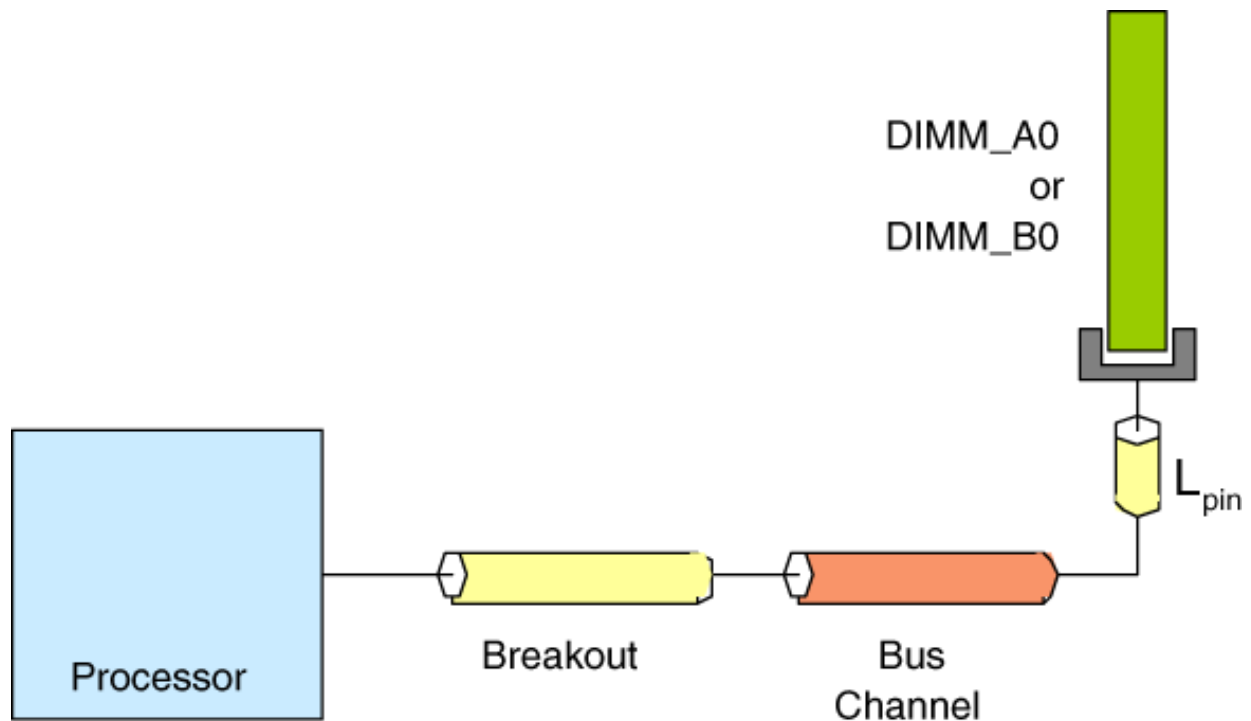


Figure 25. Data and DM Routing and Termination

Table 13 on page 63 defines DATA and DM routing parameters. Impedance target for DATA and DM is $60 \Omega \pm 10\%$.

Table 13. Parameters for Data and DM Trace Width and Separation for a Board With Five-Mil Inner Dielectric

Definition	Breakout ^{2, 3, 4}	Bus Channel ⁵	DIMM Via Field	Units
Trace Width	4.0	4.0	4.0	mils
Spacing to other nets ¹	6.0	8.0	8.0 ⁶	mils

Notes:

1. If the net is serpentine, then the space to itself should be 12 mils or more.
2. The length of the nets in the breakout area should be less than 500 mils.
3. If three nets are routed between two processor pins, then they may be routed as 4-mil traces with 4-mil spacing.
4. Maximize the space between the nets.
5. If the lead-in length of a trace is longer than 4 inches, then 4 mils should be added to the spacing to other traces in the bus channel.
6. The maximum length of parallel nets with less than the designated trace spacing is 100 mils.

5.3.3 Clock Termination and Layout Guidelines

This section describes the assignment and termination for the clock signals. Each SO-DIMM socket receives two differential clock pairs.

Figure 26 shows the unbuffered mobile DDR2 clock configuration recommendations.

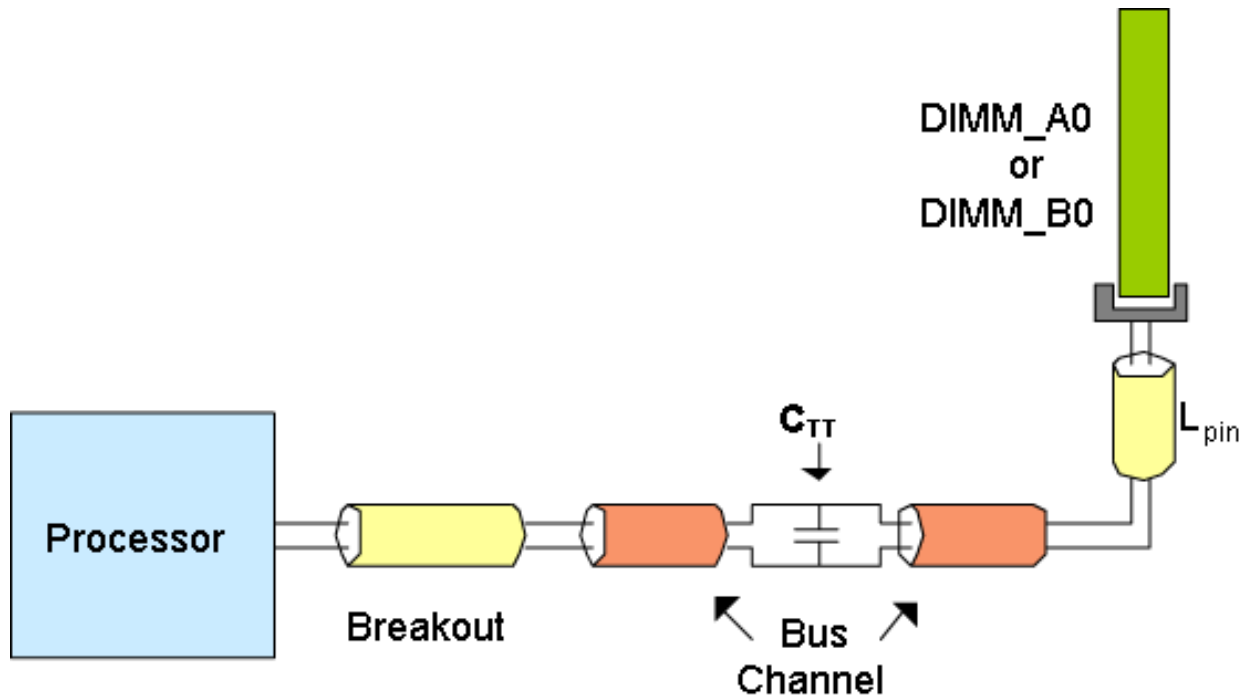


Figure 26. Clock Routing and Termination

Differential C_{TT} termination is required for CLK_H and CLK_L. The maximum length for the C_{TT} stub is 150 mils (PCB trace + via height), with a PCB trace length of less than 100 mils. Table 14 lists the termination value for the unbuffered SO-DIMM clocks.

Table 14. CLK_H and CLK_L Termination

Parameter	Value	Tolerance
C_{TT}	1.5 pF	±10%

A general depiction of the overall CLK_H and CLK_L routing is shown in Figure 27 on page 65.

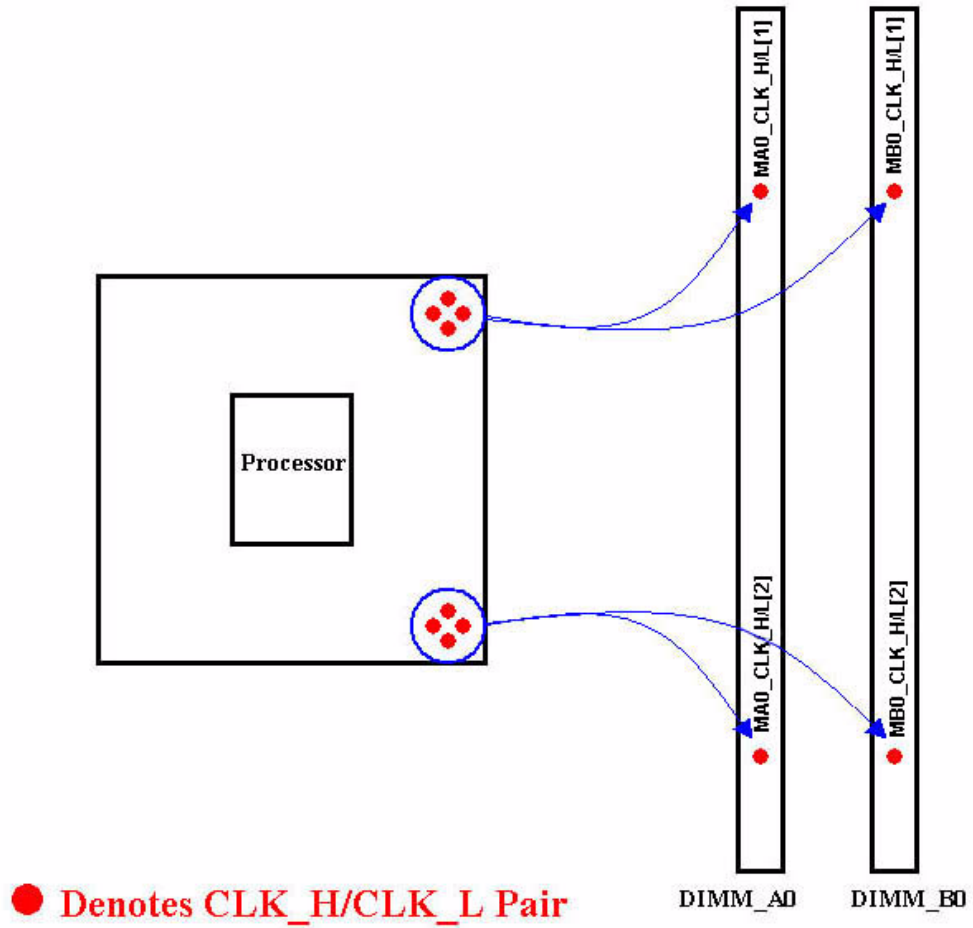


Figure 27. Clock Routing and Termination

Table 15 defines CLK routing parameters. In the bus channel the differential impedance target for CLK nets is 72 Ω with 10% tolerance, including manufacturing tolerances.

Table 15. Parameters for CLK Trace Width and Separation for a Board With Five-Mil Inner Dielectric

Definition	Breakout ^{2, 3, 4}	Bus Channel ⁵	SO-DIMM Via Field	Units
CLK Width	4.0	8.0	4.0	mils
CLK Spacing to Other Signals ¹	6.0	16.0	8.0 ⁶	mils
Differential Spacing for CLK pairs	4.0	4.0	4.0	Mils

Notes:

1. If the net is serpentine, then the space to itself should be 16 mils or more.
2. The length of the nets in the breakout area should be less than 500 mils.
3. If three nets are routed between two processor pins, then they may be routed as 4-mil traces with 4-mil spacing.
4. Maximize the space between the nets.
5. If the lead-in length of a trace is longer than 4 inches, then 4 mils should be added to the spacing to other traces in the bus channel.
6. The maximum length of parallel nets with less than the designated trace spacing is 100 mils.

5.4 SO-DIMM Routing with Five-Mil Traces

If desired, a system based on the processor and two SO-DIMMs may be routed using 5-mil traces (stripline routing on a board with 6-mil inner dielectric). In this case, three internal layers are needed for DDR2 routing. The following trace geometry should be used for 5-mil traces:

- All non-differential nets should be routed as 5-mil traces with 10-mil spacing to the adjacent traces—this includes Address, Command, Control, Data, and DM bits. In the bus channel, the target impedance for these signals is 60 Ω with 10% tolerance, including manufacturing tolerances.
- The DQS nets should be routed as 5-mil traces with 5-mil differential spacing and 15-mil spacing to any other net (15/5/5/5/15). In the bus channel, the target impedance for this group of signals is 93- Ω differential impedance with 10% tolerance, including manufacturing tolerances.
- The CLK nets should be routed as 10-mil traces with 5-mil differential spacing and 20-mil spacing to any other net (20/10/5/10/20). In the bus channel, the target impedance for this group of signals is 72- Ω differential impedance with 10% tolerance, including manufacturing tolerances.

5.5 Basic Idea for SO-DIMM DDR2 Memory Routing

The basic idea for DDR2 routing is to match all signal lengths, which go to a particular SO-DIMM, to within a given tolerance. Both data and address/command are source-synchronous buses: Address/command propagates relative to clock, and data propagates relative to strobe. A fairly loose length-matching tolerance is required for address/command and all 64-bits of data to each SO-DIMM. The tight length-matching tolerances are for data relative to strobe *within a data-group*. A step-by-step routing procedure is described in Sections 5.5.1 and 5.5.2.

These routing guidelines are based on a design in which the lower 64-bit SO-DIMM are placed before the upper 64-bit SO-DIMM. The designer is free to choose the SO-DIMM spacing, but AMD recommends a spacing of 600 to 800 mils for the most compact memory footprint.

5.5.1 Routing Procedure

1. Route MB_DATA[0] and MB_DATA[63] to SO-DIMM_B0. Choose the longest of the two nets and measure the lead-in length from the processor pin to the SO-DIMM_B0 pin. Adjust the processor placement relative to the SO-DIMMs until this length is shorter than 5.0 inches.
2. Route MA_DATA[0] and MA_DATA[63] to SO-DIMM_A0. Choose the longest of the two nets and measure the lead-in length from the processor pin to the SO-DIMM_A0 pin. Adjust the processor placement relative to the SO-DIMMs until this length is shorter than 5.0 inches.
3. Route all the remaining nets to SO-DIMM_B0. The shortest MB_DATA[63:0] net can be up to 2.00 inches shorter than the longest. The MB0_CLK[2:1]_H and MB0_CLK[2:1]_L lengths should be fixed at the median distance between the longest and shortest net to SO-DIMM_B0.
4. Route all the remaining nets to SO-DIMM_A0. The shortest MA_DATA[63:0] net can be up to 2.00 inches shorter than the longest. The MA0_CLK[2:1]_H and MA0_CLK[2:1]_L lengths should be fixed at the median distance between the longest and shortest net to SO-DIMM_A0.

5.5.2 Routing and Length-Matching Rules

In Sections 5.5.2.1, 5.5.2.2, and 5.5.2.3 the design guide rules are reduced to the most basic information. The nets routed on the top or bottom layer of the board—where a piece of the net must change from stripline to microstrip to be routed to the SO-DIMM pads—should be less than 300 mils.

5.5.2.1 CLK_H and CLK_L Signals

- All CLK signals are routed over the VSS planes.
- True-to-complement differential pairs are length matched within 50 mils.
- The two CLK pairs to a particular SO-DIMM have to be length matched within 100 mils.
- Target CLK length to SO-DIMM_A0 is $X_{\text{DIMM_A0}}$ and is a length between 2000 to 4000 mils.
- Target CLK length to SO-DIMM_B0 is $X_{\text{DIMM_B0}}$ and is a length between 2000 to 4000 mils.

- The Clocks to a SO-DIMM are routed on the same layer as used for Address/Command routing to the same SO-DIMM.
- The net connecting C_{TT} termination capacitors to the processor pins should be shorter than 1200 mils.

5.5.2.2 Data, DQS, DM Signals

- All data/DQS/DM signals are referenced to the VSS plane and do not cross a plane split.
- All Data and DM nets in each data group—consisting of one DQS_H, one DQS_L, eight Data, and one DM signal—are length matched to within 100 mils of the average of the DQS net lengths on that Data group.
- For DQS nets, true and complement mismatch tolerance is 50 mils.
- All data/DQS/DM signals are length matched to within 1000 mils of X_{DIMM_A0} and X_{DIMM_B0} , which means that the longest and shortest data/DQS/DM nets cannot differ by more than 2000 mils.

5.5.2.3 ADD, BANK, RAS, CAS, WE, CS, CKE, ODT

- All Address/Command signals are referenced to VDDIO and do not cross any power plane split.
- All address and command nets to each SO-DIMM should be in the range of the CLKs and 1000 mils shorter than the CLKs.
- The length of the traces from the SO-DIMM pads to RTT should be less than 1.9 inches.
- The routing length from each SO-DIMM to RTT is matched within 600 mils.

Chapter 6 Design Guidelines for Four Unbuffered DDR2 DIMMs

This chapter describes the bus, layout and routing rules, termination resistors, and stackup for unbuffered four-DIMM designs.

Note: A two-DIMM design can also be implemented using the four-DIMM design guidelines.

6.1 DDR2 SDRAM

The processor contains an integrated 128-bit dual-channel DDR2 SDRAM memory controller, which supports up to four, unbuffered DDR2 SDRAM DIMMs. All DIMMs must operate at the same frequency as the slowest DIMM.

6.1.1 Signal Names and Descriptions

There are four signal groups for the DDR2 bus interface—a data group, a control group, an address/command group, and a clock group. There are two independent channels, labeled A and B.

- Data A—This group consists of the data bits MA_DATA[63:0], MA_DM[8:0], the differential strobe signals MA_DQS_H[8:0] and MA_DQS_L[8:0], and the check bits MA_CHECK[7:0].
- Control A—This group consists of the chip select signals MA0_CS_L[1:0] and MA1_CS_L[1:0], MA_CKE[1:0], and the ODT (on-die termination) signals MA0_ODT[0] and MA1_ODT[0].

Notes: The Control A group follows the design guidelines for the address/command A group.

- Address/Command A—This group consists of MA_ADD[15:0], MA_BANK[2:0], MA_RAS_L, MA_CAS_L, and MA_WE_L.
- Clock A—This group consists of MA0_CLK_H[2:0], MA0_CLK_L[2:0], MA1_CLK_H[2:0], and MA1_CLK_L[2:0]. Unbuffered DIMMs require three differential clock pairs to each DIMM.
- Data B—This group consists of the data bits MB_DATA[63:0], MB_DM[8:0], the differential strobe signals MB_DQS_H[8:0] and MB_DQS_L[8:0], and the check bits MB_CHECK[7:0].
- Control B—This group consists of chip select signals MB0_CS_L[1:0] and MB1_CS_L[1:0], MB_CKE[1:0], and the ODT signals MB0_ODT[0] and MB1_ODT[0].

Notes: The Control B group follows the design guidelines for the address/command B group.

- Address/Command B—This group consists of MB_ADD[15:0], MB_BANK[2:0], MB_RAS_L, MB_CAS_L, and MB_WE_L.

- Clock B—This group consists of MB0_CLK_H[2:0], MB0_CLK_L[2:0], MB1_CLK_H[2:0], and MB1_CLK_L[2:0]. Unbuffered DIMMs require three differential clock pairs to each DIMM.

The four signal groups—the data group, address/command group, control group, and clock group—are treated differently in the motherboard layout. Each group has a particular set of layout and routing guidelines.

The signal pin names for the processor system controller and the DDR2 SDRAM DIMM are cross-referenced in Table 16.

Table 16. DDR2 Signal Descriptions

Processor Signal Name	DDR2 DIMM Signal Name	Signal Description	Direction ¹
MA_DATA[63:0], MB_DATA[63:0]	DQ[63:0]	Memory Data	BI
MA_DM[8:0], MB_DM[8:0]	DM[8:0]	Memory Datamask	OUT
MA_ADD[15:0], MB_ADD[15:0]	A[15:0]	Memory Address	OUT
MA_BANK[2:0], MB_BANK[2:0]	BA[2:0]	Bank Address	OUT
MA_CHECK[7:0], MB_CHECK[7:0]	ECC[7:0]	Memory Error Correction	BI
MA_DQS_L[8:0], MB_DQS_L[8:0]	DQS[8:0]#	Data Strobe	BI
MA_DQS_H[8:0], MB_DQS_H[8:0]	DQS[8:0]	Data Strobe	BI
MA_RAS_L, MB_RAS_L	RAS#	Row Address Select	OUT
MA_CAS_L, MB_CAS_L	CAS#	Column Address Select	OUT
MA_WE_L, MB_WE_L	WE#	Write Enable	OUT
MA_CKE[1:0], MB_CKE[1:0]	CKE	Clock Enable	OUT
MA0_ODT[0], MB0_ODT[0]	ODT0	DRAM Termination	OUT
MA1_ODT[0], MB1_ODT[0]	ODT0	DRAM Termination	OUT
MA0_CS_L[1:0], MB0_CS_L[1:0]	S1# / S0#	Chip Select	OUT
MA1_CS_L[1:0], MB1_CS_L[1:0]	S1# / S0#	Chip Select	OUT
MA0_CLK_H[2:0], MB0_CLK_H[2:0]	CK[2:0]	Differential Clock	OUT
MA0_CLK_L[2:0], MB0_CLK_L[2:0]	CK[2:0]#	Differential Clock	OUT
MA1_CLK_H[2:0], MB1_CLK_H[2:0]	CK[2:0]	Differential Clock	OUT
MA1_CLK_L[2:0], MB1_CLK_L[2:0]	CK[2:0]#	Differential Clock	OUT
Note:			
1. Direction is defined relative to the processor.			

6.1.2 Four Layer Assignments

As a minimum requirement, a 4-layer board stackup is needed for the socket AM2 processor. All data signals are referenced to VSS. All address/command signals are referenced to VDDIO, and all clock signals are referenced to either VDDIO or VSS depending on their position on the DIMM. Routing net segments too close to plane splits can cause signal integrity issues. For a dielectric height, H, a net

segment in the bus channel that is less than 3H from a power-plane split must not exceed 100 mils in length for that spacing. The layout is easier to implement if the plane splits are finalized before routing the signal layers.

The DDR2 signal assignments are shown in Figure 28 on page 73. Slightly overlapping the VSS and VDDIO planes can help to make the placement of the decoupling capacitors easier.

Board Cross-Section and Layer Assignment




Figure 28. Board Stackup Overview

A motherboard design based on the socket AM2 processor may be implemented on a board with either 4-mil or 4.5-mil outer layer dielectric thickness. The stackup parameters are shown in Figure 29, and the numerical values for each parameter are listed in Table 17 on page 75.




Figure 29. Board Stackup Parameters for a Four-Layer Board

Table 17. Parameters for Stackup on a Four-Layer Board

Parameter	Definition	4.0-Mil Dielectric	4.5-Mil Dielectric	Units
H ₁	Outer Dielectric Thickness	4.0	4.5	mils
Th ₁	Outer Signal Layer Metal Thickness	2.0	2.0	mils
Th ₂	Plane Layer Metal Thickness	1.4	1.4	mils
Er	Dielectric Constant (at 1 GHz)	4.0	4.0	

6.1.3 Six Layer Assignments

A six-layer board stackup may be implemented in two ways—(1)with DDR2 signals routed as stripline only or (2)with DDR2 signals routed as stripline and microstrip. The DDR2 signal assignments on a stripline-only implementation are shown in the Figure 30 on page 76. The DDR2 signal assignments on a all-layer implementation are shown in the Figure 31 on page 77. For both implementations, data and clock signals are referenced to VSS and all address/command signals are referenced to VDDIO. Routing net segments too close to plane splits can cause signal integrity issues. For a dielectric height, H, a net segment in the bus channel that is less than 3H from a power-plane split must not exceed 100 mils in length for that spacing. The layout is easier to implement if the plane splits are finalized before routing the signal layers.

Note: These two six-layer implementations are provided for reference only. They have not been build nor tested by AMD.

6.1.3.1 Stripline Only Routing

The stack-up for the stripline-only implementation should be as follows—S1, P1, S2, S3, P2, and S4, where

- S1 not used for DDR
- P1 is Power and GND split
- S2 is inner-layer 1 routing (controller B)
- S3 is inner-layer 2 routing (controller A)
- P2 is GND
- S4 not used for DDR

The advantages of stripline-only routing are as follows:

- All routes are buried, so this makes room for more surface-mount components.

- The design is more compact, so the DDR routing can be compressed into a smaller space. This compactness allows for closer DIMMs, as well as possibly another sub-system being placed closer to the processor.

The DDR2 signal assignments on a stripline-only implementation are shown in the Figure 30.

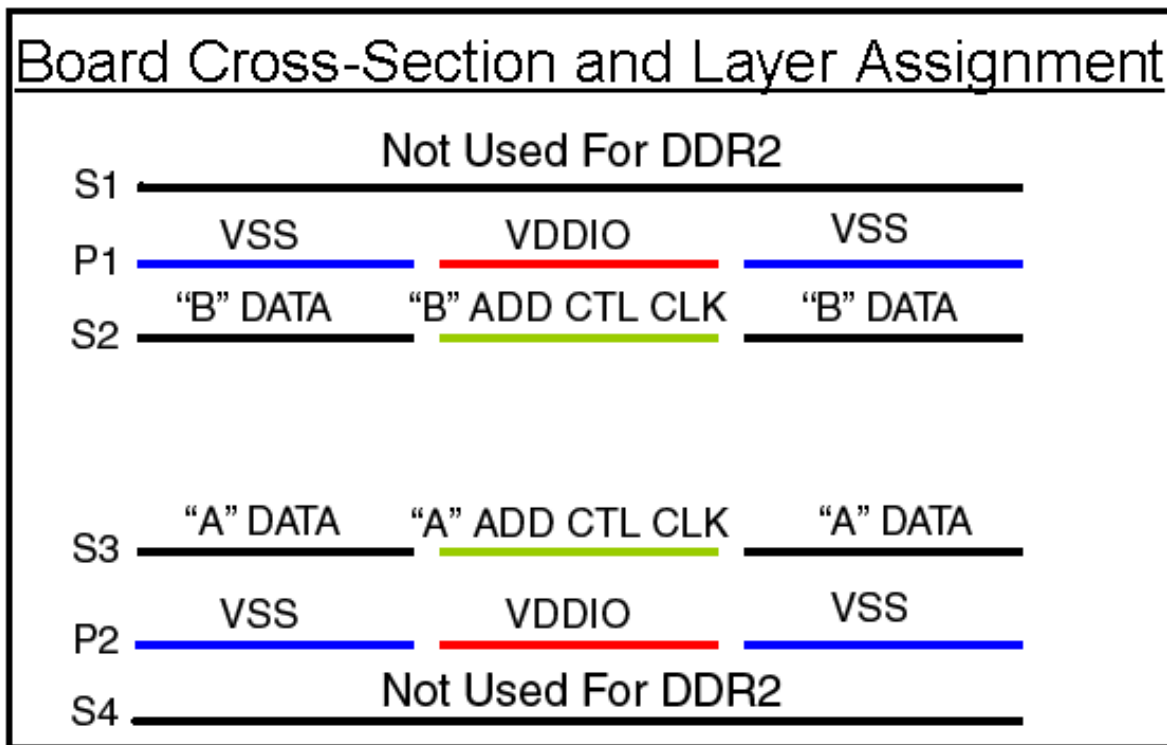


Figure 30. Six-Layer Board Stackup Overview—Stripline Only

6.1.3.2 Stripline and Microstrip Routing

The stack-up for the all-layer implementation should be as follows—S1, P1, S2, S3, P2, and S4, where

- S1 is top-layer routing (controller B)
- P1 is Power and GND split
- S2 is inner-layer 1 routing (controller A)
- S3 is inner-layer 2 routing (controller A)
- P2 is GND
- S4 is bottom-layer routing (controller B)

The advantages of using the all-layer implementation are as follows:

- The DDR routing can be spread out more to help reduce crosstalk.
- Wider traces can be used to better match load impedance.
- This stackup allows for a solid plane on P2 in the DDR routing area.

The DDR2 signal assignments on a all-layer implementation are shown in the Figure 31.

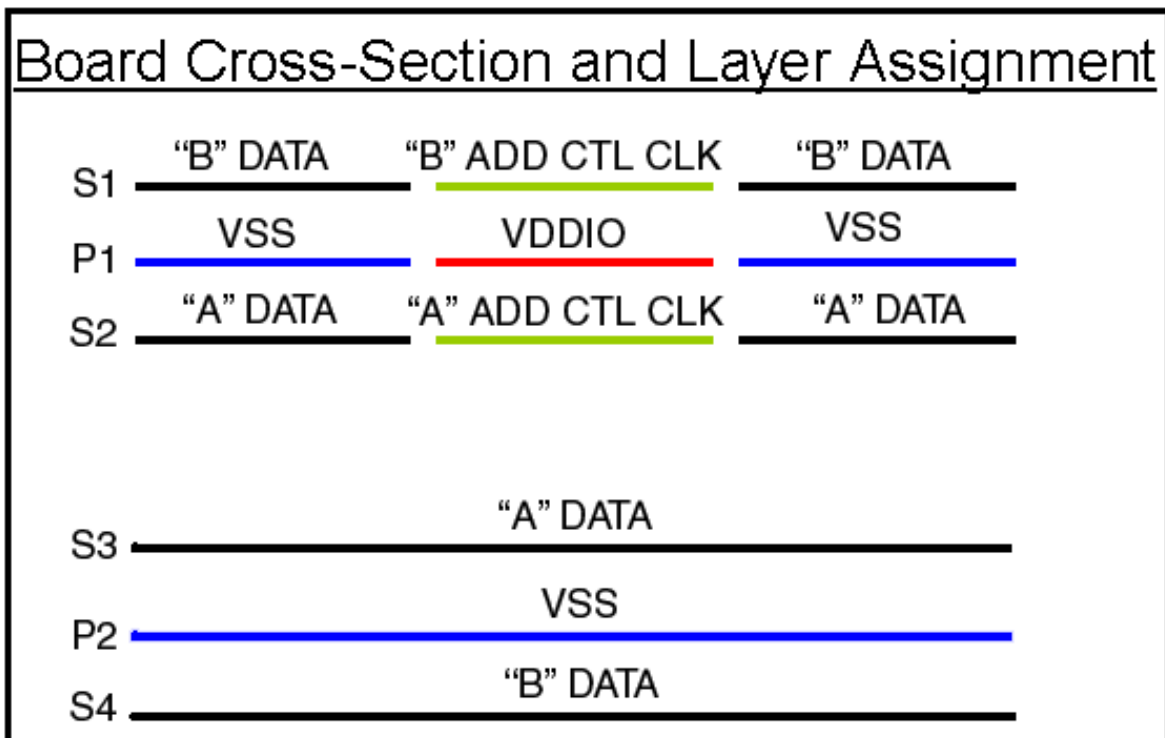


Figure 31. Six-Layer Board Stackup Overview—All Layers Used

6.2 Overview of DDR2 SDRAM Interface

The following sections provide an overview of the processor-memory interface and interconnect.

6.2.1 Interface Overview and Block Diagram

Figure 32 on page 78 shows a block diagram for a 4-DIMM unbuffered DDR2 interface design.

Note: Each DIMM receives three differential clock pairs and a copy of the address/command bus.

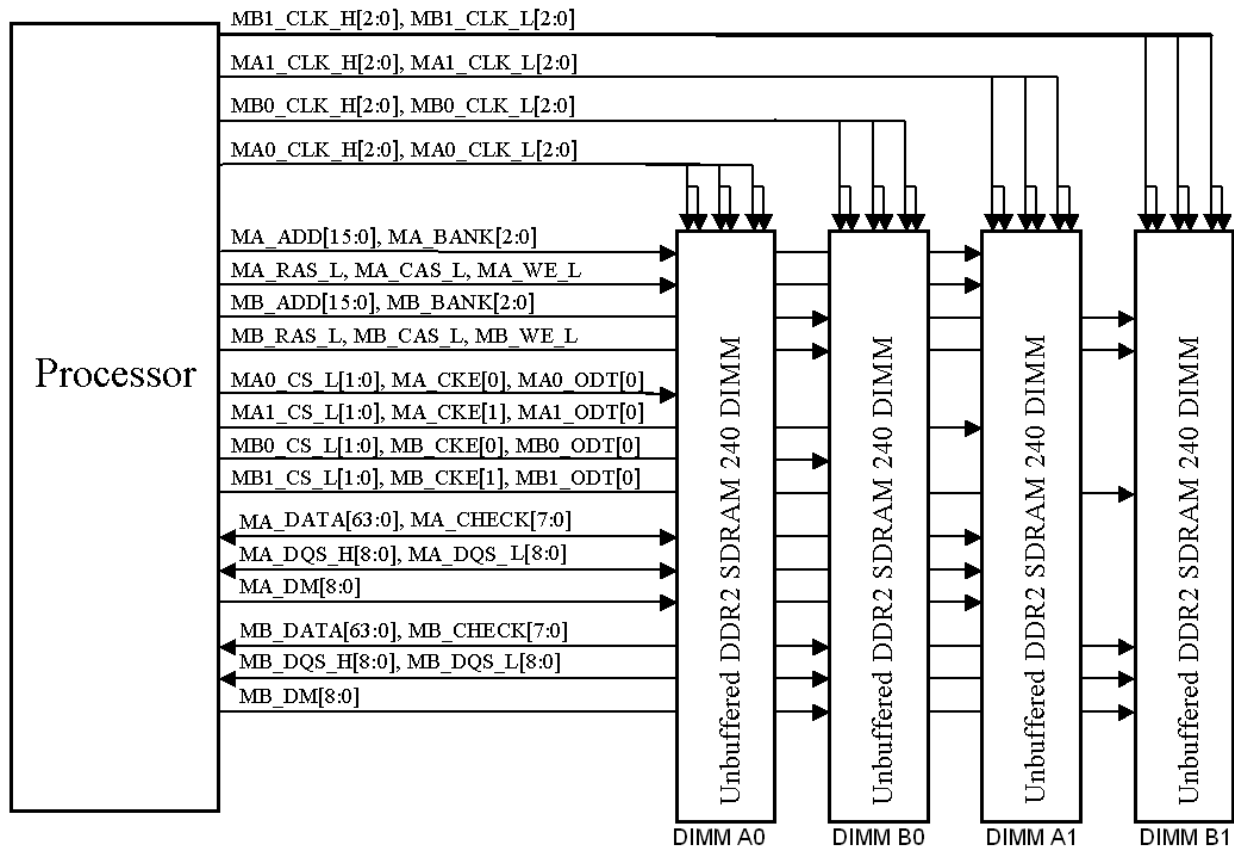


Figure 32. Unbuffered 4-DIMM Block Diagram

Figure 33 shows a block diagram for a 2-DIMM unbuffered DDR2 interface design.

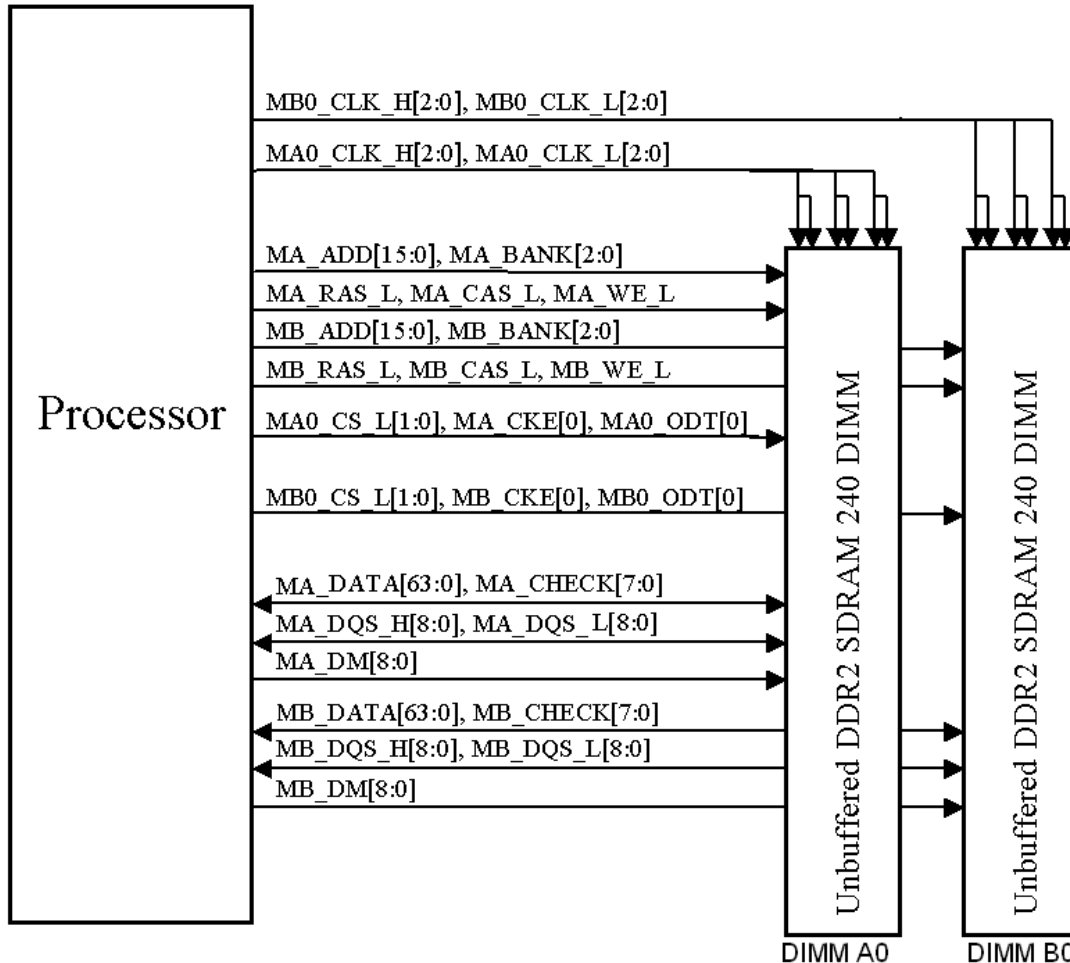


Figure 33. Unbuffered 2-DIMM Block Diagram

6.2.2 Valid DIMM Configurations – Four Unbuffered DIMMs

The socket AM2 processor can support up to four unbuffered DIMM slots. Refer to the *AMD BIOS and Kernel Developers Guide for AMD NPT Family 0Fh Processors*, order# 32559, for specific memory-size configurations.

6.2.3 Frequency

For the most up-to-date processor DIMM support documentation, refer to the *AMD BIOS and Kernel Developers Guide for AMD NPT Family 0Fh Processors*, order# 32559.

6.3 System Configuration Notes

DIMM population is always uniform from back to front. For details refer to the *AMD BIOS and Kernel Developers Guide for AMD NPT Family 0Fh Processors*, order# 32559.

- Population option for 64-bit mode is DIMM_A1.
- Population order for 128-bit mode is first DIMMs_A1/B1 and then DIMMs_A0/B0.

6.4 Unbuffered Four-DIMM DDR2 Layout Guidelines

The following sections describe the recommended termination, placement and routing, as well as the trace width and separation guidelines for the unbuffered DDR2 memory interface design. Figure 34 shows the general layout of a DDR2 bus with four unbuffered DIMM slots. It shows the recommended placement of termination for the clock, data, control, and address/command groups. Bulk, mid-frequency, and high-frequency decoupling capacitor placements are also shown.

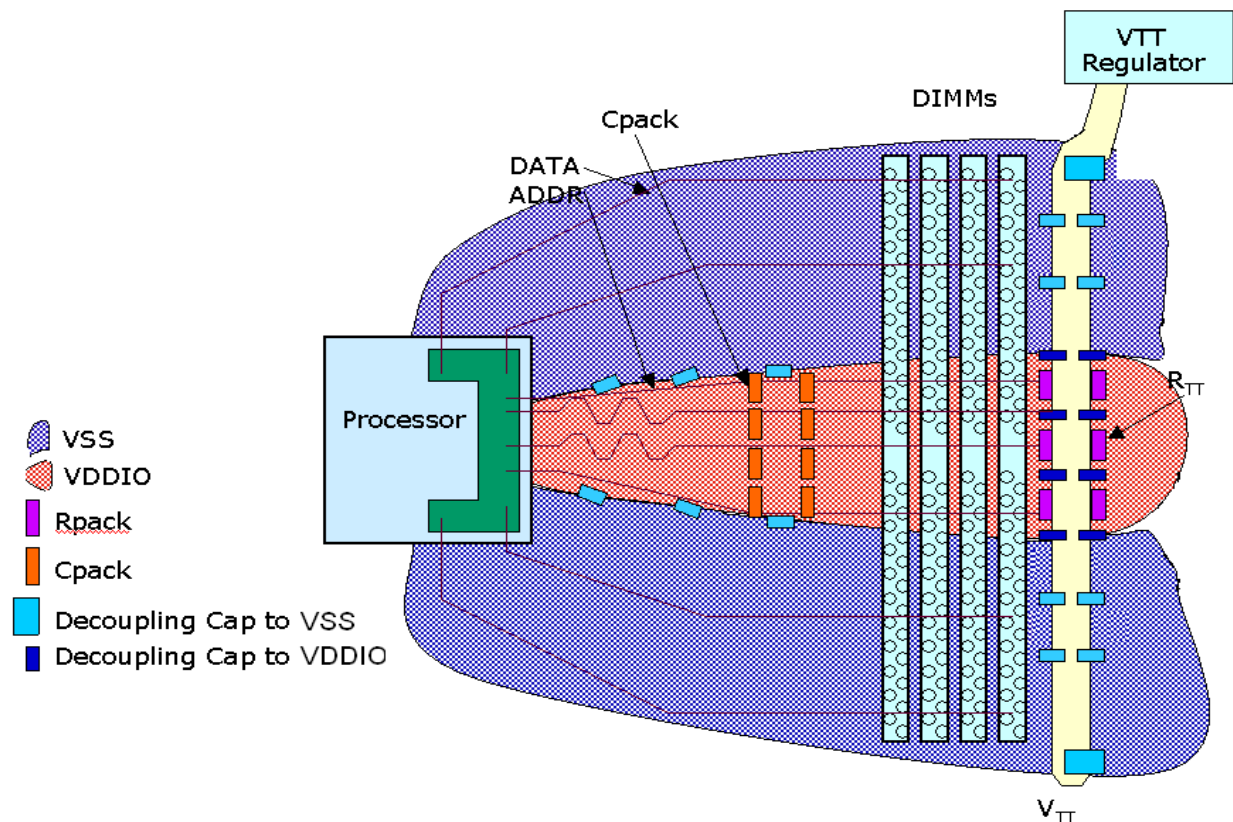


Figure 34. Unbuffered 4-DIMM General Layout Picture

The basic idea for DDR2 routing is to match all signal lengths of a particular DIMM to within a given tolerance. Both data and address/command signals are source-synchronous buses. Address/command signals propagate relative to clock signals, and data signals propagate relative to strobe signals. A fairly loose length-matching tolerance is required for address/command signals and all 64-bits of data signals to each DIMM. The tight length-matching tolerances are for data signals that are relative to strobe *within a data-group*. A step-by-step routing procedure is described in Section 6.5.1 on page 93.

AMD recommends that the two 64-bit halves of the 128-bit DDR2 memory interface be routed as a physical interleave to minimize lead-in length disparity. The designer is free to choose the DIMM spacing, but AMD recommends a spacing between 350 to 400 mils.

General routing rules are as follows:

1. Channel B should be routed on the top layer, and Channel A should be routed on the bottom layer.
2. Address/command group Cp termination is located near DIMM_A0.
3. Address/command group RTT termination is located near DIMM_B1.
4. To reduce cross talk from the pins to the nets, designers should avoid adding length to the nets under the processor as a means to achieve length matching. Additionally, nets should not be routed around processor pins. Refer to to Figure 35 for an illustration of what *not* to do.




Figure 35. Undesirable Net Routing

Figure 36 illustrates the lead-in length requirements for memory routing—from the processor to all the DIMMs.

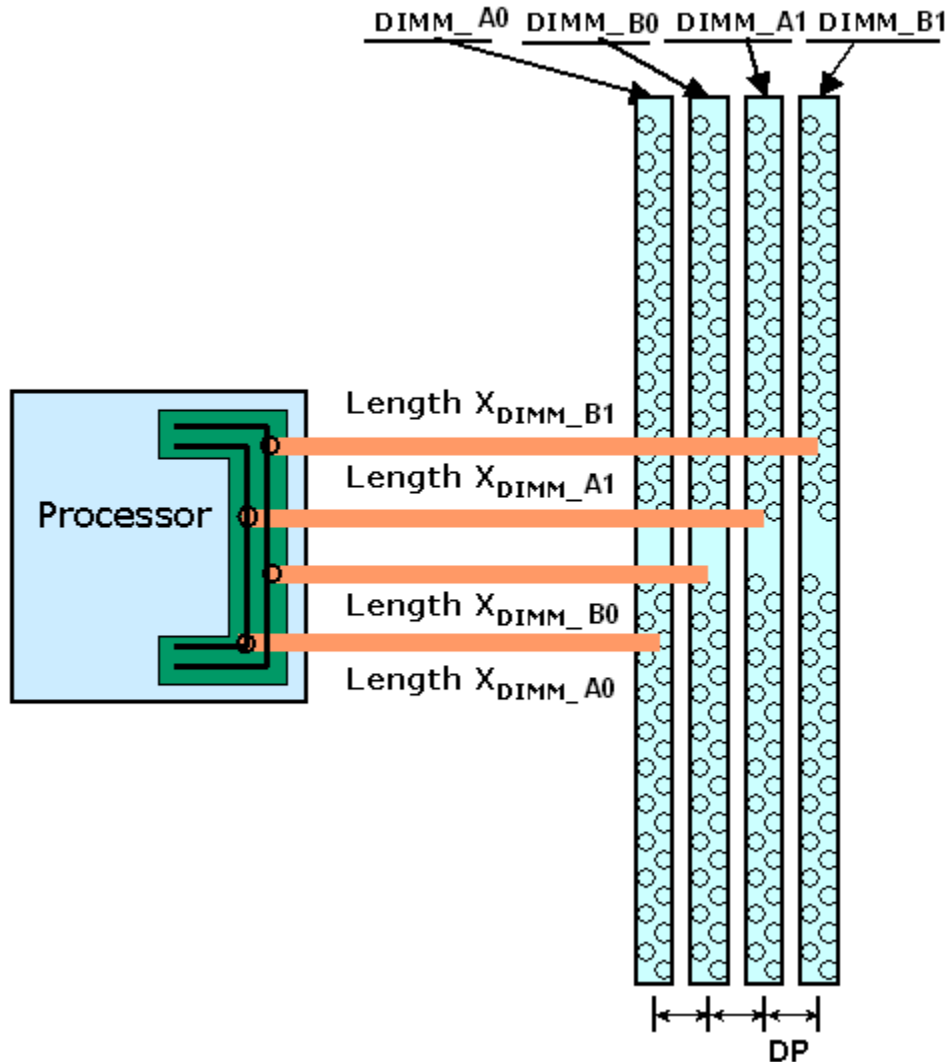


Figure 36. Lead-In Length for Memory Routing

DP is the DIMM Pitch, which refers to the center-to-center spacing between two adjacent memory DIMMs. Target lead-in lengths for each DIMM are set to help define the routing rules. These target lengths are represented by four variables— X_{DIMM_A0} , X_{DIMM_A1} , X_{DIMM_B0} , X_{DIMM_B1} . Every routing to each DIMM is defined relative to the target length to that DIMM.

To define the target lead-in length for each DIMM, use the following steps:

- Route the longest and shortest data nets to each DIMM.
- The average of the longest and shortest data net— $(\text{longest} + \text{shortest}) / 2$ —to each DIMM is the target lead-in length for that DIMM.
- For DIMM placement, DP is recommended to be between 350 to 400 mils.
- The target lead-in for DIMM_A0 and DIMM_B0 should not differ by more than DP.
- The target lead-in for DIMM_A1 and DIMM_B1 should not differ by more than DP.

6.4.1 Address/Command Layout Guidelines

The following sections provide Address/Command Layout Guidelines.

6.4.1.1 Address/Command and Control Termination

As seen in Figure 37 and Figure 38 on page 84, the address/command group terminations consist of both a capacitor C_p to VDDIO, which is located between the processor and the DIMM sockets, and a resistor R_{TT} at the end of the bus, which is pulled to VTT.

Recommended termination values are listed in Table 18.

Table 18. Termination Values for Unbuffered DIMM Address/Command Signals

Parameter	Value	Tolerance	Pullup Voltage
C_p	22 pF	$\pm 10\%$	VDDIO
R_{TT}	47 Ω	$\pm 5\%$	VTT

6.4.1.2 Address/Command Group Layout Guidelines

The following sections describe placement and routing, termination, as well as trace width and separation recommendations for the address/command group. There are two address/command channels. The “A” channel should be routed to DIMM_A0 and DIMM_A1, and the “B” channel to DIMM_B0, and DIMM_B1.

Figure 37 and Figure 38 on page 84 show the DDR2 address/command termination and layout recommendations for channel A and B.




Figure 37. Channel A Command and Address Routing




Figure 38. Channel B Command and Address Routing

6.4.1.3 CKE Guidelines

This section describes the termination recommendations for the MA_CKE[1:0] and MB_CKE[1:0] signals. There are two CKE signals per channel. Each CKE should be routed to one DIMM only and connected to two pins on the same DIMM.

Figure 39 shows the DDR2 CKE termination recommendations. RTT for CKE nets is 47 Ω .

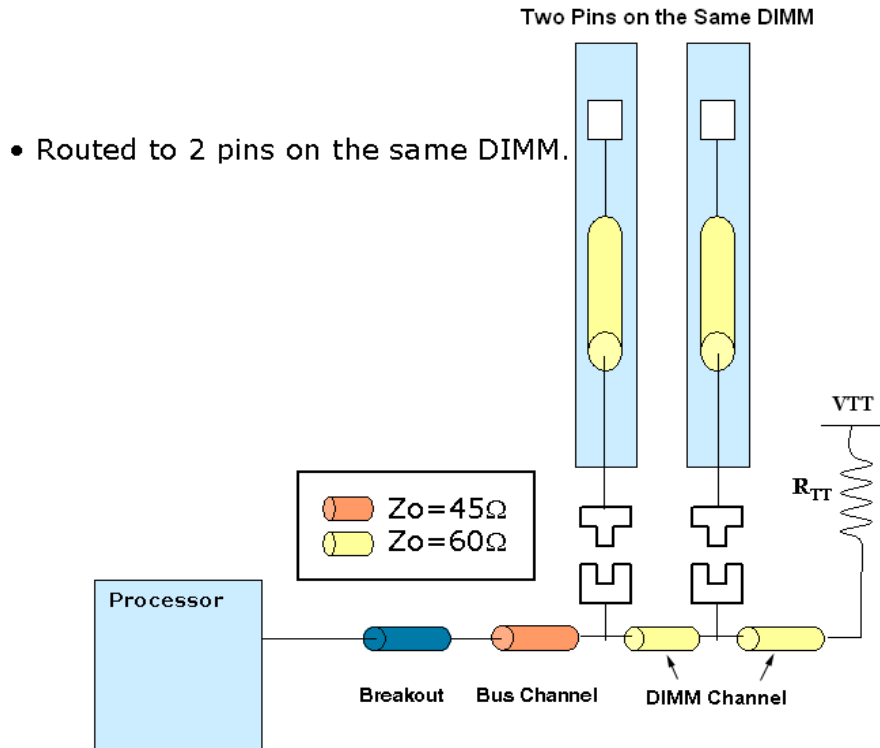


Figure 39. CKE Routing and Termination

6.4.1.4 CS and ODT Guidelines

This section describes termination recommendations for CS and ODT signals. Two CS signals and a single ODT signal are routed to each DIMM. There are a total of eight CS and four ODT signals. RTT termination for these signals is 47 Ω .

Figure 40 on page 86 shows the DDR2 CS and ODT termination for all DIMMs. Please note that each DIMM has a second ODT pin (ODT1) which should be connected to VSS.




Figure 40. CS and ODT Routing and Termination

6.4.1.5 Address/Command and Control Trace Width and Separation

Table 19 on page 87 defines routing parameters. In the bus channel, the impedance target for the address/command and control nets is $45\ \Omega$ with 10% tolerance, including manufacturing tolerances. For all other portions of the routing (excluding the breakout area) the impedance target is $60\ \Omega$ with 15% tolerance, including manufacturing tolerances.

Table 19. Parameters for Address/Command and Control Trace Width and Space

Definition	Breakout ^{3, 4, 5}	Bus Channel ⁶ (4.0-Mil Dielectric)	Bus Channel ⁷ (4.5-Mil Dielectric)	DIMM Channel	Units
Trace Width	4.0	8.0	10.0	4.0	mils
Trace Spacing ^{1, 2}	6.0	12.0	10.0	5.0	mils

Note:

1. If the net is serpentine, then the space to itself should be 20 mils.
2. In the DIMM channel, the maximum length of parallel nets with less than 6-mil spacing is 100 mils.
3. The length of the nets in the breakout area should be less than 500 mils.
4. If three nets are routed between two processor pins, then they may be routed as 4-mil traces with 4-mil spacing.
5. Maximize the space between the nets.
6. If the lead-in length of a trace is longer than 4 inches, then 4 mils should be added to the spacing to other nets in the bus channel.
7. If the lead-in length of a trace is longer than 4 inches, then 5 mils should be added to the spacing to other nets in the bus channel.

6.4.2 Data Group Guidelines

This section defines the data groups and describes their connection. There are sixteen data groups and two ECC groups, with each ECC group consisting of eight data bits and one differential data strobe signal and one datamask.

Ex: The ECC for data group “A” consists of (MA_CHECK[7:0], MA_DQS_H[8], MA_DQS_L[8], MA_DM[8]).

All data groups in a particular channel should be routed on the same layer. Figure 41 on page 88 illustrates routing and termination for the data group. Figure 42 on page 89 depicts routing for a DQS pair.

Table 20 on page 88 shows the DDR2 data groups for a single channel. Both channels have identical signals. When ganged in 128-bit mode, channel A serves as the lower bits and channel B as the upper bits.

Table 20. DDR2 Data Groups

Group	Data	Check	DQS	DM
0	[7:0]	–	0	0
1	[15:8]	–	1	1
2	[23:16]	–	2	2
3	[31:24]	–	3	3
4	[39:32]	–	4	4
5	[47:40]	–	5	5
6	[55:48]	–	6	6
7	[63:56]	–	7	7
8	–	[7:0]	8	8

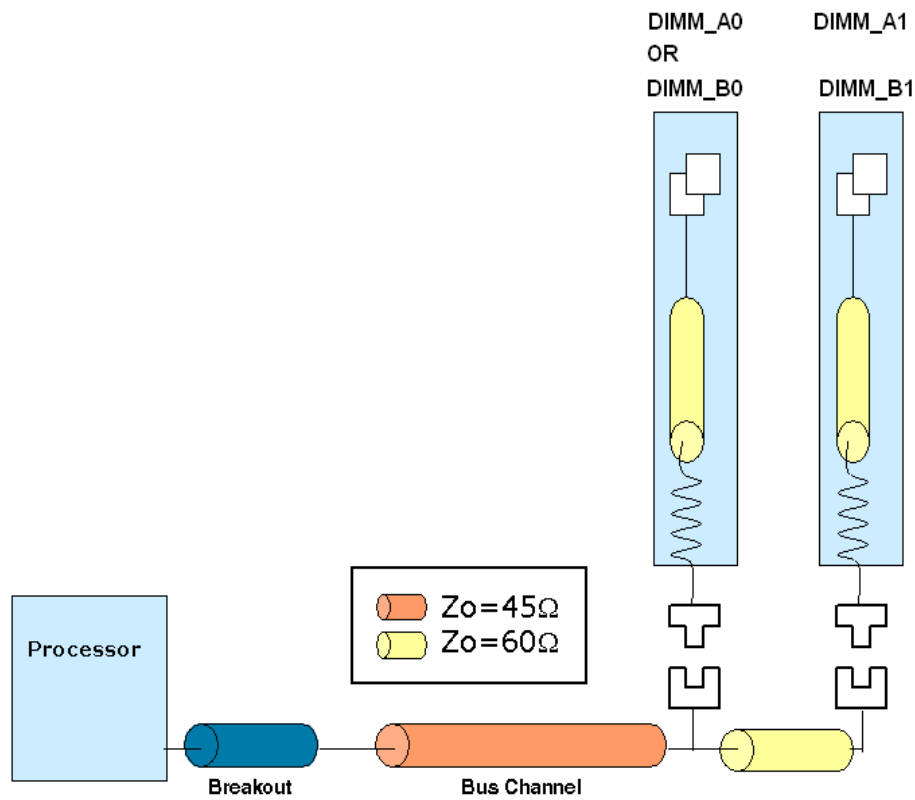


Figure 41. Data Group Routing and Termination




Figure 42. Routing for DQS Pair

6.4.2.1 Data Group Termination

No on-board termination is required.

6.4.2.2 Data Bus Trace Width and Separation

Table 21 on page 90 defines the routing parameters for DQS routing. In the bus channel the differential impedance target for DQS nets is $72\ \Omega$ with 10% tolerance, including manufacturing tolerances.

Table 21. Parameters for DQS Trace Width and Space

Definition	Breakout ^{3, 4, 5}	Bus Channel ⁶ (4.0-Mil Dielectric)	Bus Channel ⁷ (4.5-Mil Dielectric)	DIMM Channel	Units
Trace Width	4.0	8.0	10.0	4.0	mils
DQS Spacing to Other Nets ^{1, 2}	6.0	12.0	15.0	5.0	mils
Differential Spacing for DQS Nets	6.0	4.0	5.0	5.0	mils

Note:

1. If the net is serpentine, then the space to itself should be 20 mils.
2. In the DIMM channel, the maximum length of parallel nets with less than 6-mil spacing is 100 mils.
3. The length of the nets in the breakout area should be less than 500 mils.
4. If three nets are routed between two processor pins, then they may be routed as 4-mil traces with 4-mil spacing.
5. Maximize the space between the nets.
6. If the lead-in length of a trace is longer than 4 inches, then 4 mils should be added to the spacing to other nets in the bus channel.
7. If the lead-in length of a trace is longer than 4 inches, then 5 mils should be added to the spacing to other nets in the bus channel.

Table 22 defines the trace geometries for the data and check nets. In the bus channel, the impedance target for the data and check nets is 45 Ω with 10% tolerance, including manufacturing tolerances. For all other portions of the routing (excluding the breakout area) the impedance target is 60 Ω with 15% tolerance, including manufacturing tolerance.

Table 22. Parameters for Data and Check Trace Width and Space

Definition	Breakout ^{3, 4, 5}	Bus Channel ⁶ (4.0-Mil Dielectric)	Bus Channel ⁷ (4.5-Mil Dielectric)	DIMM Channel	Units
Trace Width	4.0	8.0	10.0	4.0	mils
Trace Spacing ^{1, 2}	6.0	12.0	10.0	5.0	mils

Note:

1. If the net is serpentine, then the space to itself should be 20 mils.
2. In the DIMM channel, the maximum length of parallel nets with less than 6 -mil spacing is 100 mils.
3. The length of the nets in the breakout area should be less than 500 mils.
4. If three nets are routed between two processor pins, then they may be routed as 4-mil traces with 4-mil spacing.
5. Maximize the space between nets.
6. If the lead-in length of a trace is longer than 4 inches, then 4 mils should be added to the spacing to other nets in the bus channel.
7. If the lead-in length of a trace is longer than 4 inches, then 5 mils should be added to the spacing to other nets in the bus channel.

6.4.3 Clock Layout Guidelines

This section describes the assignment, termination, placement and routing recommendations for the clock signals. Each DIMM socket receives three differential clock pairs driven from the processor.

Figure 43 on page 91 shows the unbuffered DDR2 clock configuration recommendations.

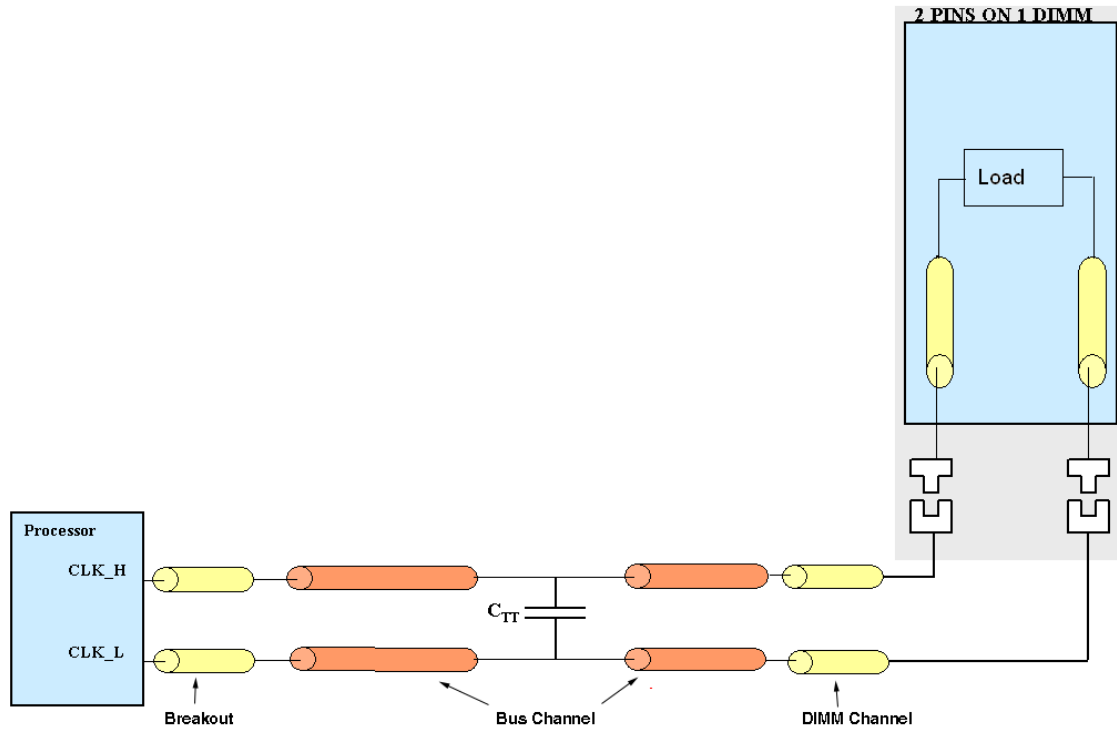


Figure 43. Clock Termination and Routing

6.4.3.1 DIMM Clock Assignment

A general depiction of overall CLK_H, CLK_L routing is shown in Figure 44.

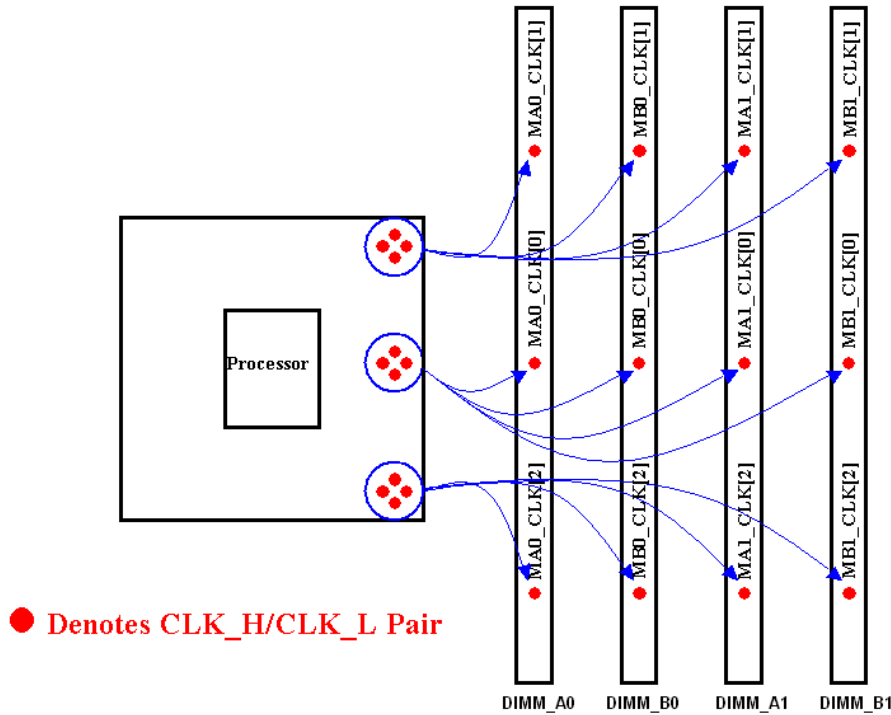


Figure 44. Clock Pair Assignment

6.4.3.2 Clock Group Termination

Differential C_{TT} termination is required for CLK_H and CLK_L. The maximum length for the C_{TT} stub is 150 mils (PCB trace + via height), with a PCB trace length of less than 100 mils. Table 23 lists the termination value for unbuffered DIMM clocks.

Table 23. CLK_H and CLK_L Termination

Parameter	Value	Tolerance
C_{TT}	1.5 pF	±10%

6.4.3.3 Clock Group Trace Width and Separation

Table 24 defines the CLK routing parameters. In the bus channel the target differential impedance for CLK nets is 72 Ω with 10% tolerance, including manufacturing tolerances.

Table 24. Parameters for CLK Trace Width and Space

Definition	Breakout ^{3, 4, 5}	Bus Channel ⁶ (4.0-Mil Dielectric)	Bus Channel ⁷ (4.5-Mil Dielectric)	DIMM Channel	Units
CLK Width	4.0	8.0	10.0	4.0	mils
CLK Spacing to Other Signals ^{1, 2}	6.0	16.0	20.0	5.0	mils
Differential Spacing for CLK Pairs	6.0	4.0	5.0	5.0	mils

Note:

1. If the net is serpentine, then the space to itself should be 20 mils.
2. In the DIMM channel, the maximum length of parallel nets with less than 6-mil spacing is 100 mils.
3. The length of the nets in the breakout area should be less than 500 mils.
4. If three nets are routed between two processor pins, then they may be routed as 4-mil traces with 4-mil spacing.
5. Maximize the space between the nets.
6. If the lead-in length of a trace is longer than 4 inches, then 4 mils should be added to the spacing to other nets in the bus channel.
7. If the lead-in length of a trace is longer than 4 inches, then 5 mils should be added to the spacing to other nets in the bus channel.

6.5 Basic Idea for Unbuffered DDR2 Memory Routing

The basic idea for DDR2 routing is to match all signal lengths, which go to a particular DIMM, to within a given tolerance. Both data and address/command are source-synchronous buses. Address/command propagates relative to clock, and data propagates relative to strobe. A fairly loose length-matching tolerance is required for address/command and all 64-bits of data to each DIMM. The tight length-matching tolerances are for data relative to strobe within a data-group. A step-by-step routing procedure is described in Sections 6.5.1 and 6.5.2.

6.5.1 Routing Procedure

Follow these steps in routing for DDR2 memory:

1. Route MB_DATA[0] and MB_DATA[63] to DIMM_B1. Choose the longest of the two nets and measure the lead-in length from the processor pin to the DIMM_B1 pin. Adjust the processor placement relative to the DIMMs until this length is shorter than 6.00 inches.

2. Route MA_DATA[0] and MA_DATA[63] to DIMM_A1. Choose the longest of the two nets and measure the lead-in length from the processor pin to the DIMM_A1 pin. Adjust the processor placement relative to the DIMMs until this length is shorter than 6.00 inches.
3. Route all the remaining nets to DIMM_B0. The shortest MB_DATA[63:0] net can be up to 2.00 inches shorter than the longest. The length of clock nets to DIMM_B0 should be fixed at the median distance between the longest and shortest net to DIMM_B0 and increased by a DIMM-pitch (DP) to subsequent DIMMs.
4. Route all the remaining nets to DIMM_A0. The shortest MA_DATA[63:0] net can be up to 2.00 inches shorter than the longest. The CLK_H and CLK_L lengths should be fixed at the median distance between the longest and shortest net to the DIMM.

6.5.2 Routing Rules

In Sections 6.5.2.1, 6.5.2.2, and 6.5.2.3 the design guide rules are reduced to only the most basic information.

6.5.2.1 CLK_H and CLK_L Signals

- True-to-complement differential pairs are length matched within 50 mils.
- The length of all CLK nets to each DIMM are matched within 100 mils tolerance.
- Net connecting C_{TT} termination capacitors to the processor pins should be shorter than 1200 mils.

6.5.2.2 Data, DQS, Check Signals

- All signals are referenced to the VSS plane and do not cross any plane splits.
- Channel A of the data bus (MA) is routed on the bottom layer.
- Channel B of the data bus (MB) is routed on the top layer.
- All the Data and DM nets in each data group—consisting of one DQS_H, one DQS_L, eight Data, and one DM signal—are length matched to within 100 mils of the average of the DQS net lengths for that Data group.
- All signals are length matched to within 1000 mils of all CLKs to a DIMM, which means that the longest and shortest Data, DQS, and DM nets cannot differ by more than 2000 mils.
- For DQS nets, True to Complement mismatch tolerance is 50 mils.
- For DIMMs A1 and B1, the longest Data, DQS, or DM signal must be shorter than 6.25 inches.

6.5.2.3 ADD, BANK, CS, CKE, RAS, CAS, WE, ODT Signals

- All signals are referenced to VDDIO and do not cross any power plane splits.
- RTT termination is placed behind the far DIMM (generally DIMM_B1).

- All nets for address group A (MA) are routed on the bottom layer.
- All nets for address group B (MB) are routed on the top layer.
- The shortest net connecting the Cp termination to DIMM_A0 or DIMM_B0 must be shorter than 1000 mils long.
- The length of all the nets connecting Cp termination to DIMM_A0 should be in the range of 600 mils (longest – shortest < 600 mils).
- The length of all the nets connecting Cp termination to DIMM_B0 should be in the range of 600 mils (longest – shortest < 600 mils).
- The command and address nets to each DIMM range from 500 to 1500 mils shorter than all CLKs to the same DIMM.
- The routing length from DIMM_A1 and DIMM_B1 to RTT is shorter than 1.90 inches.
- The routing length from DIMM_A1 and DIMM_B1 to RTT is matched to within 600 mils.

6.5.3 Unbuffered 4-DIMM DDR2 Routing Example

A general view of DDR2 component placement and routing may be seen on Figure 45.

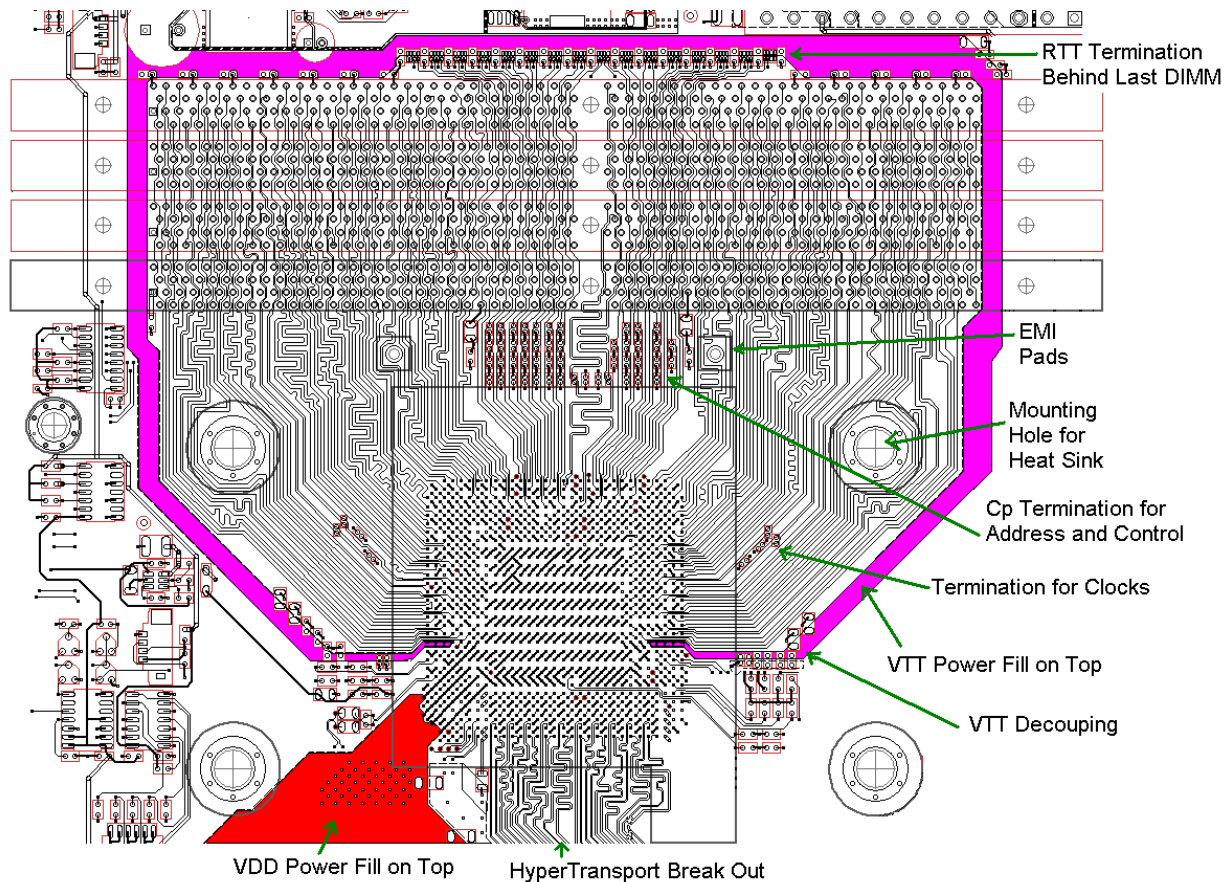


Figure 45. General DDR2 Component Placement and Routing for the Socket AM2 Processor

The processor pinout is designed to be routed with a 4-mil breakout. This configuration allows three nets to be routed between each of the two processor pads. Figure 46 shows an example of three nets that have been routed between two processor pads.




Figure 46. Example of Three Nets Routing Between the Two Processor Pads

Figure 47 shows a top-layer breakout from the processor and illustrates the recommended routing from the processor pins. This figure shows VDD power fill, VTT delivery, HyperTransport™ breakout, 4-mil nets used for processor breakout, and width change on the processor nets after breakout. After the nets breakout of the processor pin field, the net width changes from 4 mils to wide trace. The breakout length should be less than 500 mils.




Figure 47. Example of a Socket AM2 Processor Breakout—Topside

Figure 48 shows a bottom-layer breakout from the processor to illustrate the recommended routing from the processor pins.




Figure 48. Example of a Socket AM2 Processor Breakout—Bottomside

Figure 49 shows a top-layer routing from the processor to the DIMMs to illustrate the general form of the recommended routing.




Figure 49. Socket AM2 Processor Four DIMM Routing Example—Topside

Figure 50 shows a bottom-layer routing from the processor to the DIMMs to illustrate the general form of the recommended routing.




Figure 50. Socket AM2 Processor Four DIMM Routing Example—Bottomside

Chapter 7 BTX Form Factor Design Guidelines for Four Unbuffered DDR2 DIMMs

This chapter describes the bus, layout and routing rules, termination resistors, and stackup for unbuffered four-DIMM designs. This chapter also includes the figures for an implementation on a BTX form factor.

7.1 DDR2 SDRAM

The processor contains an integrated 128-bit dual-channel DDR2 SDRAM memory controller, which supports up to four, unbuffered DDR2 SDRAM DIMMs. All DIMMs must operate at the same frequency as the slowest DIMM.

7.1.1 DDR Routing on BTX Form Factor

For DDR2 signal definition, termination and routing there is no difference between a BTX and any other form factor. All the information on the following topics (see Chapter 6 for complete details) applies to the DDR interface for BTX form factors.

- Signals Names and Description (Section 6.1.1 on page 71)
- Layer Assignment (Section 6.1.2 on page 72)
- Interface Overview and Block Diagram (Section 6.2.1 on page 77)
- System Configuration Notes (Section 6.3 on page 80)
- Address/Command and Control Termination (Section 6.4.1.1 on page 83)

- DDR2 Data Group Definition (Table 20 on page 88)
- Data Group Termination (Section 6.4.2.1 on page 89)
- DIMM Clock Assignment (Section 6.4.3.1 on page 92)
- Clock Group Termination (Section 6.4.3.2 on page 92)

7.2 Unbuffered Four-DIMM DDR2 Layout Guidelines

The following sections describe placement and routing, as well as the trace width and separation guidelines for the unbuffered DDR2 memory interface design implemented on a BTX board. Layout is easier to implement if the plane splits are finalized before routing the signal layers. Figure 51 shows the general layout of a DDR2 bus with four unbuffered DIMM slots. It shows the recommended placement of termination for the clock, data, control, and address/command groups. Bulk, mid-frequency, and high-frequency decoupling capacitor placements are also shown.

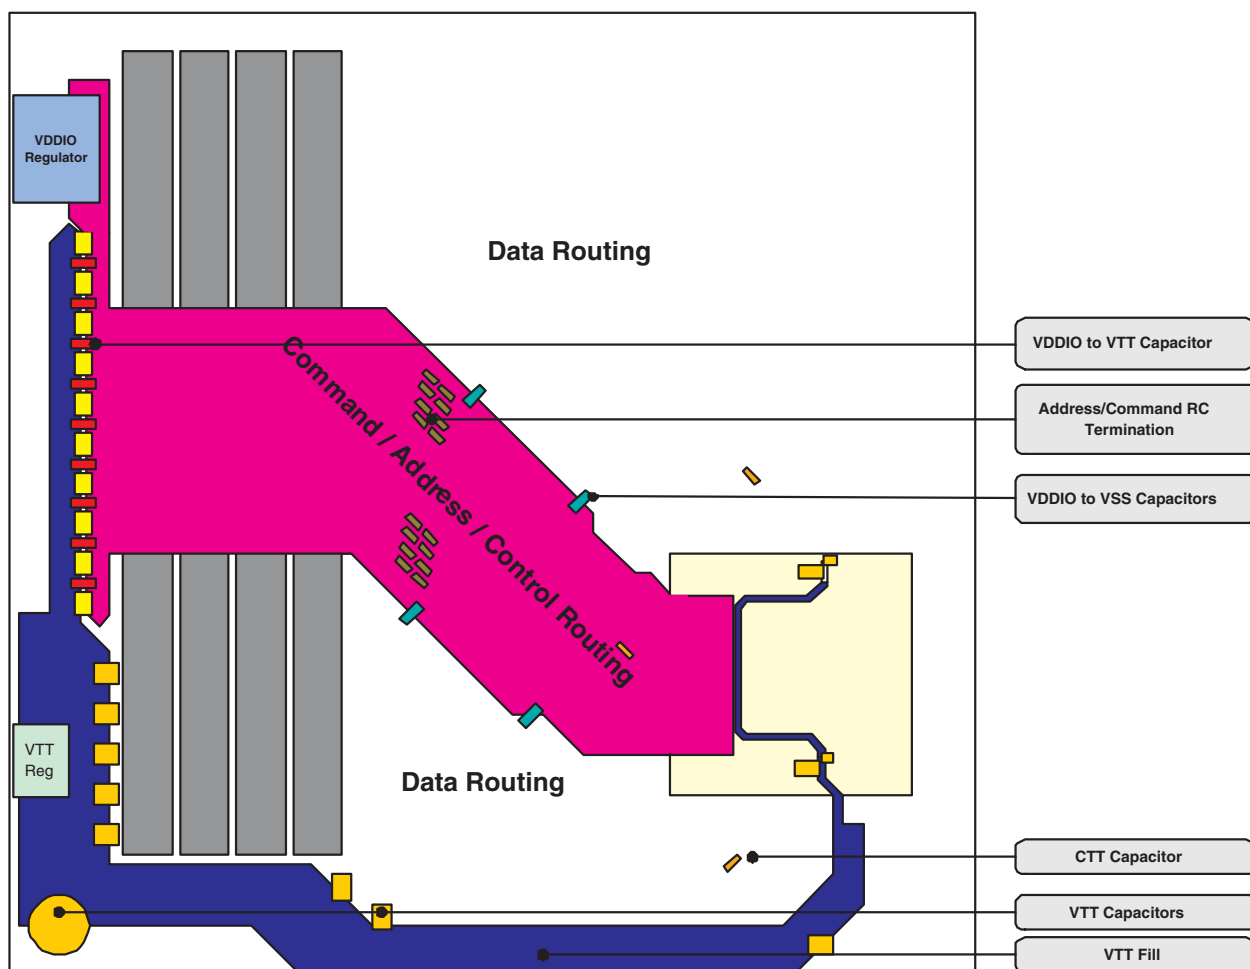


Figure 51. Unbuffered 4-DIMM General Layout Picture for BTX Boards

The basic goal for DDR2 routing is to match all signal lengths of a particular DIMM to within a given tolerance. Both data and address/command signals are source-synchronous buses. Address/command signals propagate relative to clock signals, and data signals propagate relative to strobe signals. A fairly loose length-matching tolerance is required for address/command signals and all 64-bits of data signals to each DIMM. The tight length-matching tolerances are for data signals that are relative to strobe *within a data-group*. A step-by-step routing procedure is described in Section 7.3.1 on page 117.

AMD recommends that the two 64-bit halves of the 128-bit DDR2 memory interface be routed as a physical interleave to minimize lead-in length disparity. The designer is free to choose the DIMM spacing, but AMD recommends a spacing between 350 to 400 mils.

General routing rules are as follows:

1. Channel B should be routed on the top layer, and Channel A should be routed on the bottom layer.
2. Address/command group C_p termination is located near DIMM_A0.
3. Address/command group R_{TT} termination is located near DIMM_B1.
4. To reduce cross talk from the pins to the nets, designers should avoid adding length to the nets under the processor as a means to achieve length matching. Additionally, nets should not be routed around processor pins. Refer to Figure 52 for an illustration of what *not* to do.




Figure 52. Undesirable Net Routing

Figure 53 illustrates the lead-in length requirements for memory routing—from the processor to all the DIMMs.

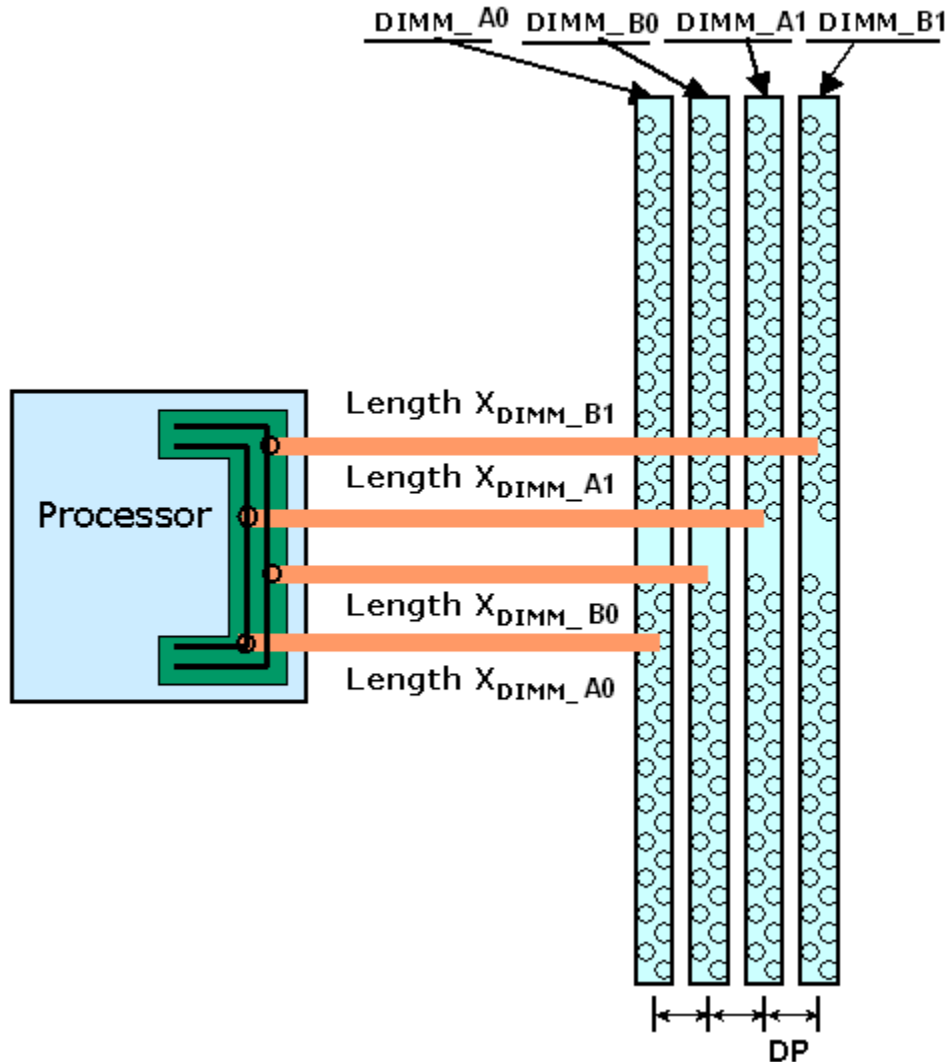


Figure 53. Lead-In Length for Memory Routing

DP is the DIMM Pitch, which refers to the center-to-center spacing between two adjacent memory DIMMs. Target lead-in lengths for each DIMM are set to help define the routing rules. These target lengths are represented by four variables— X_{DIMM_A0} , X_{DIMM_A1} , X_{DIMM_B0} , X_{DIMM_B1} . Every routing to each DIMM is defined relative to the target length to that DIMM.

To define the target lead-in length for each DIMM, use the following steps:

- Route the longest and shortest data nets to each DIMM.
- The average of the longest and shortest data net— $(\text{longest} + \text{shortest}) / 2$ —to each DIMM is the target lead-in length for that DIMM.
- For DIMM placement, DP is recommended to be between 350 to 400 mils.
- The target lead in for DIMM_A0 and DIMM_B0 should not differ by more than DP.
- The target lead in for DIMM_A1 and DIMM_B1 should not differ by more than DP

7.2.1 Address/Command Group Layout Guidelines

The following sections describe placement and routing, as well as trace width and separation recommendations for the address/command group. There are two address/command channels. The “A” channel should be routed to DIMM_A0 and DIMM_A1, and the “B” channel should be routed to DIMM_B0, and DIMM_B1.

Figure 54 and Figure 55 on page 108 show the DDR2 address/command termination and layout recommendations for channel A and B.




Figure 54. Channel A Command and Address Routing

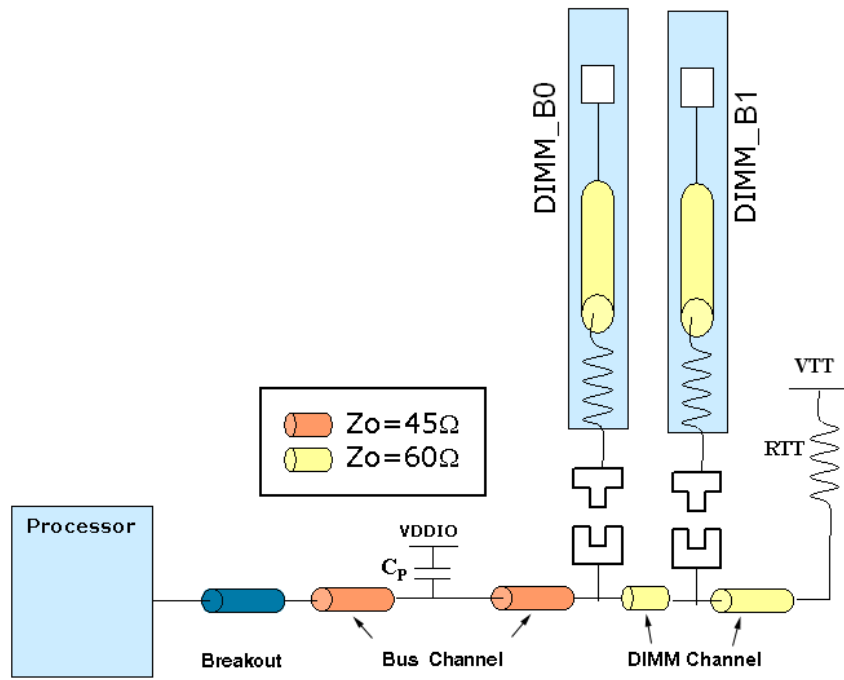


Figure 55. Channel B Command and Address Routing

7.2.1.1 CKE Guidelines

This section describes the termination recommendations for the MA_CKE[1:0] and MB_CKE[1:0] signals. There are two CKE signals per channel. Each CKE should be routed to one DIMM only and connected to two pins on the same DIMM.

Figure 56 shows the DDR2 CKE termination recommendations.

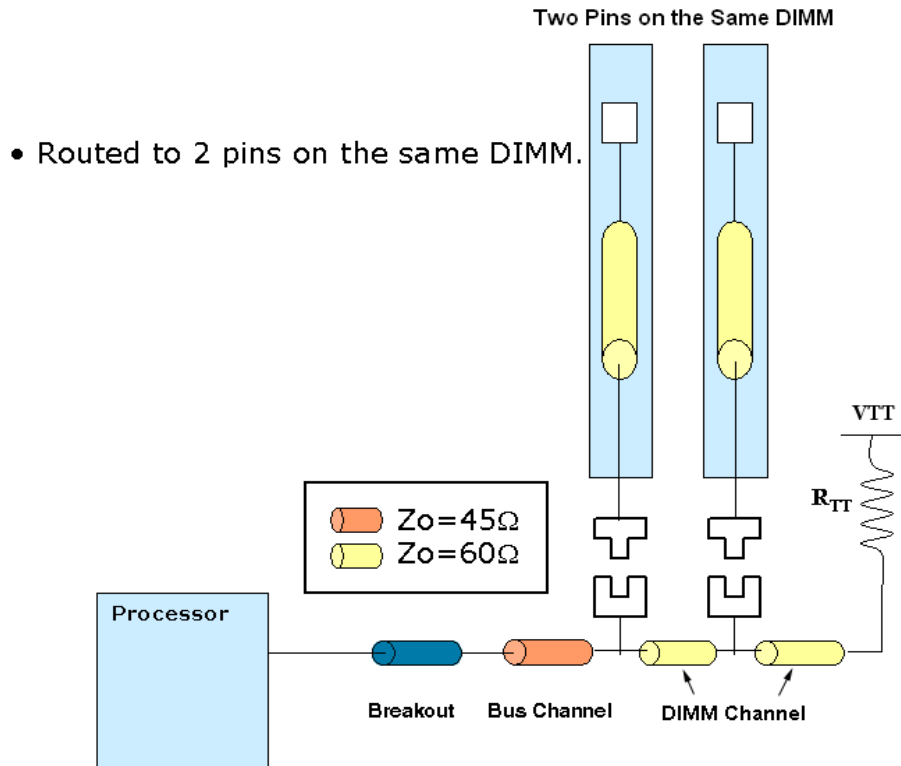


Figure 56. CKE Routing and Termination

7.2.1.2 CS and ODT Guidelines

Two CS signals and a single ODT signal are routed to each DIMM. There are a total of eight CS and four ODT signals.

Figure 57 shows the DDR2 CS and ODT termination for all DIMMs. Please note that each DIMM has a second ODT pin (ODT1) which should be connected to VSS.




Figure 57. CS and ODT Routing and Termination

7.2.1.3 Address/Command and Control Trace Width and Separation

Table 25 defines routing parameters. In the bus channel the impedance target for the address/command and control nets is 45 Ω with 10% tolerance, including manufacturing tolerances. For all other portions of the routing (excluding the breakout area) the impedance target is 60 Ω with 15% tolerance, including manufacturing tolerances.

Table 25. Parameters for Address/Command and Control Trace Width and Space

Definition	Breakout ^{3, 4, 5}	Bus Channel ⁶ (4.0-Mil Dielectric)	Bus Channel ⁷ (4.5-Mil Dielectric)	DIMM Channel	Units
Trace Width	4.0	8.0	10.0	4.0	mils
Trace Spacing ^{1, 2}	6.0	12.0	10.0	5.0	mils

Note:

1. If the net is serpentine, then the space to itself should be 20 mils.
2. In the DIMM channel, the maximum length of parallel nets with less than 6-mil spacing is 100 mils.
3. The length of the nets in the breakout area should be less than 500 mils.
4. If three nets are routed between two processor pins, then they may be routed as 4-mil traces with 4-mil spacing.
5. Maximize the space between the nets.
6. If the lead-in length of a trace is longer than 4 inches, then 4 mils should be added to the spacing to other nets in the bus channel.
7. If the lead-in length of a trace is longer than 4 inches, then 5 mils should be added to the spacing to other nets in the bus channel.

7.2.2 Data Group Guidelines

This section defines the data groups and describes their connection. There are sixteen data groups and two ECC groups, with each ECC group consisting of eight data bits, one differential data strobe signal, and one datamask.

Example: The ECC for data group A consists of (MA_CHECK[7:0], MA_DQS_H[8], MA_DQS_L[8], MA_DM[8]).

All data groups in a particular channel should be routed on the same layer. Figure 58 on page 112 illustrates routing and termination for the data group. Figure 59 on page 113 depicts routing for a DQS pair.

Table 26 on page 112 shows the DDR2 data groups for a single channel. Both channels have identical signals. When ganged in 128-bit mode, channel A serves as the lower bits and channel B as the upper bits.

Due to the offset position of the processor in relation to the DIMMs, implementing a BTX layout is easier without ECC routed on the board.

Table 26. DDR2 Data Groups

Group	Data	Check	DQS	DM
0	[7:0]	–	0	0
1	[15:8]	–	1	1
2	[23:16]	–	2	2
3	[31:24]	–	3	3
4	[39:32]	–	4	4
5	[47:40]	–	5	5
6	[55:48]	–	6	6
7	[63:56]	–	7	7
8	–	[7:0]	8	8

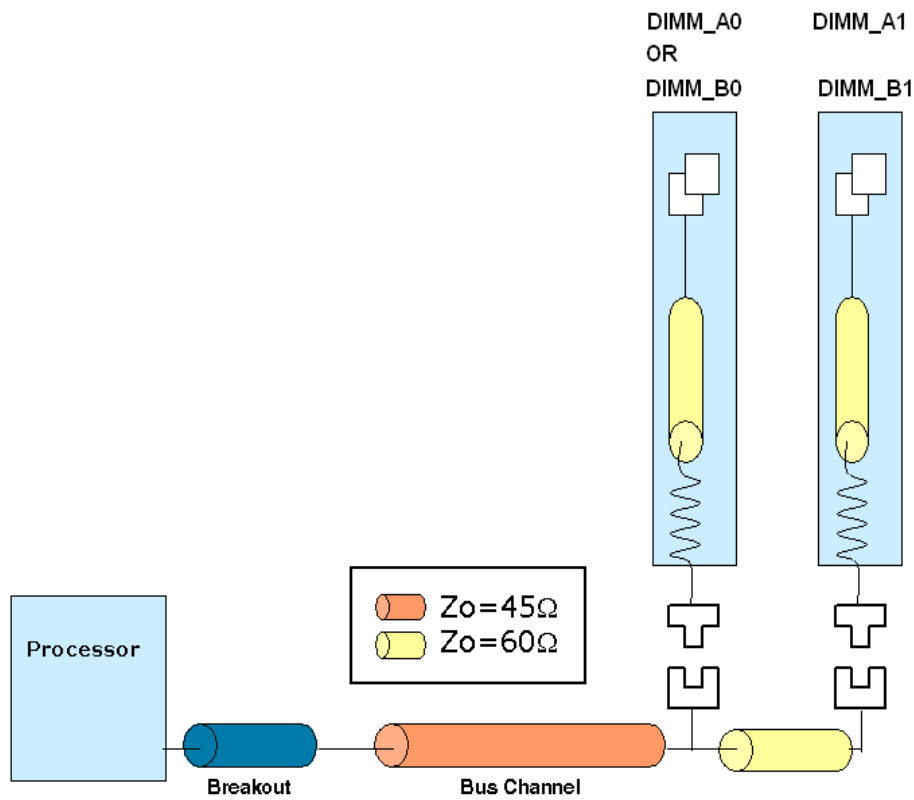


Figure 58. Data Group Routing and Termination




Figure 59. Routing for DQS Pair

7.2.2.1 Data Bus Trace Width and Separation

Table 27 defines the routing parameters for DQS routing. In the bus channel the differential impedance target for DQS nets is $72\ \Omega$ with 10% tolerance, including manufacturing tolerances.

Table 27. Parameters for DQS Trace Width and Space

Definition	Breakout ^{3, 4, 5}	Bus Channel ⁶ (4.0-Mil Dielectric)	Bus Channel ⁷ (4.5-Mil Dielectric)	DIMM Channel	Units
Trace Width	4.0	8.0	10.0	4.0	mils
DQS Spacing to Other Nets ^{1, 2}	6.0	12.0	15.0	5.0	mils
Differential Spacing for DQS Nets	6.0	4.0	5.0	5.0	mils

Note:

1. If the net is serpentine, then the space to itself should be 20 mils.
2. In the DIMM channel, the maximum length of parallel nets with less than 6-mil spacing is 100 mils.
3. The length of the nets in the breakout area should be less than 500 mils.
4. If three nets are routed between two processor pins, then they may be routed as 4-mil traces with 4-mil spacing.
5. Maximize the space between the nets.
6. If the lead-in length of a trace is longer than 4 inches, then 4 mils should be added to the spacing to other nets in the bus channel.
7. If the lead-in length of a trace is longer than 4 inches, then 5 mils should be added to the spacing to other nets in the bus channel.

Table 28 defines the trace geometries for the data and check nets. In the bus channel the impedance target for the data and check nets is $45\ \Omega$ with 10% tolerance, including manufacturing tolerances. For all other portions of the routing (excluding the breakout area) the impedance target is $60\ \Omega$ with 15% tolerance, including manufacturing tolerances.

Table 28. Parameters for Data and Check Trace Width and Space

Definition	Breakout ^{3, 4, 5}	Bus Channel ⁶ (4.0-Mil Dielectric)	Bus Channel ⁷ (4.5-Mil Dielectric)	DIMM Channel	Units
Trace Width	4.0	8.0	10.0	4.0	mils
Trace Spacing ^{1, 2}	6.0	12.0	10.0	5.0	mils
Note:					
<ol style="list-style-type: none"> 1. If the net is serpentine, then the space to itself should be 20 mils. 2. In the DIMM channel, the maximum length of parallel nets with less than 6 -mil spacing is 100 mils. 3. The length of the nets in the breakout area should be less than 500 mils. 4. If three nets are routed between two processor pins, then they may be routed as 4-mil traces with 4-mil spacing. 5. Maximize the space between nets. 6. If the lead-in length of a trace is longer than 4 inches, then 4 mils should be added to the spacing to other nets in the bus channel. 7. If the lead-in length of a trace is longer than 4 inches, then 5 mils should be added to the spacing to other nets in the bus channel. 					

7.2.3 Clock Layout Guidelines

This section describes routing recommendations for the clock signals. Each DIMM socket receives three differential clock pairs driven from the processor.

Figure 60 on page 116 shows the unbuffered DDR2 clock configuration recommendations.

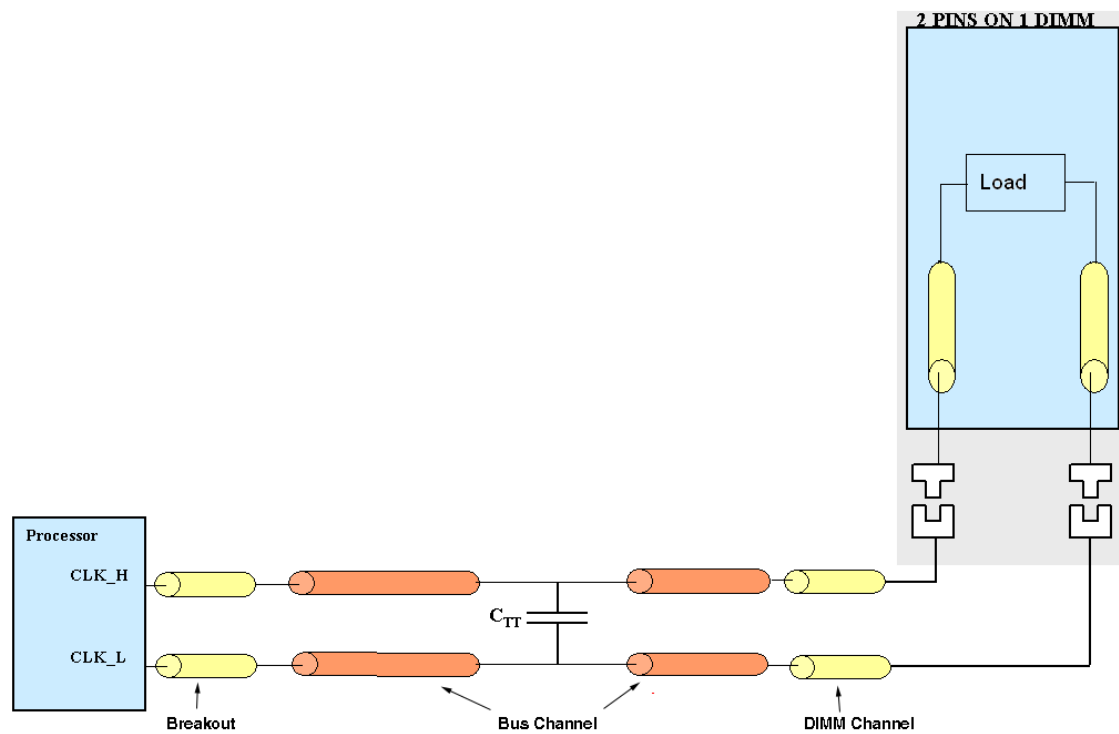


Figure 60. Clock Termination and Routing

7.2.3.1 Clock Group Trace Width and Separation

Table 29 defines the CLK routing parameters. In the bus channel the target differential impedance for CLK nets is 72 Ω with 10% tolerance, including manufacturing tolerances.

Table 29. Parameters for CLK Trace Width and Space

Definition	Breakout ^{4, 5, 6}	Bus Channel ⁷ (4.0-Mil Dielectric)	Bus Channel ⁸ (4.5-Mil Dielectric)	DIMM Channel	Units
CLK Width	4.0	8.0	10.0	4.0	mils
CLK Spacing to Other Signals ^{1, 2, 3}	6.0	16.0	20.0	5.0	mils
Differential Spacing for CLK Pairs	6.0	4.0	5.0	5.0	mils

Note:

1. If the net is serpentine, then the space to itself should be 20 mils.
2. In the DIMM channel, the maximum length of parallel nets with less than 6-mil spacing is 100 mils.
3. In the bus channel, the spacing from CLK to other traces may be reduced from the designated spacing to 5 mils. The aggregate length of the altered trace must be less than 100 mils.
4. The length of the nets in the breakout area should be less than 500 mils.
5. If three nets are routed between two processor pins, then they may be routed as 4-mil traces with 4-mil spacing.
6. Maximize the space between the nets.
7. If the lead-in length of a trace is longer than 4 inches, then 4 mils should be added to the spacing to other nets in the bus channel.
8. If the lead-in length of a trace is longer than 4 inches, then 5 mils should be added to the spacing to other nets in the bus channel.

7.3 Basic Idea for Unbuffered DDR2 Memory Routing

The basic idea for DDR2 routing is to match all signal lengths, which go to a particular DIMM, to within a given tolerance. Both data and address/command are source-synchronous buses. Address/command propagates relative to clock, and data propagates relative to strobe. A fairly loose length-matching tolerance is required for address/command and all 64-bits of data to each DIMM. The tight length-matching tolerances are for data relative to strobe within a data-group. A step-by-step routing procedure is described in Sections 7.3.1 and 7.3.2.

7.3.1 Routing Procedure

Follow these steps in routing for DDR2 memory:

1. Route MB_DATA[0] and MB_DATA[63] to DIMM_B1. Choose the longest of the two nets and measure the lead-in length from the processor pin to the DIMM_B1 pin. Adjust the processor placement relative to the DIMMs until this length is shorter than 6.00 inches.

2. Route MA_DATA[0] and MA_DATA[63] to DIMM_A1. Choose the longest of the two nets and measure the lead-in length from the processor pin to the DIMM_A1 pin. Adjust the processor placement relative to the DIMMs until this length is shorter than 6.00 inches.
3. Route all the remaining nets to DIMM_B0. The shortest MB_DATA[63:0] net can be up to 2.00 inches shorter than the longest. The length of clock nets to DIMM_B0 should be fixed at the median distance between the longest and shortest net to DIMM_B0 and increased by a DIMM-pitch (DP) to subsequent DIMMs.
4. Route all the remaining nets to DIMM_A0. The shortest MA_DATA[63:0] net can be up to 2.00 inches shorter than the longest. The CLK_H and CLK_L lengths should be fixed at the median distance between the longest and shortest net to the DIMM.

7.3.2 Routing Rules

In Sections 7.3.2.1, 7.3.2.2, and 7.3.2.3, the design guide rules are reduced to only the most basic information.

7.3.2.1 CLK_H and CLK_L Signals

- True-to-complement differential pairs are length matched to within 50 mils.
- The length of all CLK nets to each DIMM are matched to within 100 mils tolerance.
- Net connecting C_{TT} termination capacitors to the processor pins should be shorter than 1200 mils.

7.3.2.2 Data, DQS, Check Signals

- All signals are referenced to the VSS plane and do not cross any plane splits.
- Channel A of the data bus (MA) is routed on the bottom layer.
- Channel B of the data bus (MB) is routed on the top layer.
- All the Data and DM nets in each data group—consisting of one DQS_H, one DQS_L, eight Data, and one DM signal—are length matched to within 100 mils of the average of the DQS net lengths for that Data group.
- All signals are length matched to within 1000 mils of all CLKs to a DIMM, which means that the longest and the shortest Data, DQS, and DM nets cannot differ by more than 2000 mils.
- For DQS nets, true to complement mismatch tolerance is 50 mils.
- For DIMMs A1 and B1, the longest Data, DQS, or DM signal must be shorter than 6.25 inches.

7.3.2.3 ADD, BANK, CS, CKE, RAS, CAS, WE, and ODT Signals

- All signals are referenced to VDDIO and do not cross any power plane splits.
- R_{TT} termination is placed behind the far DIMM (generally DIMM_B1).

- All nets for address group A (MA) are routed on the bottom layer.
- All nets for address group B (MB) are routed on the top layer.
- The shortest net connecting the Cp termination to DIMM_A0 or DIMM_B0 must be shorter than 1000 mils.
- The length of all the nets connecting Cp termination to DIMM_A0 should be in the range of 600 mils (longest – shortest < 600 mils).
- The length of all the nets connecting Cp termination to DIMM_B0 should be in the range of 600 mils (longest – shortest < 600 mils).
- The command and address nets to each DIMM range from 500 to 1500 mils shorter than all CLKs to the same DIMM.
- The routing length from DIMM_A1 and DIMM_B1 to R_{TT} is shorter than 1.90 inches.
- The routing length from DIMM_A1 and DIMM_B1 to R_{TT} is matched to within 600 mils.

7.3.3 Unbuffered 4-DIMM DDR2 Routing Examples on a BTX Board

A general view of DDR2 component placement and routing on a BTX board may be seen in Figure 61.




Figure 61. General DDR2 Component Placement and Routing for a BTX Form Factor

The processor pinout is designed to be routed with a 4-mil breakout. This configuration allows three nets to be routed between each of the two processor pads. Figure 62 on page 121 shows an example of three nets that have been routed between two processor pads.




Figure 62. Example of Three Nets Routing Between the Two Processor Pads

Figure 63 shows a top-layer breakout from the processor and illustrates the recommended routing from the processor pins. This figure shows the VDD power fill, VTT delivery, HyperTransport™ breakout, the 4-mil nets used for processor breakout, and the width change on the processor nets after breakout. After the nets break out of the processor pin field, the net width changes from 4 mils to wide trace. The breakout length should be less than 500 mils.




Figure 63. Example of a Socket AM2 Processor Breakout on a BTX Board—Topside

Figure 64 shows a bottom-layer breakout from the processor to illustrate the recommended routing from the processor pins.





Figure 64. Example of a Socket AM2 Processor Breakout on a BTX Board—Bottomside

Figure 65 shows a top-layer routing from the processor to the DIMMs to illustrate the general form of the recommended routing.



**Figure 65. Socket AM2 Processor Four DIMM Routing Example on a BTX Board—
Topside**

Figure 66 shows a bottom-layer routing from the processor to the DIMMs to illustrate the general form of the recommended routing.




Figure 66. Socket AM2 Processor Four DIMM Routing Example on a BTX Board—Bottomside

Chapter 8 Design Guidelines for Miscellaneous Signals

This chapter describes how to properly design the circuits that accept the various processor inputs and outputs.

Table 30 on page 128 is a quick reference guide for these signals.

The signal groups are defined as follows:

- Clock Related Signals
- Power Related Signals
- HDT/JTAG Related Signals
- Driver Strength Signals
- Global Signals
- Thermal Related Signals
- No Connect and Test Point Signals
- Pullup, Pulldown, and Required Terminations

System voltage terminology is discussed in Section 2.3.1 on page 22.

Table 30. Miscellaneous Signals Quick Reference

	Signal Name	Type	Connection	Termination	Voltage
Clock Related Signals	CLKIN_H	I	Routed to clock generator		
	CLKIN_L	I	Routed to clock generator		
Power Related Signals	VID [5:0] ^{1, 2, 3}	O	Routed to VDD regulator	300 Ω on VID[1]	VDDIO
	PSI_L	O	Routed to power supply control circuit		
	VDD_FB_H	A	Routed to VDD regulator		
	VDD_FB_L	A	Routed to VDD regulator		
	VDDIO_FB_H	A	Routed to VDDIO regulator		
	VDDIO_FB_L	A	Routed to VDDIO regulator		
	VTT_SENSE	A	Routed to VTT regulator		
	CPU_PRESENT_L	A	Pulled up to a voltage	Onboard pullup	
HDT Related Signals	TMS	I	Routed to HDT header	Internal pullup	
	TCK	I	Routed to HDT header	Internal pullup	
	TRST_L	I	Routed to HDT header	Internal pullup	
	TDI	I	Routed to HDT header	Internal pullup	
	TDO	O	Routed to HDT header	Internal pullup	
	DBRDY	O	Routed to HDT header	Internal pullup	
	DBREQ_L	I	Routed to HDT header	Internal pullup	
Driver Strength Signals	HTREF0	A	Pulled down	44.2 Ω	VSS
	HTREF1	A	Pulled up	44.2 Ω	VLDT
	M_ZN	A	Pulled up	39.2 Ω	VDDIO
	M_ZP	A	Pulled down	39.2 Ω	VSS
Global Signals	PWROK ^{1, 4}	I	Routed to PWROK circuit		
	RESET_L ^{1, 4}	I	Routed to I/O Hub		
	LDTSTOP_L ^{1, 4}	I	Routed to I/O Hub		

Notes:

1. These pins may need a voltage translation circuit depending on the voltage level of the source or destination. Refer to the data sheet for edge-rate requirements.
2. Only VID[1] needs pullup to VDDIO.
3. If a 5-bit regulator is used, then only VID[4:0] should be routed to the regulator.
4. PWROK, LDTSTOP_L, and RESET_L need 300- Ω pullups on the motherboard if driven by Open Drain driver.
5. TEST29_L and TEST29_H should be routed as a differential pair with 80- Ω differential impedance to an 80.6- Ω termination resistor.

Table 30. Miscellaneous Signals Quick Reference (Continued)

	Signal Name	Type	Connection	Termination	Voltage
Thermal Related Signals	PROCHOT_L ¹	I/O	Route to a circuit for thermal monitoring and pullup	300 Ω	VDDIO
	THERMTRIP_L ¹	O	Routed to I/O Hub and pullup	300 Ω	VDDIO
	SIC ¹	I	Strap to VDDIO or VSS but not to both	300 Ω	VDDIO
	SID ¹	I/O	Can be left unconnected		
	THERMDA	A	Routed to temperature monitoring device		
	THERMDC	A	Routed to temperature monitoring device		
No Connect Pins	TEST2, TEST3, TEST6, TEST7, TEST8, TEST10, TEST26, TEST27, TEST28_H, TEST28_L		No Connect on Board		
Test Points	TEST12, TEST14, TEST15, TEST16, TEST17, TEST18, TEST19, TEST20, TEST21, TEST22, TEST23, TEST24		Need Test Points on Board		
Pullups	TEST25_H	I	Pulled Up	510 Ω	VDDIO
Pulldowns	TEST25_L	I	Pulled Down	510 Ω	VSS
Tie to VSS	TEST9, TEST13	I	Tie to VSS		VSS
Differential Termination	TEST29_L, TEST29_H ⁵		80.6 Ω differential termination between TEST29_L and TEST29_H	80.6 Ω	
Reserved pins	RSVD		No Connect on Board		

Notes:

1. These pins may need a voltage translation circuit depending on the voltage level of the source or destination. Refer to the data sheet for edge-rate requirements.
2. Only VID[1] needs pullup to VDDIO.
3. If a 5-bit regulator is used, then only VID[4:0] should be routed to the regulator.
4. PWROK, LDTSTOP_L, and RESET_L need 300- Ω pullups on the motherboard if driven by Open Drain driver.
5. TEST29_L and TEST29_H should be routed as a differential pair with 80- Ω differential impedance to an 80.6- Ω termination resistor.

8.1 CLKIN_H and CLKIN_L Differential Input

Refer to Section 4.1.1.1 on page 27 for more information on input clock termination.

8.2 Auto-Compensation

It is crucial to maintain a matched system to prevent signal reflections. To achieve a matched system, the impedance of the driver, the motherboard trace, and the receiver must be impedance matched. This task is accomplished by utilizing auto-compensated drivers. Auto-compensation is a feedback system contained within the control logic for the bus. This feedback constantly adjusts the system to maintain constant impedance. This is accomplished by combining variable arrays of transistors that produce different resistive networks. The equivalent impedance of the array is constantly monitored and adjusted to achieve the desired value. This method provides correction for process, temperature, and voltage variations.

8.2.1 HyperTransport™ I/O Compensation

The HyperTransport™ links utilize auto-compensation to implement the transmission termination (R_{ON}) and the receiver termination (RTT). HTREF0 is one of the two HyperTransport technology driver compensation inputs. This signal should be pulled down to VSS with a 44.2- Ω ($\pm 1\%$) resistor. HTREF1 is the second HyperTransport technology driver compensation input. This signal should be pulled up to VLDT with a 44.2- Ω ($\pm 1\%$) resistor. The resistors should be placed within 1500 mils of the processor pin.

These signals should be routed as indicated in Table 31.

Table 31. Compensation Resistor Routing Rules for HyperTransport™ Technology

Length (in)	Space (mil)	Breakout Width (mil)	Width (mil)
Less than 1.5	10	4	5

8.2.2 Memory I/O Compensation

The drive strength of the processor memory controller drivers is configurable through the M_ZN and M_ZP pins on the processor. For details regarding configuration, refer to the *AMD BIOS and Kernel Developers Guide for AMD NPT Family 0Fh Processors*, order# 32559.

These signals should be routed as indicated in Table 32.

Table 32. Compensation Resistor Routing Rules for M_ZP and M_ZN

Length (in)	Space (mil)	Breakout Width (mil)	Width (mil)
Less than 1.5	10	4	5

8.3 Hardware Debug Tool

The AMD Hardware Debug Tool (HDT) feature is a tool system designers can use to control and monitor the internal workings of the processor. This tool aids in debugging the vast network of interconnects between the integrated circuits. HDT is accessed with the standard Test Access Port (TAP) feature of the JTAG interface. The socket AM2 processor supports both standard Joint Test Action Group (JTAG) instructions as well as AMD proprietary commands specific for monitoring and controlling the processor. AMD recommends the SAMTEC ASP-68200 family of connectors for HDT connectors.

AMD strongly recommends including HDT connectors on the top layer of all motherboard designs. These connectors provide the ability to debug the system and BIOS. This ability is necessary if AMD is to consult on specific motherboard-processor compatibility issues. If these connectors are not included, it is recommended that the signals be routed to pads on the motherboard, so that the connectors may be added later if necessary.

Additionally, a warm reset switch physically mounted to the motherboard must be provided to enable AMD debug capabilities. In production, this switch could be no-popped, but the site must never be removed. This switch functionality cannot be wired to a chassis-mounted switch. A warm reset is defined as the condition when RESET_L is asserted at the processor while PWROK continues to be asserted (RESET_L goes low and PWROK stays high). A two pin header in parallel to the warm reset switch is also required. This two pin header is used in AMD's automated validation tool. The absence of this two pin header will slow down the debug and validation process.

8.3.1 Single-Processor HDT

The SAMTEC ASP-68200-07 connector is keyed specifically to be used as a single-processor HDT header. It is *highly* recommended that the HDT signals be routed to surface pads even if the connector is not being installed. This preparation allows the connectors to be added later, if necessary for board debugging purposes. Refer to Figure 67 on page 132 for HDT connector and pad illustrations. Refer to Table 33 on page 132 for HDT connector pin-out information.

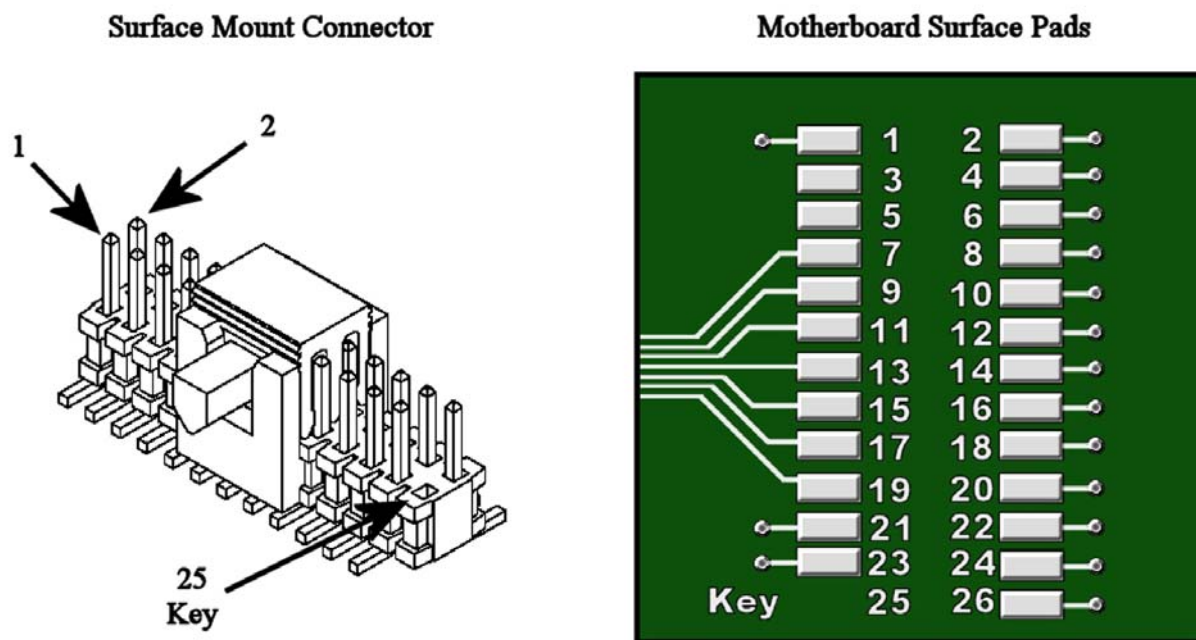


Figure 67. HDT Header

Table 33. HDT Header Pin-Out

HDT Header Connections				
Connector Side 1			Connector Side 2	
Pin#	Signal		Pin#	Signal
1	GND		2	GND
3	Reserved		4	GND
5	Reserved		6	GND
7	DBREQ_L		8	GND
9	DBRDY		10	GND
11	TCK		12	GND
13	TMS		14	GND
15	TDI		16	GND
17	TRST_L		18	GND
19	TDO		20	GND
21	VDDIO		22	GND
23	VDDIO		24	RESET_L
25	Key		26	GND

8.3.2 TMS

The TMS input is the test-mode-select input to the JTAG controller. TMS is used to enable test access to port controller operations. This signal should be routed to pin 13 of the HDT header.

8.3.3 TCK

The TCK input is the test clock input to the JTAG controller. TCK is used as the main clock for the JTAG control logic and is externally driven. This signal should be routed to pin 11 of the HDT header.

8.3.4 TRST_L

The TRST_L input is the test access port reset input to the JTAG controller. This active-Low signal initializes the test access port controller when asserted. This signal should be routed to pin 17 of the HDT header.

8.3.5 TDI

The TDI input is the test data input to the JTAG controller. TDI is used as a serial input line to write to the HDT serial shift register. This signal should be routed to pin 15 of the HDT header.

8.3.6 TDO

The TDO output is the test data output. TDO is a serial push-pull output from the HDT serial shift register and provides a means of viewing data from within the processor. Under operation, this signal remains tristated except when in the Shift-DR or Shift-IR controller states. This signal should be routed to pin 19 of the HDT header.

8.3.7 DBREQ_L

The DBREQ_L input is the debug request input to the JTAG controller. This active-Low input is used to synchronize events between external JTAG signals and processor signals. This synchronization is necessary because the two devices are driven by different clock sources. This signal should be routed to pin 7 of the HDT header. If this signal is not routed, then it has an internal pullup.

8.3.8 DBRDY

The DBRDY output is the debug ready output signal. This signal indicates that the JTAG internal logic registers are synchronized and ready to accept input or drive output. This push-pull signal should be routed to pin 9 of the HDT header.

8.3.9 RESET_L

The RESET_L is the processor's reset and would be monitored by HDT tools. RESET_L should be level shifted to 3.3 V and routed to pin 24 of the HDT header.

8.4 Voltage Regulator Signals

The socket AM2 processor has several pins dedicated to sensing and controlling the various voltage sources. Required voltages are VDD, VDDIO, and VTT. The basic control requirements for these supplies are described in the following subsections.

8.4.1 VID [5:0]

VDD is the source that provides power for the core logic circuitry within the processor. The target voltage level of this source is dynamically adjusted by the processor through manipulation of the VID (voltage identification) outputs. A 5-bit or 6-bit controller may be used in the regulator implementation. The processor does not drive (11111)b or (01111)b as a valid VID code. The following describes the usage of 5-bit and 6-bit controllers.

5-bit Implementation:

- VID[4:0] connects to VID[4:0] of the regulator (with appropriate level translation)
- VID[5] is left unconnected

6-bit Implementation:

- VID[5:0] connects to VID[5:0] of the regulator (with appropriate level translation)

The VID signals should be routed to the processor's VDD regulator and may need voltage level translation.

For compatibility with future processors, VID[1] should be pulled up to VDDIO through a 300- Ω resistor.

8.4.2 PSI_L

PSI_L is a Power Status Indicator signal. This signal is asserted when the processor is in a low power state. PSI_L should be connected to the power supply controller, if the controller supports “skip mode, or diode emulation mode”. PSI_L is asserted by the processor during the C3 and S1 states.

8.4.3 VDD_FB_H and VDD_FB_L

VDD_FB_H and VDD_FB_L are the processor output pins, which are dedicated to sensing the core voltage level at the processor die. They are used for differential feedback schemes. VDD_FB_H and VDD_FB_L should be routed as 10-mil traces with 10-mil spacing (10/10).

8.4.4 VDDIO_FB_H AND VDDIO_FB_L

The VDDIO voltage source provides power to the DDR2 output drivers and other miscellaneous functions within the processor. VDDIO_FB_H and VDDIO_FB_L are the processor differential output pins dedicated to sensing the VDDIO voltage level at the processor die.

VDDIO_FB_H should be used for systems which do not support differential voltage feedback. In this case, VDDIO_FB_L is not used and VDDIO_FB_H should be routed as a 10-mil trace with 10-mil spacing (10/10).

8.4.5 VTT_SENSE

The VTT voltage source provides power to the DDR2 termination resistors and the DDR2 output drivers. Special consideration must be made when designing the control strategy of the VTT voltage regulator. VTT voltage must closely track the VDDIO changes. To accommodate this task, the VTT_SENSE pin is designed to enable sensing of the VTT voltage level at the die. VTT_SENSE, VDDIO_FB_H, and VDDIO_FB_L pins can be used to control the output level of the VTT regulator and make sure VTT tracks VDDIO changes. Some regulators used for VTT may not have a feedback connection. In these cases, the processor's VTT_SENSE pin may be left unconnected. VTT_SENSE should be routed as a 10-mil trace with 10 mils spacing.

8.4.6 CPU_PRESENT_L

This pin is internally connected to VSS. To detect if a processor is present on a system, pull up this pin to VDDIO (or any other voltage) using a resistor. If a processor is present, voltage on this pin will be low (VSS). If a processor is not present, then voltage will be high (VDDIO or any voltage used for pullup).

8.5 Global Signals

The three signals in Sections 8.5.1, 8.5.2, and 8.5.3 are shared between all HyperTransport™ devices in a system.

8.5.1 PWROK

PWROK is a processor input signal used to indicate when the processor can attempt to lock the internal PLL. This signal should not be asserted until all processor power planes are active and the system clock generators are powered and allowed to run stably for at least one millisecond. The time a clock generation device takes from when power is applied to the time when output clocks become stable varies. Please refer to the clock generator specifications for specified timing. The POWER_GOOD input to the I/O hub should also be applied at the same time as PWROK is applied to the processor. The I/O hub then asserts the global reset signal, which initializes a system wide reset.

To achieve acceptable signal quality, the PWROK signal may need to be buffered and redriven to the various loads. The buffer is also beneficial because it provides level translation to the various input levels of the different loads.

There are many different ways to implement control of these signals. It is crucial to coordinate the application of these signals at the proper point in the power plane start-up sequence.

8.5.2 RESET_L

RESET_L is an active-Low input signal that resets the processor when asserted. Control of this signal originates from the I/O hub. This reset signal coincides with the global reset that is distributed over the entire motherboard to all the processors and various ICs. Timing coordination is critical for proper operation of the RESET_L signal.

To achieve acceptable signal quality, the reset output from the I/O hub may need to be buffered and redriven to the various loads. The buffer is also beneficial because it provides level translation to the various input levels of the different loads.

8.5.3 LDTSTOP_L

The LDTSTOP_L input to the processor controls entry and exit from the sleep and power states of the HyperTransport links. This pin should be routed to the source of LDTSTOP_L, which normally is the I/O hub.

8.6 Thermal-Related Miscellaneous Signals

This section explains the various thermal monitoring and thermal alarm output signals. The processor has a thermal sensor with which the on-chip temperature can be monitored. The processor also has thermal monitoring logic incorporated within the processor that trips an output to shutdown the system if a thermal alarm temperature is reached.

8.6.1 THERMTRIP_L

This pin is a thermal alarm output that is used to power down the system and prevent processor damage due to overheating. THERMTRIP_L is an open-drain processor output signal and requires an external pullup to VDDIO through a 300- Ω resistor. This signal may need voltage level translation if a device receiving this signal is not VDDIO compatible.

The system must power off the processor within a specified time after THERMTRIP_L is asserted. This time limit is specified in the *AMD NPT Family 0Fh Processor Electrical Data Sheet*, order# 31119. Some chipsets may have this functionality built in. If the chipset does not have the capability, Figure 68 on page 137 depicts an example of how to properly shut down the power supply when THERMTRIP_L is asserted.

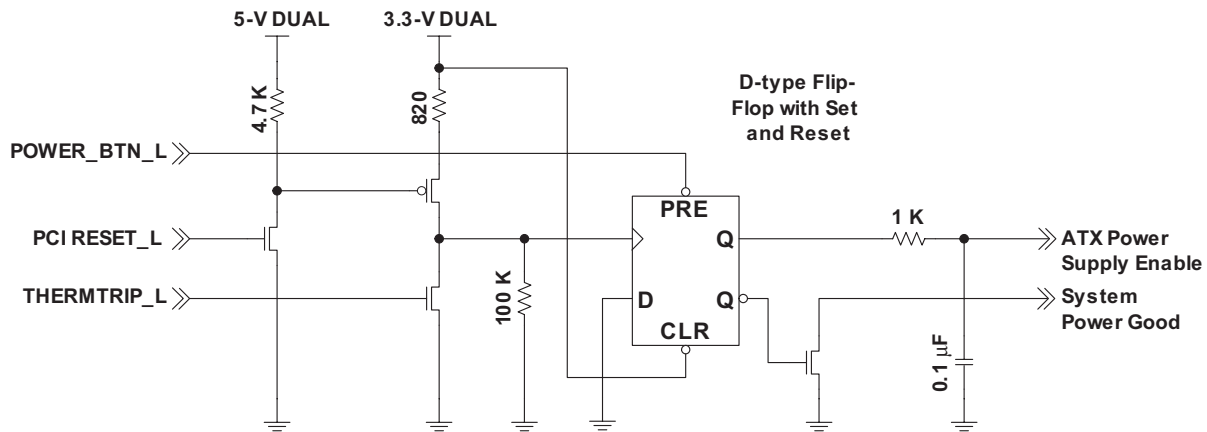


Figure 68. Thermtrip Circuit

8.6.2 PROCHOT_L

PROCHOT_L is a recommended but optional feature that designers may choose to use to help their platforms maintain the best operating conditions.

PROCHOT_L is an active-Low signal which indicates that the processor has reached a preset maximum operating temperature. The processor asserts PROCHOT_L when the HTC (hardware thermal control) is active. External hardware can assert PROCHOT_L to reduce the processor's power consumption by forcing HTC activation. For example, if the VDD power supply is getting near its maximum allowed temperature, then it can assert PROCHOT_L, which will force HTC active and reduce the processor's power, thus reducing the load on the VDD supply and help it remain within specification. This signal should be pulled up to VDDIO through a 300- Ω resistor.

Enabling PROCHOT_L may require action by the firmware. See the *AMD BIOS and Kernel Developers Guide for AMD NPT Family 0Fh Processors*, order# 32559, for more details regarding configuration of PROCHOT_L.

For an example schematic, see Figure 69 on page 138.




Figure 69. Example PROCHOT_L Schematic

8.6.3 SIC and SID

The SIC pin should be strapped to either VDDIO or VSS (but not both) using a 300- Ω ($\pm 5\%$) resistor. The SID pin can be left unconnected.

8.6.4 THERMDA and THERMDC

An on-die thermal diode is provided on the processor. The thermal diode can be used for thermal-based system management. An external temperature sensor using only two sourcing currents is required to measure the temperature of the thermal diode.

This section provides information on routing traces from a temperature sensor to the thermal diode. The thermal diode specifications are located in the *AMD NPT Family 0Fh Processor Electrical Data Sheet*, order# 31119. The *AMD BIOS and Kernel Developers Guide for AMD NPT Family 0Fh Processors*, order# 32559, defines how to access the thermal diode offset in the THERMTRIP Status Register.

Note: The thermal diode is separate from the circuit used to control the THERMTRIP_L feature for thermal shutdown. The thermal diode cannot be used to characterize the behavior of THERMTRIP_L.

8.6.4.1 Temperature Sensor Layout Recommendations

The following guidelines are recommended for temperature sensors. Adhering to these guidelines minimizes noise issues associated with routing from the on-die thermal diode to the external temperature sensor. Temperature sensors use small differences in voltage to measure temperature. A 1-mV variation results in a 5°C change in measured temperature.

Contact the temperature sensor manufacturer for specific considerations on layout and usage of the temperature sensor.

General layout guidelines are as follows:

- THERMDA and THERMDC should be routed as a differential pair and should reference VSS.
- Total trace resistance should be less than 0.5 Ω . This resistance may be achieved by using a minimal thickness of 1 oz. copper traces, 10-mils in width, for trace lengths up to 10 inches. Longer traces may be compensated for by increasing the trace width as long as the trace resistance remains below 0.5 Ω . For traces up to 10 inches long, 20/10/10/10/20, (sp/tr/sp/tr/sp) differential trace geometry is recommended.
- THERMDA and THERMDC nets should have at least 20-mil spacing to adjacent nets.
- AMD recommends adding pads for a 0402 capacitor on THERMDA and THERMDC nets. Place the pads as close to the processor as possible.
- Avoid routing THERMDA and THERMDC above or below high-speed signals (on different layers). If THERMDA or THERMDC must cross over or under high-speed traces, then route them perpendicular (90°) to the high-speed traces. Do not route THERMDA and THERMDC parallel to high-speed nets in different layers.

8.7 No Connects

Any pin named RSVD should be left unconnected on the board. These pins may be used for future enhancements or AMD internal debug purposes.

Chapter 9 Socket AM2 Processor Power Requirements

This chapter covers power requirements (including all power supplies) for supporting a socket AM2 processor. These power requirements allow the motherboard designer to properly layout and manage the power supplies that are necessary to support a socket AM2 processor. However, this chapter is not intended to replace the processor data sheet. Refer to the *AMD NPT Family 0Fh Engineering Sample Processor Power and Thermal Information*, order# 33951, for specific voltage and current numbers.

High-frequency decoupling capacitors for all motherboard supplies are listed in this chapter. However, other motherboard structures that need decoupling capacitors are listed in other chapters. All the decoupling capacitors are necessary for a correct operation of the processor.

Designers may replace a capacitor with another capacitor of the same dielectric and capacitance but with a smaller body size. Decoupling capacitors should not be substituted with another capacitor with a larger body size.

9.1 High-Frequency Motherboard Decoupling

High-frequency, board-level capacitor placement is critical to the proper functionality of the power supply. These capacitors must be able to respond to short duration and high frequency load changes. As such, it is necessary to place the high-frequency capacitors as close as possible to the final load to minimize the interconnect inductance. The optimum location for minimum interconnect inductance is the area directly under the socket on the backside of the motherboard. To place the maximum number of capacitors in this area use the 0805 or 0603 component, body-style capacitor, which offers a good combination of capacitance values per available board area.

Once all of the available area under the socket is filled with capacitors, the next best location to place additional capacitors is on the topside of the motherboard. They should be placed as close to the socket as possible without interfering with the routing.

Typical high-frequency capacitors are multi-layered ceramic (MLC) surface-mount chips. Due to their construction, they have a low parasitic inductance. This feature is important, because the most critical factor limiting the effectiveness in minimizing the power and ground noise is the inductance associated with the capacitor connection to the board.

Capacitors with side escapes are recommended over those with end escapes. With side escapes the full-current loop length is significantly reduced, allowing the best possible performance from the capacitor. For a comparison of the difference in length, see Figure 70 on page 142. Additionally, multiple-plated through hole (PTH) vias are recommended to further reduce the capacitor's connection impedance to the power and ground planes on the motherboard.

Inductive Loop Comparison of Capacitor Connection




Figure 70. High-Frequency Decoupling Routing Comparison

Place the high-frequency capacitors with the lowest values closest to the socket. Since these components typically have the smallest body size, they are least likely to interfere with the dense routing in the area. Higher-valued capacitors should be placed as close to the socket as possible, but due to their larger body size, they should be placed further away from the socket than the smaller (lowest value) capacitors.

AMD recommends that the traces from the capacitor pads to the PTH vias be as wide as possible to minimize lead inductance. Using copper pours or mini-planes on the external (top and bottom) signal layers will further establish a solid, low-inductance connection. These mini-planes should not interfere with signal routing, because it is not necessary to route signals through the center of the socket.

9.2 Power Generation and Distribution Guidelines

This section documents all of the processor's power-supply requirements—including placement, routing, and decoupling recommendations for the processor power supplies.

During the S3 power state, the power supplies (VTT, VDDIO, and VLDT) meet the sequencing requirements that are specified in Refer to the *AMD NPT Family 0Fh Engineering Sample Processor Power and Thermal Information*, order# 33951.

The various power structures—VDD, VDDIO, and VSS—are shown in the Figure 71.

VDDA, VTT, VLDT, and M_VREF (not show in Figure 71) are delivered on signal layers.




Figure 71. General View of the Processor Power Plane Cuts

9.2.1 Socket AM2 Processor Power and Ground Overview

Figure 71 shows a cut-away of an actual motherboard to illustrate the recommended layout guidelines for the power planes underneath the socket AM2 processor. The goal of the power distribution scheme is to utilize all layers of the motherboard in such a way as to reduce the inductance of the planes. In addition to providing better power profiles, solid power and ground planes provide ideal references for high-speed signals.

9.2.2 VDD Power Requirements

VDD provides power to the core logic within the processor.

A power plane should provide the processor's core power. It is recommended that the separation between the power-plane and ground-plane pair be kept to a minimum. Care should also be taken to avoid heavy perforations, which are the result of either an excessive placement of components or of vias that are not devoted to decoupling of the core power. The VDD power plane should always be at least as wide as the processor.

The processor pinout accommodates a power fill on both the top and bottom signal layers for VDD power delivery. The VDD power fill should utilize all available space on the top and bottom signal layers to connect the power supply to the pin field. Figure 72 on page 145 illustrates of how this may be done on a four-layer board.




Figure 72. Example of VDD Power Fill on the Top Layer

The high-frequency capacitors are located as close to the processor as possible, on the backside of the board. All the capacitors listed in this chapter that are placed between the processor and the VRMs should be placed as close to the processor as possible.

Table 34 shows the recommendations for the decoupling capacitors. These recommendations assume that appropriate layout recommendations have been implemented.

Table 34. VDD Decoupling Recommendations

Qty	Capacitance	Type	Location
15	22 μ F	1206 or 0805 X7R Ceramic	Under socket on the bottomside
3	0.22 μ F	0603 X7R Ceramic	Under socket on the bottomside
1	0.01 μ F	0603 X7R Ceramic	Under socket on the bottomside
1	180 pF	0402 COG or NPO Ceramic	Under socket on the bottomside
	Bulk caps as required by power supply design		Close to the regulator ¹
Note:			
1. Power-supply bulk capacitors should be placed in a way that does not block airflow to the regulator components. Air flow generated by the processor fan should be directed to the VDD power supply area with the heatsink fins.			

Figure 73 on page 147 shows the recommended placement of the high-frequency decoupling capacitors on the bottomside of the board if a backplate is used. Figure 74 on page 148 shows the recommended placement of the high-frequency decoupling capacitors on the bottomside of the board for a system with no backplate. The number of capacitors that can be placed on the bottomside under the socket varies depending on the via placement and power-plane cut.

Note: Not all of the capacitors are shown in Figure 73 and Figure 74 on page 148. The 1206 capacitors require via pads of 20 mils or smaller. Vias connecting to capacitors should have solder masks up to the via hole. For both Figure 73 and Figure 74 VDD is red, VDDIO is pink, VSS is green, and backplate windows are black.

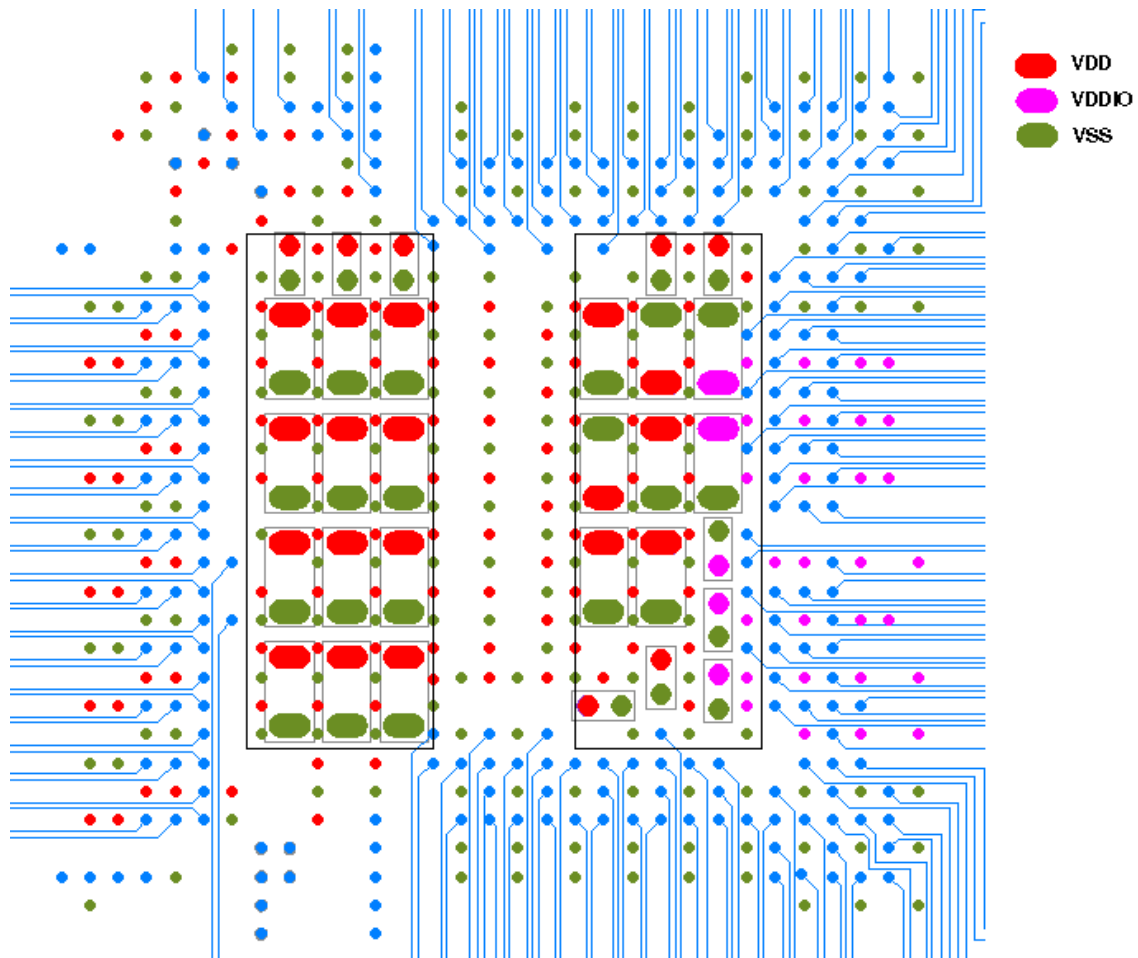


Figure 73. VDD Decoupling Placement With Backplate Windows—Bottom Signal Layer

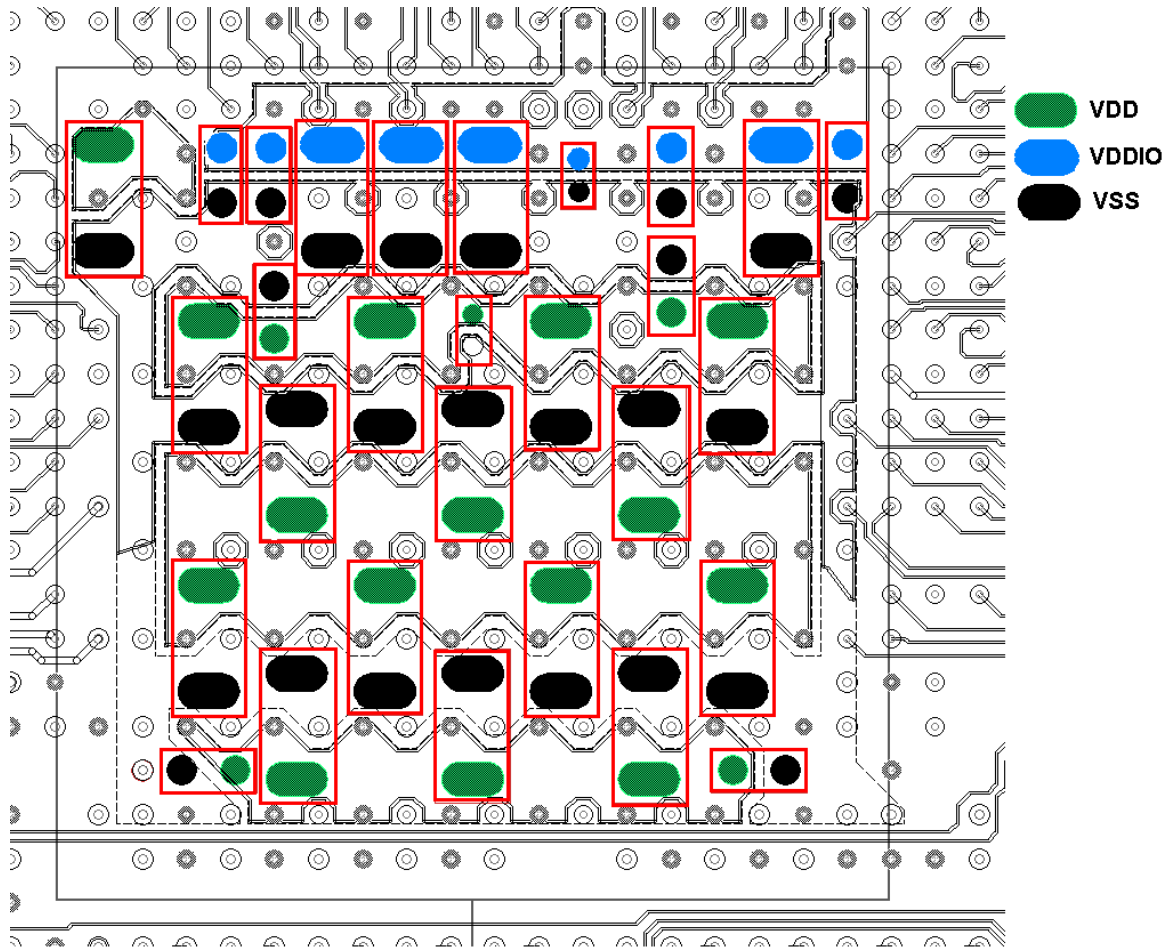


Figure 74. Decoupling Placement Strategy on Bottom Signal Layer—Without Backplate

9.2.3 VDDIO, VTT, and M_VREF Routing and Decoupling

Refer to the *AMD NPT Family 0Fh Engineering Sample Processor Power and Thermal Information*, order# 33951, for the exact current and tolerance requirements for VDDIO, VTT, and M_VREF.

9.2.3.1 VDDIO Decoupling

VDDIO should be a plane that runs beneath the DIMMs and partially under the processor socket, because VDDIO is the main power supply for the DIMMs and also an auxiliary supply for the DDR2 section of the processor. When placing the decoupling capacitors, start by positioning the smallest-valued (and smaller body size) capacitors near the socket, followed by larger capacitors. Since VDDIO provides power to the DIMMs as well, additional decoupling capacitors must be placed near each DIMM connector.

Figure 75 and Figure 77 on page 153 illustrate the recommended decoupling and routing. Table 35 on page 150 documents the VDDIO decoupling recommendations.

Note: Not all of the capacitors are shown in Figure 75.

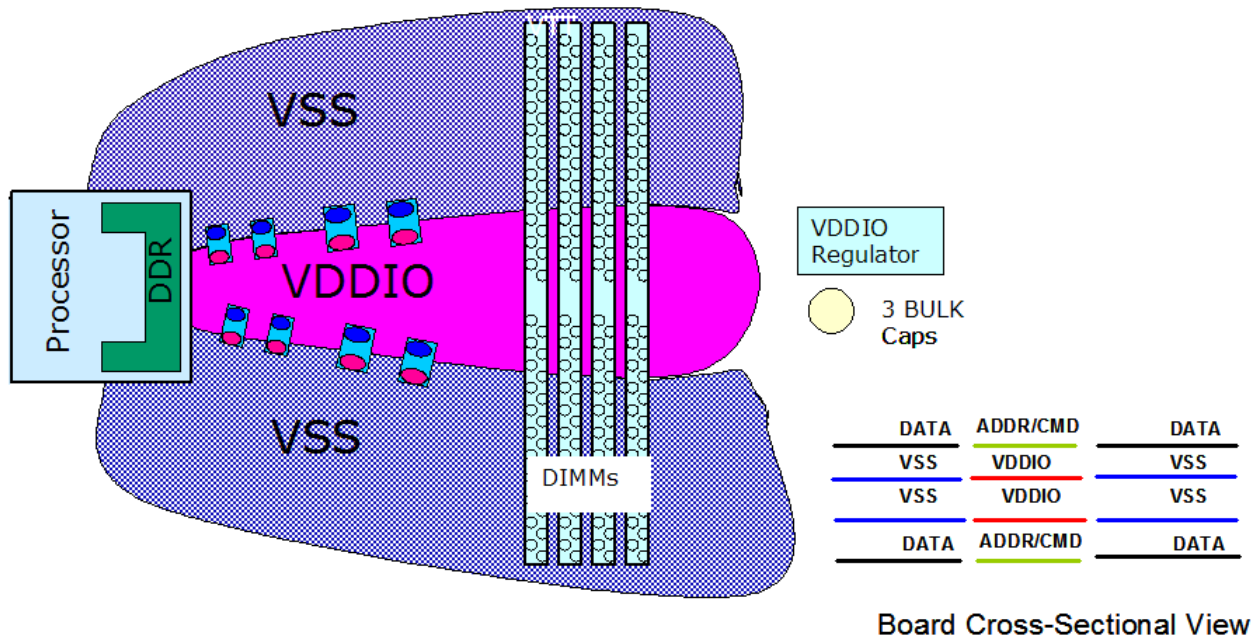


Figure 75. VDDIO Decoupling Placement Between Processor and DIMMs

Table 35. VDDIO Decoupling Recommendations

Qty	Capacitance	Type	Location
2	22 μ F	1206 X7R Ceramic	Under the socket on the bottomside
1	0.22 μ F	0603 X7R Ceramic	Under the socket on the bottomside
1	0.01 μ F	0603 X7R Ceramic	Under the socket on the bottomside
1	180 pF	0402 COG or NPO Ceramic	Under the socket on the bottomside
See note	180 pF	0402 COG or NPO Ceramic	Evenly spaced along the VDDIO/VSS plane split
4	4.7 μ F	X7R Ceramic	Between the socket and DIMMs, as close as possible to the socket
4	0.22 μ F	0603 X7R Ceramic	Between the socket and DIMMs, as close as possible to the socket
	Bulk capacitors as required by power supply design		Close to the regulator

Note:

If the VSS plane is cut in order to create a VDDIO plane, additional 180-pF capacitors need to be added along both sides of the VDDIO and VSS plane split. The number of capacitors required varies with the length of the plane split along VDDIO and VSS.

- *If the distance between the processor keep-out and first DIMM's keep-out is less than 2.5 inches, one 180-pF capacitor on each side of the plane split is needed (two total). The capacitors should be placed halfway between the CPU and the first DIMM in order to halve the distance between the processor and DIMM socket.*
- *If the distance between the processor keep-out and first DIMM's keep-out is greater than 2.5 inches, two 180-pF capacitors on each side of the plane split are needed (four total). The capacitors need to be evenly spaced between the CPU and the first DIMM.*

9.2.3.2 VTT Power Delivery and Decoupling

VTT is the termination voltage for the memory-interface signals.

VTT can be provided in one of two ways—as a plane or as a fill. VTT can be a plane beneath the DIMMs and partially under the processor. VTT also can be a copper pour on a signal layer between the processor and the DIMMs. If VTT is provided as a copper pour, it is called the “VTT fill”.

There must also be an area fill behind the DIMMs, referred to as the “VTT island”.

The VTT fill is routed on a single layer and can be either a single copper pour (at least 200 mils wide) to one side of the processor or it can be two copper pours (at least 100 mils wide) with one pour on each side of the processor. The VTT fill can narrow to 50 mils wide for the segment within 1 inch of the processor. The VTT island must be at least 100 mils wide and connected to the VTT plane or the VTT fill between the processor and the DIMMs. VTT is decoupled to VSS close to the socket. Behind the DIMMs, VTT is decoupled to VDDIO and VSS.

Table 36 on page 151 lists the VTT decoupling recommendations.

Table 36. VTT Decoupling Recommendations

Qty	Capacitance	Type	Location
4	4.7 μ F	0805 X7R Ceramic	VTT to VSS: Capacitors are placed between the socket and DIMMs, as close as possible to the socket
4	0.22 μ F	0603 X7R Ceramic	VTT to VSS: Capacitors are placed between the socket and DIMMs, as close as possible to the socket
4	1 nF	0603 or 0402 X7R Ceramic	VTT to VSS: Capacitors are placed between the socket and DIMMs, as close as possible to the socket
1 capacitor for each 1Rpacks	0.1 μ F	0603 X7R Ceramic	VTT to VDDIO: Capacitors are placed behind the DIMMs and follow the Rpacks.
1 capacitor for each 1 Rpacks	0.1 μ F	0603 X7R Ceramic	VTT to VSS: Capacitors are placed behind the DIMMs.
4	180 pF	0402 COG or NPO Ceramic	VTT to VSS: Capacitors are placed between the socket and DIMMs, as close as possible to the socket
	Bulk caps as required by power supply design		VTT to VSS: Close to the regulator

The processor has two sets of VTT connections, one on either side. In cases where the VTT fill is delivered as a single copper pour, both sides must be connected on the motherboard. Figure 76 on page 152 depicts an implementation that connects both sets of VTT pins together.

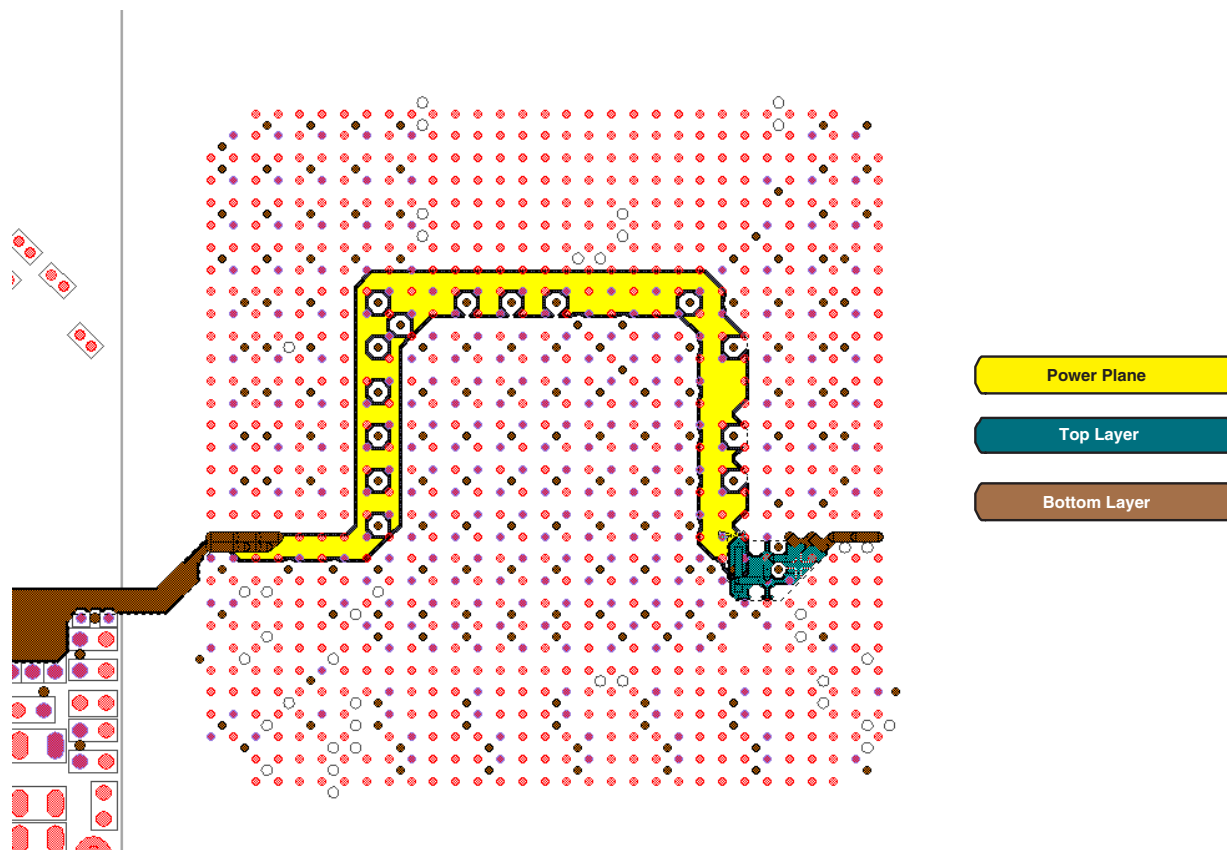


Figure 76. VTT Delivery on Single Side

Figure 77 on page 153 illustrates the VTT caps and routing.

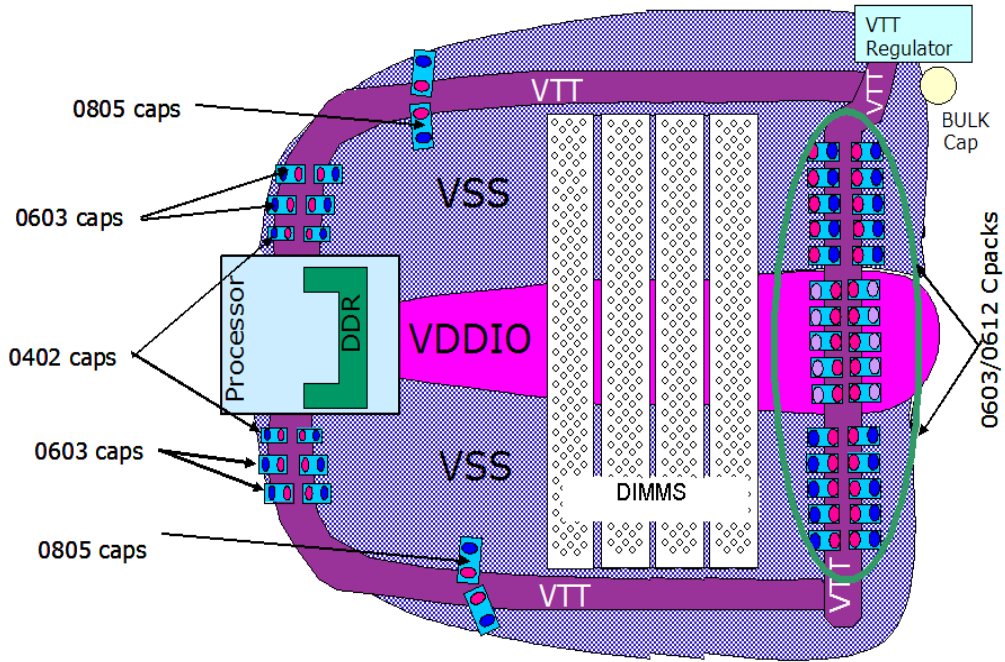


Figure 77. VTT Decoupling Placement Between Processor and DIMMs

9.2.3.3 M_VREF Decoupling and Routing Guidelines

M_VREF is the reference voltage for the memory interface. There are two M_VREF signals. One signal is generated at the DIMMs for the DDR2 DIMMs, and it provides power for the DIMM's reference. The other M_VREF signal is generated close to the processor from a VDDIO divider, which should consist of two 15- Ω resistors. For required decoupling capacitors, refer to Table 37.

Table 37. M_VREF-to-Processor Decoupling Recommendations

Qty	Capacitance	Type	Location
1	0.1 μ F	0603 X7R Ceramic	Between M_VREF and VSS, less than 1000 mils from the pin
1	1 nF	0603 or 0402 X7R Ceramic	Between M_VREF and VSS, less than 1000 mils from the pin

M_VREF to the processor should be routed as a 15-mil wide trace with 20-mil spacing and should be shorter than 6 inches. See Figure 78 for M_VREF topology and decoupling.

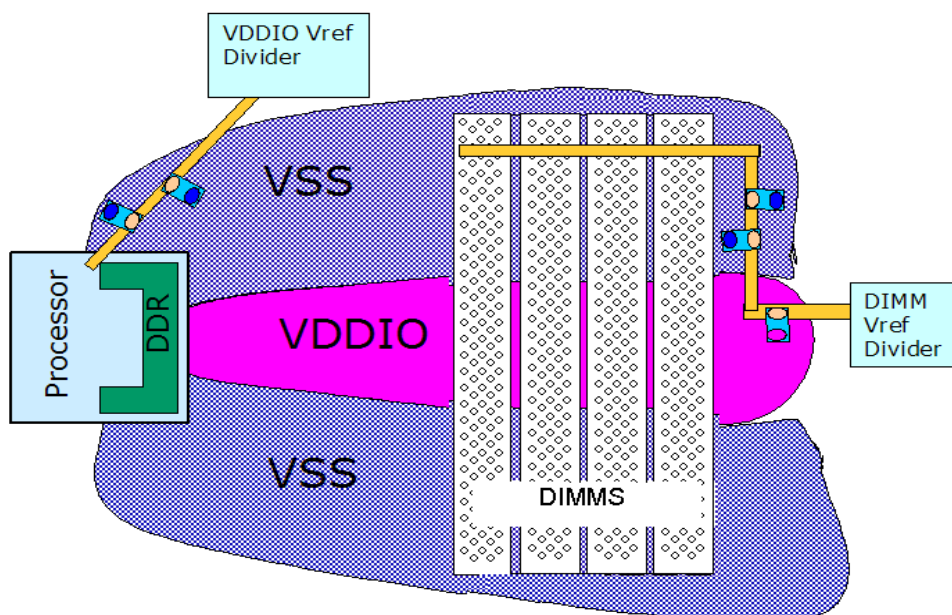


Figure 78. M_VREF Decoupling Placement

9.2.3.4 VDDSPD

If the design is an SO-DIMM implementation, the VDDSPD input of the SO-DIMM socket should be connected to 3.3V to accommodate the SO-DIMM temperature sensor.

9.2.4 VDDA Power Requirements

VDDA supplies voltage for the phase lock loop (PLL) in the processor. This supply must be isolated from the other supplies to keep the clock PLL locked under extreme power supply loading or bus activity conditions, which could cause AC transients and/or noise on the other voltage planes. Great care must be taken to ensure that the VDDA power generation is as quiet as possible. It is recommended that the power from the regulator be delivered to the processor with a 15-mil wide trace that is referenced to VSS. This trace should be kept at least 20 mils away from all other signals.

The filter circuit on the VDDA power supply should include an inductor (or ferrite bead) and several capacitors. Figure 79 on page 155 illustrates a VDDA power-delivery circuit to the processor. The inductor and capacitors create a filter that is used to keep the VDDA pin on the processor as quiet as possible at frequencies up to 10 MHz. For circuitry and component values, refer to Figure 79 on page 155 and Table 38 on page 156.

The inductor or ferrite bead must meet the following requirements:

- DC Resistance $\leq 40 \text{ m}\Omega$
- $I_{\text{max}} \geq 500 \text{ mA}$
- Impedance is approximately 35Ω

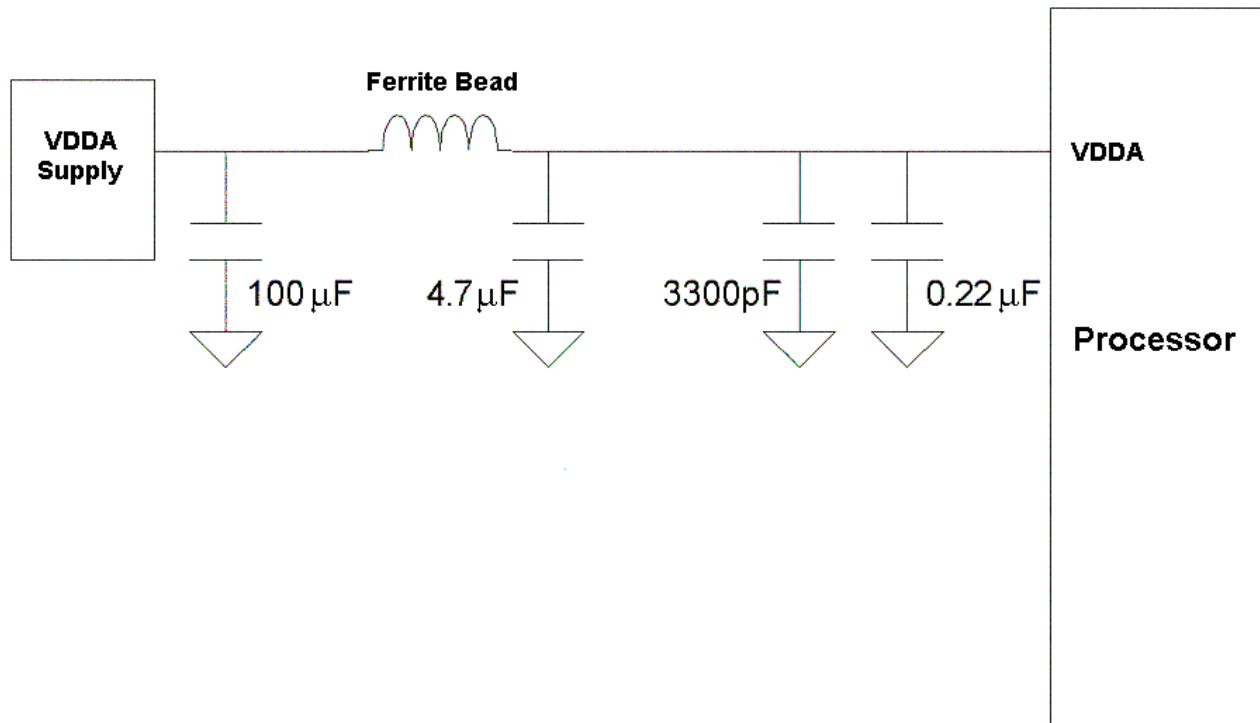


Figure 79. VDDA Filtering Circuit

Table 38 gives the VDDA filter recommendations.

Table 38. VDDA Filter Recommendations

Qty	Components	Type	Location
1	0.22 μ F	X7R Ceramic	1000 mils or less from VDDA pins
1	3300 pF	X7R Ceramic	1000 mils or less from VDDA pins
1	4.7 μ F	X7R Ceramic	As close as possible to the ferrite bead
1	30-300 nH	Ferrite bead	At the power supply output
1	100 μ F	X5R Ceramic	At the power supply output

9.2.5 VLDT Power Requirements

VLDT is the power supply for the HyperTransport™ I/O cells. Since HyperTransport technology has a very high data-rate, careful design practices must be exercised to ensure a low DC and AC impedance path from the VLDT power supply to the die.

This voltage is OFF during the S3 state.

VLDT is connected with a topside pour from the VLDT power-supply output to the processor.

When routing VLDT as either a trace or a pour referenced to VSS, the following recommendations should be followed:

- VLDT may be routed to either the VLDT_A or VLDT_B pins. Connecting VLDT to both sets of pins is not necessary.
- If routed as a trace, VLDT (from the VLDT power supply to VLDT_A or VLDT_B) should be routed at least 200 mils wide, until it reaches the pin field; in the pin field, VLDT should be routed as 20-mil traces. It should reference VSS, when possible.
- The high-frequency capacitors should be placed as close as possible to the VLDT power pins, but they should be placed in such a way as to maximize the copper connection to the VLDT pins.
- One mid-frequency decoupling capacitor (recommended value of 4.7- μ F 1206 or 0805 X7R) is connected between the processor pins and the capacitor with a 100-mil wide pour (the VLDT_A or VLDT_B pins not directly connected to the VLDT power supply). In other words, if VLDT_A pins are connected to the VLDT power supply on the motherboard, then VLDT_B pins should be connected to a single 4.7- μ F capacitor routed with a 100-mil wide trace.
- VLDT should be kept at least 20 mils away from any signal.

Note: If HyperTransport signals cross over plane splits, AMD recommends that designers use 0603 0.01 μ F stitching capacitors (one per every four pair of signals) to bridge the planes.

Table 39 provides the VLDT decoupling recommendations.

Table 39. VLDT Decoupling Recommendations

Qty	Capacitance)	Type	Location
2	4.7 μ F	0805 X7R Ceramic	Near the socket on the VLDT pour
2	0.22 μ F	0603 X7R Ceramic	Near the socket on the VLDT pour
2	180 pF	0603 or 0402 COG or NPO Ceramic	Near the socket on the VLDT pour
1	4.7 μ F	0805 X7R Ceramic	Near the socket on unused VLDT pins when VLDT is routed as a pour to one set of VLDT_A or VLDT_B pins. If VLDT is routed as a plane, place the capacitor between the VLDT power supply and the processor VLDT.
	Bulk caps as required by power supply design		Close to the regulator

Chapter 10 EMI Design Guidelines

Motherboard designers should follow these design recommendations to minimize system electromagnetic interference (EMI). The emphasis is on implementing the necessary electromagnetic compatibility (EMC) design in the early stages to eliminate costly and time-consuming debug and design changes at the end of the development cycle. In addition to following these EMC design guidelines, it is recommended that an EMC design engineer complete EMC schematic and layout board reviews on each hardware revision level. Also, early level hardware should be EMC tested to identify and then implement necessary design changes on subsequent levels of hardware.

Computer systems using AMD processors must meet various EMC agency specifications before they can be offered for sale. In the US and Canada, the FCC CFR 47 part 15 Class A and B limits apply. Throughout Europe and Asia, the CE standards (or similar EMC standards) apply.

10.1 I/O Signal Partitioning and Separation

To minimize system level EMI emissions, I/O cables must be kept free of EMI. To reduce coupled EMI onto I/O signals, do the following:

1. Partition high-frequency signals such that they are separate from I/O signals. Of highest priority is the separation between I/O signals and the high-frequency periodic (clock) signals.
2. Place I/O logic as close as possible to its associated I/O connector. This reduces the I/O trace routing on the motherboard and the potential EMI noise coupling. If it is necessary to locate I/O logic more than a few inches from the I/O connector, take care to choose the shortest and “least noisy” path for routing the I/O signals to the connector.
3. Route the I/O traces in logical groups to each I/O connector and maintain a minimum 0.1 inch separation between I/O to any non-I/O signals.

10.2 Decoupling, Bypass, Stitching, and Filtering Capacitors

The following section discusses capacitor selection and application as it pertains to EMI reduction.

10.2.1 Voltage Plane Decoupling Capacitors

A motherboard voltage decoupling analysis should be completed using Spice, other simulation tools, or by following vendor data sheet information. Once the decoupling requirements are identified, place the smaller-value SMT decoupling capacitors as close as possible to each component’s voltage pins. Use a low-impedance connection to place this capacitor. In practice this means using a very short

trace with a minimum 20-mil width to connect the decoupling capacitor terminals to voltage and ground.

10.2.2 Voltage Filtering Requirements

Specialty voltages, such as Analog, PLL, and especially Clock voltage supplies, require high-frequency decoupling and filtering to reduce EMI. Closely follow each vendor's data sheet design specification.

See Chapter 9 for additional information on this topic.

10.2.3 High-Frequency Bypass and Stitching Capacitors

The motherboard voltage and ground planes serve as high-frequency signal returns. For this reason, it is important to maintain a low impedance at higher frequencies between all voltage and ground planes. In practice, this is accomplished by placing distributed bypass capacitors between all voltage and ground planes, and also across the voltage-plane isolation breaks. These bypass and stitching capacitors are low impedance at high frequencies, thus offering high-speed signal-return currents the shortest and lowest impedance return paths.

Voltage-plane EMC bypass capacitors should be placed through-out the entire motherboard on a 1-inch to 1.5-inch grid on every voltage plane.

Note: *Smaller value ($< 0.1\mu\text{F}$) surface mount technology (SMT) voltage decoupling capacitors can satisfy this EMC bypass capacitor requirement if present in a region of the motherboard.*

Voltage-plane stitching capacitors are required in the event that high-speed signals must cross a reference plane break. To mitigate this crossing, place a stitching capacitor across the break no further than 0.5 inches from the signal crossing.

Note: *The more restrictive stitching capacitor placement requirements for HyperTransport™ signals are in Section 4.2.2.6.*

Failure to add these capacitors can have undesired consequences, such as the following:

- Signals that cross a reference-plane split without a stitching capacitor close by have a higher impedance and different propagation rate than expected.
- The signal return path may become significantly large. This large return loop area contributes to a rise in emissions from this signal.

Values for stitching and bypass capacitors typically range from 100 pF to 0.1 μF . However, the exact capacitor value is of secondary importance compared to the use of a low-impedance connection from the capacitor terminals to the motherboard and using multi-layer ceramic small package SMT capacitors.

10.2.4 DDR VDDIO to VSS Stitching Capacitor Requirements

To achieve proper DDR2 signal referencing and length matching on four layer motherboards, a VDDIO voltage island is cut on both internal power and ground planes beneath the DDR routing.

Doing this requires placement of a small SMT stitching capacitor across the VDDIO/VSS plan break. See Section 9.2.3 for additional details.

10.3 Motherboard Grounding

Proper motherboard grounding is essential for reduced system EMI emissions. The following section provides proper motherboard grounding techniques.

10.3.1 Motherboard-to-Chassis Grounding

Under most circumstances, optimal system-level EMI results are obtained when all motherboard mounting locations are connected directly to the chassis.

Create a low-impedance mounting hole to ground the connection with eight via connections as shown in Figure 80.

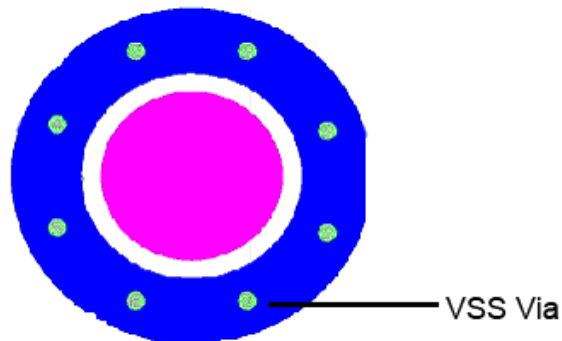


Figure 80. PCB Mounting Hole

10.3.2 I/O Filtering and Shield Grounding

All system I/O signals should include EMI filtering located close to the I/O connector. This filtering should consist of series impedance (resistor, ferrite bead, or inductor) and/or provisions for a shunt capacitor-to-ground on each signal, if functionality permits. If an I/O signal exits the system on an unshielded cable, provisions for a common mode choke filter are advised.

Additionally each I/O cable connector shield tab should connect directly to the motherboard ground or to the I/O isolated ground as described below.

To further reduce common-mode EMI currents on I/O cables, it is an acceptable practice to isolate the ground plane just before the I/O connectors. If this is done, care needs to be taken to not have high-speed signals cross the I/O ground isolation break. Normally this technique is implemented on slower I/O ports and may only isolate the I/O connector shield tabs of higher speed ports. Always, include jumper pad locations that bridge the I/O ground isolation break close to each I/O connector shield tab. In most cases, these 0- Ω jumpers are populated and would only be removed if necessary during EMI problem resolution.

10.3.3 Processor Heatsink Grounding

In some cases, the processor and memory EMI harmonics can be reduced by grounding the processor heatsink in a low-impedance fashion. Include provisions for heatsink grounding by including the motherboard grounding pads, which are illustrated in Chapter 11.

10.4 Clocks and EMI

High-frequency periodic (clock) circuits are often the main source of EMI in a computer system. The following sections provide guidelines for reducing and containing clock EMI.

10.4.1 Clock Circuit Placement and Wiring

To reduce EMI, place crystal circuits 0.5 inches or greater from the edge of the motherboard and in a manner as to reduce the clock-trace lengths. To minimize the coupling of clock EMI to victim circuits, place crystal circuits 1.0 inch or greater from the internal cable connectors and 1.5 inches or greater from system I/O cable connectors. Always route clock nets as short as possible and with a minimum of layer changes while maintaining 0.1-inch or greater separation from any I/O signals and the edge of the circuit board. Also, do not route a clock trace over an adjacent layer reference plane break.

See Section 10.2.3 on page 160 for additional information on this topic.

10.4.2 Clock Signal Termination

To improve signal integrity and reduce EMI, it is important to properly terminate high-speed clock signals. Correct signal termination can reduce ringing and reflections, and thus EMI, without adversely affecting timing or slew rate. For single-ended clock signals, locate the series terminating resistor as close as possible to the output driver. Choose a resistor value that, when combined with the driver output impedance, matches the trace impedance on the motherboard. Additionally, load-end terminations, such as a shunt R/C to ground, can be used to minimize incident wave reflections.

10.4.3 Spread Spectrum Clocking

Spread spectrum clocking (SSC) should always be enabled on the processor and memory subsystems to significantly reduce EMI levels. All AMD processor clock-generator modules include the 0.5% down spread capability, which should be utilized on all motherboard designs. Configure the processor clock-generator module with spread spectrum on as the default. Please refer to the vendor's clock generator data sheet for a complete circuit implementation.

10.4.4 Special Considerations for Clock Generator Modules

Below are some design guidelines to reduce EMI from the clock generator modules:

Note: Please refer to the vendor's clock generator data sheet for additional circuit implementation details.

- If possible, route all clock generator clocks referencing the motherboard ground layer. On a four-layer stackup, this would be the bottom layer.
- On the top layer, directly beneath the clock generator module, create a VSS pour connecting to all clock generator ground pins. See Figure 81.
- If the solid layer beneath the clock generator is power, cut at the island into it to form the filtered voltage for the clock generator module. See Figure 82 on page 164.)
- Place high-frequency decoupling capacitors as close as physically possible to each clock generator power pin with short connecting traces. See Figure 81.

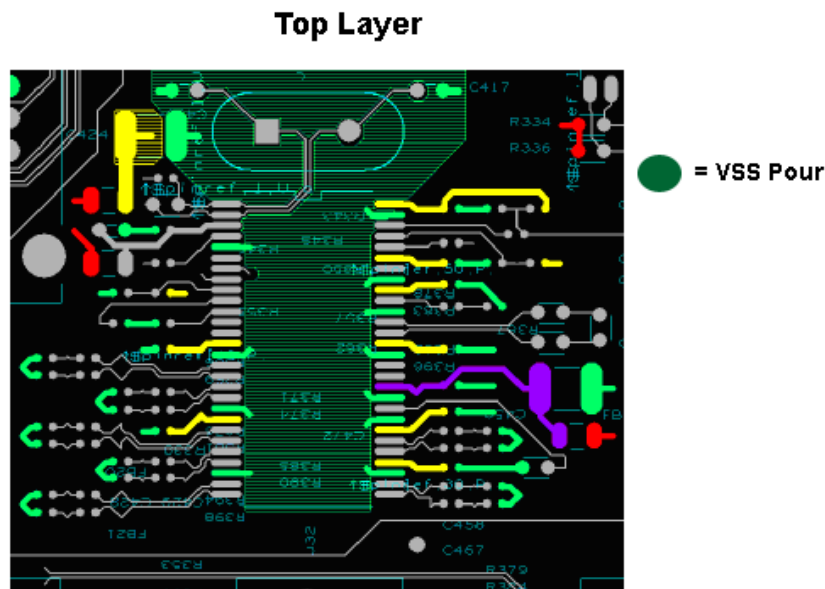


Figure 81. Top Layer Beneath the Clock Generator Module

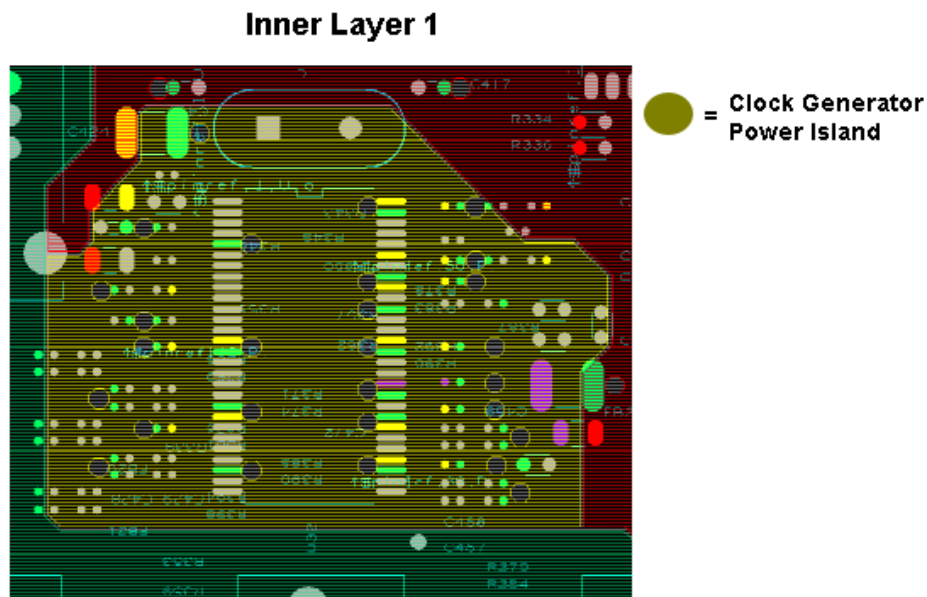


Figure 82. Inner Layer Beneath the Clock Generator Module

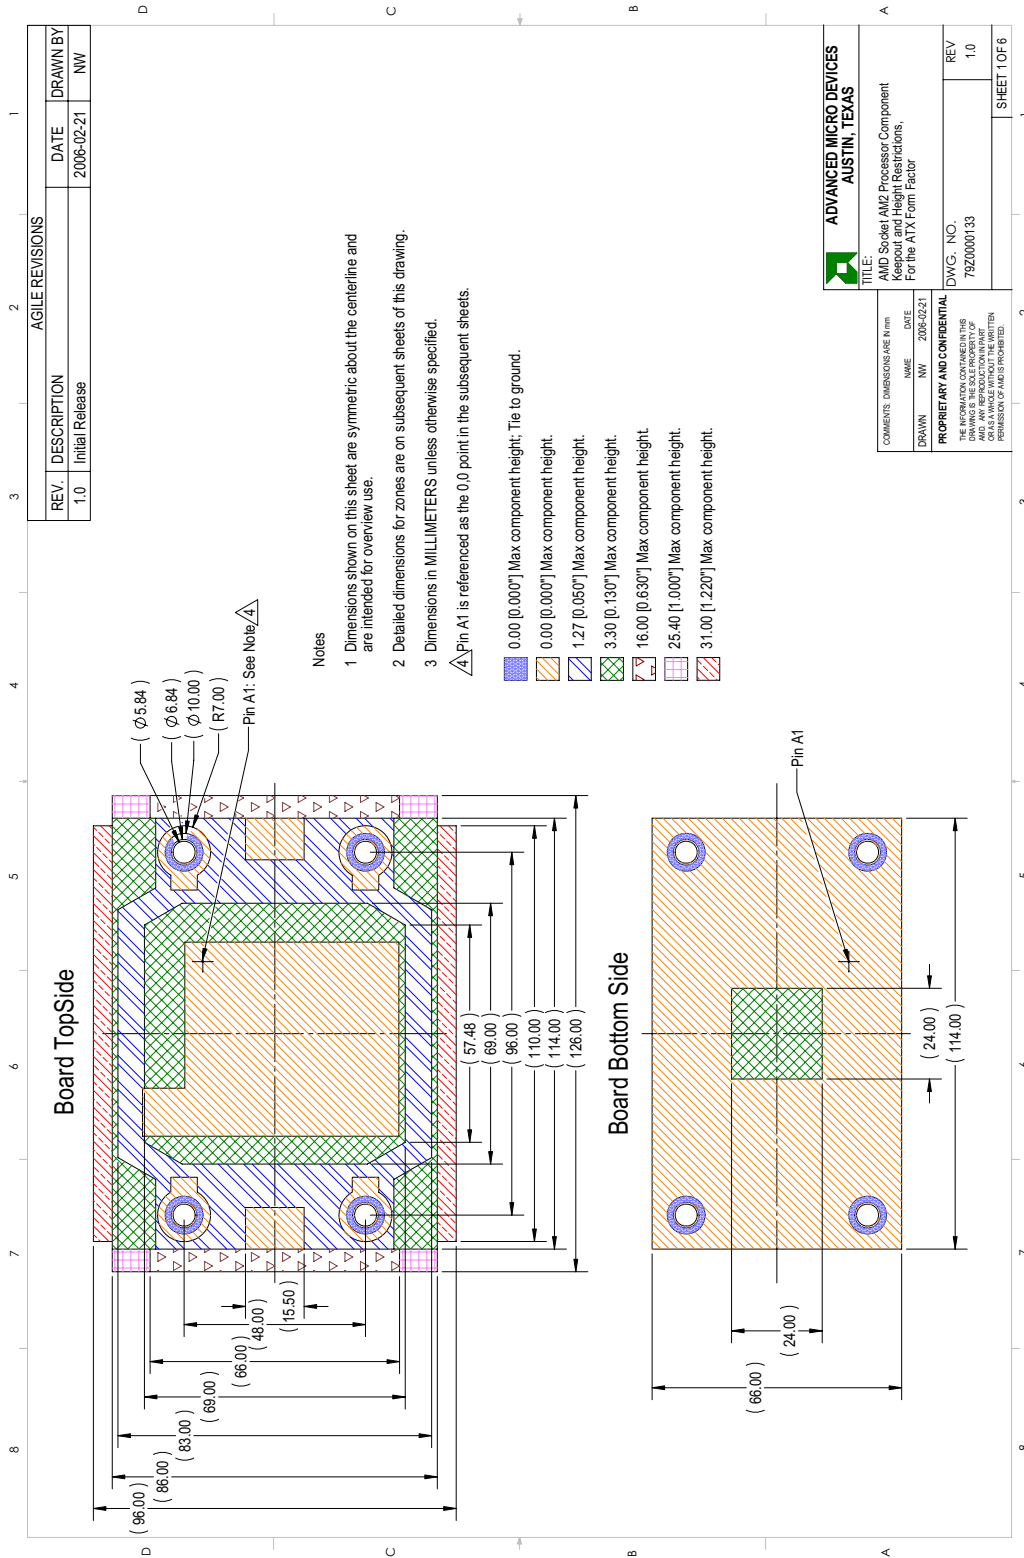
10.4.5 Disable Unused Clock Outputs


Disable all unused clock outputs via hardware and/or BIOS settings. This includes, but is not limited to, unused memory clocks, unused PCI clocks, and unused clock-generator clocks, such as processor, USB, Ethernet, and 14.318-MHz system clocks.


Chapter 11 Thermal and Mechanical Motherboard Requirements

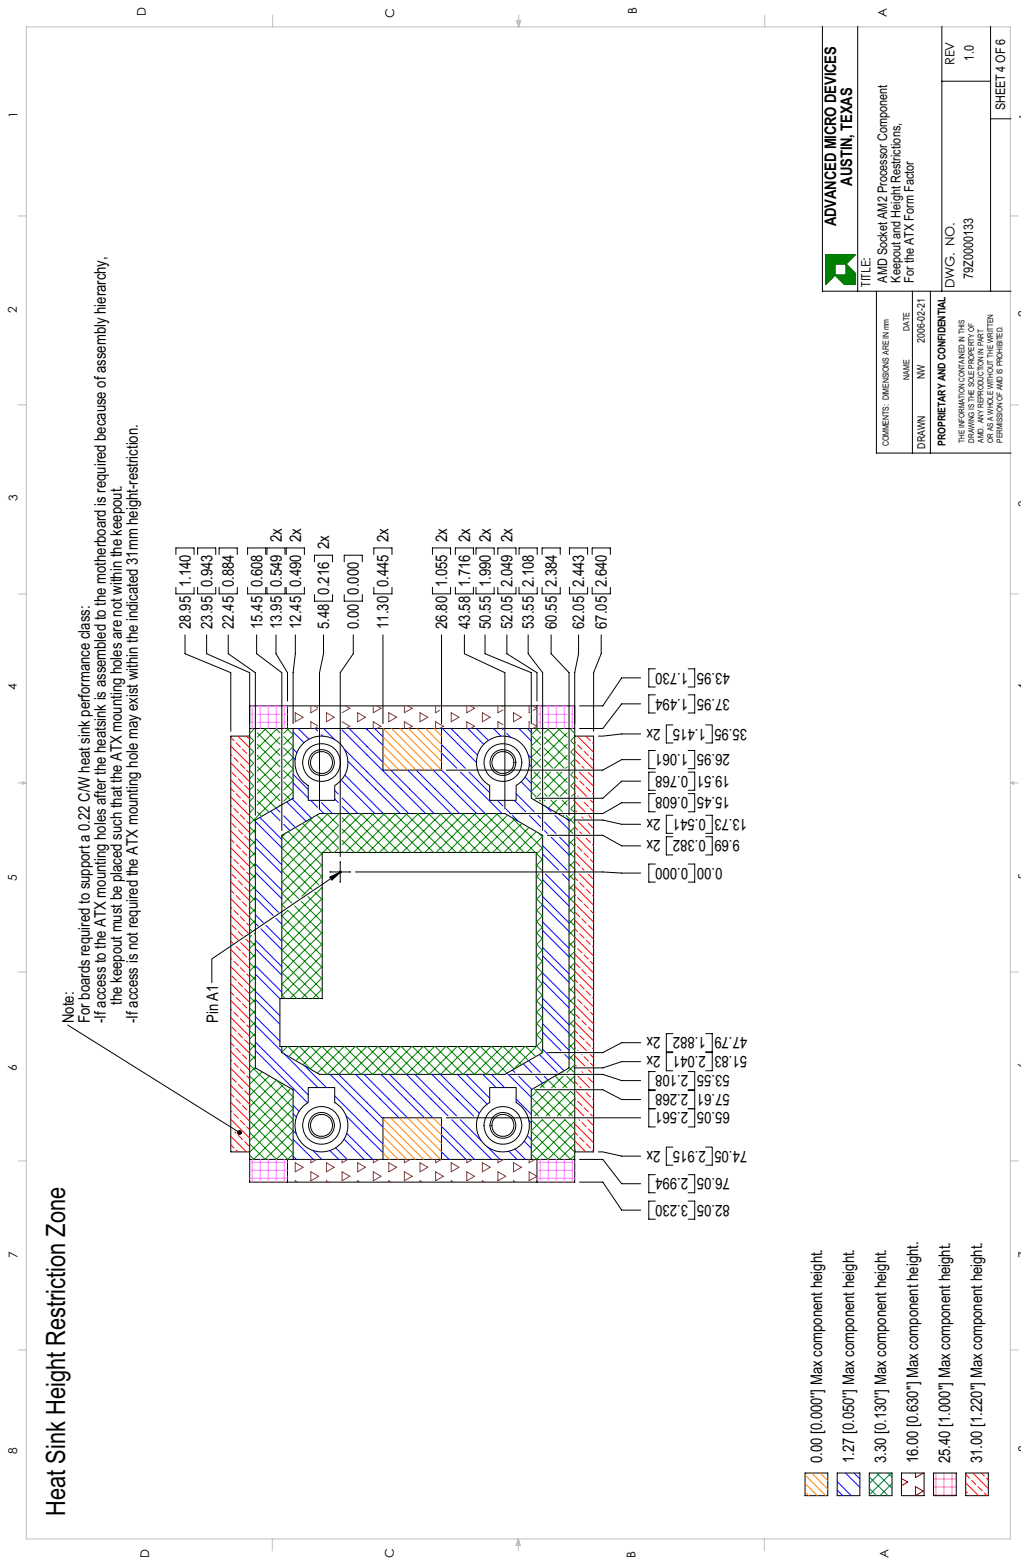
Information and figures are preliminary and will be updated in future revisions.


11.1 Socket AM2 ATX Processor Component Keepout and Height Restrictions

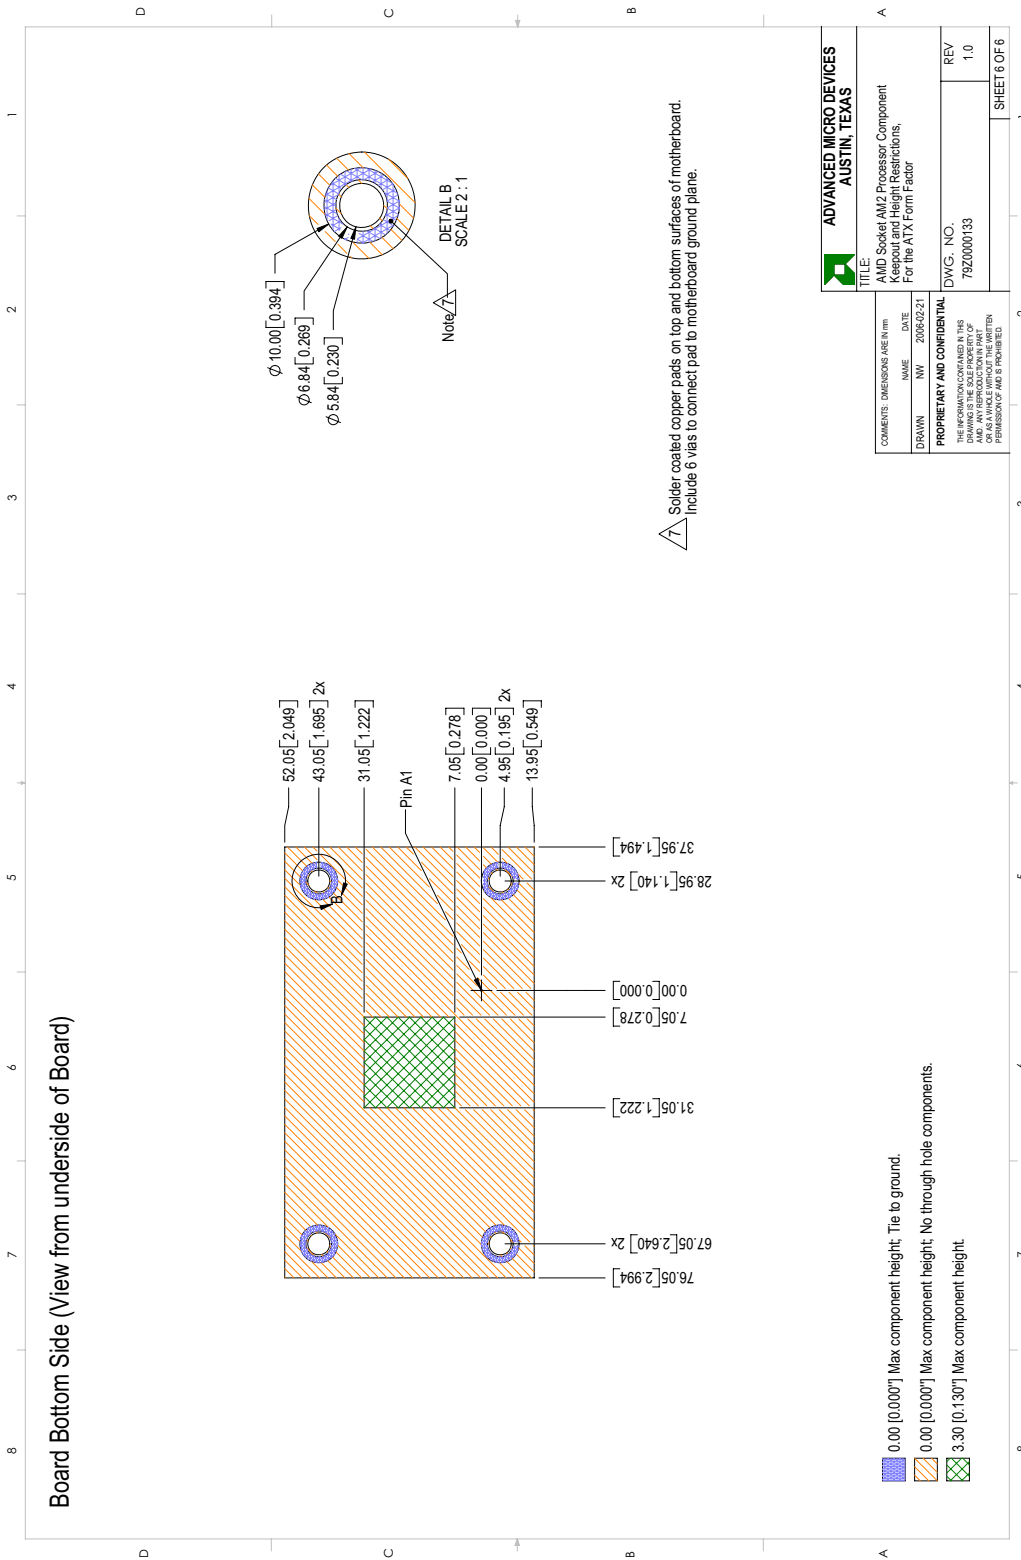




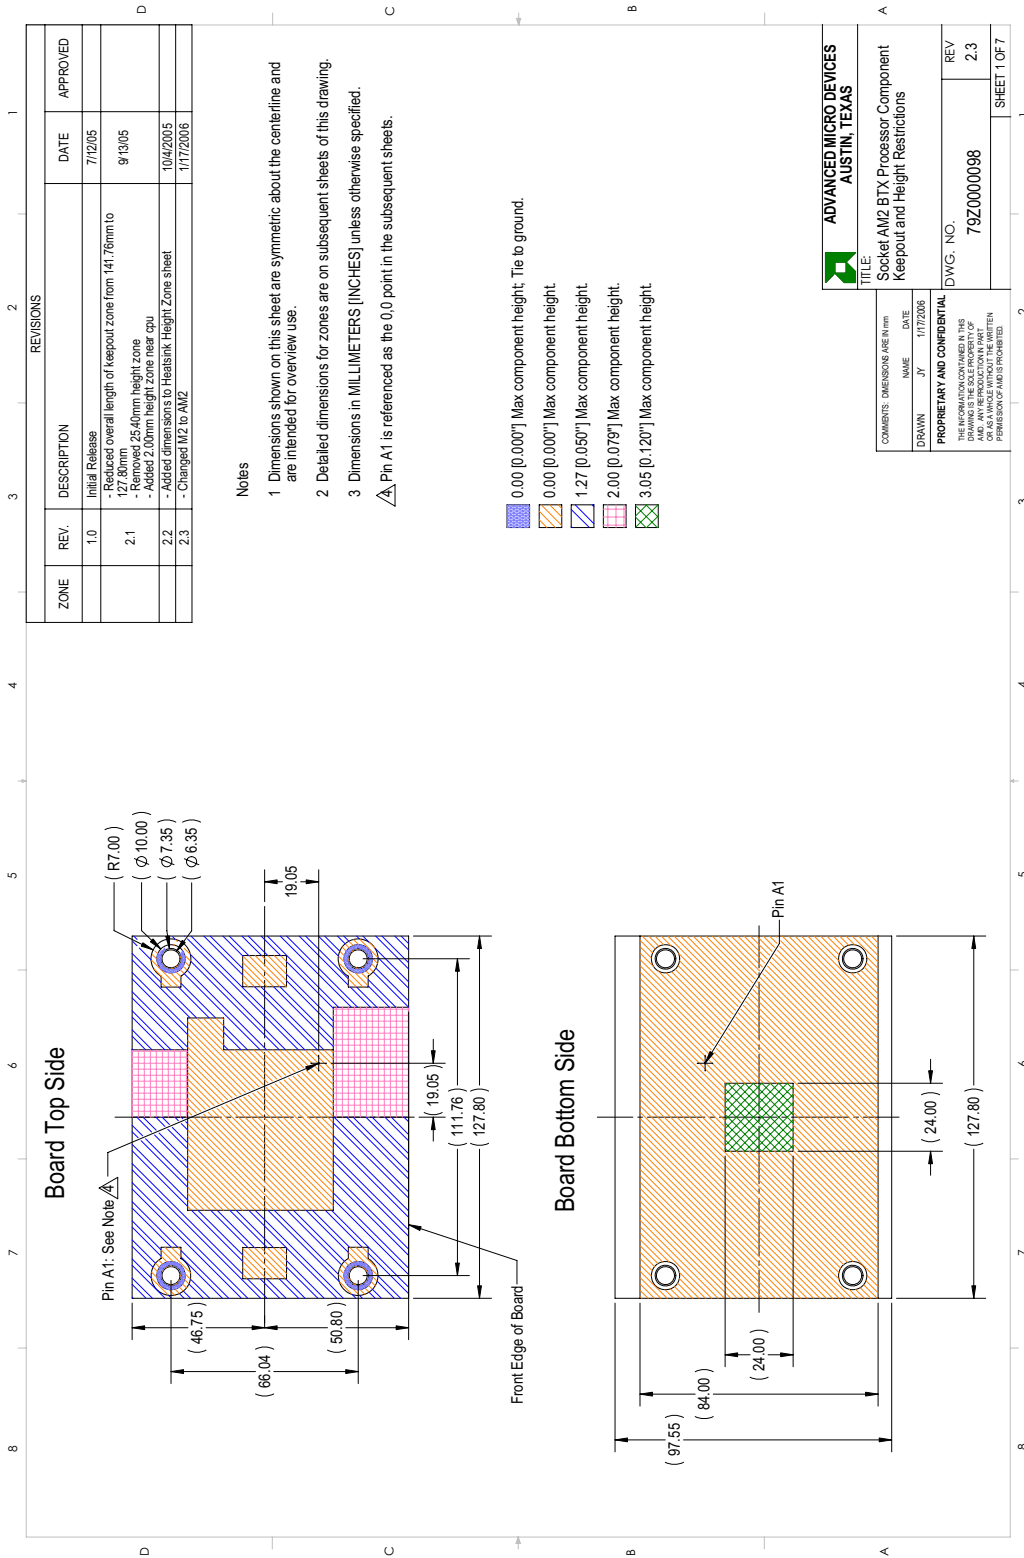











11.2 Socket AM2 BTX Processor Component Keepout and Height Restrictions

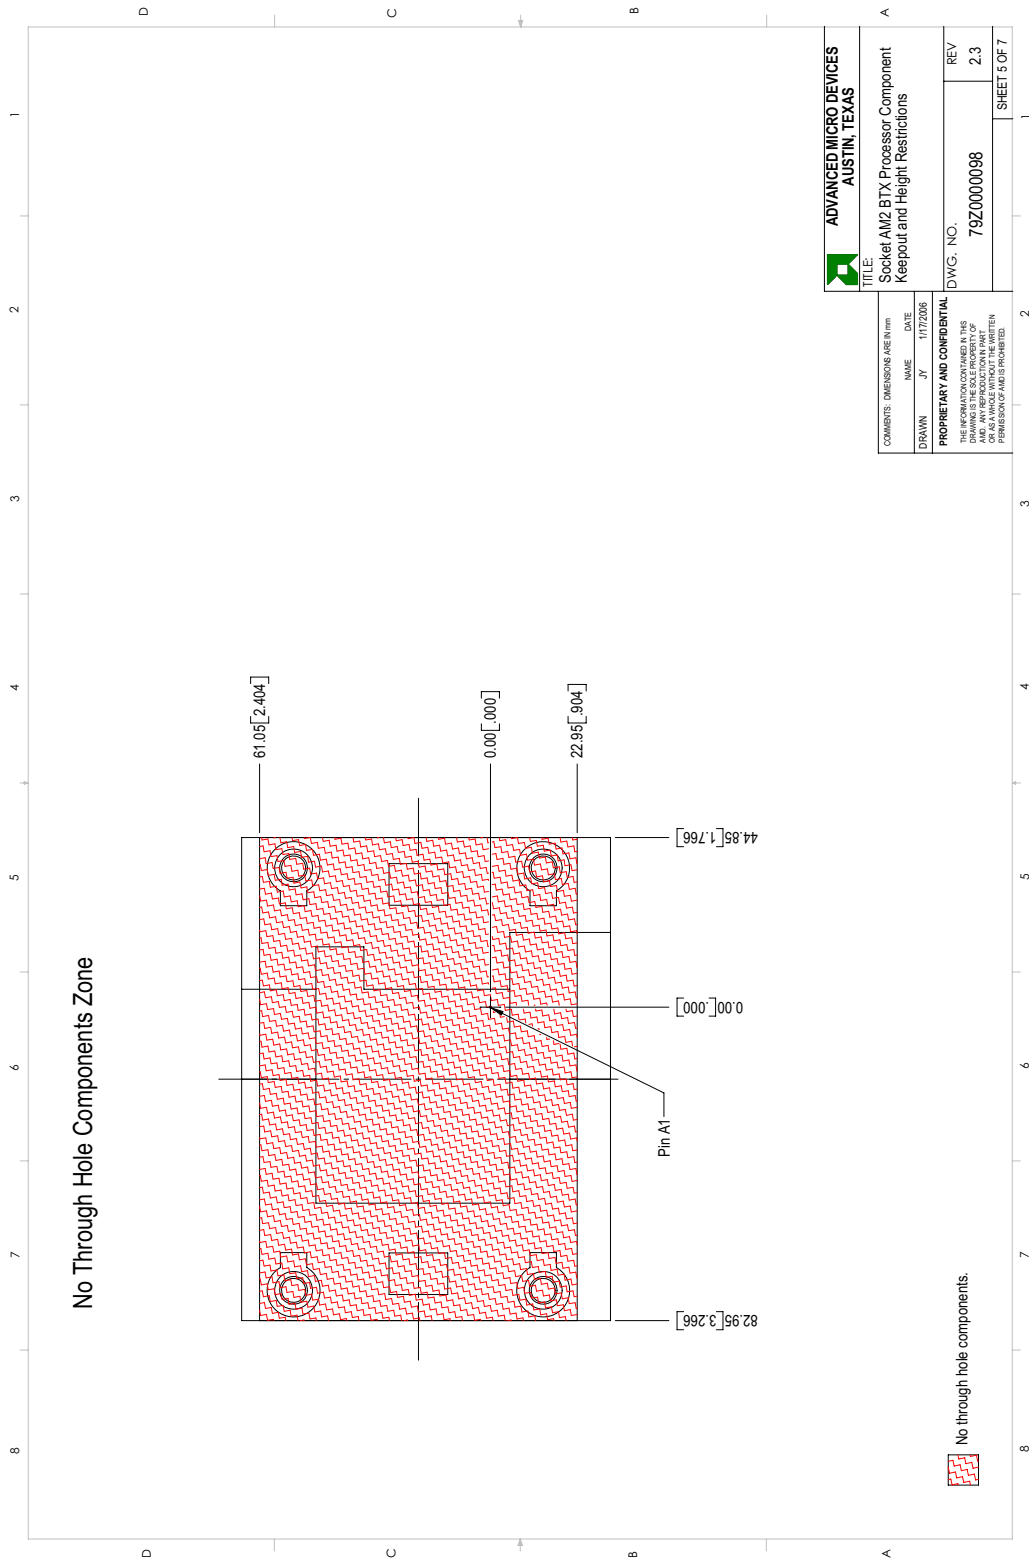






ADVANCED MICRO DEVICES AUSTIN, TEXAS	
TITLE: Socket AM2 BTX Processor Component Keepout and Height Restrictions	
COMMENTS: DIMENSIONS ARE IN mm	DATE: 1/17/2005
DRAWN: JY	NAME:
PROPRIETARY AND CONFIDENTIAL THIS INFORMATION IS UNCLASSIFIED AND UNCONTROLLED UNLESS INDICATED OTHERWISE. ANY REPRODUCTION, PARTIAL OR COMPLETE, WITHOUT THE EXPRESS WRITTEN PERMISSION OF AMD IS PROHIBITED.	
DWG. NO.: 79Z0000098	REV: 2.3
SHEET 3 OF 7	





ADVANCED MICRO DEVICES AUSTIN, TEXAS	
TITLE: Socket AM2 BTX Processor Component Keepout and Height Restrictions	
COMMENTS: DIMENSIONS ARE IN MM	DATE: 11/7/2005
DRAWN: JY	DATE:
PROPRIETARY AND CONFIDENTIAL THIS INFORMATION IS THE PROPERTY OF ADVANCED MICRO DEVICES. IT IS TO BE USED ONLY FOR THE PURPOSES SPECIFIED HEREIN. NO PART OF THIS DOCUMENT OR ANY INFORMATION CONTAINED HEREIN IS TO BE REPRODUCED OR TRANSMITTED IN ANY FORM OR BY ANY MEANS WITHOUT THE WRITTEN PERMISSION OF AMD IS PROHIBITED.	
DWG. NO.: 7920000098	REV: 2.3
SHEET 5 OF 7	

