IBM PowerPC 970FX RISC Microprocessor and CPC945 Bridge and Memory Controller

Design Guide

SA14-970FXDG-03

October 31, 2007

Printed in the United States of America October 2007

IBM Systems and Technology Group
2070 Route 52, Bldg. 330
Hopewell Junction, NY 12533-6351

The IBM home page can be found at **ibm.com®**

The IBM semiconductor solutions home page can be found at **ibm.com**/chips

SA14-970FXDG-03
October 31, 2007

Contents

List of Figures

List of Tables

Revision Log

Revision Date	Description
October 31, 2007	SA14-970FXDG-03 • Edited the document for clarity and consistency.
March 19, 2007	Version 1.2 • Added *Section 1.5 PowerPC 970FX and CPC945 Bridge System Bring Up* on page 23.
August 31, 2006	First Issue

About this Book

This design guide contains design recommendations for systems based on the IBM PowerPC® 970FX Microprocessor[1] and the CPC945 Bridge and Memory Controller[2]. The design considerations, board schematics, and debug recommendations provided in this document and in referenced documents were developed to ensure flexibility for board designers and to help reduce the risk of board design problems. The guidelines presented here are based on experience, simulation, and platform design performed by IBM.

Board designers can also use the associated IBM PowerPC 970FX/CPC945 evaluation board schematics for reference. The board is supplied for the evaluation of the PPC970FX processor and the CPC945 bridge in a potentially wide variety of applications. Therefore, the board circuitry design in many instances is more complex than equivalent circuitry for a specific application might require. In particular, the layout enables probing and analysis. The board is suitable for software development, for benchmarking, and for detailed study of the hardware. However, it is not intended as a true hardware reference design.

The recommendations in this document are subject to change. Verify with your IBM representative that you have the latest versions of all documents before finalizing any designs. All the recommendations are intended to help design a functional system. However, they are only guidelines and do not take the place of design-specific results obtained from signal integrity modeling. IBM provides both Input/Output Buffer Information Specification (IBIS) models and encrypted HSPICE models for the processor interface bus on the PowerPC 970FX processor. IBM also provides encrypted HSPICE models for the CPC945 processor interface, and for the double data rate (DDR) memory, HyperTransport, and PCI Express (PCI-E) interfaces.

Note: Although IBIS can provide a quick analysis of the layout, use of a layout and analysis tool that is HSPICE-capable is recommended. HSPICE gives a more accurate view of the actual signal integrity.

Who Should Read This Book

The design guide is intended for system software and hardware developers who plan to develop their products using the PowerPC 970FX processor and the CPC945 bridge. It is assumed that the reader understands operating systems, microprocessor system design, basic principles of reduced instruction set computer (RISC) processing, and the details of Power Architecture™.

Related Documents

This design guide refers to several other technical documents. Most are available through IBM Customer Connect at www.ibm.com/chips/techlib/techlib.nsf/products/PowerPC.

Note: Check this IBM Web site to ensure that you have the latest documentation before starting a design.

When you register as an IBM Customer Connect user, you can select documents and areas of interest. You will receive e-mail notification whenever updates to selected documents are posted. A weekly e-mail notifies you of newly published documents in your areas of interest. For more information about Customer Connect, visit www.ibm.com/technologyconnect/.

1. In this document, the PowerPC 970FX Microprocessor is abbreviated as PowerPC 970FX processor or PPC970FX.
2. In this document, the CPC945 Bridge and Memory Controller is abbreviated as CPC945 bridge or CPC945.

Note: If you cannot find a listed document on the Web site, it might require a confidential disclosure agreement (CDA). Contact your IBM representative to obtain access to confidential documents through the IBM Customer Connect portal.

IBM PPC970FX Documentation

- *IBM PowerPC 970FX RISC Microprocessor User's Manual*
- *IBM PowerPC 970FX RISC Microprocessor Datasheet*
- *IBM PowerPC 970FX RISC Microprocessor Errata List for DD3.X*
- *IBM PowerPC 970FX Power On Reset Application Note*
- *IBM PowerPC 970FX Boundary Scan Application Note*
- *IBM PowerPC 970MP/CPC945 Automated PI Tuning Application Note*
- *Manual Collection of PowerPC 970FX Processor Thermal Calibration Values Application Note*
- *On Chip Voltage Monitoring on the PowerPC970FX Application Note*
- *PowerPC970FX Thermal Diode Application Note*
- *PowerPC970 Debug Notes*
- *Improving BGA to PCB Thermo-Mechanical Integrity White Paper*
- *IBIS Model for the PPC970FX*
- *PowerPC 970FX HSpice Models*

IBM CPC945 Documentation

- *CPC945 Bridge and Memory Controller User's Manual*
- *CPC945 Bridge and Memory Controller Datasheet*
- *IBM PowerPC CPC945 Errata List for DD1.X (confidential)*
- *CPC945 Chip Thermal Considerations Application Note*
- *CPC945 Memory Signal Delay Tuning Application Note*
- *CPC945 BSDL file*

The hardware documentation for the PowerPC 970FX/CPC945 evaluation board is included in the CD-ROM delivered with every board. The firmware source and documentation for the PowerPC 970FX and the AMCC PowerPC 405GPr service processor is included in the CD-ROM delivered with every board. It can also be downloaded from the IBM developerWorks® Web site: http://www.ibm.com/developerworks/power/pibs/

Conventions and Notations Used in This Book

The use of overbars designates signals that are active low or the complement of differential signals. For example, $\overline{\text{DDEL\_OUT}}$.

Acronyms and Abbreviations

AD	address/data
ADIN	processor interface inbound
ADOUT	processor interface outbound
ALU	arithmetic logic unit
ATA	advanced technology attachment
AV_{DD}	analog phase-locked loop (PLL) supply voltage
BA	bank address
BGA	ball grid array
BOM	bill of materials
BSDL	boundary scan description language
CAS	column enable
CBGA	ceramic ball grid array
CDA	confidential disclosure agreement
CK	double data rate two (DDR2) dynamic random access memory (DRAM) clock
CKE	clock enable
CS	chip select
DDR	double data rate
DDR2	double data rate two
DIMM	dual inline memory modules
DM	data mask
DQ	data line
DQS	strobe line
DRAM	dynamic random access memory
ECC	error checking and correction
EMF	electromagnetic fields
EMI	electromagnetic interference
ENET	ethernet
EPOS	IBM Embedded PowerPC Operating System

ESD	electrostatic discharge
FCC	Federal Communications Commission
FRAM	ferroelectric random access memory (RAM)
GPIO	general-purpose input/output
HSTL	high-speed transceiver logic
HT	HyperTransport
I^2C	interintegrated circuit
IAP	initial alignment procedure
IBIS	Input/Output Buffer Information Specification
IDE	integrated drive electronics
IEEE	Institute of Electrical and Electronics Engineers
JTAG	Joint Test Action Group
L2	level 2 cache
ODT	on-die termination
OV_{DD}	I/O supply voltage
PCB	printed circuit board
PCI	peripheral component interconnect
PCI-E	peripheral component interconnect express
PCI-X	peripheral component interconnect extended
PI	processor interface
PIBS	PowerPC initialization boot software
PLL	phase-locked loop
POR	power-on reset
RAM	random access memory
RAS	row enable
RISC	reduced instruction set computer
ROM	read-only memory
SDRAM	synchronous dynamic random access memory
SIMD	single-instruction, multiple-data

SMP	symmetric multiprocessor
SOC	system-on-a-chip
SPD	serial presence detect
SPU	service processor unit
SRB	snoop response bus
SSB	source-synchronous buses
SSTL	stub series terminated logic
SYSCLK	system clock
TH	transfer handshake bus
USB	universal serial bus
V_{DD}	core voltage
VFC	V_{DD} fuse code
VPU	vector processing unit
V_{REF}	DDR2 reference voltage
VRM	voltage regulation modules
V_{TT}	termination voltage
WE	write enable

1. System Overview

This section summarizes the key features of the IBM PowerPC 970FX Microprocessor and the CPC945 Bridge and Memory Controller.

1.1 PowerPC 970FX Processor

The PowerPC 970FX is a 64-bit PowerPC microprocessor with vector processing unit (VPU) extensions—the single-instruction, multiple-data (SIMD) operations that accelerate data intensive processing tasks. This processor is designed to support multiple system configurations ranging from desktop and low-end server applications and uniprocessors, up through 4-way symmetric multiprocessor (SMP) configurations.

1.1.1 Architectural Features

- 64-bit implementation of the PowerPC Architecture (version 2.0.1)
- 1.0 to 2.2 GHz core frequency operation, supporting high-speed processor interface data-transfer speeds up to 1.1 GHz for original equipment manufacturer (OEM) applications
- 10 execution units:
  - Two integer
  - Two floating point (single or double precision)
  - Two load/store
  - Two vector/SIMD (one combined arithmetic logic unit [ALU], one permute)
  - Condition unit
  - Branch units
- 64 KB direct-mapped instruction cache (I-Cache), 32 KB 2-way, set-associative data cache (D-Cache), both with parity protection
- 512 KB level 2 (L2) cache with error checking and correction (ECC)
- Thermal diode (application-specific calibration necessary)
- Support for multiple thermal management modes:
  - Static power management
    - Software initiated doze, nap, and deep nap modes. (See the *IBM PowerPC 970FX RISC Microprocessor User's Manual* for more information.)
  - Dynamic power management
    - Certain sections of the design stop their hardware-initiated clocks when not in use. (See the *IBM PowerPC 970FX RISC Microprocessor User's Manual* for more information.)
  - Power tuning
    - Software-initiated slow down of the processor; selectable to half of the nominal operating frequency. (See the *IBM PowerPC 970FX RISC Microprocessor User's Manual* for more information.)

1.1.2 Packaging and Power

- Available in a 576-pin ceramic ball grid array (CBGA), 25 × 25 mm square, 1.0 mm pitch package.
- Packaging options include reduced lead CBGA for Restriction of Hazardous Substances Directive (RoHS) compliance.
- Core voltage (V_{DD}) operation depends on part number and application conditions, and can range from 1.0 V to 1.25 V. (See the *IBM PowerPC 970FX RISC Microprocessor Datasheet* for more information.)
- I/O supply voltage (OV_{DD}) at 1.50 V.
- Analog phase-locked loop (PLL) supply voltage (AV_{DD}) at 2.80 V.

For additional details on packaging and power, see the *IBM PowerPC 970FX RISC Microprocessor Datasheet*.

1.1.3 Power Tuning and Power Management

The power tuning engine controls the power management modes and on-chip and off-chip clock frequency, and supports voltage scaling for the PowerPC 970FX processor. To implement power tuning as described here, your design must use a bridge or system controller that supports this feature. In the PPC970FX/CPC945 evaluation board design, all processors and processor interfaces in the CPC945 bridge change the power tuning mode concurrently for transitions between full-power and half-power mode. Either processor can request the mode change. This information is then transmitted to the CPC945 bridge through the processor interface bus as a special request. The CPC945 bridge grants the requests, mirrors this special request to all processors, and waits for all processors to signal that they have quiesced the processor interface bus and are ready to switch mode. The CPC945 bridge then triggers the mode switch, which is completed within 200 nanoseconds for most bus ratios. The frequency scaling on the processor interface bus requires changing certain processor interface timing parameters. Since the I/O voltage is not changed, an initial alignment procedure (IAP) is not required. The new parameters are sent along with the power tuning command and overwrite the old parameters when the frequency switch occurs. For additional details on power tuning, see the *IBM PowerPC 970FX RISC Microprocessor User's Manual*.

1.1.4 Thermal Management Using a Thermal Diode Control Application

The PowerPC 970FX processor features an on-board temperature sensing diode, located in the hottest portion of the chip, connected to pins AA1 (DIODENEG) and Y1 (DIODEPOS). For information on this diode, see the *IBM PowerPC 970FX RISC Microprocessor Datasheet*. For information about manual calibration techniques and examples, see the *Manual Collection of PowerPC 970FX Processor Thermal Calibration Values* and *Thermal Considerations: PowerPC 970FX* application note.

Other temperature monitoring hardware can also be implemented with the PowerPC 970FX processor and mounted as close to the PowerPC 970FX processor as practical. The PPC970FX/CPC945 evaluation board provides one example of external temperature-sensing hardware; see the PPC970FX/CPC945 evaluation board documentation for details. The documentation describes a programmable system-on-a-chip (SOC)-based circuit that monitors the PowerPC 970FX thermal diode, and the service processor firmware source code that queries the monitoring device to determine if an unsafe operating temperature has been detected.

1.2 CPC945 Bridge and Memory Controller

The CPC945 Bridge and Memory Controller is a frontside bus controller that is compatible with single core (PowerPC 970FX) and dual core (PowerPC 970MP) PowerPC 970 microprocessors. It provides a 5-way interconnection among two PowerPC 970 processor interfaces, a double data rate (DDR) synchronous dynamic random access memory (SDRAM) subsystem, a Peripheral Component Interconnect (PCI)-Express root complex, and a HyperTransport host bridge.

1.2.1 Architectural Features

- Dual 36-bit PowerPC 970FX processor interface buses with cache coherency and snooping, running up to 550 MHz (1000 megatransfers per second [MTps])
- PCI Express interface
  - x1, x4, x8, x16
- 128-bit, 533-MTps DDR SDRAM controller and interface with ECC
- One slave and two master interintegrated circuit (I^2C) interfaces
- 16-bit HyperTransport host bridge interface, operating at up to 1000 GTps (1 gigatransfer per second)
- Interrupt controller with support for 8 internal and up to 120 external interrupts (routed over the HyperTransport interface)

See the *CPC945 Bridge and Memory Controller User's Manual* for additional information.

1.2.2 Packaging and Power

- Available in a 1182-pin, flip-chip, plastic ball grid array (BGA), 37.5 × 37.5 mm square, 1.0 mm pitch package.

For additional details on packaging and power, see the *CPC945 Bridge and Memory Controller Datasheet*.

1.2.3 Power Tuning and Power Management

To optimize the electrical power consumption and thermal performance of systems built with the CPC945 Bridge and Memory Controller, power management logic in the CPC945 bridge switches off clocks to various parts of the chip when those parts are not needed. Logic also controls the speeds at which different interfaces operate, allowing additional power savings and configuration control.

The CPC945 Clock Control Register provides software access to manage power by turning off unused clocks and PLLs. All clocks and PLLs are stopped and started cleanly without spikes or short cycles, so logic does not have to be reset after stopping and then restarting clocks.

At the top of the power management hierarchy, the CPC945 bridge controls the idling of all buses and internal operations as the system goes in and out of the system sleep state. This is done in tandem with the service processor unit (SPU), an external circuit or microprocessor that manages power-up and hardware resets and their relationships to the suspend and power management functions of the rest of the system. The CPC945 bridge and the SPU have a 2-wire handshake interface consisting of the suspend request ($\overline{\text{SUSPENDREQ}}$) and suspend acknowledgment ($\overline{\text{SUSPENDACK}}$) signals. Using these signals and the control bits in the CPC945 System Power Management Register, the CPC945 bridge safely suspends all internal operations, idles all external buses, and places the memory subsystem in the self-refresh state.

1.2.4 Thermal Management Using a Thermal Diode Control Application

The CPC945 bridge features an on-board, conventional thermal monitor. The thermal monitor consists of a vertical PNP bipolar junction transistor, with the emitter connected to pin AP30 (SYS_THDIOD_D) and the base connected to pin AP29 (SYS_THDIO_G). The collector is grounded. The bipolar is parasitic; the element of interest in this circuit is the emitter-base diode. The emitter-base diode is connected such that the voltage on AP30 must be greater than the voltage on AP29 to forward bias the diode. The forward-bias voltage AP30 - AP29 should be between 0.50 - 0.70 V. Because of the electrostatic discharge (ESD) protection diodes, the voltages on both AP30 and AP29 should be between 0 and V_{DD}. If this guideline is followed, the ESD protection diodes will not conduct measurable current. The chip temperature can be determined with standard, off-the-shelf, thermal monitoring devices.

1.3 Peak Bandwidth Summary

- The maximum core frequency supported for PowerPC 970FX general market availability is 2.2 GHz.
- The fastest processor interface transfer rate for a 2.2 GHz core frequency is attained with a bus ratio of 2:1 and is 550 MHz DDR or 1.1 GTps.
- The CPC945 memory interface supports DDR2 operational data rates of 400 MTps and 533 MTps. The memory interface can be configured as either a 64-bit or a 128-bit (72-bit or 144-bit with ECC) wide interface, with a maximum bandwidth of 8.53 GBps.
- The CPC945 bridge supports a 16-bit PCI Express port
- The CPC945 HyperTransport interface supports a 16-bit wide link that operates at a default data rate of 200 MHz for 400 MTps × 2 bytes = 800 MBps per direction (1.6 GBps in total). The CPC945 DD1.0 hardware supports data rates up to 500 MHz for 1 GTps × 2 bytes = 2 GBps maximum bandwidth per direction (4 GBps in total).

1.4 PPC970FX/CPC945 Evaluation Board System Configuration

As *Figure 1-1* on page 22 shows, the PPC970FX/CPC945 evaluation board includes the following major system components:

- Dual PowerPC 970FX reduced instruction set (RISC) processors
- AMCC PowerPC 405GPr service and control processor
- CPC945 bridge and memory controller
- 16-bit bus between the CPC945 bridge and the first HyperTransport tunnel to peripheral component interconnect express (PCI-X) bridge (HT7520)
- 8-bit buses between the HyperTransport tunnel to PCI-X bridge and the AMD-8111 HyperTransport hub (Southbridge)
- Dual 10/100/1000 (gigabit ethernet [GigE]) ethernet ports off the PCI-X bus to the I/O area RJ-45 plugs
- Intel® GD31244 Quad Serial ATA disk peripheral
- Single 10/100 ethernet port off the Southbridge
- Two universal serial bus (USB) 2.0 ports and two serial ports
- Advanced technology attachment (ATA) 133-capable integrated drive electronics (IDE) interface (master and slave)

- · Expansion slots
  - Two 3.3 V 64-bit PCI-X slots that run at up to 133 MHz
  - One 3.3 V 32-bit PCI slot running at 33 MHz
  - PCI Express connector

The following sections describe board layout and manufacturing guidelines that are specific to this particular design. This information should be used as a reference. Other designs are possible but must be verified through both simulation and operation.

Figure 1-1. PPC970FX/CPC945 Evaluation Board with Two Processor Cards

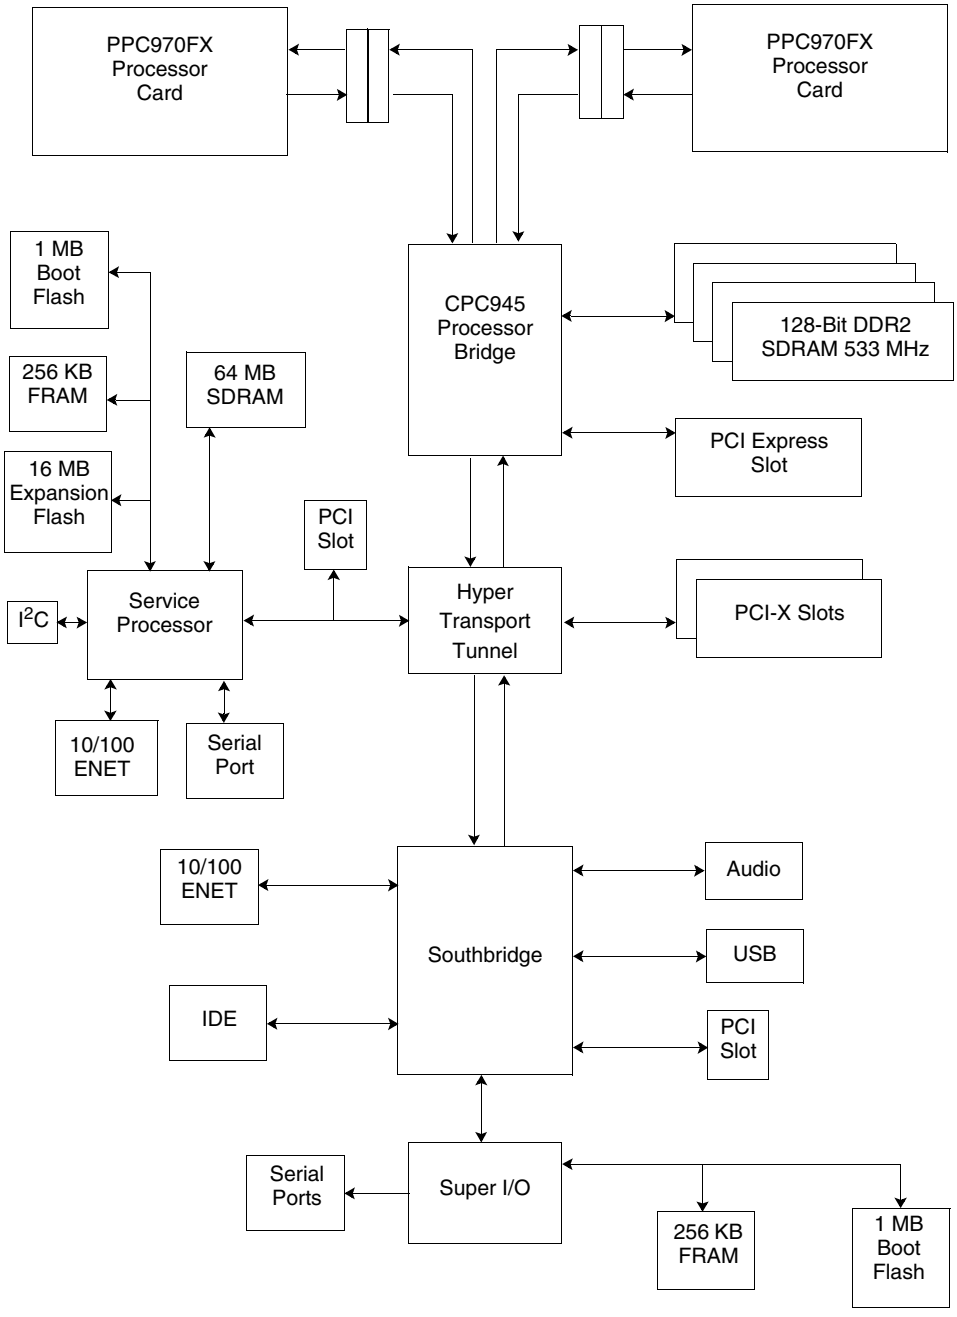

DDR2	Double data rate two
ENET	Ethernet
FRAM	Ferroelectric random access memory (RAM)
I^2C	Interintegrated circuit
IDE	Integrated drive electronics
PCI	Peripheral component interconnect
PCI-X	Peripheral component interconnect-X
SDRAM	Synchronous dynamic random access memory
USB	Universal serial bus

Note: The processor cards are separate assemblies that plug onto 400-pin MEG-array connectors.

1.5 PowerPC 970FX and CPC945 Bridge System Bring Up

Because these processors are configurable, parameters must be set to define the configuration. The *PowerPC 970FX Power-On Reset Application Note* provides the configuration details. However, it does not describe the actual bring up sequence because that is specific to each design. The *CPC945 Bridge and Memory Controller User Manual* has sections on initialization for the processor interface, the DDR controller, and the HyperTransport interface.

PowerPC 970FX designs use a service processor to initialize the system board. A significant portion of this initialization is done using the I^2C bus. Whether by setting general purpose I/O bits or using the I^2C interface, the service processor runs a start up script or program for initialization.

The service processor used on the PPC970FX/CPC945 evaluation board is an AMCC PowerPC 405GPr. The actual script used to initialize the PPC970FX/CPC945 evaluation board is provided as an example or starting point for each customer's specific design. This script brings up the voltages on the board, sets the clocks, and configures the processor interface, the DDR memory interface, the HyperTransport interface, and the PCI Express interface. The script steps the processor through the power-on reset (POR) stages described in the *IBM PowerPC 970FX Power On Reset Application Note*. The example scripts can be found on the IBM developerWorks Web site in the Power Architecture area under downloads and products (http://www.ibm.com/developerworks/power/pibs/). The scripts are part of IBM PowerPC 970/CPC945 Evaluation Kit Software.

At this point, the system has been initialized with configuration values for the processor interface that allow the system to run but might not be optimized for the specific processor interface design on the customer board. To optimize the processor interface configuration values for both the PowerPC 970FX processor and the CPC945 bridge, see the *IBM PowerPC 970MP/CPC945 Automated PI Tuning Application Note*. Some utility code is provided with the application note to help select the optimum configuration parameters. This utility code runs on the service processor or a host that has I^2C access to system. The code consists of Perl scripts that must be modified for the target system. Both the processor interface tuning information and the scripts can be obtained from Customer Connect (see *Related Documents* on page 11).

The configuration values for the DDR2 memory interface also must be tuned for each design. The *CPC945 Memory Signal Delay Tuning Application Note* describes how to properly tune the DDR2 memory interface. This note describes techniques for obtaining required measurements with lab equipment and provides scripts to complete the procedure. Also, as noted in the *CPC945 Bridge and Memory Controller Datasheet* and the *Memory Signal Delay Tuning Application Note*, CPC945 parts used for the initial tuning phase should cover the entire manufacturing process window. The *CPC945 Bridge and Memory Controller Datasheet* describes how to get this range of parts.

When the interfaces have been tuned, update the startup script with the optimized processor interface parameter values for both the processor and the CPC945 bridge, and update the DDR2 memory parameters for the CPC945 bridge.

2. Component Dimensions and Physical Layout

For information about packaging options and dimensions, see the mechanical packaging section of the *IBM PowerPC 970FX RISC Microprocessor Datasheet*.

Figure 2-1, *Figure 2-2*, and *Figure 2-3* show the signal assignments for the PowerPC 970FX processor and the CPC945 bridge, with color coding to indicate signal groups as defined in *Table 2-1* and *Table 2-2*.

Table 2-1. Signal Groupings for the PowerPC 970FX Microprocessor (see *Figure 2-1* on page 26 and *Figure 2-2* on page 27)

Signal Group	Color
Processor interface inbound (ADIN) signals	
Processor interface inbound snoop signals	
Processor interface inbound bus clock signals	
Processor interface outbound (ADOUT) signals	
Processor interface outbound snoop signals	
Processor interface outbound bus clock signals	
Processor interface status and control signals	
Clock control signals	
Interrupts and resets	
Joint Test Action Group (JTAG) debug and interintegrated circuit (I^2C) signals	
Thermal diodes and Kelvin pins	

Table 2-2. Signal Groupings for the CPC945 Bridge and Memory Controller (see *Figure 2-3* on page 28)

Signal Group	Color
Processor interface input signals	
Processor interface output signals	
HyperTransport interface signals	
Double data rate (DDR) interface signals	
Peripheral Component Interconnect (PCI) Express interface signals	

Figure 2-1. IBM PowerPC 970FX Microprocessor Ball Placement (Top View)

AD24 GND	AD23 OVDD	AD22 TMS	AD21 TCK	AD20 GND	AD19 OVDD	AD18 MCP	AD17 TBEN	AD16 GND	AD15 OVDD	AD14 PSYNC_ OUT	AD13 TDO	AD12 ATTEN- TION	AD11 LSSD_S TOP_EN ABLE	AD10 GND	AD9 OVDD	AD8 LSSD_S TOPC2_ ENABLE	AD7 LSSD_S TOPC2S TAR_EN ABLE	AD6 GND	AD5 VDD	AD4 GND	AD3 VDD	AD2 GND	AD1 OVDD	AD
AC24 BI_MOD E	AC23 GND	AC22 VDD	AC21 GND	AC20 VDD	AC19 BUS_CF G1	AC18 VDD	AC17 GND	AC16 C1_UND _GLOBA L	AC15 C2_UND _GLOBA L	AC14 VDD	AC13 GND	AC12 VDD	AC11 GND	AC10 PULSE_ SEL2	AC9 PULSE_ SEL0	AC8 VDD	AC7 GND	AC6 VDD	AC5 GND	AC4 VDD	AC3 GND	AC2 VDD	AC1 GND	AC
AB24 SYNC_E NABLE	AB23 VDD	AB22 GND	AB21 TDI	AB20 GND	AB19 INT	AB18 GND	AB17 VDD	AB16 BUS_CF G2	AB15 VDD	AB14 GND	AB13 VDD	AB12 QREQ	AB11 PULSE_ SEL1	AB10 GND	AB9 VDD	AB8 GND	AB7 PLL_RA NGE0	AB6 RAMSTO P_ENAB LE	AB5 LSS- DMODE	AB4 SRESET	AB3 VDD	AB2 GND	AB1 VDD	AB
AA24 OVDD	AA23 GND	AA22 GPUL- DBG	AA21 GND	AA20 I2CCK	AA19 BUS_CF G0	AA18 VDD	AA17 GND	AA16 VDD	AA15 GND	AA14 CKTERMI _DIS	AA13 SPARE	AA12 AFN	AA11 GND	AA10 PSYNC	AA9 PLL_RA NGE1	AA8 PLL_MU LT	AA7 GND	AA6 OVDD	AA5 RI	AA4 VDD	AA3 GND	AA2 VDD	AA1 DIODEN EG	AA
Y24 GND	Y23 VDD	Y22 VDD	Y21 I2CDT	Y20 GND	Y19 VDD	Y18 GND	Y17 OVDD	Y16 GND	Y15 VDD	Y14 GND	Y13 OVDD	Y12 GND	Y11 VDD	Y10 GND	Y9 VDD	Y8 GND	Y7 VDD	Y6 GND	Y5 VDD	Y4 GND	Y3 VDD	Y2 GND	Y1 DIODE- POS	Y
W24 OVDD	W23 AVP_RE SET	W22 PLLTEST	W21 GND	W20 TRST	W19 GND	W18 VDD	W17 GND	W16 VDD	W15 GND	W14 VDD	W13 GND	W12 VDD	W11 GND	W10 VDD	W9 GND	W8 VDD	W7 GND	W6 VDD	W5 GND	W4 SPARE2	W3 GND	W2 VDD	W1 GND	W
V24 BYPASS	V23 PSRO0	V22 THERM_I NT	V21 QACK	V20 HRESET	V19 VDD	V18 GND	V17 VDD	V16 GND	V15 VDD	V14 GND	V13 VDD	V12 GND	V11 VDD	V10 GND	V9 VDD	V8 GND	V7 VDD	V6 GND	V5 PSRO_ Enable	V4 GND	V3 VDD	V2 GND	V1 VDD	V
U24 DI2	U23 GND	U22 VDD	U21 GND	U20 VDD	U19 LSSD_S CAN_EN ABLE	U18 VDD	U17 GND	U16 VDD	U15 GND	U14 VDD	U13 GND	U12 VDD	U11 GND	U10 VDD	U9 GND	U8 VDD	U7 GND	U6 VDD	U5 GND	U4 VDD	U3 GND	U2 VDD	U1 GND	U
T24 GND	T23 VDD	T22 SYSCLK	T21 OVDD	T20 PLL_LOC K	T19 PLLT- ESTOUT	T18 GND	T17 VDD	T16 GND	T15 OVDD	T14 GND	T13 VDD	T12 GND	T11 OVDD	T10 GND	T9 VDD	T8 GND	T7 VDD	T6 GND	T5 VDD	T4 GND	T3 VDD	T2 KVPRB- GND	T1 VDD	T
R24 ANALOG _GND	R23 GND	R22 SYSCLK	R21 GND	R20 CHK- STOP	R19 GND	R18 VDD	R17 GND	R16 VDD	R15 GND	R14 VDD	R13 GND	R12 VDD	R11 GND	R10 VDD	R9 GND	R8 VDD	R7 GND	R6 VDD	R5 GND	R4 VDD	R3 GND	R2 KVPRB- VDD	R1 Z_SENS E	R
P24 AVDD	P23 VDD	P22 GND	P21 VDD	P20 EI_DISA BLE	P19 VDD	P18 GND	P17 OVDD	P16 GND	P15 VDD	P14 GND	P13 VDD	P12 GND	P11 VDD	P10 GND	P9 VDD	P8 GND	P7 VDD	P6 GND	P5 VDD	P4 GND	P3 VDD	P2 Z_OUT	P1 VDD	P
N24 OVDD	N23 GND	N22 I2CGO	N21 TRIG- GERIN	N20 VDD	N19 TRIG- GEROUT	N18 VDD	N17 GND	N16 VDD	N15 GND	N14 VDD	N13 GND	N12 VDD	N11 GND	N10 VDD	N9 GND	N8 VDD	N7 GND	N6 VDD	N5 GND	N4 VDD	N3 ADOUT0	N2 VDD	N1 SPARE_ GND	N
M24 GND	M23 VDD	M22 GND	M21 VDD	M20 GND	M19 PROCID 1	M18 PROCID 2	M17 VDD	M16 GND	M15 VDD	M14 GND	M13 VDD	M12 GND	M11 VDD	M10 GND	M9 OVDD	M8 GND	M7 VDD	M6 GND	M5 VDD	M4 GND	M3 ADOUT4	M2 GND	M1 OVDD	M
L24 SRIN0	L23 GND	L22 SRIN1	L21 SRIN1	L20 VDD	L19 PROCID 0	L18 VDD	L17 GND	L16 OVDD	L15 GND	L14 VDD	L13 GND	L12 VDD	L11 GND	L10 OVDD	L9 GND	L8 VDD	L7 GND	L6 VDD	L5 GND	L4 VDD	L3 SROUT0	L2 SROUT0	L1 ADOUT3	L
K24 SRIN0	K23 VDD	K22 ADIN6	K21 OVDD	K20 GND	K19 VDD	K18 GND	K17 OVDD	K16 GND	K15 VDD	K14 GND	K13 VDD	K12 GND	K11 VDD	K10 GND	K9 VDD	K8 GND	K7 VDD	K6 GND	K5 OVDD	K4 ADOUT5	K3 ADOUT2	K2 ADOUT6	K1 GND	K
J24 ADIN8	J23 GND	J22 ADIN3	J21 ADIN1	J20 VDD	J19 GND	J18 OVDD	J17 GND	J16 VDD	J15 GND	J14 OVDD	J13 GND	J12 VDD	J11 GND	J10 OVDD	J9 GND	J8 OVDD	J7 GND	J6 VDD	J5 GND	J4 VDD	J3 GND	J2 VDD	J1 OVDD	J
H24 OVDD	H23 ADIN7	H22 ADIN2	H21 ADIN0	H20 GND	H19 VDD	H18 GND	H17 VDD	H16 GND	H15 VDD	H14 GND	H13 VDD	H12 GND	H11 VDD	H10 GND	H9 VDD	H8 GND	H7 VDD	H6 GND	H5 VDD	H4 GND	H3 ADOUT7	H2 ADOUT1	H1 ADOUT8	H
G24 ADIN13	G23 GND	G22 VDD	G21 ADIN11	G20 ADIN9	G19 ADIN14	G18 VDD	G17 GND	G16 VDD	G15 GND	G14 VDD	G13 GND	G12 VDD	G11 GND	G10 VDD	G9 GND	G8 VDD	G7 GND	G6 VDD	G5 GND	G4 ADOUT9	G3 ADOUT1 3	G2 GND	G1 SROUT1	G
F24 GND	F23 ADIN10	F22 GND	F21 ADIN25	F20 GND	F19 OVDD	F18 GND	F17 VDD	F16 GND	F15 VDD	F14 GND	F13 VDD	F12 GND	F11 VDD	F10 GND	F9 VDD	F8 GND	F7 OVDD	F6 GND	F5 VDD	F4 ADOUT1 1	F3 OVDD	F2 ADOUT1 0	F1 SROUT1	F
E24 CLKIN	E23 GND	E22 VDD	E21 ADIN22	E20 ADIN21	E19 GND	E18 VDD	E17 GND	E16 VDD	E15 GND	E14 VDD	E13 GND	E12 ADOUT1 6	E11 GND	E10 VDD	E9 GND	E8 VDD	E7 GND	E6 VDD	E5 GND	E4 VDD	E3 CLK- OUT	E2 ADOUT1 2	E1 OVDD	E
D24 CLKIN	D23 VDD	D22 ADIN12	D21 VDD	D20 ADIN32	D19 VDD	D18 ADIN30	D17 VDD	D16 GND	D15 ADIN18	D14 GND	D13 VDD	D12 GND	D11 ADOUT1 5	D10 GND	D9 OVDD	D8 ADOUT2 7	D7 VDD	D6 ADOUT2 3	D5 VDD	D4 GND	D3 CLK- OUT	D2 ADOUT2 6	D1 GND	D
C24 OVDD	C23 GND	C22 ADIN20	C21 GND	C20 VDD	C19 ADIN34	C18 ADIN35	C17 ADIN29	C16 ADIN42	C15 ADIN17	C14 ADIN28	C13 ADIN4	C12 ADOUT2 8	C11 ADOUT1 9	C10 ADOUT4 1	C9 ADOUT3 9	C8 ADOUT3 5	C7 ADOUT3 3	C6 ADOUT3 6	C5 ADOUT2 1	C4 ADOUT3 2	C3 GND	C2 OVDD	C1 ADOUT2 0	C
B24 ADIN24	B23 ADIN23	B22 GND	B21 ADIN31	B20 OVDD	B19 ADIN27	B18 GND	B17 ADIN26	B16 OVDD	B15 ADIN15	B14 GND	B13 VDD	B12 GND	B11 OVDD	B10 ADOUT1 8	B9 GND	B8 ADOUT1 4	B7 OVDD	B6 ADOUT3 0	B5 GND	B4 ADOUT3 1	B3 VDD	B2 ADOUT2 2	B1 GND	B
A24 OVDD	A23 ADIN37	A22 ADIN33	A21 ADIN36	A20 ADIN38	A19 ADIN43	A18 ADIN39	A17 ADIN41	A16 ADIN19	A15 ADIN40	A14 ADIN16	A13 ADIN5	A12 ADOUT2 9	A11 ADOUT1 7	A10 ADOUT4 0	A9 ADOUT3 8	A8 ADOUT4 2	A7 ADOUT3 4	A6 ADOUT4 3	A5 ADOUT2 4	A4 ADOUT3 7	A3 GND	A2 ADOUT2 5	A1 OVDD	A
24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	

Note: This diagram is oriented as if looking down through the IBM PowerPC 970FX RISC Microprocessor with it placed and soldered on the system board - a top view.

Figure 2-2. IBM PowerPC 970FX Microprocessor Ball Placement (Bottom View)

AD	AD1 OVDD	AD2 GND	AD3 VDD	AD4 GND	AD5 VDD	AD6 GND	AD7 LSSD_S TOPC2S TAR_EN ABLE	AD8 LSSD_S TOPC2_ ENABLE	AD9 OVDD	AD10 GND	AD11 LSSD_S TOP_EN ABLE	AD12 ATTEN- TION	AD13 TDO	AD14 PSYNC_ OUT	AD15 OVDD	AD16 GND	AD17 TBEN	AD18 MCP	AD19 OVDD	AD20 GND	AD21 TCK	AD22 TMS	AD23 OVDD	AD24 GND
AC	AC1 GND	AC2 VDD	AC3 GND	AC4 VDD	AC5 GND	AC6 VDD	AC7 GND	AC8 VDD	AC9 PULSE_ SEL0	AC10 PULSE_ SEL2	AC11 GND	AC12 VDD	AC13 GND	AC14 VDD	AC15 C2_UND _GLOBA L	AC16 C1_UND _GLOBA L	AC17 GND	AC18 VDD	AC19 BUS_CF G1	AC20 VDD	AC21 GND	AC22 VDD	AC23 GND	AC24 BI_MOD E
AB	AB1 VDD	AB2 GND	AB3 VDD	AB4 SRESET	AB5 LSS- DMODE	AB6 RAMSTO P_ENAB LE	AB7 PLL_RA NGE0	AB8 GND	AB9 VDD	AB10 GND	AB11 PULSE_ SEL1	AB12 QREQ	AB13 VDD	AB14 GND	AB15 VDD	AB16 BUS_CF G2	AB17 VDD	AB18 GND	AB19 INT	AB20 GND	AB21 TDI	AB22 GND	AB23 VDD	AB24 SYNC_E NABLE
AA	AA1 DIODEN EG	AA2 VDD	AA3 GND	AA4 VDD	AA5 RI	AA6 OVDD	AA7 GND	AA8 PLL_MU LT	AA9 PLL_RA NGE1	AA10 PSYNC	AA11 GND	AA12 AFN	AA13 SPARE	AA14 CKTERM _DIS	AA15 GND	AA16 VDD	AA17 GND	AA18 VDD	AA19 BUS_CF G0	AA20 I2CCK	AA21 GND	AA22 GPUL- DBG	AA23 GND	AA24 OVDD
Y	Y1 DIODE- POS	Y2 GND	Y3 VDD	Y4 GND	Y5 VDD	Y6 GND	Y7 VDD	Y8 GND	Y9 VDD	Y10 GND	Y11 VDD	Y12 GND	Y13 OVDD	Y14 GND	Y15 VDD	Y16 GND	Y17 OVDD	Y18 GND	Y19 VDD	Y20 GND	Y21 I2CDT	Y22 VDD	Y23 VDD	Y24 GND
W	W1 GND	W2 VDD	W3 GND	W4 SPARE2	W5 GND	W6 VDD	W7 GND	W8 VDD	W9 GND	W10 VDD	W11 GND	W12 VDD	W13 GND	W14 VDD	W15 GND	W16 VDD	W17 GND	W18 VDD	W19 GND	W20 TRST	W21 GND	W22 PLLTEST	W23 AVP_RE SET	W24 OVDD
V	V1 VDD	V2 GND	V3 VDD	V4 GND	V5 PSRO_ Enable	V6 GND	V7 VDD	V8 GND	V9 VDD	V10 GND	V11 VDD	V12 GND	V13 VDD	V14 GND	V15 VDD	V16 GND	V17 VDD	V18 GND	V19 VDD	V20 HRESET	V21 QACK	V22 THERM_I NT	V23 PSRO0	V24 BYPASS
U	U1 GND	U2 VDD	U3 GND	U4 VDD	U5 GND	U6 VDD	U7 GND	U8 VDD	U9 GND	U10 VDD	U11 GND	U12 VDD	U13 GND	U14 VDD	U15 GND	U16 VDD	U17 GND	U18 VDD	U19 LSSD_S CAN_EN ABLE	U20 VDD	U21 GND	U22 VDD	U23 GND	U24 DI2
T	T1 VDD	T2 KVPRB- GND	T3 VDD	T4 GND	T5 VDD	T6 GND	T7 VDD	T8 GND	T9 VDD	T10 GND	T11 OVDD	T12 GND	T13 VDD	T14 GND	T15 OVDD	T16 GND	T17 VDD	T18 GND	T19 PLLT- ESTOUT	T20 PLL_LOC K	T21 OVDD	T22 SYSCLK	T23 VDD	T24 GND
R	R1 Z_SENS E	R2 KVPRB- VDD	R3 GND	R4 VDD	R5 GND	R6 VDD	R7 GND	R8 VDD	R9 GND	R10 VDD	R11 GND	R12 VDD	R13 GND	R14 VDD	R15 GND	R16 VDD	R17 GND	R18 VDD	R19 GND	R20 CHK- STOP	R21 GND	R22 SYSCLK	R23 GND	R24 ANALOG _GND
P	P1 VDD	P2 Z_OUT	P3 VDD	P4 GND	P5 VDD	P6 GND	P7 VDD	P8 GND	P9 VDD	P10 GND	P11 VDD	P12 GND	P13 VDD	P14 GND	P15 VDD	P16 GND	P17 OVDD	P18 GND	P19 VDD	P20 EI_DISA BLE	P21 VDD	P22 GND	P23 VDD	P24 AVDD
N	N1 SPARE_ GND	N2 VDD	N3 ADOUT0	N4 VDD	N5 GND	N6 VDD	N7 GND	N8 VDD	N9 GND	N10 VDD	N11 GND	N12 VDD	N13 GND	N14 VDD	N15 GND	N16 VDD	N17 GND	N18 VDD	N19 TRIG- GEROUT	N20 VDD	N21 TRIG- GERIN	N22 I2CGO	N23 GND	N24 OVDD
M	M1 OVDD	M2 GND	M3 ADOUT4	M4 GND	M5 VDD	M6 GND	M7 VDD	M8 GND	M9 OVDD	M10 GND	M11 VDD	M12 GND	M13 VDD	M14 GND	M15 VDD	M16 GND	M17 VDD	M18 PROCID 2	M19 PROCID 1	M20 GND	M21 VDD	M22 GND	M23 VDD	M24 GND
L	L1 ADOUT3	L2 SROUT0	L3 SROUT0	L4 VDD	L5 GND	L6 VDD	L7 GND	L8 VDD	L9 GND	L10 OVDD	L11 GND	L12 VDD	L13 GND	L14 VDD	L15 GND	L16 OVDD	L17 GND	L18 VDD	L19 PROCID 0	L20 VDD	L21 SRIN1	L22 SRIN1	L23 GND	L24 SRIN0
K	K1 GND	K2 ADOUT6	K3 ADOUT2	K4 ADOUT5	K5 OVDD	K6 GND	K7 VDD	K8 GND	K9 VDD	K10 GND	K11 VDD	K12 GND	K13 VDD	K14 GND	K15 VDD	K16 GND	K17 OVDD	K18 GND	K19 VDD	K20 GND	K21 OVDD	K22 ADIN6	K23 VDD	K24 SRIN0
J	J1 OVDD	J2 VDD	J3 GND	J4 VDD	J5 GND	J6 VDD	J7 GND	J8 OVDD	J9 GND	J10 OVDD	J11 GND	J12 VDD	J13 GND	J14 OVDD	J15 GND	J16 VDD	J17 GND	J18 OVDD	J19 GND	J20 VDD	J21 ADIN1	J22 ADIN3	J23 GND	J24 ADIN8
H	H1 ADOUT8	H2 ADOUT1	H3 ADOUT7	H4 GND	H5 VDD	H6 GND	H7 VDD	H8 GND	H9 VDD	H10 GND	H11 VDD	H12 GND	H13 VDD	H14 GND	H15 VDD	H16 GND	H17 VDD	H18 GND	H19 VDD	H20 GND	H21 ADIN0	H22 ADIN2	H23 ADIN7	H24 OVDD
G	G1 SROUT1	G2 GND	G3 ADOUT1 3	G4 ADOUT9	G5 GND	G6 VDD	G7 GND	G8 VDD	G9 GND	G10 VDD	G11 GND	G12 VDD	G13 GND	G14 VDD	G15 GND	G16 VDD	G17 GND	G18 VDD	G19 ADIN14	G20 ADIN9	G21 ADIN11	G22 VDD	G23 GND	G24 ADIN13
F	F1 SROUT1	F2 ADOUT1 0	F3 OVDD	F4 ADOUT1 1	F5 VDD	F6 GND	F7 OVDD	F8 GND	F9 VDD	F10 GND	F11 VDD	F12 GND	F13 VDD	F14 GND	F15 VDD	F16 GND	F17 VDD	F18 GND	F19 OVDD	F20 GND	F21 ADIN25	F22 GND	F23 ADIN10	F24 GND
E	E1 OVDD	E2 ADOUT1 2	E3 CLK- OUT	E4 VDD	E5 GND	E6 VDD	E7 GND	E8 VDD	E9 GND	E10 VDD	E11 GND	E12 ADOUT1 6	E13 GND	E14 VDD	E15 GND	E16 VDD	E17 GND	E18 VDD	E19 GND	E20 ADIN21	E21 ADIN22	E22 VDD	E23 GND	E24 CLKIN
D	D1 GND	D2 ADOUT2 6	D3 CLK- OUT	D4 GND	D5 VDD	D6 ADOUT2 3	D7 VDD	D8 ADOUT2 7	D9 OVDD	D10 GND	D11 ADOUT1 5	D12 GND	D13 VDD	D14 GND	D15 ADIN18	D16 GND	D17 VDD	D18 ADIN30	D19 VDD	D20 ADIN32	D21 VDD	D22 ADIN12	D23 VDD	D24 CLKIN
C	C1 ADOUT2 0	C2 OVDD	C3 GND	C4 ADOUT3 2	C5 ADOUT2 1	C6 ADOUT3 6	C7 ADOUT3 3	C8 ADOUT3 5	C9 ADOUT3 9	C10 ADOUT4 1	C11 ADOUT1 9	C12 ADOUT2 8	C13 ADIN4	C14 ADIN28	C15 ADIN17	C16 ADIN42	C17 ADIN29	C18 ADIN35	C19 ADIN34	C20 VDD	C21 GND	C22 ADIN20	C23 GND	C24 OVDD
B	B1 GND	B2 ADOUT2 2	B3 VDD	B4 ADOUT3 1	B5 GND	B6 ADOUT3 0	B7 OVDD	B8 ADOUT1 4	B9 GND	B10 ADOUT1 8	B11 OVDD	B12 GND	B13 VDD	B14 GND	B15 ADIN15	B16 OVDD	B17 ADIN26	B18 GND	B19 ADIN27	B20 OVDD	B21 ADIN31	B22 GND	B23 ADIN23	B24 ADIN24
A	A1 OVDD	A2 ADOUT2 5	A3 GND	A4 ADOUT3 7	A5 ADOUT2 4	A6 ADOUT4 3	A7 ADOUT3 4	A8 ADOUT4 2	A9 ADOUT3 8	A10 ADOUT4 0	A11 ADOUT1 7	A12 ADOUT2 9	A13 ADIN5	A14 ADIN16	A15 ADIN40	A16 ADIN19	A17 ADIN41	A18 ADIN39	A19 ADIN43	A20 ADIN38	A21 ADIN36	A22 ADIN33	A23 ADIN37	A24 OVDD
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24

Figure 2-3. CPC945 Bridge and Memory Controller Pinout Drawing (Top View)

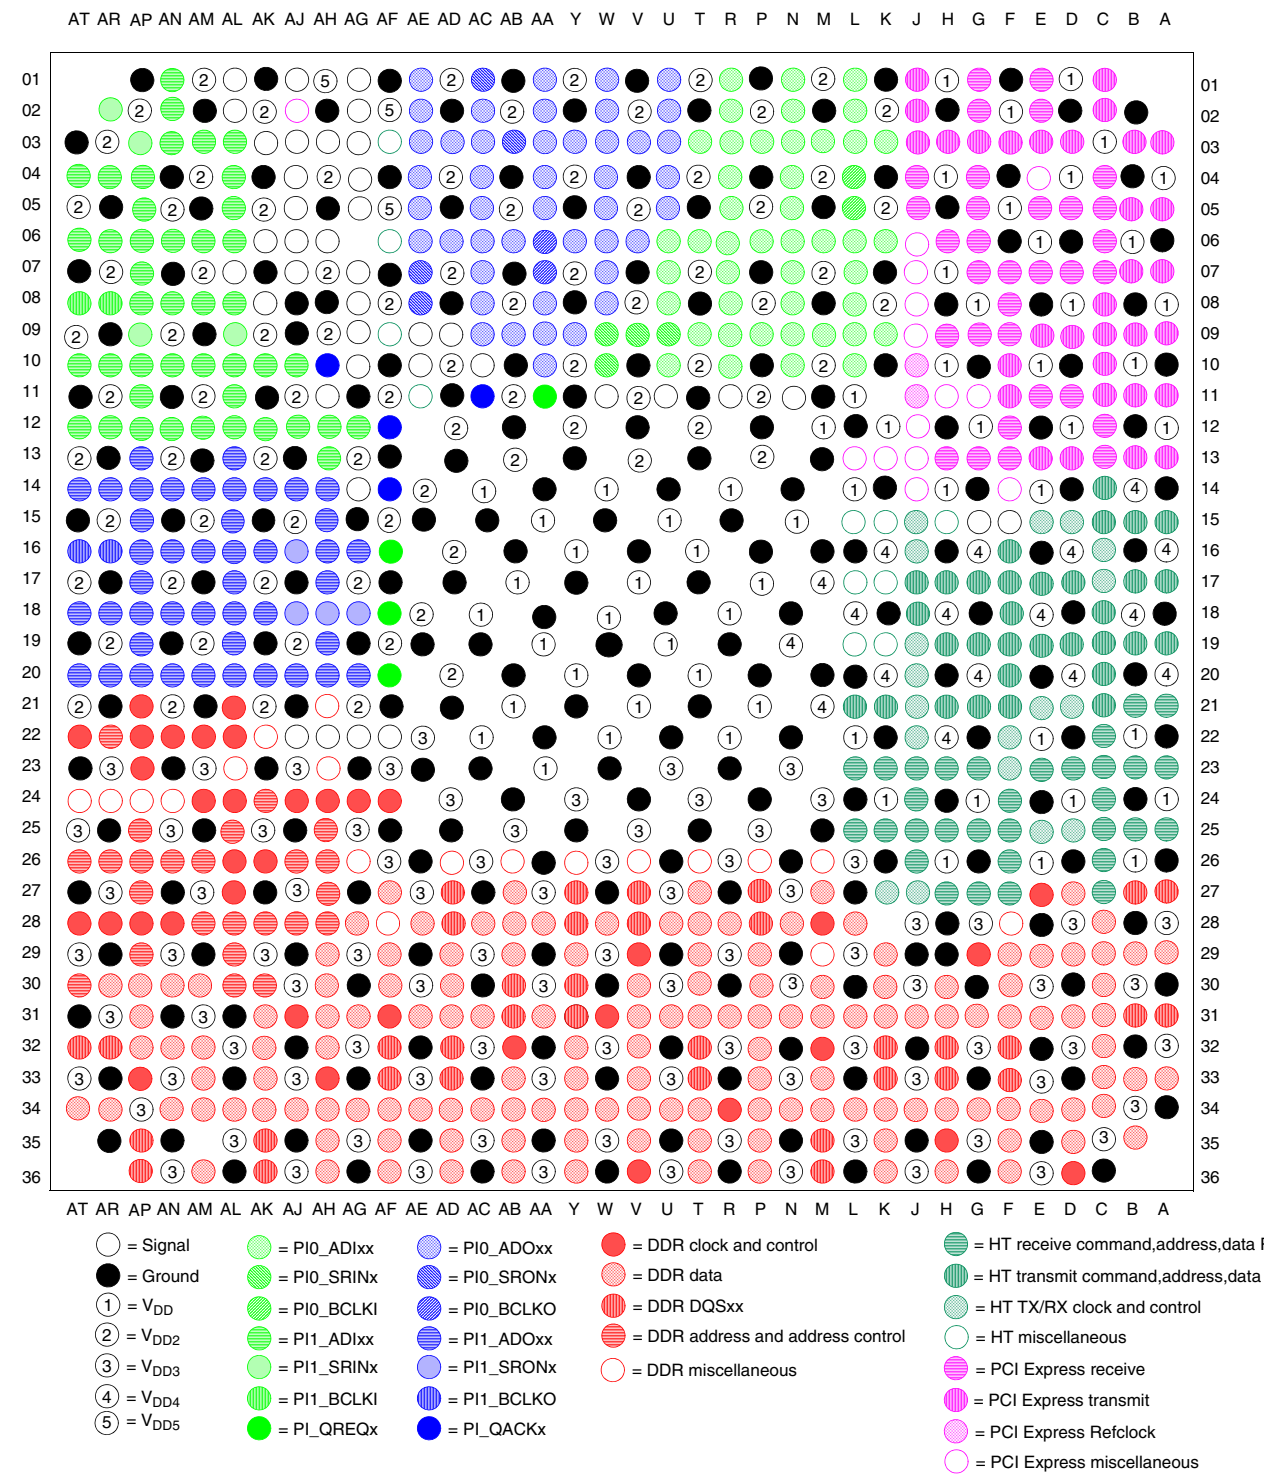

3. Thermal Design Guidelines

The PowerPC 970FX processor is a high-performance microprocessor that requires a focus on thermal management. Attention must be paid to thermal environment factors such as power dissipation, ambient temperatures, and maximum airflow. The core of the PowerPC 970FX processor contains an internal thermal diode that provides the only accurate way to determine the temperature of the die.

The PowerPC 970FX processor is a controlled collapse chip connection (C4) flip chip that attaches to a ceramic ball grid array (BGA). The back of the silicon die is the thermal interface. The die of the PowerPC 970FX processor is approximately 66.2 mm^2 (0.103 in^2), which means that, in addition to the power dissipation, the power density (heat flux) needs to be considered. Typically 2% or less of the chip power is dissipated through the substrate, so essentially all the heat is dissipated through the thermal solution attached to the back of the die.

The relatively small, 7.07 mm × 9.36 mm, size of the PowerPC 970FX die requires taking special care in the selection of such crucial parts of the cooling system as the heat sink, the thermal interface material, and the mounting method. One key constraint affecting the cooling system is the heat-sink-to-die mounting pressure. The maximum mounting force that should be applied to the die and substrate is detailed in the *IBM PowerPC 970FX RISC Microprocessor Datasheet*. For thermal solutions that require higher package mounting forces, an underfill material can be used. See the *Improving BGA to PCB Thermo-Mechanical Integrity White Paper*, which discusses the benefits of underfill material for additional mounting force and solder joint reliability due to thermal cycling stress. It is the customer's responsibility to investigate and characterize the use of epoxy underfill in their design.

3.1 Thermal Interface Materials

Depending on the PowerPC 970FX processor chosen, the power dissipation can range from 14 - 76 W. At the higher wattages, and the resultant heat flux because of the small die, selection of the thermal interface material is as significant as the heat sink. The thermal interface material fills the voids and surface imperfections of the die-to-heat-sink interface and facilitates an efficient transfer of heat from the die to the heat sink. There are many different types of thermal interface materials. However, given the heat flux especially of the higher wattage parts, the die-to-heat-sink interface requires the use of high-quality thermal transfer materials.

Two primary categories of thermal interface material are used in typical designs: thermal greases and phase-change materials. Other products, such as graphite-based materials, might work, depending on the given application.

Phase change material starts as a thin wax-based material loaded with thermally conducting powders. The phase change materials melt into final position when the microprocessor reaches operating temperature.

Thermal greases are basically built from a heavy oil carrier loaded with thermally conducting powders. Thermal greases loaded with silver offer excellent heat transfer and the lowest thermal impedance, for processors in the higher wattage ranges. These silver-loaded thermal greases are not only highly conductive thermally, but also electrically conductive. The package drawing in the *IBM PowerPC 970FX RISC Microprocessor Datasheet* indicates that there are decoupling capacitors located on the top of the substrate near the die. Therefore, when using electrically conductive greases, care must be taken to avoid shorting the capacitors.

Note: To assure the feasibility of the thermal solution design and the lowest assembly cost, it is important to involve manufacturing engineering early in the thermal solution design, particularly in decisions involving the application of the thermal interface material.

3.2 Heat Sink Selection

Depending on the power dissipation, a variety of heat sinks can be used to cool the PowerPC 970FX processor. Technologies range from basic types, such as solid metal and heat-pipe-based, to liquid cooled and even more exotic solutions. Constraints of the specific application determine the required thermal solution and associated cost of each solution. For instance, conventional solid metal heat sinks can be used for a wide variety of applications provided the power dissipation is low enough, adequate space and airflow are available, and the temperature differential between the operating ambient and maximum chip temperature is large enough.

The performance of the heat sink is determined by factors such as a ambient air temperature, air flow volume available to the heat sink, obstructions to the cooling air flow, the number and spacing of fins, and the proximity of other heat generating devices. In applications where space is constrained, such as a server blade, and at higher wattages, more expensive solutions such as heat-pipes and liquid cooling might be required. Experienced heat sink vendors can help determine the most cost-effective solution for a particular application.

3.3 Physical Considerations

Placement of the PowerPC 970FX processor on the printed circuit board (PCB) requires consideration of several thermal factors. The heat sink and its attendant mounting method must be determined before the layout is started to ensure enough clearance around the processor. This includes not only the physical placement of other parts, but ensuring that the heat sink and its mounting hardware do not constrain the routing of the processor buses.

The mounting forces listed in the *IBM PowerPC 970FX RISC Microprocessor Datasheet* include both long-term static forces and manufacturing-assembly dynamic forces. Note that there can be no long-term tensile stress on the die.

The flip-chip die and substrate are robust, but not indestructible. Care must be take in the assembly process to ensure that the stresses applied are as planar as possible to the die surface to avoid the possibility of edge or corner chipping. Mechanical samples can be provided to assist in the manufacturing characterization of the thermal solution assembly process.

3.4 Thermal Diode

The PowerPC 970FX processor contains a thermal diode that, when used in conjunction with external circuitry, can monitor the temperature of the processor die. Most importantly, this diode is located at the hot spot of the die. The thermal diode value is therefore used to determine the maximum temperature specification of the part. The *IBM PowerPC 970FX RISC Microprocessor Datasheet* specifies the maximum temperature as T_{DIODE}, not T_J. The diode is accompanied by some electrostatic discharge (ESD) diodes, and is driven with a 100 μA constant current source. The voltage across the p-n junction is measured and converted to a temperature. The datasheet contains some information about the use of the thermal diode. Each thermal diode is factory calibrated and the information is stored in the microprocessor's fuse ring. The *IBM PowerPC 970FX Power On Reset Application Note* describes how to read this information.

Because of the unique structure and biasing requirements of the diode, commercial parts for reading thermal diodes that use the two current method cannot be used with the PowerPC 970FX processor. Several application notes describe the thermal diode, how to drive it, and how to interpret the information (see *Related Documents* on page 11 for more information).

3.4.1 Thermal Diode Monitoring Circuit

The PPC970FX/CPC945 evaluation board contains circuitry that is used to drive and read the thermal diode of the PowerPC 970FX core. The thermal diode is properly biased in accordance with the *IBM PowerPC 970FX RISC Microprocessor Datasheet.* A Cypress Programmable System-On-Chip (PSoC) microcontroller is used to read the voltage of the processor thermal diode and send the digitized voltage values to the service processor unit (SPU). The SPU then converts the voltage values into a temperature, using calibration information contained in the PowerPC 970FX processor's mode ring. The PSoC is configured to communicate over the I^2C bus to the SPU.

In general, the voltage across the thermal diode varies approximately 1.5 mV per 1°C. This is a small change that can be distorted easily by on-board electromagnetic interference (EMI) and electromagnetic field (EMF) noise if the signal lines are not routed carefully. Route the thermal diode signal lines to avoid coupling noise to the analog signals. The voltage drop across the diode can vary approximately between 0.80 V down to 0.60 V. The final voltage range is chip dependent, and the characteristics of the diode are stored in the PowerPC 970FX fuse ring.

4. Processor Interface Routing Guidelines

The PowerPC 970FX processor interface uses high-speed source-synchronous buses (SSBs) to transfer data between the PowerPC 970FX processor and the CPC945 bridge or system controller chips, and to support the cache-coherency snooping protocols for multiprocessor configurations. The SSBs are unidirectional point-to-point connections between a drive side (D) and a receive side (R). SSBs, and the 1-signal snoop response buses (SRBs), are paired to form a bidirectional channel between a processor and a bridge chip as shown in *Figure 4-1*.

Figure 4-1. Processor Interfaces

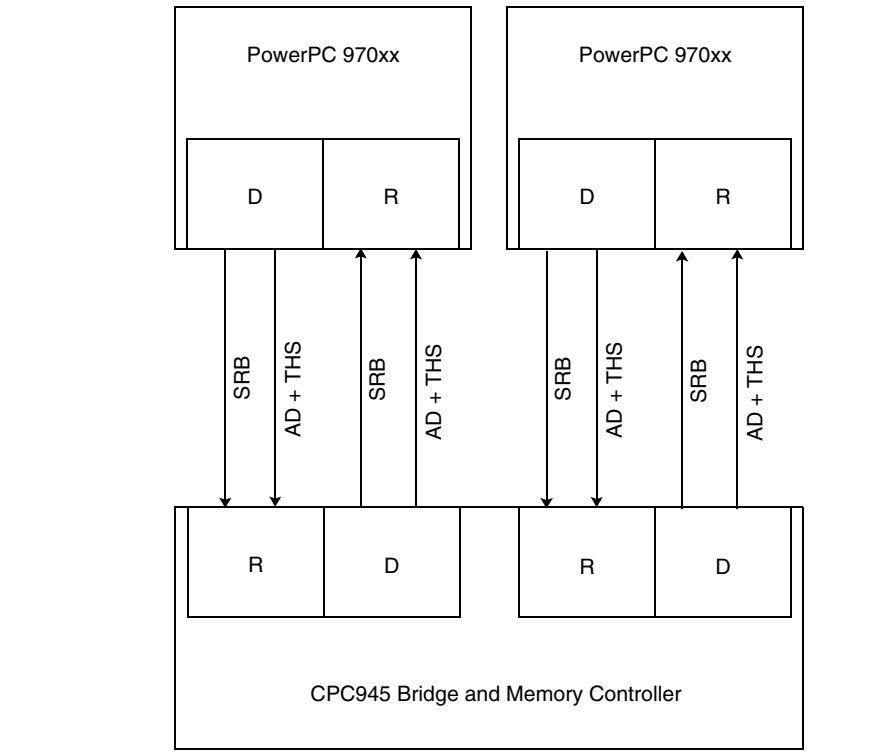

AD: Address/data channel
D: Drive side
R: Receive side
SRB: Snoop response bus
TH: Transfer handshake bus

SSB data is transferred on every bus-clock edge; that is, at double the data rate (DDR) of the bus-clock frequency. There are 50 signal lines per SSB. Two lines are used for the differential bus clock, 44 signal lines are used to communicate 36 bits of logical data and error checking, and four signal lines are used for the differential snoop response bus. The 36 data bits consist of 35 bits of the address and data (AD) channel and a single bit for the transfer handshake bus (TH).

The SSBs achieve high-speed operation using low-cost packaging solutions by exploiting four features:

1. Source synchronous signalling. The differential bus clocks are bundled with the single-ended data signals. The reference for the single-ended signals is approximately $OV_{DD}/2$ as described in the *IBM PowerPC 970FX RISC Microprocessor Datasheet*.

2. Far-end (parallel) termination. The single-ended data signals use parallel termination, located in the receiver on the processor, at the far end of the signal line to absorb signal reflections and maintain a quasi-constant current loading for each data signal line.
3. Error checking. Two types of error checking are available on the bus: error checking and correction (ECC) mode and a balanced coding method. Both methods are described in the *IBM PowerPC 970FX RISC Microprocessor User's Manual*. The application of balanced coding to the SSB maintains a quasi-constant current loading across the entire SSB interface. Within the SSB, there is no net current flow across the power planes. This dramatically reduces noise problems that are due to power-supply rail compression (that is, Ldi/dt noise) and current voltage offsets between the chips.
4. Point-to-point unidirectional signalling. Restricting the signal fan-out to a single point and keeping the signal flow unidirectional mitigates problems associated with high-frequency signal attenuation.

For additional design information about the processor interface electrical and physical characteristics, see the *IBM PowerPC 970FX RISC Microprocessor User's Manual* and the *IBM PowerPC 970FX RISC Microprocessor Datasheet*.

4.1 General Guidelines

The PPC970FX/CPC945 evaluation board is designed to achieve the maximum supported processor and processor interface frequency. Other configurations are possible; however, appropriate modeling and simulation, and careful layout and placement, are required to produce a reliable design that deviates from the following recommendations.

Note: Some of this information comes from the *IBM PowerPC 970FX RISC Microprocessor Datasheet*. It is repeated here for convenience, but review the processor interface section of the datasheet before starting a design. If there is a discrepancy between this document and the datasheet, the datasheet value takes precedence.

- Adhere to the suggested maximum printed circuit board (PCB) trace lengths for data and clocks. These are 18 cm for transfer rates up to 1.5 gigatransfers per second (GTps) and 22.5 cm for transfer rates up to 1.0 GTps for a PCB of typical Flame Resistant 4 (FR-4) material. It is always best to minimize the length of traces that have high data rates. These maximum guidelines are not intended to encourage or license the use of arbitrarily long processor interface traces.
- Route all traces on inside PCB layers using stripline construction (that is, locate each signal trace on a layer that is between a power and ground plane).
- Match the length of all processor interface inputs to within 150 ps.
- Match the length of all processor interface outputs to within 150 ps.
- For a multiprocessor system, match the length of the inputs to processor A to the length of inputs to processor B and the length of the outputs from processor A to the length of the outputs from processor B. If these do not match, the longer one determines the target time of the interface.
- Ensure that the clock delay is longer than the longest data delay particularly for bus transfer rates above 1.1 GTps.
- Ensure that the characteristic impedance of the processor interface bus traces is 50 Ω. An impedance of 50 Ω is dependent on the line width, line spacing, and board stack-up. To escape the processor or companion part, the line widths and spacing selected to achieve a characteristic impedance of 50 Ω, might have to vary somewhat for a short distance.

- Route the clocks as differential pairs, with very accurate length matching (< 0.254 mm [10 mils] delta). The routes should be internal to the board except for short regions in the ball escape area (these should be .508 cm [200 mils] or less). These traces should also have a characteristic impedance of 50 Ω.
- Physically and electrically isolate each high-speed bus. Isolate processor bus inputs from outputs. If economically feasible, keep input and output signals on separate layers. If input and output signals must exist on the same layer, only the input signals should be routed amongst each other as a group; similarly, only the output signals should be routed amongst each other as a group.
- Match traces in terms of delay, not just trace length. Extra vias on routes should be considered to add approximately 20 ps per via. The processor interface routes should minimize vias. If possible, target one at the system controller and bridge, one at the processor, and one en route. There is always a trade-off between the length matching and the route density.
- Avoid using right angle bends on the processor interface; it is better to use 45 degree angles on corners. Review tight serpentine traces. It is better not to have excessive corners on a net; they have more capacitance and tend to slow the route to some degree. Fewer long bends are better than many short bends for flight time matching.
- Do not cross reference plane splits

All these recommendations are intended to help design a functional system. However, they are only guidelines and do not take the place of design-specific results obtained from signal integrity modeling. For designs of this sophistication, time spent modeling actually saves time in bring up and debug. IBM provides both Input/Output Buffer Information Specification (IBIS) and encrypted HSPICE models for the processor interface bus on the PowerPC 970FX processor.

Note: Although IBIS can provide a quick analysis of the layout, we recommend using a layout and analysis tool that is HSPICE-capable. HSPICE gives a more accurate view of the actual signal integrity.

4.2 Example System Implementations

Three possible system implementations are shown in *Figure 4-2 System Implementations* on page 36. The system A configuration, which has been implemented in hardware, breaks the processor interface bus with a connector between the processor and the CPC945 bridge. This configuration supports a system design that can support a single or dual processor configuration based whether one or two processor cards are plugged into a system board. If a connector is used, it must be of high quality and designed for high-speed signal use. The bill of materials (BOM) available with the PPC970FX/CPC945 evaluation board lists the connectors that can be used. The connector pair should be impedance matched to the board traces as closely as possible. System B avoids breaking the processor interface bus with a connector. System C is a single-chip version of system B.

Figure 4-2. System Implementations

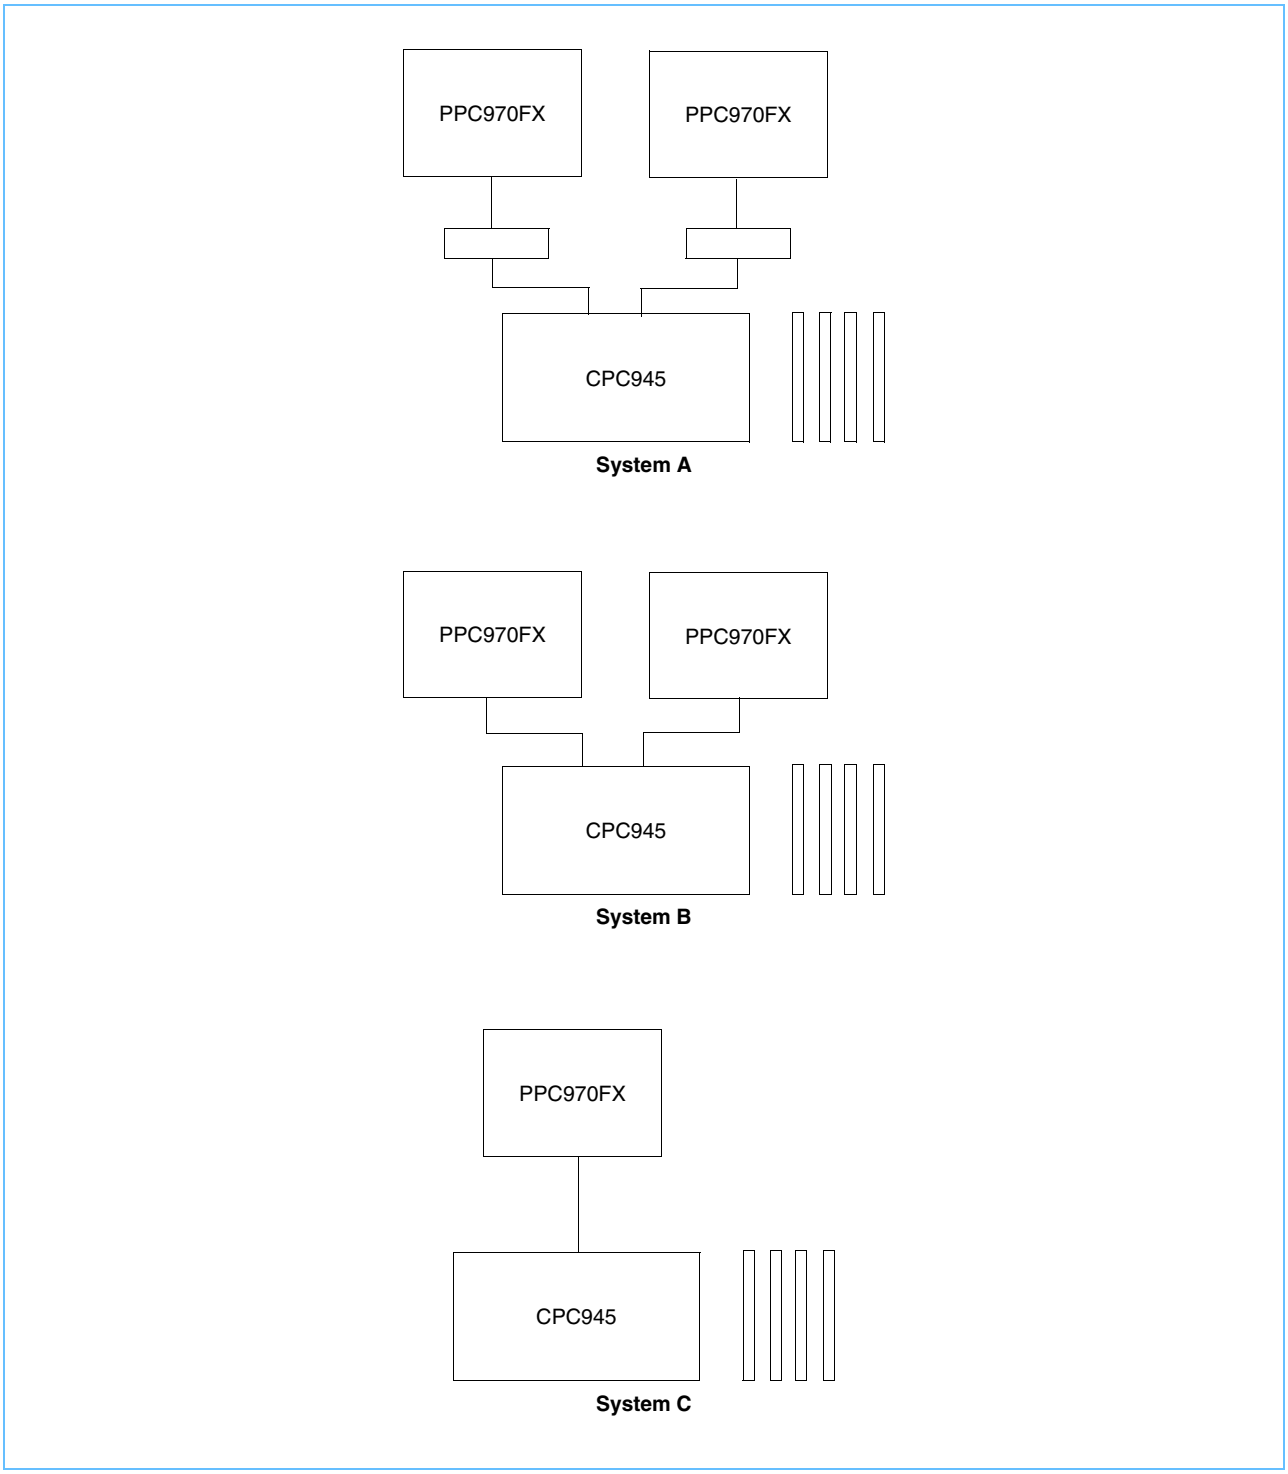

4.3 System A Layout Example

A sample board layout for the processor card is shown in *Figure 4-3* through *Figure 4-5*. *Figure 4-3* shows all the processor interface traces from the processor to the connector. These traces are approximately 3.3 cm (1.3 in.) long and are matched to 0.0762 cm (30 mils). As shown in *Figure 4-3* through *Figure 4-5*, the whole processor interface bus is routed on three layers with only a few traces on the third layer. The bottom layer is used for escaping the components only and shows that the escapes are essentially just dog bones (surface pad to via connected by a short trace). Note that for the processor, via-in-pad technology is used. Although the guidelines recommend isolating processor bus inputs from outputs, it was not efficient to route the input bus on separate layers from the output bus in this case. However, as illustrated by the white line in *Figure 4-3* and *Figure 4-4*, the inputs (above the line) and outputs (below the line) are spaced apart from each other.

The routings pictured are actually for a PowerPC 970MP pinout. The PowerPC 970FX pinout is different but has similar signal groupings. The processor board used as part of the PPC970FX/CPC945 evaluation board has a PowerPC 970MP footprint, but accepts PowerPC 970FX parts with a small interposer card soldered between the processor chip and the card.

Figure 4-3. Processor Interface Bus In and Out, Layer S03

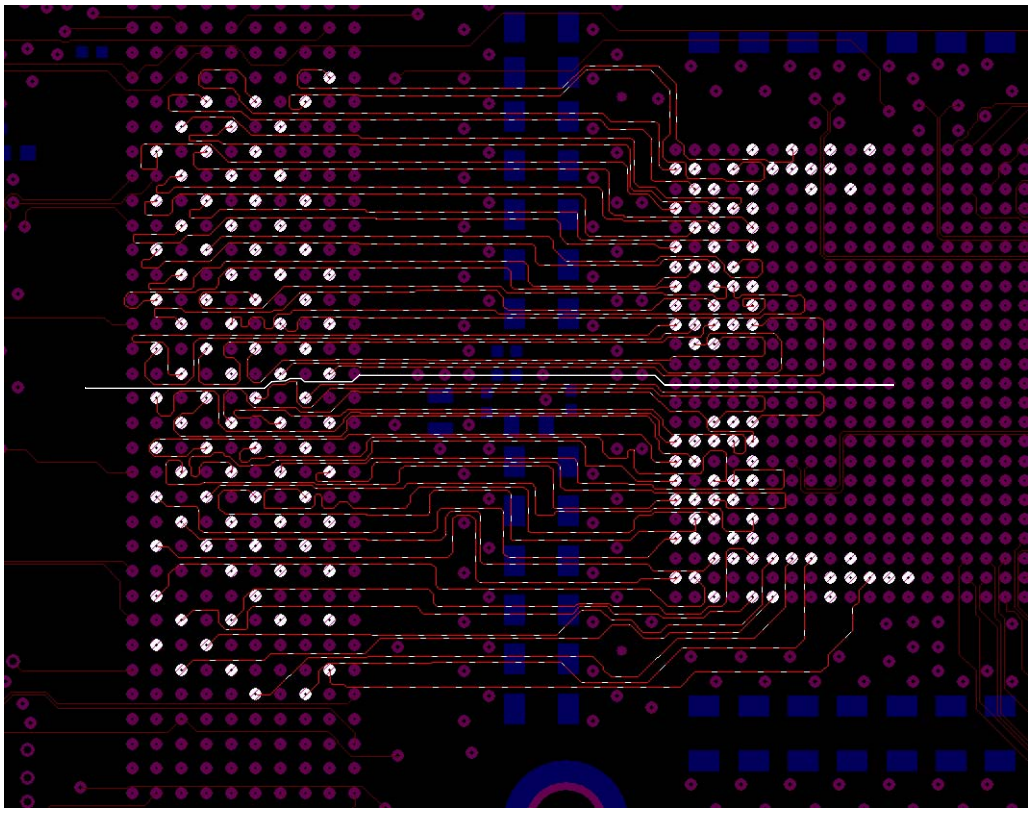

Figure 4-4. Processor Interface Bus In and Out, Layer S05

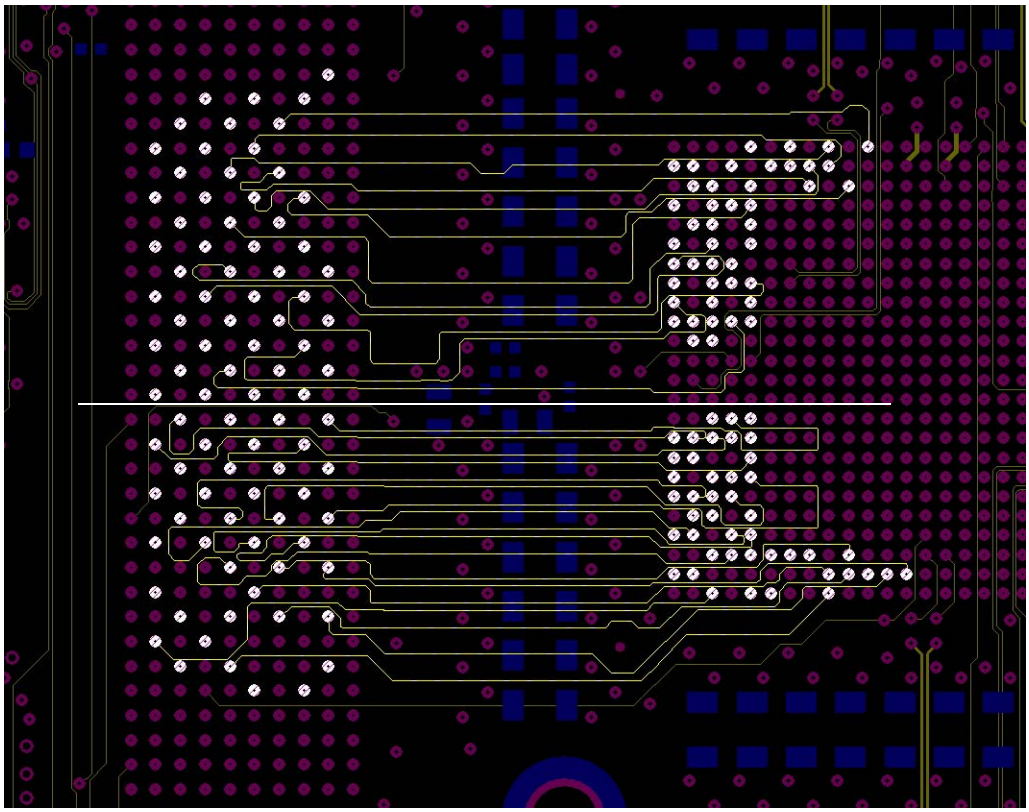

Figure 4-5. Processor Interface Bus In and Out, Layer S12

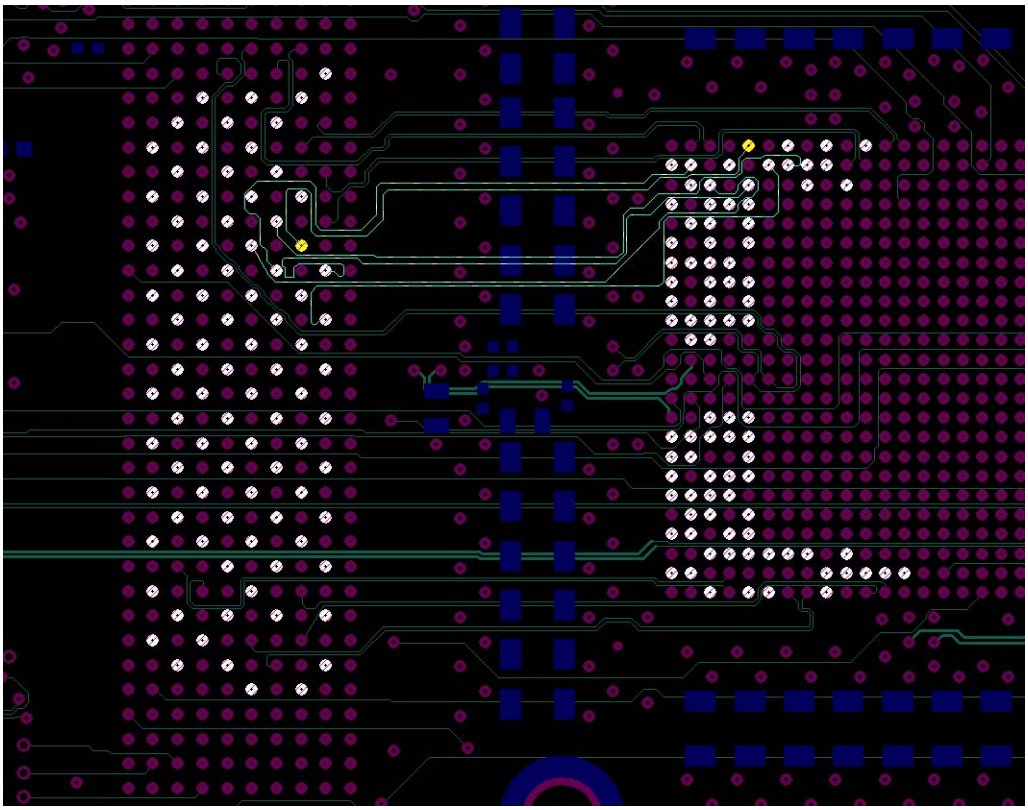

Figure 4-6 through *Figure 4-8* show that two complete processor interface buses can be routed on only three layers. These traces are approximately 11.2 cm (4.41 in.) long and are matched to 0.5 cm (200 mils).

Figure 4-6. Processor Interface Bus A and B In and Out, Layer S03

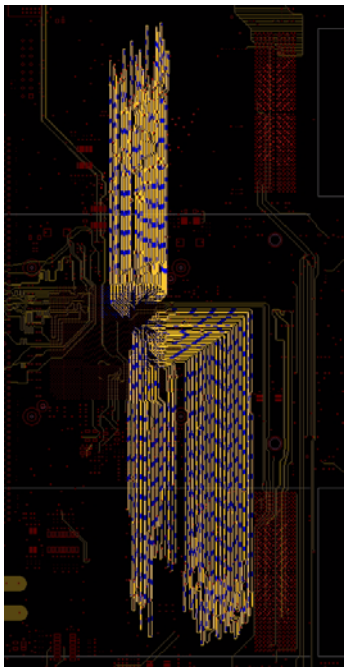

Figure 4-7. Processor Interface Bus A and B In and Out, Layer S05

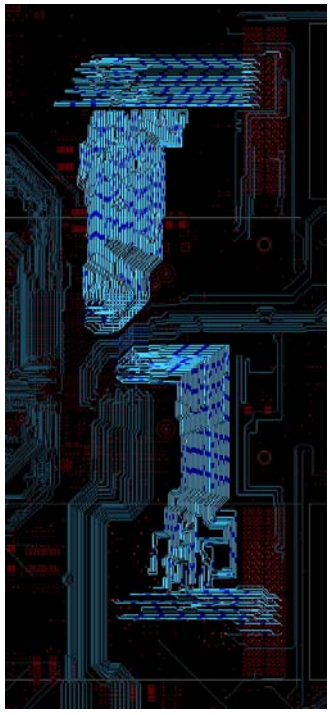

Figure 4-8. Processor Interface Bus A and B In and Out, Layer S07

5. Clocking

As described in the *IBM PowerPC 970FX RISC Microprocessor Datasheet*, the primary input clock for the PowerPC 970FX processor is a differential pair: SYSCLK and $\overline{\text{SYSCLK}}$. This system clock (SYSCLK) is either 1/8 or 1/12 of the core frequency, based on the chosen processor-interface bus clock ratio. The *IBM PowerPC 970FX RISC Microprocessor User's Manual* provides the functional description of the PowerPC 970FX processor clocking. This design guide focuses on the implementation of the clocking scheme for the system.

5.1 Processor and Companion Part Clock Balancing

The choice of companion part, or system bridge chip and memory controller, determines if the system can be designed with a single processor (either single or dual core) or dual processors. The SYSCLK to each part, whether a processor or companion part, must be matched in frequency and phase. This can be done through layout exclusively or through a combination of layout and output skewing, a technique that is supported by some clocking chips. Trace length matching of the differential clocks to each part is a good start, but might not be sufficient. The signal propagation velocity in a printed circuit board (PCB) can vary depending on a number of layout factors; therefore, the goal is to use the simulation capability of the board design system, along with the models provided by IBM, to match the clock propagation times.

5.2 Processor and Companion Part Synchronization

The *Processor-Clock Timing Relationship Between PSYNC and SYSCLK* figure in the *IBM PowerPC 970FX RISC Microprocessor Datasheet* shows that the PSYNC signal is provided for system synchronization. (In fact, the PSYNC signal is provided as an input to the PowerPC 970FX processor; the APSYNC signal is provided as an input to the CPC945 bridge.) This signal is a 1-SYSCLK wide pulse that is typically generated every 24 SYSCLKs. Since PSYNC is used to establish the concept of Time0 for the processors and companion parts, the traces for PSYNC should be propagation time matched between the parts. The key is to ensure that the same relative core clock samples the PSYNC signal. Matching to 1.25 cm or approximately 100 ps is a good target.

5.3 Spread Spectrum

Spread spectrum is a technique whereby the input SYSCLK frequency is slowly varied according to a modulation profile. This technique is used to detune harmonic frequency peaks, which spreads out the radiated electromagnetic interference (EMI) energy. This makes it easier to pass Federal Communications Commission (FCC) emission standards in the United States and the radiated emission standards of other countries. High-end board layout systems can minimize radiated emission by rounding signal corners and by using other known techniques, but implementation of spread spectrum has the largest impact. The *IBM PowerPC 970FX RISC Microprocessor Datasheet* details the expected modulation profile, the maximum modulation frequency, and the maximum amount of SYSCLK frequency down spread considered acceptable for the PowerPC 970FX processor.

5.4 Component Recommendation

The LCK4013 (from LSI) clock part supports the clocking features described in this design guide. The LCK4013 has three low voltage differential outputs that can be individually skewed to balance propagation time in a system with up to two processors and a companion part. In addition, the LCK4013 can generate multiple individually skewable PSYNCs. Spread spectrum and frequency slewing are also supported. For simple systems, a less sophisticated clock generation part would suffice. Read the SYSCLK requirements in the *IBM PowerPC 970FX RISC Microprocessor Datasheet*. The processor requires moderate to high frequency (>100 MHz), relatively low jitter, low voltage (OV_{DD}), controlled rise and fall time, and differential high-speed transceiver logic (HSTL) signals.

5.5 Implementation

Figure 5-1 shows one example of the SYSCLK circuit to a processor. In this example, the clocks are ac coupled to the processor. The dc bias is established by the SYSCLK termination. The PowerPC 970FX processor has on-chip termination for the differential SYSCLK, which can be enabled or disabled using the CKTERM_DIS signal. The termination network shown in blue is an alternative to using the on-chip termination. Some designers place these components, but do not populate them unless they cannot achieve the correct signal integrity with the on-chip termination. The ac coupling is not a requirement, as indicated in the schematics of the PPC970FX/CPC945 evaluation board.

Figure 5-1. SYSCLK Circuit Example

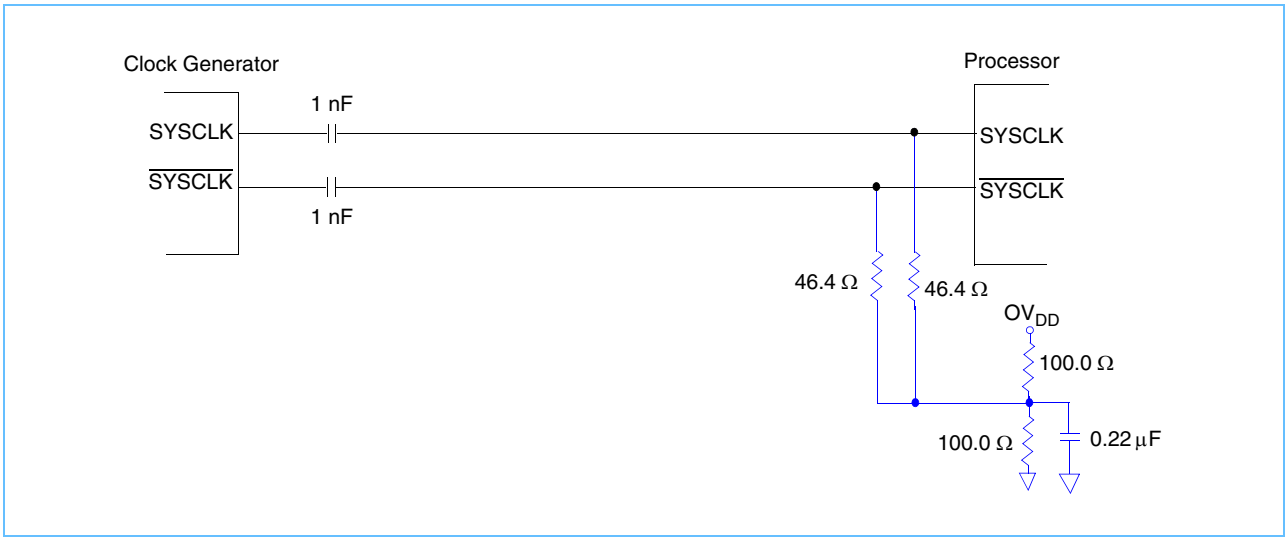

5.6 Layout Considerations

The two differential clock traces must be carefully managed. SYSCLK and $\overline{\text{SYSCLK}}$ should be routed on the same layer within 0.5 mm (20 mils) of each other (maximum separation), matched in length to within 0.5 mm (20 mils), and must be separated from other signals by at least 0.5 mm (20 mils). The intent is that the two signals see the same conditions to maximize the ability of differential signals to reject common mode error sources.

Figure 5-2 shows the traces for the example system referred to as System A in *Section 4.2 Example System Implementations* on page 35. The differential SYSCLK traces are paired with 0.102 mm (4 mil) separation, and 0.635 mm (25 mil) clearance to the nearest nonclock trace. They are length matched to within ±0.127 mm (5 mil) on the processor card.

Figure 5-2. Processor Card SYSCLK Pairs

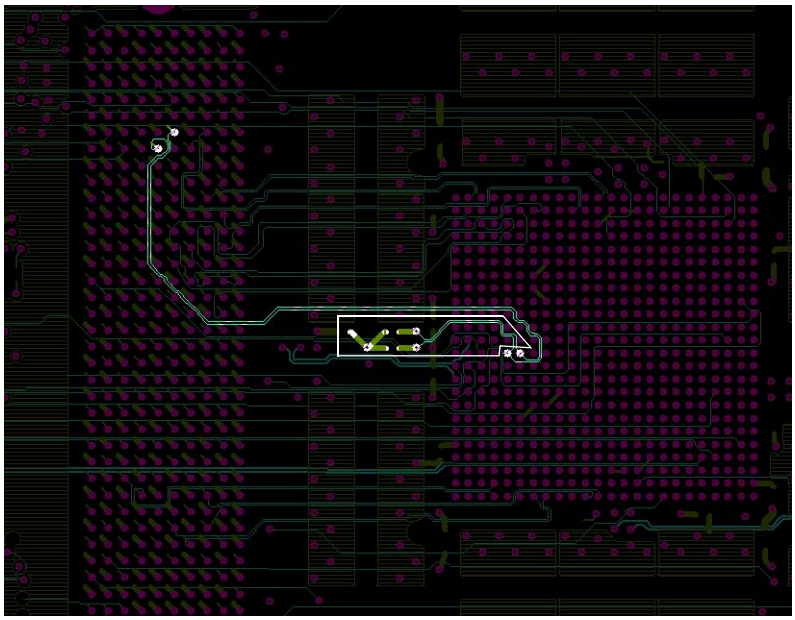

This card design was developed as an evaluation and characterization test bed for the processor. The external termination shown in blue in *Figure 5-1 SYSCLK Circuit Example* on page 44 is part of the layout. The extra trace and component pads for the external termination are indicated by the white polygon. Using the IBIS and HSPICE models provided for the PowerPC 970FX, the signal integrity analysis done on the clock traces in most cases indicates that the on-die SYSCLK termination is sufficient. The SYSCLK traces are edited in *Figure 5-3* on page 46 to show a possibly cleaner routing of the SYSCLK differential pair if external termination can be avoided.

Figure 5-3. Removing External Termination

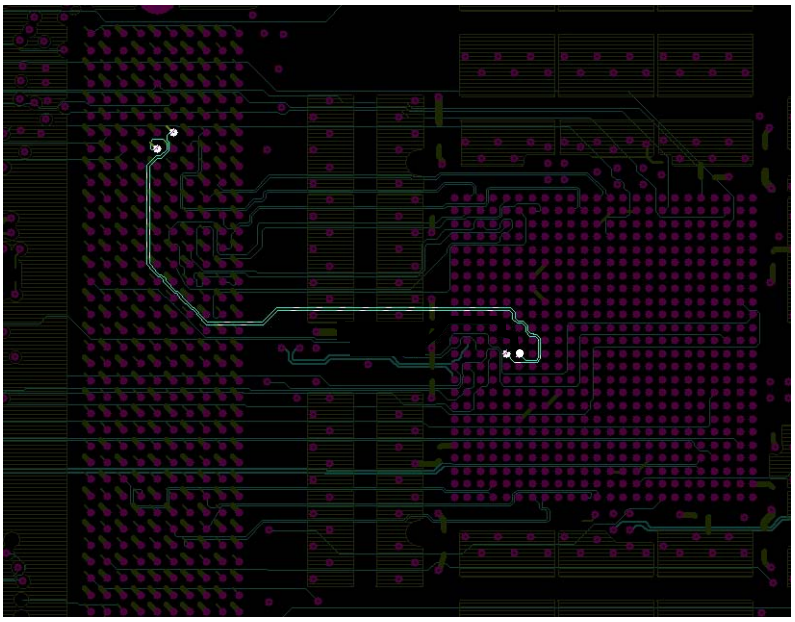

Figure 5-4 shows the traces for the example system referred to as System A in *Section 4.2 Example System Implementations* on page 35. The differential SYSCLK traces are paired with 0.102 mm (4 mils) separation, and 1.02 mm (40 mils) clearance to the nearest nonclock trace. They are length matched to within ±0.38 mm (15 mils) on the system board; however, it was not possible to keep the whole trace on the same layer as shown by the different layer colors. As specified in *Section 5.1 Processor and Companion Part Clock Balancing*, the trace lengths for the SYSCLK pairs are matched within 1.02 cm (400 mils). This is evident in the large serpentine routing on the blue layer shown in *Figure 5-4* on page 47.

Figure 5-4. System Board Processor A and Processor B SYSCLK Pairs

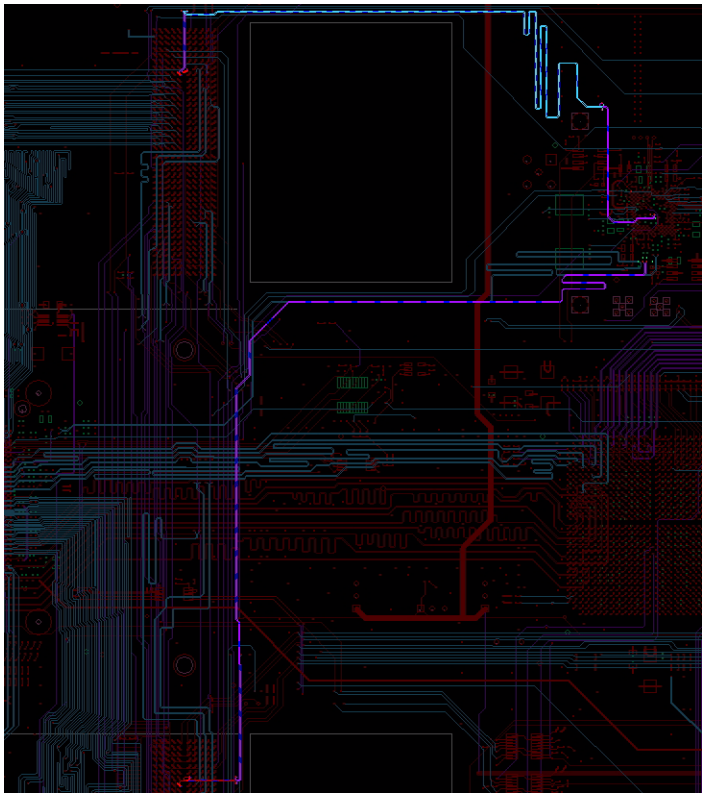

6. CPC945 Memory Interface and Routing Guidelines

6.1 Memory Summary

The CPC945 Bridge and Memory Controller has the following features:

- Double data rate two (DDR2), 533 MHz and 400 MHz.
- The maximum data rate of up to 8.5 GBps with a 266 MHz dynamic random access memory (DRAM) clock (533 million transfers per second of 128 bits).
- 144-bit data bus (with error correction and checking [ECC]) or 128-bit data bus (without ECC) or 72-bit data bus (with ECC) or 64-bit data bus (without ECC) (See the memory section of the *CPC945 Bridge and Memory Controller User's Manual* for full details on the bus width configurations.)
- Support for JEDEC DDR2 SDRAM Specification (JESD79-2) 256 Mb, 512 Mb, 1 Gb, and 2 Gb chip sizes and x4, x8, and x16 organizations. The x8 and x16 chips can be mixed. If x4 chips are used, all the chips must be x4.
- Support for 8 ranks. Eight dual inline memory modules (DIMMs) maximum (4 double-sided pairs).
- 64 GB maximum capacity using four pairs of 8 GB double-sided DIMMs. (Only 62 GB are addressable because of the 2 GB gap in the I/O portion of the memory map.)
- ECC support for single-symbol (4-bit) correction and double-symbol error detection; chip kill correction with x4 chips.

Note:

- The CPC945 bridge supports up to four memory DIMMs (two pairs) directly attached, or eight memory DIMMs (four pairs) with the use of external data multiplexing switch devices on the memory bus. Review the restrictions on DIMM characteristics in the memory section of the *CPC945 Bridge and Memory Controller User's Manual*.
- Always review the *IBM PowerPC CPC945 Errata List for DD1.X* document, including the design notes section, as part of the design information.

6.2 CPC945 DDR2 Interface Overview

In general, a DDR2 SDRAM interface uses source synchronous transfers for all signaling. The address, control, chip select (CS), and clock enable (CKE) signals are timed relative to the DDR2 DRAM clock (CK). The address, bank address (BA0 - BAX), row enable (RAS), column enable (CAS), write enable (WE), and clock (CK) signals are always sourced from the memory controller. Data (DQ) lines are timed relative to a strobe (DQS) in a source synchronous manner, and they are bidirectional.

Data (DQ) signals use source synchronous transfers that employ a strobe (DQS), per byte lane, that travels with the DQ signals over what should optimally be length-matched routing paths. The intrabyte lane routes (for the bits within each byte) must also be length matched. The DQS strobe latches the DQ signals. The DDR2 SDRAM DIMMs employ differential clocks (CK). The CPC945 bridge provides two differential clock pairs. The DDR2 specification defines a tight relationship between CK and DQS transitions; these signals effectively are required to align within ±1/2 bit time.

For the CPC945 bridge, all address lines including BA0 - BA2, RAS, CAS, WE, CKE, CS are buffered. CK comes from two differential pairs on the CPC945 bridge. Each pair is designed to drive two DIMMs, so a designer who wants to implement a larger configuration must buffer or fan out the CK clocks with external logic such as a zero-delay stub series terminated logic (SSTL) buffer.

The CK to DQS relationship dictates the DRAM clocking scheme. The CPC945 memory controller clock outputs (CK) can track with the DQS outputs to maintain the ±1/2 bit time relationship over process, voltage, and temperature variations. For the memory tuning process, when the CK and CS relationship is properly set, the user then should select DQ and DQS delays such that DQS arrival time at the DRAM device matches CK.

On writes, the strobe (DQS) is sent with the required phase relationship relative to data lines (DQ). SDRAM devices do not adjust the strobe placement. The CPC945 bridge has internal delay lines that assist in aligning the DQS and DQ pair for each byte lane, such that the DQS-to-CK relationship can be met without excessive inter-byte-lane route matching. The CK output pairs timing can be adjusted within the CPC945 Clock Delay Adjustment Register (x‘F800 2520’).

Note: The CPC945 bridge does not provide data mask (DM) signals.

For writes, DQS should transition near the clock edges of CK and be centered in the write data window. The CPC945 bridge has 32 delay lines that tune the strobe (DQS) placement relative to data (DQ)—two for each of the 16 byte lanes to allow the rising and falling strobes for each byte lane to be tuned independently. These are found in the DQS Data Delay Registers. For these specific registers, the rising edge is controlled by WrRDelayOffset, and the falling edge is controlled by WrFDelayOffset.

In addition to the DQS-to-DQ adjustments, the CPC945 memory controller can adjust the clock phase that DQ uses relative to its internal core clock phase. This allows the different byte lanes to be launched earlier or later relative to each other to tune out flight time variations from byte lane to byte lane. Tuning out the flight time variations across the byte lanes facilitates the arrival of all of the DQS lines within ±1/2 bit time of the single CK per DIMM. There are 16 delay lines that control the timing when data (DQ) in each byte lane is sent, one per byte lane, for the write data delay offset (WrDataDelayOffset). So, DQS timing is delayed in total by the sum of a DQ delay plus a DQS-to-DQ delay.

On reads, the DRAM sources strobe is at the same phase as data, so the memory controller must delay the strobe to assure data eyes are sampled in the center of the valid eye.

For reads, the DQS is used to latch the data values provided by the DIMMs. The DQS from the DIMM is aligned with the clock and skewed automatically by 1/2 bit time by logic within the CPC945 bridge. Half of the DQS Delay Adjustment Registers allow fine tuning adjustments to be made to the delay parameter ResMuxDel to help adjust read timings to compensate for any skew between byte lanes.

The DQS load varies with the device load capacitance variation and with the device count variation; therefore, a DIMM might have one or two loads per DQS line. Mixing DIMM sizes, such as using some single bank versions with dual bank ones affects timing. Such a configuration is not recommended. For detailed recommendations, review the restrictions on DIMM characteristics in the memory section of the *CPC945 Bridge and Memory Controller User’s Manual*. HSPICE simulation is strongly advised, as good design practice, to ensure that your design topology can support the data rates required for the maximum loading.

In addition to tuning out the flight time skew between byte lanes, the delay mechanisms can also have enough range to use for timing margin tests during bring-up. Align DQS and CK by using the same settings as the DQ and DQS delays, with an offset for loading variance, possible CK buffering variance, and flight-time issues. The buffer output times and flight times on boards should be set up to reasonably match in the CK and DQ and DQS paths.

6.3 DDR2 Channel Impedance, Termination, and General Routing Recommendations

To control trace impedance and manage signal integrity for these high-speed signals, it is important to reference the address and command lines to a solid ground or power plane. The data lines should be referenced to a solid ground plane. Most of the SDRAM routing of the evaluation board was done as stripline traces, where the signal layer has a reference plane, either ground or power, above and below the signal layer. Address and command signals should be routed separately from the data group signals, from the memory controller to the first DIMM. Address and command signals are latched at the DIMMs using the clock signals; therefore, they must maintain a closely matched length relationship to the clock signals at the DIMM.

With DDR2, the termination for the data signals is provided by the on-die termination (ODT) of the SDRAM parts. The address and control signals still require parallel termination to the termination voltage (V_{TT}). Series resistors might also be required on the address and controls signals.

Place all serpentine routing between the memory controller and the first DIMM, and keep the series resistor close to the driver. Make wiring lengths between the first DIMM, the second DIMM, and the final terminating resistor the same across all bits and keep them as short as practical. HSPICE simulations are recommended to verify the design. Use a length match rule between DIMMs and have a maximum length rule from the last DIMM to V_{TT} termination resistors, to ensure compliance in layout. The terminating resistor voltage plane must have sufficient bypassing.

Keep V_{TT} noise constrained (100 mV is suggested as a target) with good decoupling and a good V_{TT} plane with a fast response linear power regulator. The terminating resistors should be decoupled. If resistor packs are used in the system, ensure that the supply pins are sufficiently decoupled. Via count on the DDR2 routes should be constrained and optimally should be the same per net in a byte lane. Use return path vias when changing a reference plane.

Registered DIMMs only require one clock pair and have on-board series termination for control and address and parallel termination for clock lines.

6.3.1 DQ and DQS Signals

The data (DQ) and strobe (DQS) lines driver impedance is nominally 35 Ω. The driver impedance should be reasonably matched with the line impedance to reduce crosstalk-related signal integrity issues. This should be verified with simulation: HSPICE models for the DDR2 I/O devices are available from IBM applications engineering. Other ways to reduce crosstalk effects include:

- Using a design rule for minimum trace pitch of 0.254 mm (10 mils)
- Shielding the data strobe signals on each side by a ground trace
- Using shielding (well tied to ground) to help to reduce transient currents

It is important that the data lines be referenced to a solid ground plane because they operate at twice the frequency of the address and control signals. The data signals require a good ground return path to avoid degradation of signal quality because of inductance in the signal return path. One way to accomplish this is to ensure that these signals are ground referenced from the CPC945 bridge to the first DIMM connectors and between each DIMM connector to provide a low-impedance current return path.

6.3.2 Clock Signals

The CK lines are differential pairs, routed together and length matched for each pair. Route CK on the surface if DQ and DQS are routed on the surface, or route it internally if DQS and DQ are routed internally, to avoid the propagation time difference between microstrip and stripline routing. If constraints affect only the propagation time, the average trace velocity on the surface to the internal layers needs to be very accurate. The trace separation between the clock traces and adjacent signals should be at least 0.254 mm (10 mils); however, more separation (up to 0.5 mm [20 mils]) is preferred.

The CPC945 bridge provides two CK, $\overline{CK}$ output differential pairs. These lines can each drive two DIMMs. Use a resistor network near the DIMMs to split the clock to the two DIMMs. A CK, $\overline{CK}$ pair should drive the upper and lower bank DIMMs for the same address space, such that either both are loaded or neither are loaded.

6.3.3 Address and Control Signals

For address and control signals, the buffering provided by registered DIMMs limits the load difference between a four DIMM, 8-bank system and a two DIMM, 2-bank system to only twice as much. For four DIMM cases, terminate the address, row enable (RAS), column enable (CAS), write enable (WE), and bank address (BA) lines to V_{TT} to obtain valid windows in 2-cycle address cycles with fully loaded configurations. Simulate a range of termination values with HSPICE to obtain the best value for your design, especially if you plan to implement a four DIMM configuration. Keep in mind that the registered DIMMs have series resistors between the address, BA, WE, RAS, and CAS pins and the register load. The address, RAS, CAS, WE, and BA lines should be impedance controlled nets with 0.254 mm (10 mils) spacing if possible.

The CS, CKE, and MUX control lines must also be impedance controlled nets. They can be parallel terminated at the end with 50 Ω to V_{TT}, or series terminated with an impedance-matching resistor near the source (defined as <125 ps from the CPC945 bridge).

6.4 PPC970FX/CPC945 Evaluation Board Routing Rules for DDR2 Interface

The rest of the section contains the routing rules used on the PPC970FX/CPC945 evaluation board. Routing constraints are specified on a group basis.

The DDR2 interface is routed on the board top layer (layer 1), and inner layers 5, 7, 10, and 12. The data signals are predominantly on layers 1, 5,and 12, with the data strobes on layers 7 and 10. The series termination resistors are placed as close to the CPC945 bridge as is feasible, while not impacting routing guidelines.

In general, follow these guidelines:

- Minimize the number of vias.
- Provide 0.254 mm (10 mils) trace-to-trace clearance; closer when fanning out of the ball grid array (BGA).
- Provide 0.381 mm (15 mils) clearance from DQS lines.
- Provide 0.381 mm (15 mils) clearance from clock lines.

6.4.1 DQS and Data Line Routing

- Lines must be no shorter than 2.54 cm (1 in.) and no longer than 17.78 cm (7 in.), matched to within 1.27 mm (50 mils) of each other.

- The finished routed length of DQSx signals is the target length for the data bits in the associated byte lane.
- Data lines must be matched to within 1.27 mm (50 mils) of each other in a byte and matched to the corresponding DQS line length, within .0635 mm (25 mils).

6.4.2 Clock Line Routing

- Make the clock line as short as reasonable; no target length—match the lengths of the CK, $\overline{CK}$ differential pair (±0.127 mm [0.005 in.]: (treat as a clock, 0.381 mm [15 mils] clearance).

6.4.3 Address and Control Line Routing

- Lines must be no shorter than 5.08 cm (2 in.) and no longer than 6.606 cm (2.6 in.), matched to within 1.27 cm (500 mils) of each other.

7. Service Processor Interface and General Layout Guidelines

7.1 Functional Overview

The PowerPC 970FX processor requires a more complicated power-on reset sequence than earlier PowerPC processors. The microprocessor initialization is performed by on-chip logic and is initiated and controlled by sequencing signals from a service processor unit (SPU) over the interintegrated circuit (I^2C) or Joint Test Action Group (JTAG) port. The SPU, usually implemented with a microcontroller, initiates and monitors the PowerPC 970FX initialization and test sequences to ensure its correct operation. The SPU role does not have to be limited to initializing the PowerPC 970FX processor. Having a microcontroller act as an SPU allows for design flexibility because it can take on additional board level tasks such as setting and monitoring the supply voltages, monitoring the processor's thermal diode, and initializing other on-board devices.

The service processor does not have to be a 32-bit microcontroller; the prime factor leading to the choice of SPU is the workload. If the SPU is only initializing the PowerPC 970FX processor and its associated bridge chip, a small 8-bit microcontroller with minimal read-only memory (ROM) is sufficient. If other board level functions are included in the SPU workload, more memory and processing power might be required. Alternatively, some card designs might already require the presence of a watchdog or low-level monitoring controller that provides connectivity such as a 10/100 ethernet port, serial communications, universal serial bus (USB), peripheral component interconnect (PCI), or other functions. If that is the case, the role of the SPU can be easily incorporated into the programming for the design's existing microcontroller.

7.2 PowerPC 970FX I^2C and JTAG Interface Design Considerations

7.2.1 Interoperability of I^2C and JTAG Interface

The primary electrical connection between the service processor and the processor is the I^2C or JTAG bus. The PowerPC 970FX I^2C and JTAG interfaces share common logic, so designers must plan accordingly. If a system supports JTAG attached debuggers, special considerations are required to prevent the I^2C and the chip-level JTAG I/O from interfering with each other. The PowerPC 970FX processor has an output semaphore pin, I2CGO, that can be used with external logic to quiesce the off-chip JTAG or I^2C slave while the other is active.

Note: The RISCWatch JTAG debugger does not sample the I2CGO pin; therefore, for systems planning to support RISCWatch, consider gating off the SPU I^2C interface with I2CGO.

See the *IBM PowerPC 970FX RISC Microprocessor User's Manual* and the *IBM PowerPC 970FX Power On Reset Application Note* for additional information about the I^2C and JTAG internal logic and operation.

7.2.2 I^2C Interface Voltage Level and Operating Speed

The voltage level on the I^2C interface is the same as the processor interface's I/O voltage level. The *IBM PowerPC 970FX RISC Microprocessor Datasheet* provides the signal levels for this interface under the heading for "NonPI input" and "NonPI output" signals in the dc Electrical Specifications table. For example, if OV_{DD} is 1.5 V, the input low voltage for the I^2C interface is 0.45 V because the NonPI input low voltage is specified as ($0.3 \times OV_{DD}$).

Additionally, the datasheet has information about correct pull-up selection. To avoid problems in determining the correct pull-up resistor value, do not wire the level-shifted PowerPC 970FX I^2C bus pins together with non-PowerPC 970FX parts in a system. Each PowerPC 970FX processor should have its own level shifter.

As noted in the *IBM PowerPC 970FX RISC Microprocessor Datasheet* section on I^2C and JTAG, the I^2C bus speed is limited to 50 KHz for the standard-mode timing specification and does not support the high-speed (Hs-mode) or fast-mode timing. Review the datasheet for information about interfacing, programming, and known limitations of these interfaces. Also, see the *IBM PowerPC 970FX RISC Microprocessor Errata List for DD3.X* for additional information about nonstandard I^2C interface operation at the beginning of the power-on reset (POR) process.

7.2.3 JTAG Interface Operation

The *IBM PowerPC 970FX RISC Microprocessor Datasheet* also contains information about the interface ac timing, including a list of the PowerPC 970FX nonstandard Institute of Electrical and Electronics Engineers (IEEE) ac timing implementations. The datasheet describes in detail the correct pull-ups and pull-downs on the JTAG pins, which vary depending on the intended use of I^2C, JTAG, or both.

The *IBM PowerPC 970FX RISC Microprocessor User's Manual* has information on JTAG programming details. The PowerPC 970FX processor does not completely conform to the IEEE Standard Test Access Port and Boundary Scan Architecture (IEEE Std. 1149.1-1990); however, this does not mean that boundary-scan functions cannot be performed. To perform the boundary-scan function and PCB-based testing, additional steps must be taken to kick off boundary-scan actions. This modified testing method relies upon setting the processor in the correct mode of operation, controlling the system clock and the JTAG clock signals, and shifting serial data streams into the PowerPC 970FX processor, which sets the level of the internal BSDL latches and therefore the I/O pins, or pushes a serial data stream out of the part to read the state of the I/O pins. To learn more, see the *IBM PowerPC 970FX Boundary Scan Application Note*.

7.3 CPC945 I^2C and JTAG Interface Design Considerations

7.3.1 I^2C Interface Voltage Level and Operating Speed

The CPC945 bridge has three I^2C interfaces; the two master I^2C interfaces are available for use with dual in-line memory module (DIMM) serial presence detect (SPD), and the slave I^2C interface is for communication with the SPU or external I^2C debuggers. See the *CPC945 Bridge and Memory Controller Datasheet* for information about operating voltage levels for the JTAG interface.

Note: The I^2C voltage levels of 2.5 V from the datasheet needs to be translated to 3.3 V to interface to the memory and the service processor.

The *CPC945 Bridge and Memory Controller Datasheet* also describes the I^2C interface operation, addressing and programming. Review the datasheet for information about the I^2C interface speed and modes supported, and especially review the *IBM PowerPC CPC945 Errata List for DD1.X* for additional information about nonstandard operation of the slave I^2C interface.

7.3.2 JTAG Interface Operation

The JTAG interface on the CPC945 bridge is separate from the I^2C interfaces. See the *CPC945 Bridge and Memory Controller Datasheet* for information about operating voltage levels for the JTAG interface. The BSDL file for the CPC945 bridge can be downloaded from IBM Customer Connect (see *Related Documents* on page 11).

7.4 Overview of PPC970FX/CPC945 Evaluation Board Service Processor Implementation

The PPC970FX/CPC945 evaluation board uses the AMCC PowerPC 405GPr microcontroller as its service processing unit (for more information, go to http://www.amcc.com/ and select embedded processors under the products). This highly integrated system-on-a-chip (SOC) device supports a high-level software system that allows it to drive and control its many on-chip resources such as Ethernet, PCI, and I^2C, and enables easy code development and testing. Many designs using the PowerPC 970FX processor do not need an SPU with this level of processing power or peripheral mix. However, the PowerPC 970FX processor uses the 405GPr service processor to take advantage of the existing PowerPC initialization boot software (PIBS) and IBM Embedded PowerPC Operating System (EPOS) firmware. (IBM provided the firmware for the PPC970FX/CPC945 evaluation board.)

The IBM PPC970FX/CPC945 evaluation kit includes software for both the PowerPC 970FX processor and the 405GPr service processor, which includes the PIBS resident in the flash memory on the board, the PIBS source code, the EPOS, sample application programs, and application development libraries and tools. Documentation includes software technical specifications and an application note that describes step-by-step how to obtain and build GNU software development tools for use with the evaluation kit software. You can download this code from the IBM developerWorks Web site at http://www.ibm.com/developerworks/power/pibs/

As an SPU, the primary purpose of the 405GPr processor is to communicate with the PowerPC 970FX processors, the selected bridge chip, and other on-board devices via the I^2C bus and general-purpose-I/O (GPIO)-based control lines. The PPC970FX/CPC945 evaluation board uses the I^2C bus to initialize and configure the PowerPC 970FX processors and CPC945 bridge chip and to communicate with various on-board I^2C-based devices and the thermal diode monitoring circuits.

7.5 Routing Guidelines

The I^2C bus and control lines must be carefully routed across the PCB to minimize crosstalk on the bus and interference with other signals. Depending upon the I^2C voltage levels involved, it might be necessary to use some type of I^2C hub and level shifters as was done on the PPC970FX/CPC945 evaluation board. The PPC970FX/CPC945 evaluation board uses the Phillips Semiconductor PCA9516 hub and PCA9517 level shifters. Designers must be aware of the input voltage requirements for these parts and ensure that the interfaces can drive signals to the correct levels. If the I^2C interface does not operate reliably, review the information on the Philips Web site, and check signal levels to make sure that the logic low voltage is low enough.

The PPC970FX/CPC945 evaluation board contains circuitry that is used to drive and read the thermal diodes of each PowerPC 970FX processor. To derive the temperature of the processor core, the thermal diode must be properly biased. Then, the sensor circuit can read the voltage across the diode and convert it into the temperature of the processor. The conversion process has separate analog and digital sections, and can be

done with either a standalone microcontroller or by the SPU (one with either internal analog-to-digital (A/D) converter hardware or one using external logic.) If done externally, account for the heat generated by the thermal diode sensing chip (see *Section 7.7 Thermal Considerations* for more information).

7.6 Electrical Considerations

The primary function of a service processor is to manage the orderly initialization of the system board. This process involves managing the voltage levels of the many different possible power supplies, initializing the PowerPC 970FX processors, and other functions. The service processor should be one of the first items to receive power when the system board is turned on, since the SPU will interface with a variety of components on the system board and must be able to reliably communicate with them. The specific requirements of the SPU are supplied by the component manufacturer.

The SPU does not have to be placed next to the PowerPC 970FX processors, but it must be able to reliably communicate with the PowerPC 970FX processors, bridge chip, and other on-board peripherals.

7.7 Thermal Considerations

Generally, the SPU is a relatively low-power part. Low-cost SPU solutions can dissipate less than a watt of power, but this power must be accounted for when developing the overall thermal solution for the end product. Some SPU solutions might dissipate more power and might even need some type of heat sink to function correctly. The thermal environment in which the SPU operates must be evaluated to ensure that the SPU is adequately cooled and does not affect other on-board devices. It might be helpful to do a computational fluid dynamics thermal model to evaluate the thermal environment including the SPU.

8. Power Delivery Guidelines

This section provides power delivery guidelines for integration of the PowerPC 970FX processor and the target system.

8.1 Overview

Depending on the system architecture and performance requirements, the power requirements can vary between 14 W, for a simple 1.0 GHz PowerPC 970FX system, to over 150 W, for a 2.2 GHz dual processor system. Designing the power delivery for the low-end systems is fairly straight forward, but becomes a reasonably significant design effort at the high end. Additionally, the designer should be aware of system requirements for any of the power-savings modes, especially those involving voltage slewing or frequency reductions (f/2 mode), and consider the impact of rapid increases and decreases in current consumption. The current draw affects the design of the printed circuit board (PCB), power supply, and decoupling requirements. Power saving modes can affect the processor core voltage requirement, frequency of operation, and processor interface operational speed. This section of the design guide describes the major issues that require attention during the design of a power supply system for the PowerPC 970FX processor. As an implementation example, this section uses the PPC970FX/CPC945 evaluation board.

8.2 PowerPC 970FX Processor Voltage Delivery

The PowerPC 970FX processor has connections for three different power supplies that deliver current to the core logic, V_{DD}, I/O drivers and receivers, OV_{DD}, and analog phase-locked loop (PLL) of the processor, AV_{DD}. The values for V_{DD} and OV_{DD} for specified frequencies are detailed in the *IBM PowerPC 970FX RISC Microprocessor Datasheet*. The analog voltage is for the on-chip PLL used as the basis of the clock generation for the chip. *Figure 8-1* is a basic block diagram of the processor and the power supplies that power the chip.

Figure 8-1. Basic Power Supplies Driving the PowerPC 970FX Processor

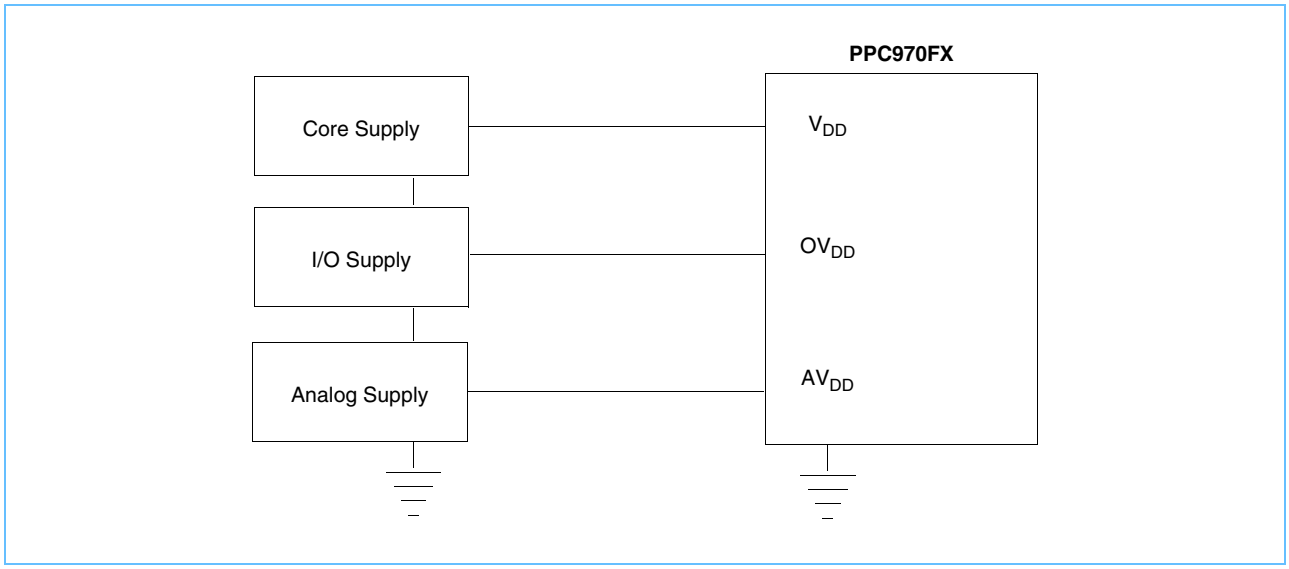

8.2.1 Board Layout Considerations

For high-end systems, handling the amount of current required by the processor is a major design issue. The processor frequency determines its operating voltage and resulting power. Given the specified power and voltage, the current requirements can be determined. With currents in the range of 14 A to over 70 A, careful consideration must be paid to power supply topology and layout. As an example, a 2.2 GHz part operating with a V_{DD} of 1.25 V and a maximum power of 76 W, draws 61 A.

Sixty one amps is a large amount of current to be routed through the thin foil of a PCB. The PCB must be carefully designed to ensure that the power is consistently delivered to the processor with the necessary levels for voltage, ripple, and noise. In the case of the PPC970FX/CPC945 evaluation board, three heavily connected 2-ounce V_{DD} layers route power from the power supply, an Artesyn VRM11 130 W dc-dc convertor, to the processor. This covers the highest power parts. For lower power requirements, modules from Artesyn, Lite-on, Linear Technologies, or other vendors can be selected to match the specific power requirements as described in the *IBM PowerPC 970FX RISC Microprocessor Datasheet*.

[[Do we need this paragraph from the 970MP design guide? “The voltage required for each part is specified by the V_{DD} fuse code (VFC) data programmed into the fuse ring of each part. The supply current must be designed to handle the range of possible current loads. For example, the 115 W, 2.5 GHz, power-optimized part has a voltage range of 1.2 - 1.35 V. The current range would then be 95.8 - 85.2 A. These values would actually be approximately 1 - 2 A lower, because that part of the power is due to OVDD, which has its own supply. For sizing purposes, the power consumption from the OV_{DD} supply will be in the 1 - 3 W range.”]]

8.3 PPC970FX/CPC945 Evaluation Board Example

The PPC970FX/CPC945 evaluation board contains several different power supplies that supply the processors, CPC945 bridge, double data rate two (DDR2) interface, and service processor. The design uses switching power supplies for flexibility and efficiency. In most systems, the service processor is the first item to receive power because it must direct the start-up sequence for the remaining power supplies. Each power supply is briefly described in the following sections.

8.3.1 Power Delivery to the PowerPC 970FX Processor

The choice of power supply design for the V_{DD} depends on the amount of current required and the preferred level of efficiency. Usually, the V_{DD} supply is provided by polyphase (multiphase), dc-to-dc switching voltage regulators. A good understanding of system power plane distribution and construction is required for high current, low voltage, high frequency (transient) processor applications. Correct copper weight, printed circuit board stackup, and proximity of the supply is required to assure a low-resistive, low-inductive path between the load (processor) and the power supply. Ensure that there is sufficient via surface area to handle the power requirements.

The V_{DD} specified in the datasheet is the required core voltage at the die. The only way to accurately determine the core voltage is by monitoring the Kelvin pins. These power and ground sense lines, designated as KVPRBVDD and KVPRBGND, are located at the high-current density areas of the die. Never tie these pins to the power and ground planes of the PCB. Use the Kelvin pins to regulate the voltage level of the power supply, but take care to avoid oscillation and other classic control issues when using feedback. See the *On Chip Voltage Monitoring on the PowerPC 970FX Application Note* for more information on the Kelvin voltage pins and their use.

It is not unreasonable to see, for high current conditions, as much as a 55 mV of drop through the package to the V_{DD} logic. This illustrates why the measurement must be done with the Kelvin pins, and why the voltage measured between planes on the PCB might not seem to conform with the specifications. If the voltage does fall below the minimum voltage for a particular frequency of operation, the processor might not operate correctly.

The I/O portion of the processor consumes less power than the core; it consumes less than 2 A. The PPC970FX/CPC945 evaluation board uses a much smaller OV_{DD} power supply than the one used for the V_{DD} logic.

The remaining power source is the analog PLL supply. This supply is low current, less than 20 mA, but must be clean and stable. According to the datasheet, the AV_{DD} supply must be filtered to ensure the stability of the internal clock. The datasheet suggests a specific AV_{DD} filter circuit, which is depicted in *Figure 8-2*. To optimize the capacitor filter noise reduction, place it as close as possible to the AV_{DD} and ANALOG_GND pins. The capacitor should have minimal inductance. The ferrite bead should supply an impedance of less than 70 Ω in the 100 - 500 MHz region. Regarding the filter circuit, check the current versions of the datasheet to be sure you are using the latest specifications and filter suggestions. Measure the AV_{DD} voltage between the AV_{DD} ball and the analog ground ball under the package. This is important to account for any drop across the filter circuit.

Note: AV_{DD} measured at the pins of the part should never be more than 50 mV lower than the AV_{DD} voltage range specified in the recommended operating conditions table in the *IBM PowerPC 970FX RISC Microprocessor Datasheet*.

Figure 8-2. PLL Power Supply Filter Circuit

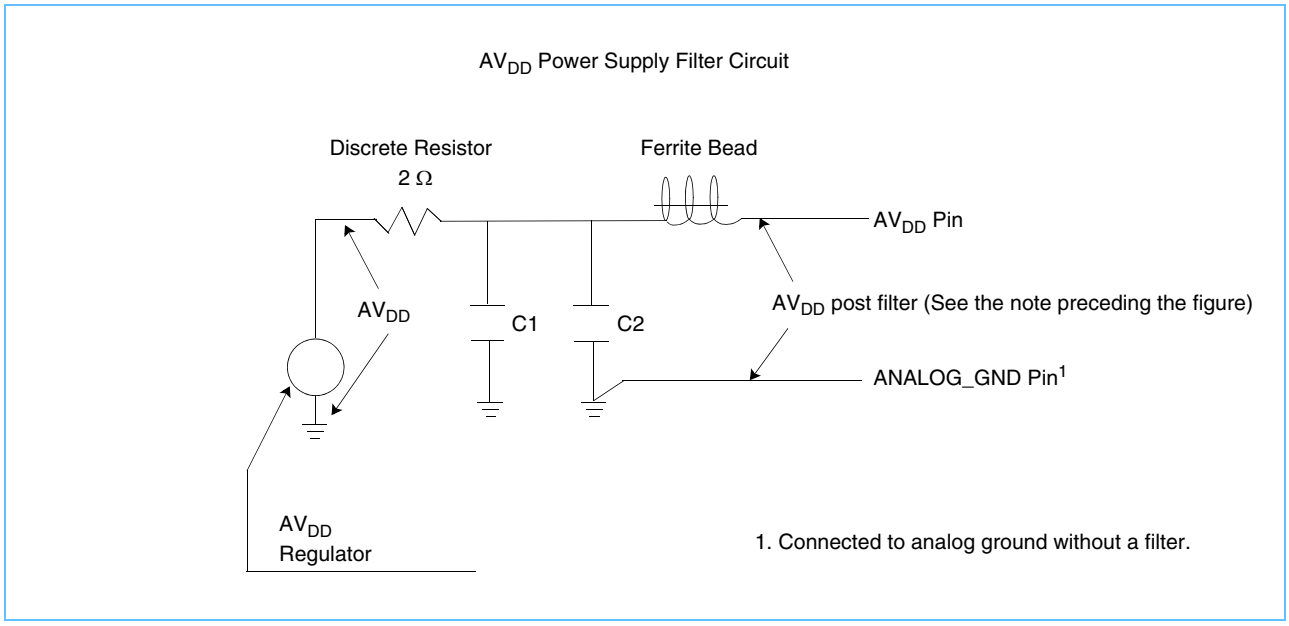

8.3.2 Using VRM Modules

Standard voltage regulation modules (VRMs) can be used for the power supply modules. Although the PowerPC 970FX processor requires a tighter tolerance on the load line than the standard VRM, vendors have been able to easily supply modified versions of their standard VRMs. The PowerPC 970FX processor does not have voltage identification (VID) bits output from the processor, but the service processor can translate the voltage specification to general purpose I/O bits that mimic VID bits to set the VRM to the required V_{DD}.

8.3.3 Power Delivery to the CPC945 Bridge

Power delivery to the CPC945 bridge is more complicated than delivery of power to the PowerPC 970FX processor because of the assortment of interfaces supported (see *Figure 8-3 PowerPC 970FX Block Diagram* on page 63). The CPC945 bridge supports a variety of system interfaces such as DDR2 memory, interintegrated circuit (I^2C), dual processor interface buses to support two PowerPC 970FX processors, HyperTransport, and PCI Express (PCI-E). The CPC945 bridge requires core, PLL, processor interface power supplies, and three ancillary supplies for the on-chip interfaces. Although some of these supplies are not directly tied to the chip, they are required to support the interface. For example, the DDR2 interface requires supplies for DDR2 terminator and reference voltage. These interfaces have their own power supply requirements. which are not covered in depth here. Before the design process is started, a thorough review of the interface specification is strongly recommended. The interface specifications can be found at the appropriate user's group or manufacturer Web sites.

The primary emphasis of this section is on the supply design for main chip power: the I/O, PLL, and core logic. The *CPC945 Bridge and Memory Controller Datasheet* table of recommended dc operating conditions lists all the required supply voltages and tolerances. Because the CPC945 bridge can support two PowerPC 970FX processors, both processor interface buses must operate at the same voltage (and frequency). Note that the I/O driver and receiver used is specified from 1.2 V nominal to 1.5 V nominal. The processor interface operating voltage is provided by supply voltage V_{DD2}, and must correspond directly to the processor interface voltage of the PowerPC 970FX processor interfaces. System designers should be aware that the slave I^2C interface, used for communication with the service processor, is powered from V_{DD5} and might require voltage translation to other I^2C devices in the system.

While it might be possible to design a power distribution where the 1.5 V supply is shared across different interfaces, depending on the target system architecture, it might not be practical to share supplies.

In the PPC970FX/CPC945 evaluation board design routing rules, there are no special requirements listed for CPC945 V_{DD} and I/O power routing, although good design practice recommends the use of planes for power and ground for correct operation.

Figure 8-3. PowerPC 970FX Block Diagram

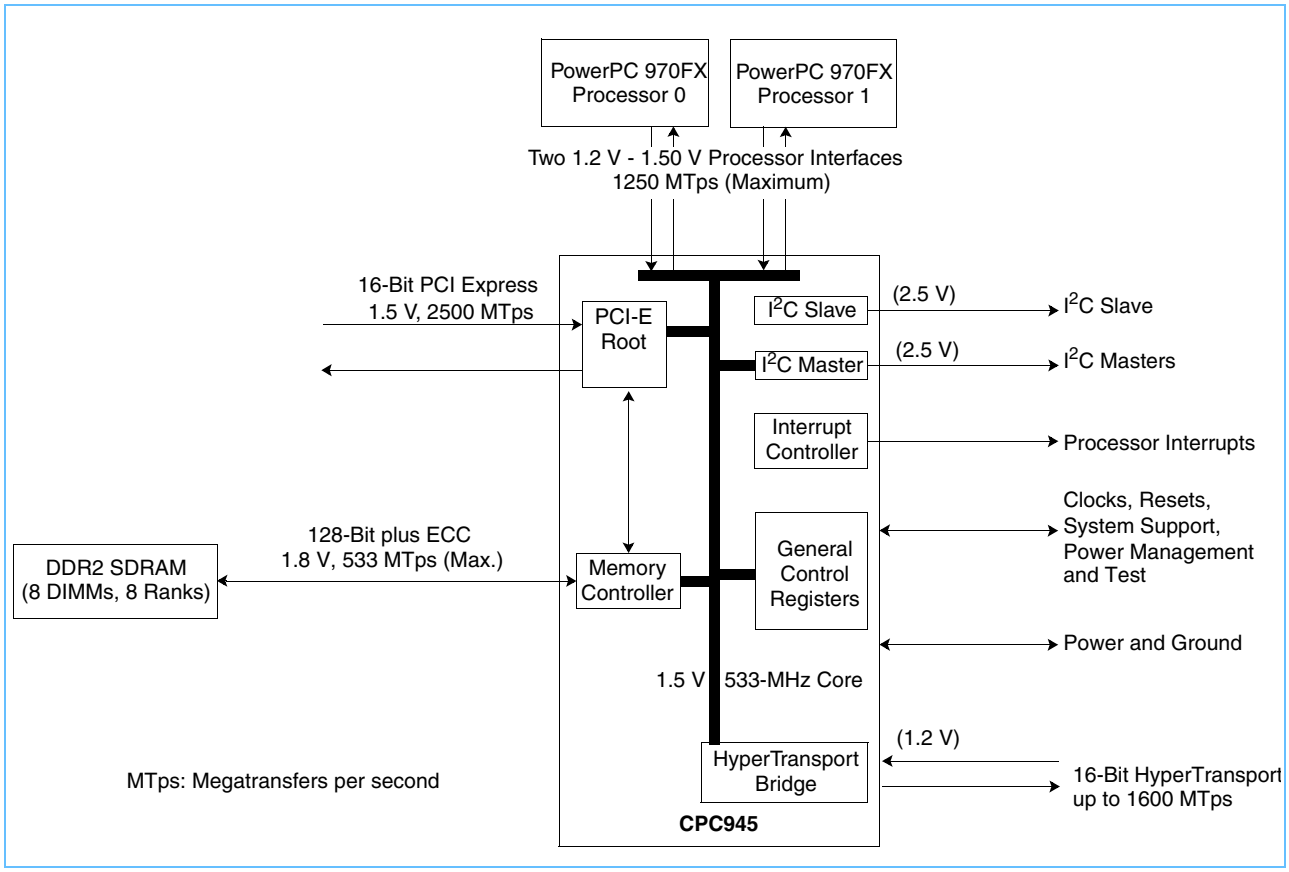

8.3.4 Power Delivery to the HyperTransport Tunnel

The HyperTransport interface logic and data I/Os require a 1.2 V supply (see *Figure 8-3*). Additional, miscellaneous I/O pins require a separate 2.5 V supply. For additional information, see the *CPC945 Bridge and Memory Controller Datasheet* and the *HyperTransport Interface Specification*.

8.3.5 Power Delivery to the DDR2 Interface

The DDR2 interface I/Os require a 1.8 V supply (see *Figure 8-3*). The master I^2C interface, for connection to the memory DIMMs, uses a 2.5 V supply. In addition, there are inputs for the DDR2 reference voltage (V_{REF}). For additional information, see the *CPC945 Bridge and Memory Controller Datasheet* and the *JEDEC DDR2 Memory Specification*. Also, be sure to obtain and read any application notes and design guideline documentation from your chosen DDR2 memory module or dual in-line memory module (DIMM) supplier.

Note: In general, while the termination voltage (V_{TT}) and V_{REF} are the same voltage, they should be handled differently. V_{REF} is a low current input that must be very stable and quiet. Do not derive the reference voltage from the V_{TT} supply. The V_{REF} divider resistors should be 1% tolerance components with low resistance values to supply a solid reference voltage that is not affected by switching currents. For all nine V_{REF} pins, the PowerPC 970FX/CPC945 evaluation design uses 499 Ω resistors with bypass capacitors.

For the V_{TT} supply, memory suppliers suggest fast-response linear regulators, capable of supplying up to several ampheres of current for many memory configurations. V_{TT} must be designed with sufficient trace area, sufficient decoupling, and low inductance to keep the noise on V_{TT} to within 100 mV in a high di/dt (noisy) environment. Pay particular attention to decoupling recommendations from your memory supplier, and understand that decoupling capacitor placement is critical; even a few millimeters can affect the decoupling efficiency.

8.3.6 Power Delivery to the PCI-Express Interface

The PCI-E interface logic and data I/Os require a 1.5 V supply (see *Figure 8-3*). For additional information, see the *CPC945 Bridge and Memory Controller Datasheet* and the *PCI Express Specification*

8.3.7 Power Delivery to the Analog Logic

The remaining CPC945 power sources are the analog PLL supplies. These supplies power the four PLLs used internally in the CPC945 bridge to generate the different interface and core-logic clock domains. These supplies must be clean and stable. According to the *CPC945 Bridge and Memory Controller Datasheet*, the AV_{DD} supply must be filtered to assure the stability of the internal clock. The datasheet suggests the specific AV_{DD} filter circuits shown in *Figure 8-4* and *Figure 8-5*. The recommended high-frequency capacitor should be of type X5R ceramic construction, size 0603, and have a 6.3 V rating. The low-frequency bulk decoupling capacitor should be a tantalum surface-mount technology (SMT) capacitor. The smaller valued capacitor should be as close as possible to the PLL power supply pins. The larger valued capacitor should be located adjacent to the device package. The analog ground should be directly connected to the filter, then to the digital ground through a single point connection. There are three filter circuits for PCI-E, two filter circuits for HyperTransport, and one each for the processor interface and the double data rate (DDR) bus. The capacitor values shown should be regarded as starting values. On the PowerPC 970FX/CPC945 evaluation design, the HyperTransport AV_{DD} filters use the values shown in *Figure 8-4*, while the PCI-E uses 22 µF and 0.1 µF capacitors.

Figure 8-4. Analog V_{DD} Filtering for the HyperTransport and PCI Express Phase-Locked Loops

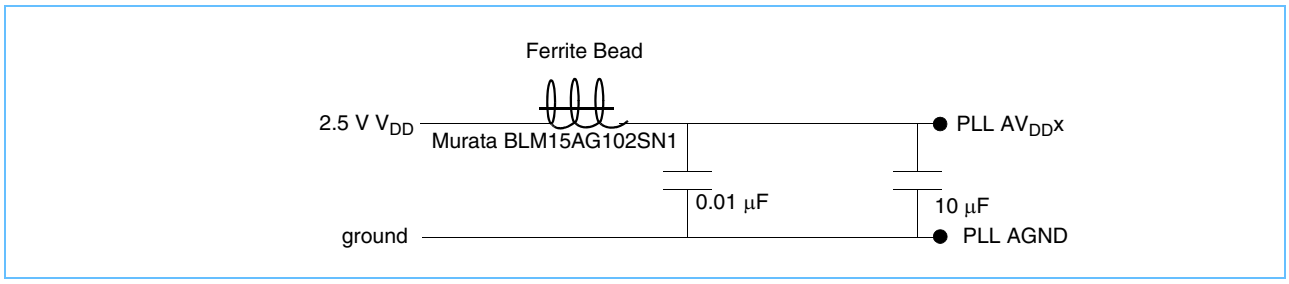

Figure 8-5. Analog V_{DD} Filtering for the Processor Interface and DDR2 Interface Phase-Locked Loops

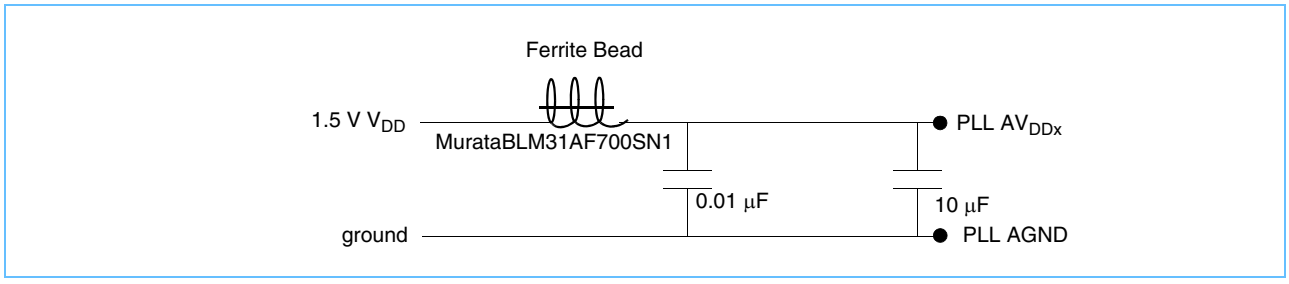

8.3.8 Processor Voltage Sequencing

The processor power supplies must be brought up in a specific sequence to ensure correct operation. The processor must not begin the power-on reset (POR) sequence until V_{DD}, OV_{DD}, and AV_{DD} voltages are within their specified tolerances. According to the *IBM PowerPC 970FX RISC Microprocessor Datasheet*, the following considerations apply:

- The power supply ramping order does not matter as long as the supplies reach their final destination in 50 ms.
- V_{DD} cannot exceed OV_{DD} by more than 0.8 V, except for 50 ms during power up or down. (During those intervals, it is allowed to be approximately 1.35 V.)
- OV_{DD} cannot exceed V_{DD} by more than 0.8 V except for 50 ms during power up or down. (During those intervals, it is allowed to be approximately 1.55 V.)
- AV_{DD} cannot exceed V_{DD} by more than 2.5 V except for 50 ms during power up or down. (During those intervals, it is allowed to be approximately 2.75 V.)

8.3.9 CPC945 Bridge Power Sequencing

The CPC945 bridge also has its own power requirements. A specific time interval recommendation is not given. The core voltage should be brought up first, followed by the I/O voltages. Note that correct I^2C slave interface operation depends on stable core, processor I/O, and DDR2 I/O voltages. No voltage should be applied to an I/O pad if the associated power supply is not on.

8.3.10 Processor Decoupling Recommendations

The decoupling recommendations contained in the datasheet call for the use of many small 0402 capacitors to provide a high-frequency, low-inductance power source for the di/dt currents generated by the processor. See the latest datasheet to review the decoupling recommendations for the applicable processors and the bridge chip. Bulk decoupling requires an understanding of the power distribution topology for the board, and an understanding of the output capacitance requirements of the dc-to-dc switching regulator design. Bulk decoupling recommendations are therefore beyond the scope of this document. The *CPC945 Bridge and Memory Controller Datasheet* contains the recommended placement of the decoupling capacitors.