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

Design Guide

SA14-970MPDG-03

September 13, 2007

© Copyright International Business Machines Corporation 2006, 2007
All Rights Reserved
Printed in the United States of America September 2007

The following are trademarks of International Business Machines Corporation in the United States, or other countries, or both.

IBM
ibm.com
IBM Logo
developerWorks
PowerPC
PowerPC Architecture

Intel is a registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.

Other company, product, and service names may be trademarks or service marks of others.

All information contained in this document is subject to change without notice. The products described in this document are NOT intended for use in applications such as implantation, life support, or other hazardous uses where malfunction could result in death, bodily injury, or catastrophic property damage. The information contained in this document does not affect or change IBM product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of IBM or third parties. All information contained in this document was obtained in specific environments, and is presented as an illustration. The results obtained in other operating environments may vary.

THE INFORMATION CONTAINED IN THIS DOCUMENT IS PROVIDED ON AN “AS IS” BASIS. In no event will IBM be liable for damages arising directly or indirectly from any use of the information contained in this document.

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

SA14-970MPDG-03
September 13, 2007

Contents

List of Figures

List of Tables

Revision Log

Revision Date	Description
September 13, 2007	SA14-970MPDG-03 • Edited the document for clarity and consistency.
March 19, 2007	Version 1.2 • Added *Section 1.5 PPC970MP/CPC945 PowerPC 970MP Processor and CPC945 Bridge System Bring Up* on page 23.
July 31, 2006	First Issue

About this Book

This design guide contains design recommendations for systems based on the IBM PowerPC® 970MP 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 970MP/CPC945 evaluation board schematics for reference. The board is supplied for the evaluation of the PPC970MP 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 970MP 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 970MP 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 details of the PowerPC Architecture™.

---

1. In this document, the IBM PowerPC 970MP Microprocessor is abbreviated as PowerPC 970MP processor or PPC970MP.
2. In this document, the CPC945 Bridge and Memory Controller is abbreviated as CPC945 bridge or CPC945.

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.

Important: Check this IBM Web site to ensure 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/.

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 PowerPC 970MP Documentation

- *IBM PowerPC 970MP RISC Microprocessor User's Manual*
- *IBM PowerPC 970MP RISC Microprocessor Datasheet*
- *IBM PowerPC 970MP RISC Microprocessor Errata List for DD1.0x, DD1.1*
- *PowerPC 970MP Power On Reset Application Note*
- *IBM PowerPC 970MP/CPC945 Automated PI Tuning Application Note*
- *Collecting Thermal Diode Calibration Values for the Dual-Core PowerPC 970MP*
- *Using the Kelvin Voltage Sense Pins in the PPC970MP*
- *Using Thermal Diodes in the PowerPC970MP® Processor*
- *PowerPC 970MP Thermal Considerations Application Note*
- *PowerPC970 Debug Notes Application Note*
- *Improving BGA to PCB Thermo-Mechanical Integrity White Paper*
- *PowerPC 970MPProcessor Interface (PI) Bus IBIS file*
- *PowerPC 970MP HSPICE Models*

IBM CPC945 Documentation

- *CPC945 Bridge and Memory Controller User Manual*
- *CPC945 Bridge and Memory Controller Datasheet*
- *IBM PowerPC CPC945 RISC Microprocessor Errata List for DD 1.2*
- *CPC945 Bridge and Memory Controller Memory Signal Delay Tuning Application Note*
- *PowerPC CPC945 Thermal Considerations Application Note*
- *CPC945 BSDL file*

The hardware documentation for the PPC970MP/CPC945 evaluation board is included in the CD-ROM delivered with every board. The firmware source and documentation for the PowerPC 970MP processor 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
VREF	DDR2 reference voltage
VRM	voltage regulation modules
VTT	termination voltage
WE	write enable

1. System Overview

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

1.1 PowerPC 970MP Processor

The PowerPC 970MP is a dual core, 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, up through 4-way symmetric multiprocessor (SMP) configurations.

1.1.1 Architectural Features

- 64-bit implementation of the PowerPC Architecture™ (version 2.0.1)
- Dual cores
- 1.2 - 2.0 GHz core frequency operation, supporting high-speed processor interface data-transfer speeds up to 1.0 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
- 1 MB 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 970MP 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 970MP 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 970MP RISC Microprocessor User’s Manual* for more information.)

1.1.2 Packaging and Power

- Available in a 575-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 0.9 V to 1.3 V. (See the *IBM PowerPC 970MP 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 2.70 V.

For additional details on packaging or power, see the *IBM PowerPC 970MP 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 970MP processor. To implement power tuning as described here, your design must use a bridge or system controller that supports this feature. In the PPC970MP/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 970MP RISC Microprocessor User's Manual*.

1.1.4 Thermal Management Using a Thermal Diode Control Application

The PowerPC 970MP processor features two on-board temperature sensing diodes, one for each core. They are located in the hottest portion of the chip. For information about these diodes, see the *IBM PowerPC 970MP RISC Microprocessor Datasheet*. For information about manual calibration techniques and examples, see the *PowerPC 970MP Thermal Considerations Application Note*.

Other temperature monitoring hardware can also be implemented with the PowerPC 970MP processor and mounted as close to the PowerPC 970MP processor as practical. The PPC970MP/CPC945 evaluation board provides one example of external temperature-sensing hardware; see the PPC970MP/CPC945 evaluation board documentation for details. The PPC970MP/CPC945 evaluation board documentation includes details of a programmable system-on-a-chip (SOC)-based circuit that monitors the PowerPC 970MP 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 (PPC970FX) and dual core (PPC970MP) 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 970MP processor interface buses with cache coherency and snooping, running up to 500 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 1 gigatransfers per second (GTps)
- Interrupt controller with support for 8 internal and up to 120 external interrupts (routed over the Hyper-Transport interface)

See the *CPC945 Bridge and Memory Controller User 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, connected to pins AP30 (SYS_THDIOD_D) and AP29 (SYS_THDIO_G). The thermal monitor consists of a vertical PNP bipolar junction transistor, with the emitter connected to pin AP30 and the base connected to pin AP29. 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 970MP general market availability is 2.0 GHz.
- The fastest processor interface transfer rate for a 2.0 GHz core frequency is attained with a bus ratio of 2:1 and is 500 MHz DDR or 1.0 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 PPC970MP/CPC945 Evaluation Board System Configuration

As *Figure 1-1* on page 22 shows, the PPC970MP/CPC945 evaluation board includes the following major system components:

- Dual IBM PowerPC 970MP reduced instruction set (RISC) processors
- AMCC PowerPC 405GPr (405GPr) service and control processor
- IBM 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. PPC970MP/CPC945 Evaluation Board with Two Processor Cards

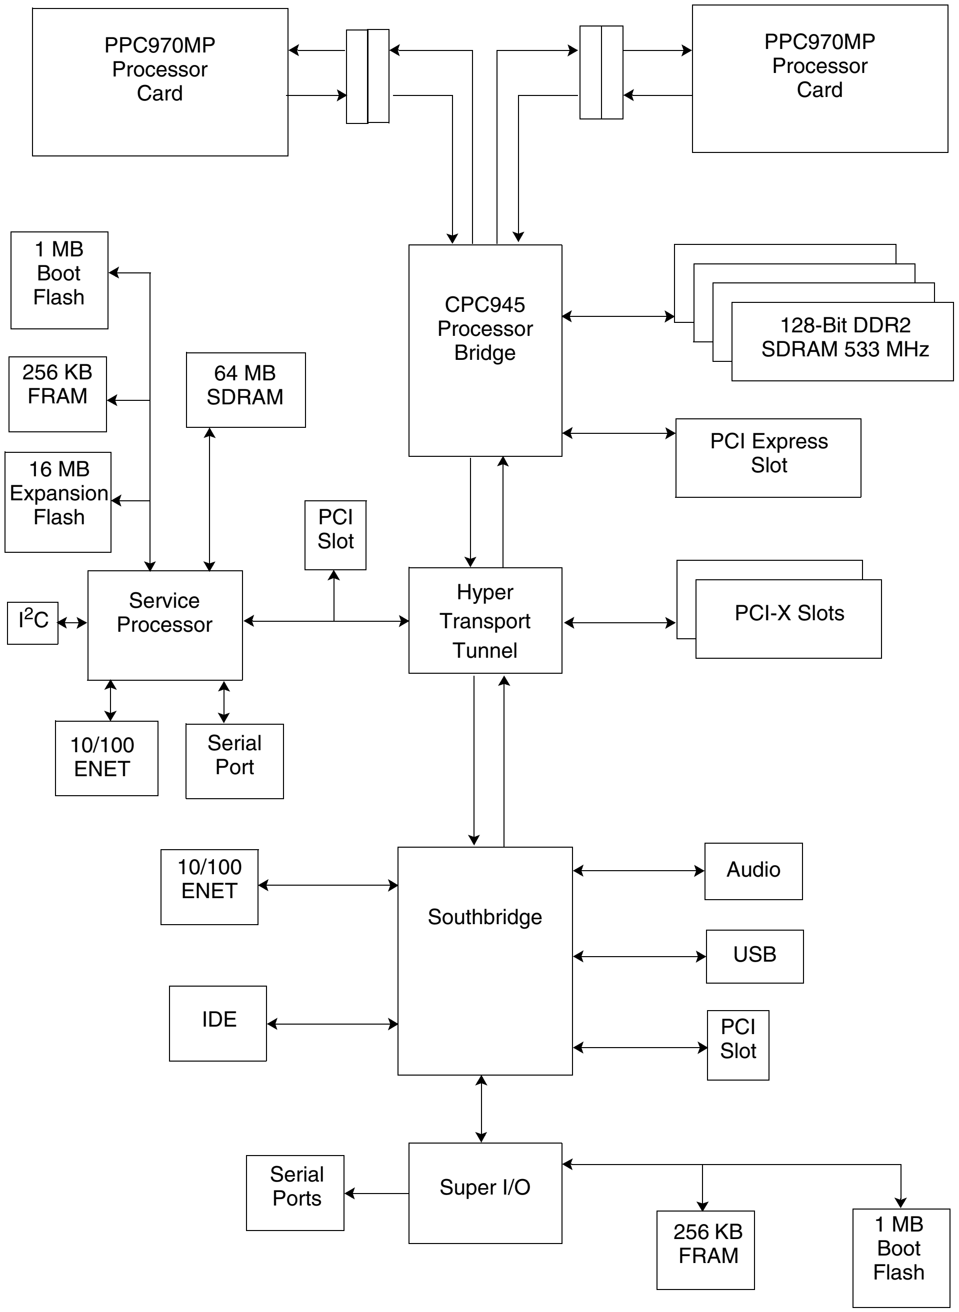

Note: The processor cards are separate assemblies that plug onto 400-pin MEG-array connectors.

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

1.5 PPC970MP/CPC945 PowerPC 970MP Processor and CPC945 Bridge System Bring Up

Because these processors are configurable, parameters must be set to define the configuration. The *PowerPC* 970MP *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 970MP designs use a service processor to initialize the system board. A significant portion of this initialization is done using the I^2C bus. The service processor used on the PPC970MP/CPC945 evaluation board is an AMCC PowerPC 405GPr (405GPr). The service processor runs a start up script or program for initialization.

The actual script used to initialize the PPC970MP/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 *PowerPC 970MP 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 970MP 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 the technical library in IBM Customer Connect (see *Related Documents* on page 12).

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 *CPC945 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 970MP RISC Microprocessor Datasheet*.

Figure 2-1, *Figure 2-2*, and *Figure 2-3* show the signal assignments for the PowerPC 970MP processor and 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 IBM PowerPC 970MP Microprocessor Ball Placement (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^{2}C) signals	🟪
Thermal diodes and Kelvin pins	🟪

Table 2-2. Signal Groupings for CPC945 Bridge and Memory Controller Pinout (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 970MP Microprocessor Ball Placement (Top View)

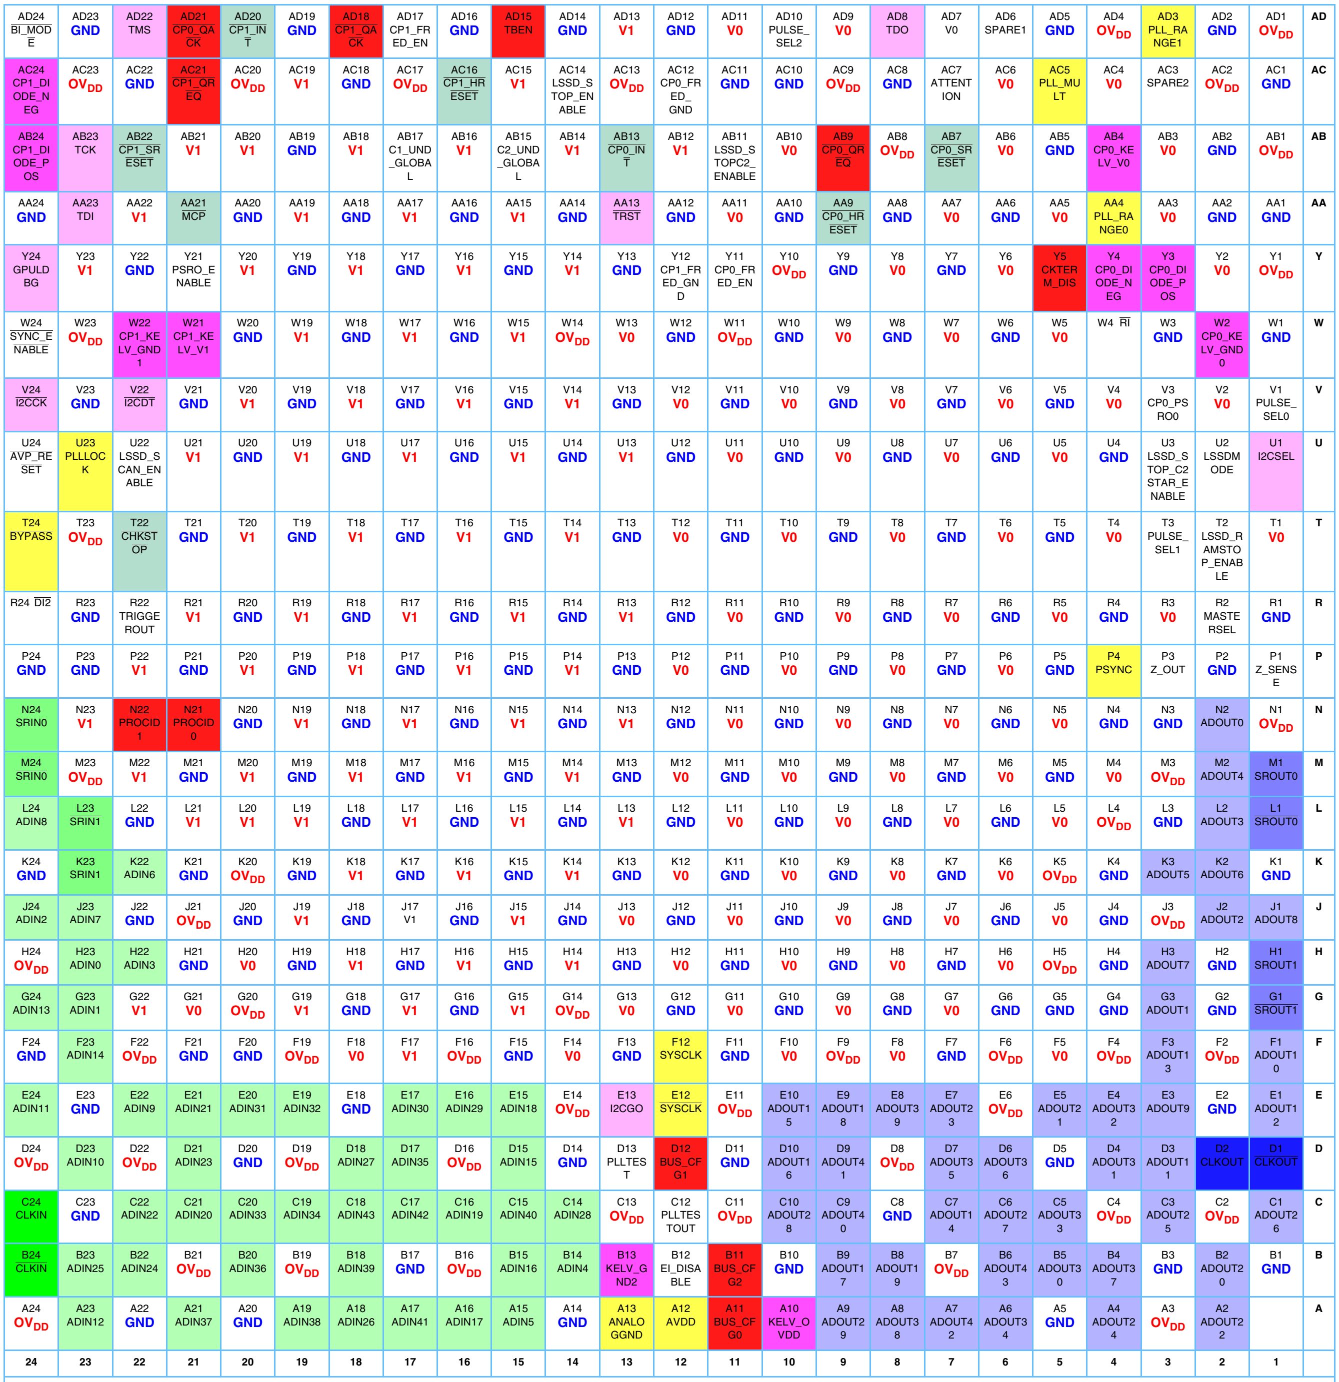

AD24 BI_MODE	AD23 GND	AD22 TMS	AD21 CP0_QACK	AD20 CP1_INT	AD19 GND	AD18 CP1_QACK	AD17 CP1_FRED_EN	AD16 GND	AD15 TBEN	AD14 GND	AD13 V1	AD12 GND	AD11 V0	AD10 PULSE_SEL2	AD9 V0	AD8 TDO	AD7 V0	AD6 SPARE1	AD5 GND	AD4 OV$_{DD}$	AD3 PLL_RANGE1	AD2 GND	AD1 OV$_{DD}$	**AD**
AC24 CP1_DIODE_NEG	AC23 OV$_{DD}$	AC22 GND	AC21 CP1_QREQ	AC20 OV$_{DD}$	AC19 V1	AC18 GND	AC17 OV$_{DD}$	AC16 CP1_HRESET	AC15 V1	AC14 LSSD_STOP_ENABLE	AC13 OV$_{DD}$	AC12 CP0_FRED_GND	AC11 GND	AC10 GND	AC9 OV$_{DD}$	AC8 GND	AC7 ATTENTION	AC6 V0	AC5 PLL_MULT	AC4 V0	AC3 SPARE2	AC2 OV$_{DD}$	AC1 GND	**AC**
AB24 CP1_DIODE_POS	AB23 TCK	AB22 CP1_SRESET	AB21 V1	AB20 V1	AB19 GND	AB18 V1	AB17 C1_UND_GLOBAL	AB16 V1	AB15 C2_UND_GLOBAL	AB14 V1	AB13 CP0_INT	AB12 V1	AB11 LSSD_STOPC2_ENABLE	AB10 V0	AB9 CP0_QREQ	AB8 OV$_{DD}$	AB7 CP0_SRESET	AB6 V0	AB5 GND	AB4 CP0_KELV_V0	AB3 V0	AB2 GND	AB1 OV$_{DD}$	**AB**
AA24 GND	AA23 TDI	AA22 V1	AA21 MCP	AA20 GND	AA19 V1	AA18 GND	AA17 V1	AA16 GND	AA15 V1	AA14 GND	AA13 TRST	AA12 GND	AA11 V0	AA10 GND	AA9 CP0_HRESET	AA8 GND	AA7 V0	AA6 GND	AA5 V0	AA4 PLL_RANGE0	AA3 V0	AA2 GND	AA1 GND	**AA**
Y24 GPULDBG	Y23 V1	Y22 GND	Y21 PSRO_ENABLE	Y20 V1	Y19 GND	Y18 V1	Y17 GND	Y16 V1	Y15 GND	Y14 V1	Y13 GND	Y12 CP1_FRED_GND	Y11 CP0_FRED_EN	Y10 OV$_{DD}$	Y9 GND	Y8 V0	Y7 GND	Y6 V0	Y5 CKTERM_DIS	Y4 CP0_DIODE_NEG	Y3 CP0_DIODE_POS	Y2 V0	Y1 OV$_{DD}$	**Y**
W24 SYNC_ENABLE	W23 OV$_{DD}$	W22 CP1_KELV_GND1	W21 CP1_KELV_V1	W20 GND	W19 V1	W18 GND	W17 V1	W16 GND	W15 V1	W14 OV$_{DD}$	W13 V0	W12 GND	W11 OV$_{DD}$	W10 GND	W9 V0	W8 GND	W7 V0	W6 GND	W5 V0	W4 RI	W3 GND	W2 CP0_KELV_GND0	W1 GND	**W**
V24 I2CCK	V23 GND	V22 I2CDT	V21 GND	V20 V1	V19 GND	V18 V1	V17 GND	V16 V1	V15 GND	V14 V1	V13 GND	V12 V0	V11 GND	V10 V0	V9 GND	V8 V0	V7 GND	V6 V0	V5 GND	V4 V0	V3 CP0_PSRO0	V2 V0	V1 PULSE_SEL0	**V**
U24 AVP_RESET	U23 PLLLOCK	U22 LSSD_SCAN_ENABLE	U21 V1	U20 GND	U19 V1	U18 GND	U17 V1	U16 GND	U15 V1	U14 GND	U13 V1	U12 GND	U11 V0	U10 GND	U9 V0	U8 GND	U7 V0	U6 GND	U5 V0	U4 GND	U3 LSSD_STOP_C2STAR_ENABLE	U2 LSSDMODE	U1 I2CSEL	**U**
T24 BYPASS	T23 OV$_{DD}$	T22 CHKSTOP	T21 GND	T20 V1	T19 GND	T18 V1	T17 GND	T16 V1	T15 GND	T14 V1	T13 GND	T12 V0	T11 GND	T10 V0	T9 GND	T8 V0	T7 GND	T6 V0	T5 GND	T4 V0	T3 PULSE_SEL1	T2 LSSD_RAMSTOP_ENABLE	T1 V0	**T**
R24 DI2	R23 GND	R22 TRIGGEROUT	R21 V1	R20 GND	R19 V1	R18 GND	R17 V1	R16 GND	R15 V1	R14 GND	R13 V1	R12 GND	R11 V0	R10 GND	R9 V0	R8 GND	R7 V0	R6 GND	R5 V0	R4 GND	R3 V0	R2 MASTERSEL	R1 GND	**R**
P24 GND	P23 GND	P22 V1	P21 GND	P20 V1	P19 GND	P18 V1	P17 GND	P16 V1	P15 GND	P14 V1	P13 GND	P12 V0	P11 GND	P10 V0	P9 GND	P8 V0	P7 GND	P6 V0	P5 GND	P4 PSYNC	P3 Z_OUT	P2 GND	P1 Z_SENSE	**P**
N24 SRIN0	N23 V1	N22 PROCID1	N21 PROCID0	N20 GND	N19 V1	N18 GND	N17 V1	N16 GND	N15 V1	N14 GND	N13 V1	N12 GND	N11 V0	N10 GND	N9 V0	N8 GND	N7 V0	N6 GND	N5 V0	N4 GND	N3 GND	N2 ADOUT0	N1 OV$_{DD}$	**N**
M24 SRIN0	M23 OV$_{DD}$	M22 V1	M21 GND	M20 V1	M19 GND	M18 V1	M17 GND	M16 V1	M15 GND	M14 V1	M13 GND	M12 V0	M11 GND	M10 V0	M9 GND	M8 V0	M7 GND	M6 V0	M5 GND	M4 V0	M3 OV$_{DD}$	M2 ADOUT4	M1 SROUT0	**M**
L24 ADIN8	L23 SRIN1	L22 GND	L21 V1	L20 V1	L19 V1	L18 GND	L17 V1	L16 GND	L15 V1	L14 GND	L13 V1	L12 GND	L11 V0	L10 GND	L9 V0	L8 GND	L7 V0	L6 GND	L5 V0	L4 OV$_{DD}$	L3 GND	L2 ADOUT3	L1 SROUT0	**L**
K24 GND	K23 SRIN1	K22 ADIN6	K21 GND	K20 OV$_{DD}$	K19 GND	K18 V1	K17 GND	K16 V1	K15 GND	K14 V1	K13 GND	K12 V0	K11 GND	K10 V0	K9 GND	K8 V0	K7 GND	K6 V0	K5 OV$_{DD}$	K4 GND	K3 ADOUT5	K2 ADOUT6	K1 GND	**K**
J24 ADIN2	J23 ADIN7	J22 GND	J21 OV$_{DD}$	J20 GND	J19 V1	J18 GND	J17 V1	J16 GND	J15 V1	J14 GND	J13 V0	J12 GND	J11 V0	J10 GND	J9 V0	J8 GND	J7 V0	J6 GND	J5 V0	J4 GND	J3 OV$_{DD}$	J2 ADOUT2	J1 ADOUT8	**J**
H24 OV$_{DD}$	H23 ADIN0	H22 ADIN3	H21 GND	H20 V0	H19 GND	H18 V1	H17 GND	H16 V1	H15 GND	H14 V1	H13 GND	H12 V0	H11 GND	H10 V0	H9 GND	H8 V0	H7 GND	H6 V0	H5 OV$_{DD}$	H4 GND	H3 ADOUT7	H2 GND	H1 SROUT1	**H**
G24 ADIN13	G23 ADIN1	G22 V1	G21 V0	G20 OV$_{DD}$	G19 V1	G18 GND	G17 V1	G16 GND	G15 V1	G14 OV$_{DD}$	G13 V0	G12 GND	G11 V0	G10 GND	G9 V0	G8 GND	G7 V0	G6 GND	G5 GND	G4 GND	G3 ADOUT1	G2 GND	G1 SROUT1	**G**
F24 GND	F23 ADIN14	F22 OV$_{DD}$	F21 GND	F20 GND	F19 OV$_{DD}$	F18 V0	F17 V1	F16 OV$_{DD}$	F15 GND	F14 V0	F13 GND	F12 SYSCLK	F11 GND	F10 V0	F9 OV$_{DD}$	F8 V0	F7 GND	F6 OV$_{DD}$	F5 V0	F4 OV$_{DD}$	F3 ADOUT13	F2 OV$_{DD}$	F1 ADOUT10	**F**
E24 ADIN11	E23 GND	E22 ADIN9	E21 ADIN21	E20 ADIN31	E19 ADIN32	E18 GND	E17 ADIN30	E16 ADIN29	E15 ADIN18	E14 OV$_{DD}$	E13 I2CGO	E12 SYSCLK	E11 OV$_{DD}$	E10 ADOUT15	E9 ADOUT18	E8 ADOUT39	E7 ADOUT23	E6 OV$_{DD}$	E5 ADOUT21	E4 ADOUT32	E3 ADOUT9	E2 GND	E1 ADOUT12	**E**
D24 OV$_{DD}$	D23 ADIN10	D22 OV$_{DD}$	D21 ADIN23	D20 GND	D19 OV$_{DD}$	D18 ADIN27	D17 ADIN35	D16 OV$_{DD}$	D15 ADIN15	D14 GND	D13 PLLTEST	D12 BUS_CFG1	D11 GND	D10 ADOUT16	D9 ADOUT41	D8 OV$_{DD}$	D7 ADOUT35	D6 ADOUT36	D5 GND	D4 ADOUT31	D3 ADOUT11	D2 CLKOUT	D1 CLKOUT	**D**
C24 CLKIN	C23 GND	C22 ADIN22	C21 ADIN20	C20 ADIN33	C19 ADIN34	C18 ADIN43	C17 ADIN42	C16 ADIN19	C15 ADIN40	C14 ADIN28	C13 OV$_{DD}$	C12 PLLTESTOUT	C11 OV$_{DD}$	C10 ADOUT28	C9 ADOUT40	C8 GND	C7 ADOUT14	C6 ADOUT27	C5 ADOUT33	C4 OV$_{DD}$	C3 ADOUT25	C2 OV$_{DD}$	C1 ADOUT26	**C**
B24 CLKIN	B23 ADIN25	B22 ADIN24	B21 OV$_{DD}$	B20 ADIN36	B19 OV$_{DD}$	B18 ADIN39	B17 GND	B16 OV$_{DD}$	B15 ADIN16	B14 ADIN4	B13 KELV_GND2	B12 EI_DISABLE	B11 BUS_CFG2	B10 GND	B9 ADOUT17	B8 ADOUT19	B7 OV$_{DD}$	B6 ADOUT43	B5 ADOUT30	B4 ADOUT37	B3 GND	B2 ADOUT20	B1 GND	**B**
A24 OV$_{DD}$	A23 ADIN12	A22 GND	A21 ADIN37	A20 GND	A19 ADIN38	A18 ADIN26	A17 ADIN41	A16 ADIN17	A15 ADIN5	A14 GND	A13 ANALOGGND	A12 AVDD	A11 BUS_CFG0	A10 KELV_OVDD	A9 ADOUT29	A8 ADOUT38	A7 ADOUT42	A6 ADOUT34	A5 GND	A4 ADOUT24	A3 OV$_{DD}$	A2 ADOUT22		**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 970MP Microprocessor with it placed and soldered on the system board

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

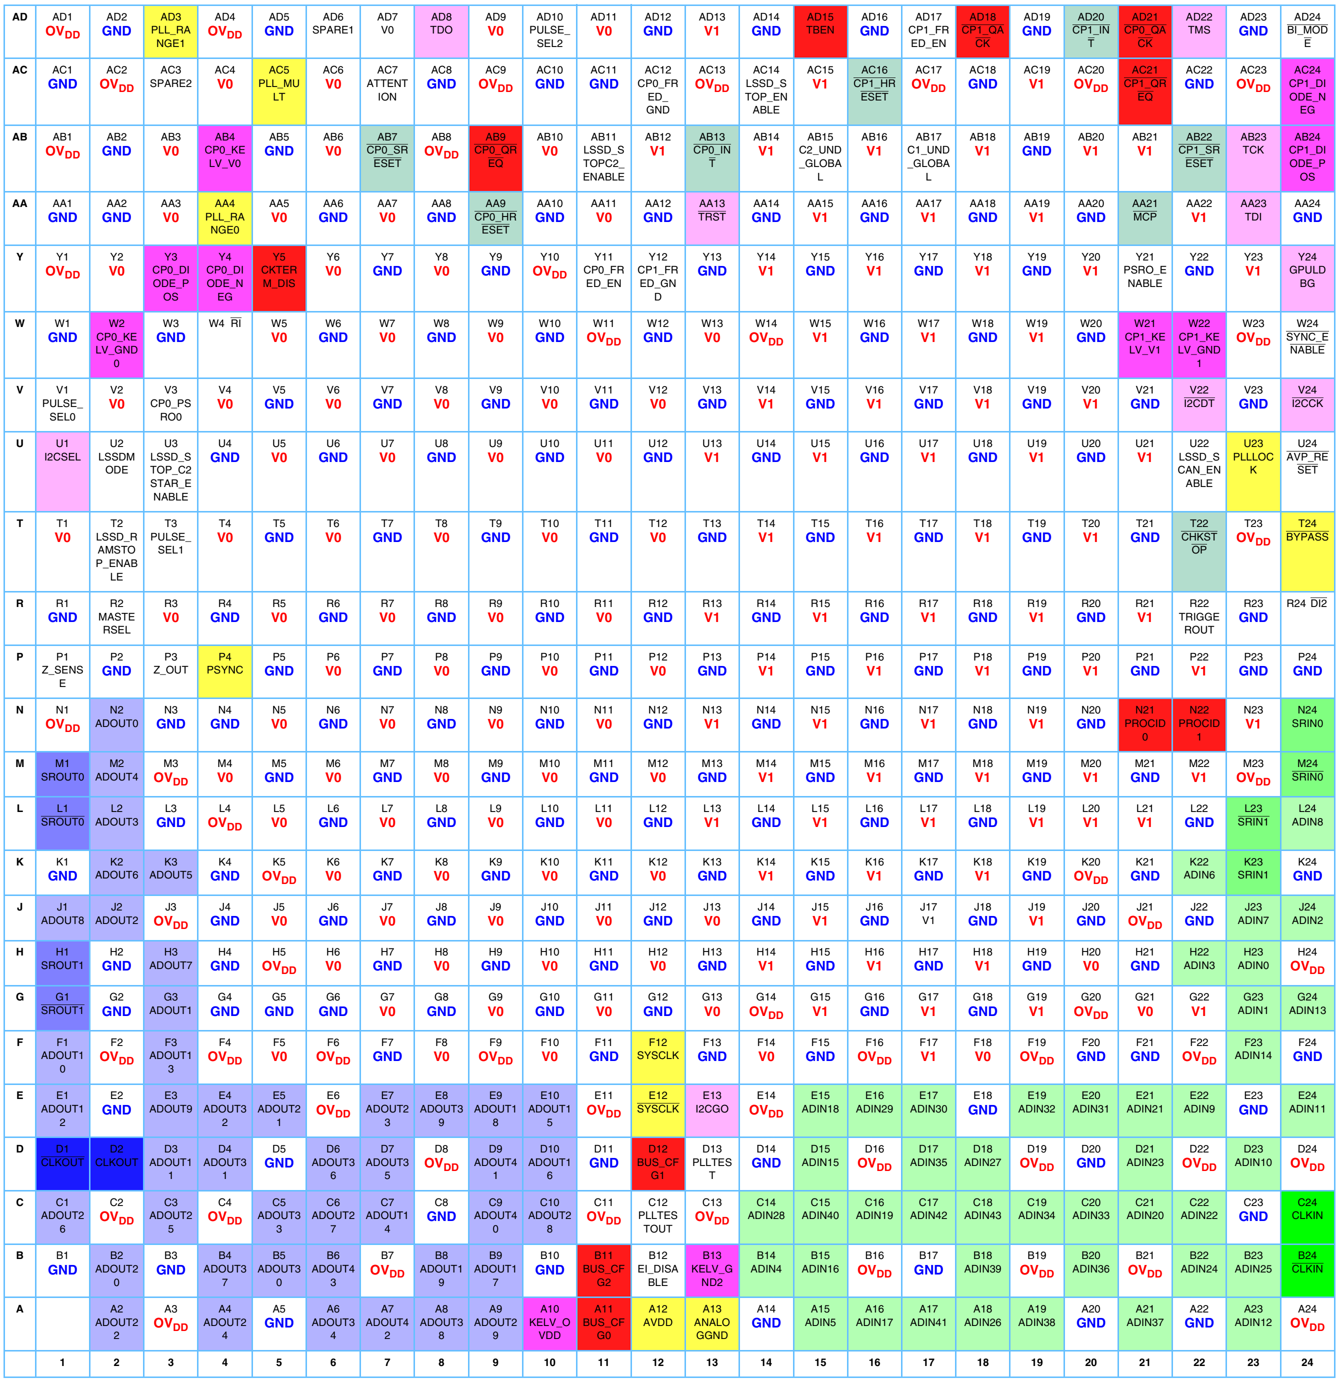

AD	AD1 OV_{DD}	AD2 GND	AD3 PLL_RANGE1	AD4 OV_{DD}	AD5 GND	AD6 SPARE1	AD7 V0	AD8 TDO	AD9 V0	AD10 PULSE_SEL2	AD11 V0	AD12 GND	AD13 V1	AD14 GND	AD15 TBEN	AD16 GND	AD17 CP1_FRED_EN	AD18 CP1_QACK	AD19 GND	AD20 CP1_INT	AD21 CP0_QACK	AD22 TMS	AD23 GND	AD24 BI_MODE
AC	AC1 GND	AC2 OV_{DD}	AC3 SPARE2	AC4 V0	AC5 PLL_MULT	AC6 V0	AC7 ATTENTION	AC8 GND	AC9 OV_{DD}	AC10 GND	AC11 GND	AC12 CP0_FRED_GND	AC13 OV_{DD}	AC14 LSSD_STOP_ENABLE	AC15 V1	AC16 CP1_HRESET	AC17 OV_{DD}	AC18 GND	AC19 V1	AC20 OV_{DD}	AC21 CP1_QREQ	AC22 GND	AC23 OV_{DD}	AC24 CP1_DIODE_NEG
AB	AB1 OV_{DD}	AB2 GND	AB3 V0	AB4 CP0_KELV_V0	AB5 GND	AB6 V0	AB7 CP0_SRESET	AB8 OV_{DD}	AB9 CP0_QREQ	AB10 V0	AB11 LSSD_STOPC2_ENABLE	AB12 V1	AB13 CP0_INT	AB14 V1	AB15 C2_UND_GLOBAL	AB16 V1	AB17 C1_UND_GLOBAL	AB18 V1	AB19 GND	AB20 V1	AB21 V1	AB22 CP1_SRESET	AB23 TCK	AB24 CP1_DIODE_POS
AA	AA1 GND	AA2 GND	AA3 V0	AA4 PLL_RANGE0	AA5 V0	AA6 GND	AA7 V0	AA8 GND	AA9 CP0_HRESET	AA10 GND	AA11 V0	AA12 GND	AA13 TRST	AA14 GND	AA15 V1	AA16 GND	AA17 V1	AA18 GND	AA19 V1	AA20 GND	AA21 MCP	AA22 V1	AA23 TDI	AA24 GND
Y	Y1 OV_{DD}	Y2 V0	Y3 CP0_DIODE_POS	Y4 CP0_DIODE_NEG	Y5 CKTERM_DIS	Y6 V0	Y7 GND	Y8 V0	Y9 GND	Y10 OV_{DD}	Y11 CP0_FRED_EN	Y12 CP1_FRED_GND	Y13 GND	Y14 V1	Y15 GND	Y16 V1	Y17 GND	Y18 V1	Y19 GND	Y20 V1	Y21 PSRO_ENABLE	Y22 GND	Y23 V1	Y24 GPULDBG
W	W1 GND	W2 CP0_KELV_GND0	W3 GND	W4 RI	W5 V0	W6 GND	W7 V0	W8 GND	W9 V0	W10 GND	W11 OV_{DD}	W12 GND	W13 V0	W14 OV_{DD}	W15 V1	W16 GND	W17 V1	W18 GND	W19 V1	W20 GND	W21 CP1_KELV_V1	W22 CP1_KELV_GND1	W23 OV_{DD}	W24 SYNC_ENABLE
V	V1 PULSE_SEL0	V2 V0	V3 CP0_PSRO0	V4 V0	V5 GND	V6 V0	V7 GND	V8 V0	V9 GND	V10 V0	V11 GND	V12 V0	V13 GND	V14 V1	V15 GND	V16 V1	V17 GND	V18 V1	V19 GND	V20 V1	V21 GND	V22 I2CDT	V23 GND	V24 I2CCK
U	U1 I2CSEL	U2 LSSDMODE	U3 LSSD_STOP_C2STAR_ENABLE	U4 GND	U5 V0	U6 GND	U7 V0	U8 GND	U9 V0	U10 GND	U11 V0	U12 GND	U13 V1	U14 GND	U15 V1	U16 GND	U17 V1	U18 GND	U19 V1	U20 GND	U21 V1	U22 LSSD_SCAN_ENABLE	U23 PLLLOCK	U24 AVP_RESET
T	T1 V0	T2 LSSD_RAMSTOP_ENABLE	T3 PULSE_SEL1	T4 V0	T5 GND	T6 V0	T7 GND	T8 V0	T9 GND	T10 V0	T11 GND	T12 V0	T13 GND	T14 V1	T15 GND	T16 V1	T17 GND	T18 V1	T19 GND	T20 V1	T21 GND	T22 CHKSTOP	T23 OV_{DD}	T24 BYPASS
R	R1 GND	R2 MASTERSEL	R3 V0	R4 GND	R5 V0	R6 GND	R7 V0	R8 GND	R9 V0	R10 GND	R11 V0	R12 GND	R13 V1	R14 GND	R15 V1	R16 GND	R17 V1	R18 GND	R19 V1	R20 GND	R21 V1	R22 TRIGGEROUT	R23 GND	R24 DI2
P	P1 Z_SENSE	P2 GND	P3 Z_OUT	P4 PSYNC	P5 GND	P6 V0	P7 GND	P8 V0	P9 GND	P10 V0	P11 GND	P12 V0	P13 GND	P14 V1	P15 GND	P16 V1	P17 GND	P18 V1	P19 GND	P20 V1	P21 GND	P22 V1	P23 GND	P24 GND
N	N1 OV_{DD}	N2 ADOUT0	N3 GND	N4 GND	N5 V0	N6 GND	N7 V0	N8 GND	N9 V0	N10 GND	N11 V0	N12 GND	N13 V1	N14 GND	N15 V1	N16 GND	N17 V1	N18 GND	N19 V1	N20 GND	N21 PROCID0	N22 PROCID1	N23 V1	N24 SRIN0
M	M1 SROUT0	M2 ADOUT4	M3 OV_{DD}	M4 V0	M5 GND	M6 V0	M7 GND	M8 V0	M9 GND	M10 V0	M11 GND	M12 V0	M13 GND	M14 V1	M15 GND	M16 V1	M17 GND	M18 V1	M19 GND	M20 V1	M21 GND	M22 V1	M23 OV_{DD}	M24 SRIN0
L	L1 SROUT0	L2 ADOUT3	L3 GND	L4 OV_{DD}	L5 V0	L6 GND	L7 V0	L8 GND	L9 V0	L10 GND	L11 V0	L12 GND	L13 V1	L14 GND	L15 V1	L16 GND	L17 V1	L18 GND	L19 V1	L20 V1	L21 V1	L22 GND	L23 SRIN1	L24 ADIN8
K	K1 GND	K2 ADOUT6	K3 ADOUT5	K4 GND	K5 OV_{DD}	K6 V0	K7 GND	K8 V0	K9 GND	K10 V0	K11 GND	K12 V0	K13 GND	K14 V1	K15 GND	K16 V1	K17 GND	K18 V1	K19 GND	K20 OV_{DD}	K21 GND	K22 ADIN6	K23 SRIN1	K24 GND
J	J1 ADOUT8	J2 ADOUT2	J3 OV_{DD}	J4 GND	J5 V0	J6 GND	J7 V0	J8 GND	J9 V0	J10 GND	J11 V0	J12 GND	J13 V0	J14 GND	J15 V1	J16 GND	J17 V1	J18 GND	J19 V1	J20 GND	J21 OV_{DD}	J22 GND	J23 ADIN7	J24 ADIN2
H	H1 SROUT1	H2 GND	H3 ADOUT7	H4 GND	H5 OV_{DD}	H6 V0	H7 GND	H8 V0	H9 GND	H10 V0	H11 GND	H12 V0	H13 GND	H14 V1	H15 GND	H16 V1	H17 GND	H18 V1	H19 GND	H20 V0	H21 GND	H22 ADIN3	H23 ADIN0	H24 OV_{DD}
G	G1 SROUT1	G2 GND	G3 ADOUT1	G4 GND	G5 GND	G6 GND	G7 V0	G8 GND	G9 V0	G10 GND	G11 V0	G12 GND	G13 V0	G14 OV_{DD}	G15 V1	G16 GND	G17 V1	G18 GND	G19 V1	G20 OV_{DD}	G21 V0	G22 V1	G23 ADIN1	G24 ADIN13
F	F1 ADOUT10	F2 OV_{DD}	F3 ADOUT13	F4 OV_{DD}	F5 V0	F6 OV_{DD}	F7 GND	F8 V0	F9 OV_{DD}	F10 V0	F11 GND	F12 SYSCLK	F13 GND	F14 V0	F15 GND	F16 OV_{DD}	F17 V1	F18 V0	F19 OV_{DD}	F20 GND	F21 GND	F22 OV_{DD}	F23 ADIN14	F24 GND
E	E1 ADOUT12	E2 GND	E3 ADOUT9	E4 ADOUT32	E5 ADOUT21	E6 OV_{DD}	E7 ADOUT23	E8 ADOUT39	E9 ADOUT18	E10 ADOUT15	E11 OV_{DD}	E12 SYSCLK	E13 I2CGO	E14 OV_{DD}	E15 ADIN18	E16 ADIN29	E17 ADIN30	E18 GND	E19 ADIN32	E20 ADIN31	E21 ADIN21	E22 ADIN9	E23 GND	E24 ADIN11
D	D1 CLKOUT	D2 CLKOUT	D3 ADOUT11	D4 ADOUT31	D5 GND	D6 ADOUT36	D7 ADOUT35	D8 OV_{DD}	D9 ADOUT41	D10 ADOUT16	D11 GND	D12 BUS_CFG1	D13 PLLTEST	D14 GND	D15 ADIN15	D16 OV_{DD}	D17 ADIN35	D18 ADIN27	D19 OV_{DD}	D20 GND	D21 ADIN23	D22 OV_{DD}	D23 ADIN10	D24 OV_{DD}
C	C1 ADOUT26	C2 OV_{DD}	C3 ADOUT25	C4 OV_{DD}	C5 ADOUT33	C6 ADOUT27	C7 ADOUT14	C8 GND	C9 ADOUT40	C10 ADOUT28	C11 OV_{DD}	C12 PLLTESTOUT	C13 OV_{DD}	C14 ADIN28	C15 ADIN40	C16 ADIN19	C17 ADIN42	C18 ADIN43	C19 ADIN34	C20 ADIN33	C21 ADIN20	C22 ADIN22	C23 GND	C24 CLKIN
B	B1 GND	B2 ADOUT20	B3 GND	B4 ADOUT37	B5 ADOUT30	B6 ADOUT43	B7 OV_{DD}	B8 ADOUT19	B9 ADOUT17	B10 GND	B11 BUS_CFG2	B12 EI_DISABLE	B13 KELV_GND2	B14 ADIN4	B15 ADIN16	B16 OV_{DD}	B17 GND	B18 ADIN39	B19 OV_{DD}	B20 ADIN36	B21 OV_{DD}	B22 ADIN24	B23 ADIN25	B24 CLKIN
A		A2 ADOUT22	A3 OV_{DD}	A4 ADOUT24	A5 GND	A6 ADOUT34	A7 ADOUT42	A8 ADOUT38	A9 ADOUT29	A10 KELV_OVDD	A11 BUS_CFG0	A12 AVDD	A13 ANALOGGND	A14 GND	A15 ADIN5	A16 ADIN17	A17 ADIN41	A18 ADIN26	A19 ADIN38	A20 GND	A21 ADIN37	A22 GND	A23 ADIN12	A24 OV_{DD}
	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)

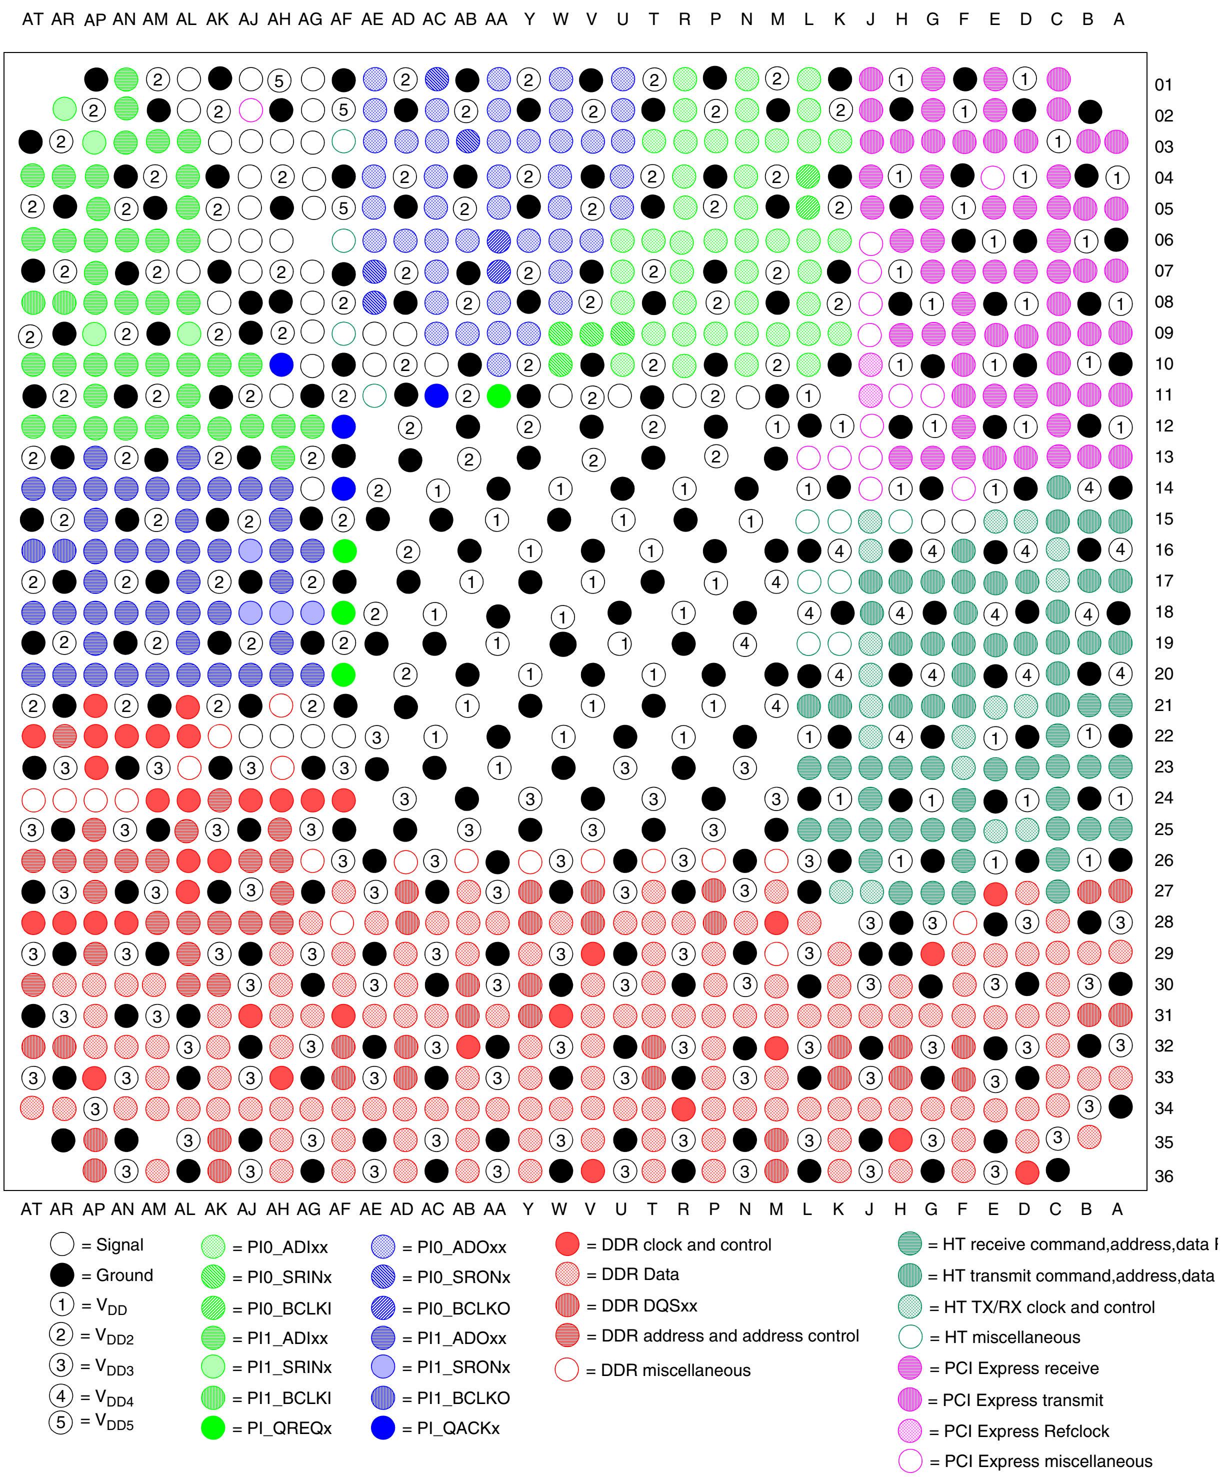

3. Thermal Design Guidelines

The PowerPC 970MP 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. For each core, the PowerPC 970MP processor contains an internal thermal diode that provides the only accurate way to determine the temperature of the die.

The PowerPC 970MP 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 970MP processor is approximately 153.8 mm^2 (0.238 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, 13.225 mm × 11.629 mm, size of the PowerPC 970MP die requires taking special care in the selection of crucial parts of the cooling system such as the heat sink, thermal interface material, and 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 970MP 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 970MP processor chosen, the power dissipation can range from 28 - 125 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; but given the heat flux, especially of the higher wattage parts, the die-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 970MP RISC Microprocessor Datasheet* indicates that there are decoupling caps located on the top of the substrate near the die. Therefore, when using electrically conductive greases, care must be taken to avoid shorting the caps.

Note: To assure the feasibility of the thermal solution design and the lowest assembly cost, it is important to involve manufacturing engineering early, 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 970MP 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 970MP 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 970MP 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 970MP processor contains two thermal diodes that, when used in conjunction with external circuitry, can monitor the temperature of the processor die. Most importantly, these diodes are 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 *PowerPC 970MP 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 *PowerPC 970MP 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 970MP processor. Several application notes describe the thermal diode, how to drive it, and how to interpret the information (see *Related Documents* on page 12 for more information).

3.4.1 Thermal Diode Monitoring Circuit

The PPC970MP/CPC945 evaluation board contains circuitry that is used to drive and read the thermal diode of each core of each PowerPC 970MP processor. The thermal diode is properly biased in accordance with the *IBM PowerPC 970MP RISC Microprocessor Datasheet*. A Cypress Programmable System-On-Chip (PSoC) microcontroller is used to read the voltage of each processor 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 970MP 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 970MP fuse ring.

4. Processor Interface Routing Guidelines

The PowerPC 970MP processor interface uses high-speed source-synchronous buses (SSBs) to transfer data between the PowerPC 970MP 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

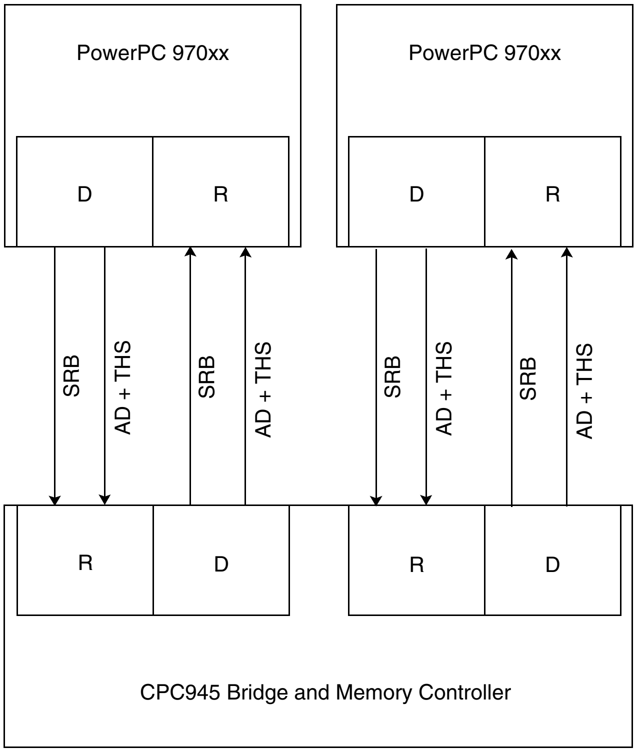

AD:	Address/data channel	SRB:	Snoop response bus
D:	Drive side	TH:	Transfer handshake bus
R:	Receive side		

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 detailed in the *IBM PowerPC 970MP 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 970MP 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 970MP RISC Microprocessor User's Manual* and the *IBM PowerPC 970MP RISC Microprocessor Datasheet*.

4.1 PowerPC 970MP Processor Interface Design Recommendations

4.1.1 General Guidelines

The PPC970MP/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 970MP 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 printed circuit board (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 inputs to processor A to the length of inputs to processor B and the length of 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 com-

panion 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 to not 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 on 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 970MP 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.

4.1.2 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 dual or quad processor configuration based whether one or two processor cards are plugged into a system board. If you choose this configuration, use a high-quality connector and designed for high-speed signal use. The bill of materials (BOM) available with the PPC970MP/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, dual core version of system B.

Figure 4-2. System Implementations

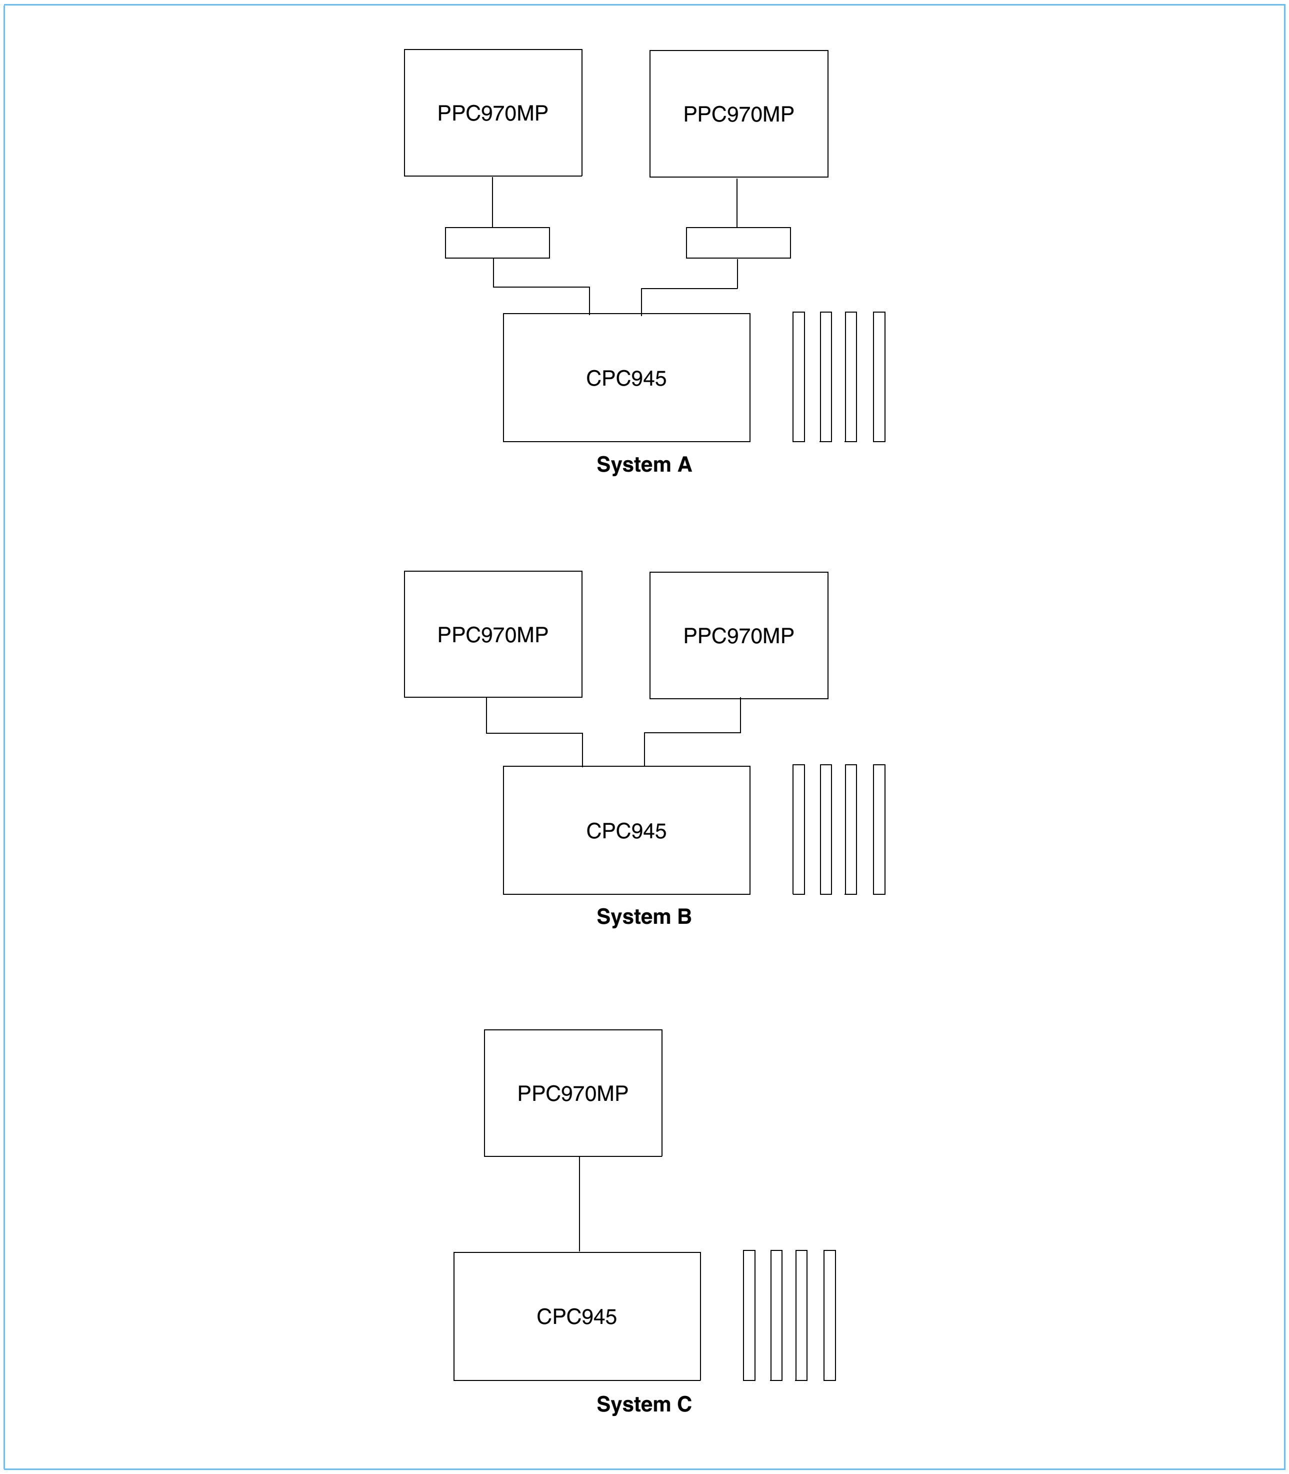

4.1.3 System A Layout Example

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.

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

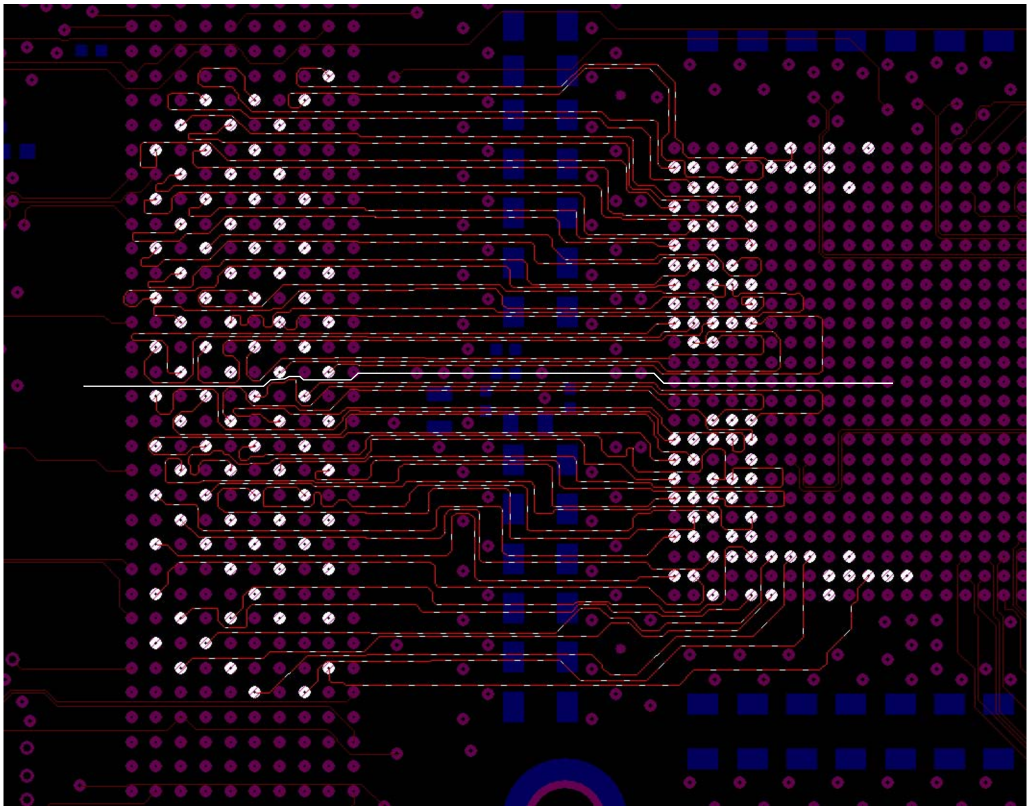

Figure 4-4. Processor Interface Bus In and Out, Layer S05

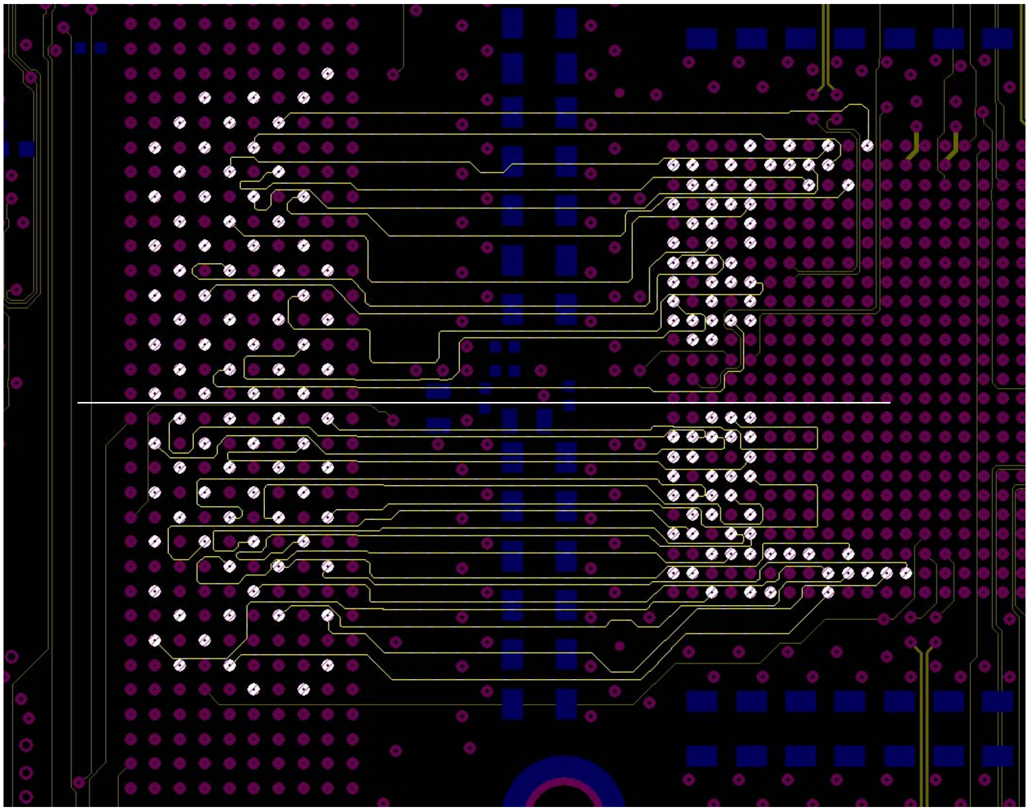

Figure 4-5. Processor Interface Bus In and Out, Layer S12

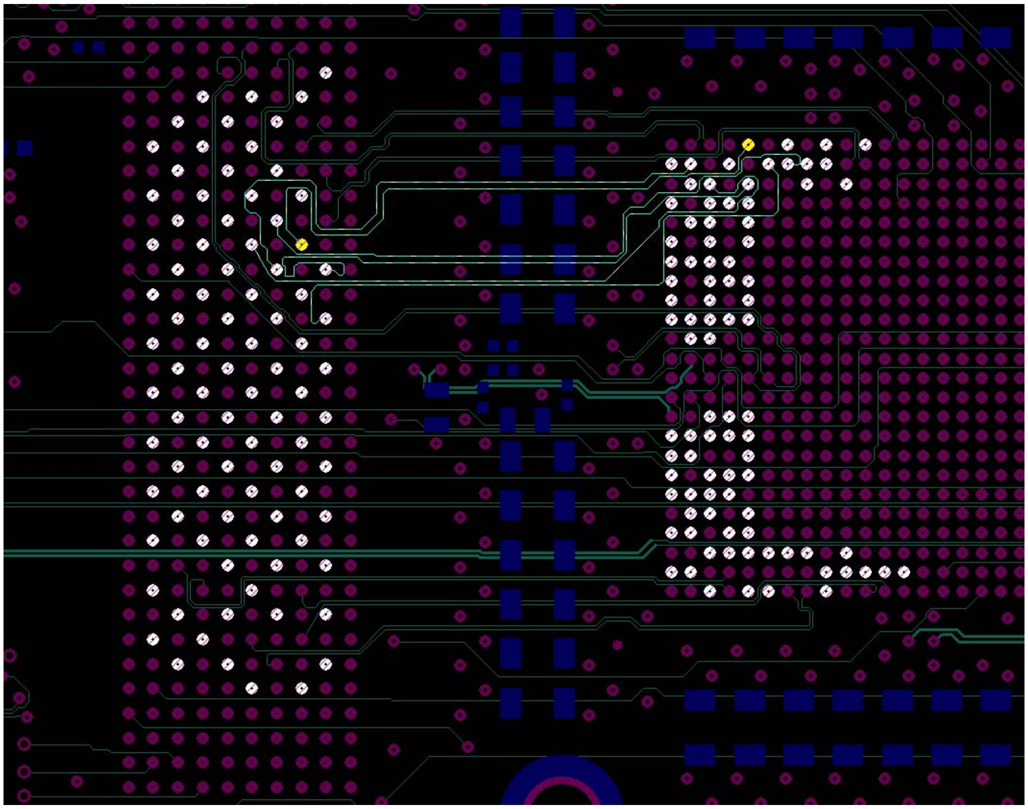

Figure 4-6 through *Figure 4-8* show that two complete processor interface buses can be routed on only three layers. These traces are approximately 11.2 cm (4.41 in.) long and are matched to 0.5 cm (200 mils)

Figure 4-6. Processor Interface Bus A and B In and Out S3

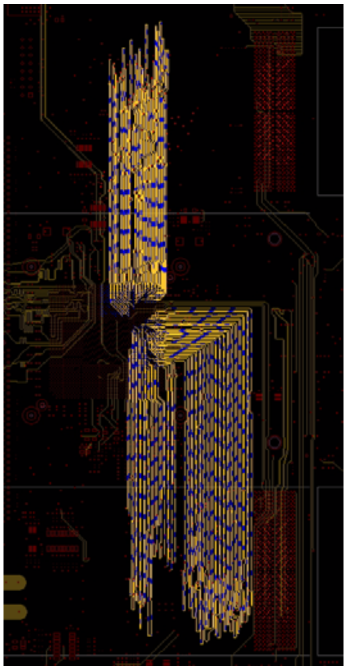

Figure 4-7. Processor Interface Bus A and B In and Out S5

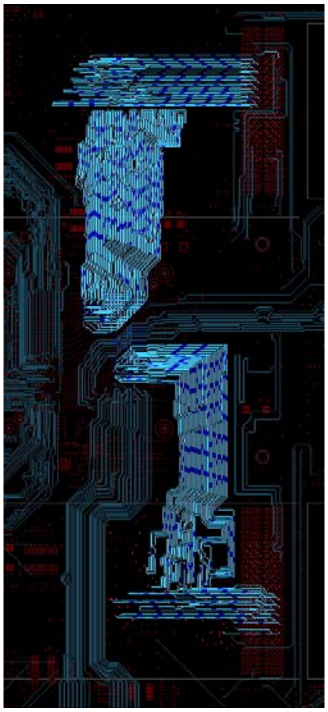

Figure 4-8. Processor Interface Bus A and B In and Out S7

5. Clocking

As described in the *IBM PowerPC 970MP RISC Microprocessor Datasheet*, the primary input clock for the PowerPC 970MP processor is provided to the processor as 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. *IBM PowerPC 970MP RISC Microprocessor User's Manual* provides the functional description of the PowerPC 970MP 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 *Core-Clock Timing Relationship Between PSYNC and SYSCLK* figure in the *IBM PowerPC 970MP 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 970MP processor; the APSYNC signal is provided as an input to the CPC945 bridge.) This signal is 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 using other known techniques, but implementation of spread spectrum has the largest impact. The *IBM PowerPC 970MP 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 970MP processor.

5.4 Frequency Slewing

Power dissipation has emerged as one of the most important topics in PowerPC 970MP design. The PowerPC 970MP processor allows SYSCLK to be gradually reduced to lower power dissipation. The processor has a number of power reduction capabilities built in, but frequency slewing provides the finest resolution in adjusting the power level. To properly use frequency slewing, all affected parts must support slewing, and the range and rate must stay within the guidelines as detailed in the *IBM PowerPC 970MP RISC Microprocessor Datasheet*. The clock generator must be able to provide a glitch-free, dynamically variable SYSCLK frequency.

5.4.1 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 carefully in the *IBM PowerPC 970MP 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 970MP processor has on-chip termination for the differential SYSCLK, which can be enabled or disable using the CKTERM_DIS signal. The termination network shown is 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 PPC970MP/CPC945 evaluation board.

Figure 5-1. SYSCLK Circuit Example

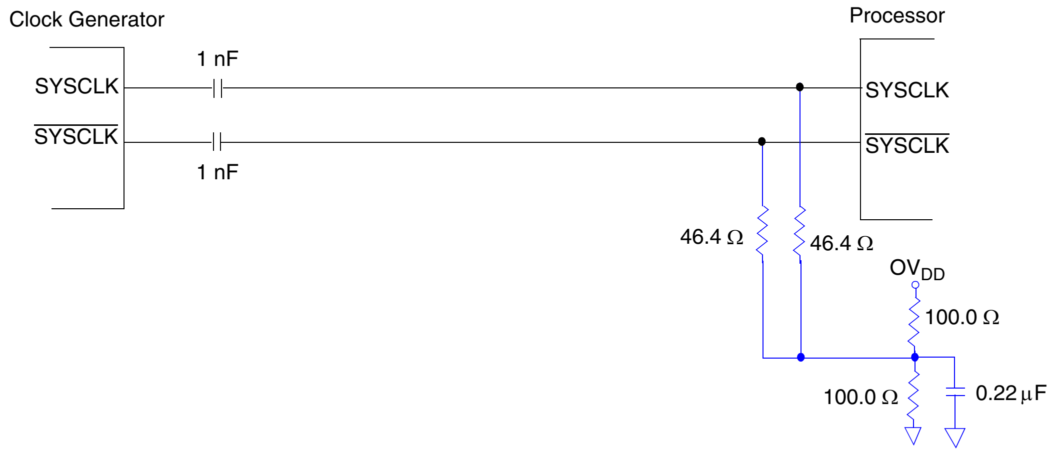

5.6 Layout Considerations

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 20 mils (0.5 mm), 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 the *Section 4.1.2 Example System Implementations* on page 35. The differential SYSCLK traces are paired with 0.102 mm (4 mils) separation, and 0.635 mm (25 mils) clearance to the nearest nonclock trace. They are length matched to within ±0.127 mm (5 mils) on the processor card.

Figure 5-2. Processor Card SYSCLK Pairs

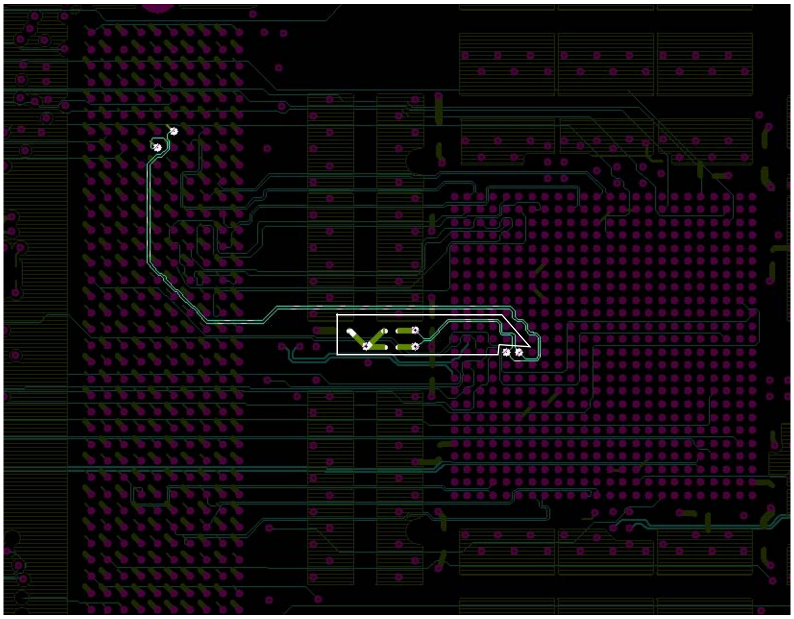

This card design was done as an evaluation and characterization test bed for the processor. The external termination shown in blue in *Figure 5-1 SYSCLK Circuit Example* on page 44 is part of the layout. The extra trace and component pads for the external termination are indicated by the white polygon. Using the IBIS and HSPICE models provided for the PowerPC 970MP processor, 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* to show the possible cleaner routing of the SYSCLK differential pair if external termination can be avoided.

Figure 5-3. Removing External Termination

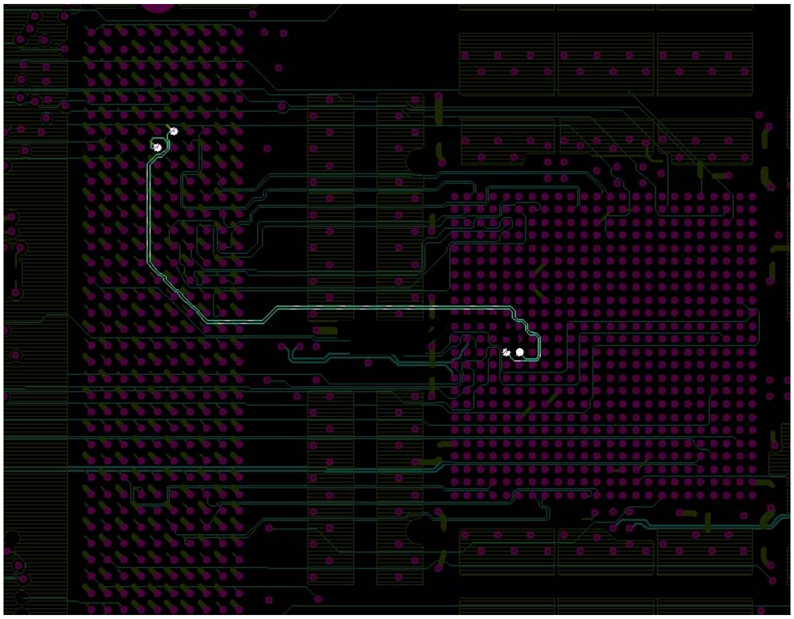

Figure 5-4 shows the traces for the example system referred to as System A in *Section 4.1.2 Example System Implementations* on page 35. The differential SYSCLK traces are paired with 0.102 mm (4 mils) separation, and 1.02 mm (40 mils) clearance to the nearest nonclock trace. They are length matched 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*.

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

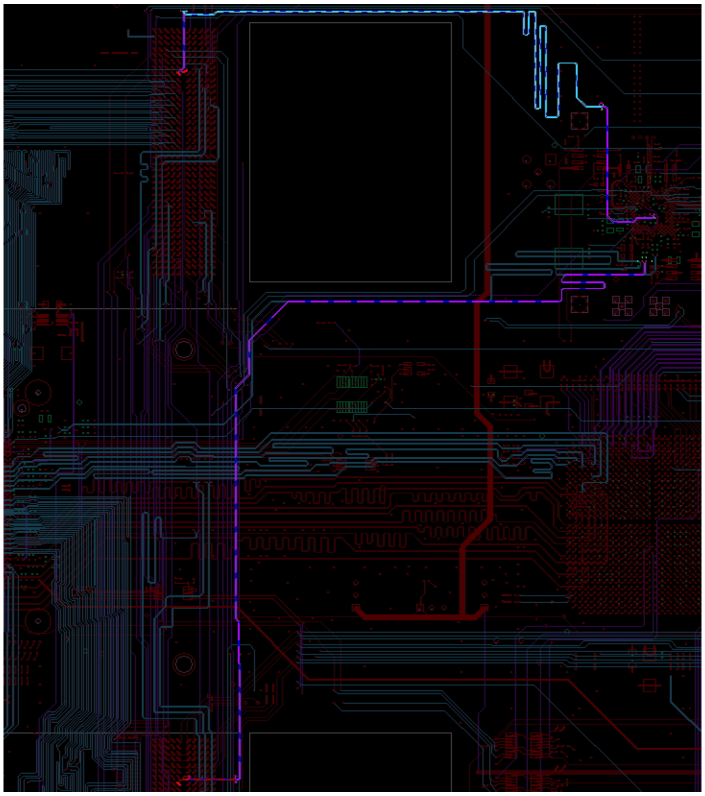

6. CPC945 Memory Interface and Routing Guidelines

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 supported is up to 8.5 GBps at 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 (ECC) or 64-bit data bus (without ECC) (See the memory section of the *CPC945 Bridge and Memory Controller User Manual* for the full details on the bus width configurations.)
- JEDEC DDR2 Synchronous Dynamic Random Access Memory (SDRAM) Specification (JESD79-2) 256 Mb, 512 Mb, 1 Gb, and 2 Gb chip sizes and x4, x8 and x16 organizations supported. The x8 and x16 chips can be mixed. If x4 chips are used, all the chips must be x4.
- Support for 8 ranks. 8 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 supports single-symbol (4-bit) correction and double-symbol error detection; chip kill correction with x4 chips.

Notes:

- 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 Manual*.
- Always review the *IBM PowerPC CPC945 RISC Microprocessor Errata List for DD 1.2*, 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-X), 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 BA0-2, RAS, CAS, WE, CKE, CS are buffered. CK comes from two differential pairs on the CPC945 bridge. Each pair was 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 1 or 2 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 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, 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 that the address and command lines are referenced 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 your 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 a 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 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, $\overline{\text{CK}}$ 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, $\overline{\text{CK}}$ 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 4 DIMM, 8-bank system and a 2 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 4 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 PPC970MP/CPC945 Evaluation Board Routing Rules for the DDR2 Interface

The rest of the section contains the routing rules used on the PPC970MP/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, $\overline{\text{CK}}$ 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 970MP 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 970MP initialization and test sequences to ensure its correct operation. The SPU role does not have to be limited to initializing the PowerPC 970MP 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 970MP 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 970MP 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 970MP I^2C and JTAG interfaces share common logic, so designers must plan accordingly. Use of the I^2C or JTAG bus is mutually exclusive and is controlled by the I2CSEL pin. If this pin is high, the I^2C bus can be used. If the pin is low, the JTAG bus can be used. Traffic on the nonselected bus (I^2C or JTAG depending on the I2CSEL pin) is ignored and should not have any side effects.

If concurrent use of both interfaces is required, the I2CSEL pin can be switched while the system is running. For correct operation, switch the I2CSEL pin only while no traffic is active on either interface to prevent misrecognition of a partial transmission. To ease this operation in debug mode (GPULDBG = 1), the I2CGO pin can be monitored or directly connected to the I2CSEL pin. The I2CGO pin switches from 0 to 1 whenever it is safe to switch I2CSEL from 0 to 1 for I^2C usage. Similarly, it switches from 1 to 0 whenever it is safe to switch I2CSEL from 1 to 0 for JTAG usage.

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 970MP RISC Microprocessor User's Manual* and the *PowerPC 970MP 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 970MP 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 as 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 970MP I^2C bus pins together with non-PowerPC 970MP parts in a system. Each PowerPC 970MP processor should have its own level shifter.

As noted in the *IBM PowerPC 970MP 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 970MP RISC Microprocessor Errata List for DD1.0x, DD1.1* for additional information about nonstandard I^2C interface operation at the beginning of the power-on reset process.

7.2.3 JTAG Interface Operation

The *IBM PowerPC 970MP RISC Microprocessor Datasheet* also contains information about the interface ac timing, including a list of the PowerPC 970MP 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.

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 volts to interface to the memory and the service processor.

The C*PC945 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 RISC Microprocessor Errata List for DD 1.2* 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 the IBM Customer Connect (see *Related Documents* on page 12).

7.4 Overview of the PPC970MP/CPC945 Evaluation Board Service Processor Implementation

The PPC970MP/CPC945 evaluation board uses the AMCC PowerPC 405GPr (405GPr) microcontroller as its service processing unit. 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 970MP processor do not need an SPU with this level of processing power or peripheral mix. However, the PowerPC 970MP 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 PPC970MP/CPC945 evaluation board.)

The IBM PPC970MP/CPC945 evaluation kit includes software for both the PowerPC 970MP 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 970MP 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 PPC970MP/CPC945 evaluation board uses the I^2C bus to initialize and configure the PowerPC 970MP 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 PPC970MP/CPC945 evaluation board. The PPC970MP/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 PPC970MP/CPC945 evaluation board contains circuitry that is used to drive and read the two thermal diodes of each PowerPC 970MP 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. Add this information to the total power budget of the end product. In the case of the PPC970MP/CPC945 evaluation board implementation, the power dissipation is minimal, but other sensing solutions might not be so low powered. In addition, be aware that the heat generated by other components on the board can affect the operation of the sensing circuit.

In general, the SPU does not require special thermal considerations unless it needs its own cooling solution. Additional thermal considerations depend upon the layout of the PCB, enclosure, and airflow specifications

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 970MP 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 970MP processors, but it must be able to reliably communicate with the PowerPC 970MP 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 970MP processor and the target system.

8.1 Overview

Depending on the system architecture and performance requirements, the power requirements can vary between 28 W, for a simple 1.2 GHz PowerPC 970MP system, to over 250 W, for a 2.5 GHz dual processor (quad core) 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 970MP processor. As an implementation example, this section uses the PPC970MP/CPC945 evaluation board.

8.2 PowerPC 970MP Processor Voltage Delivery

The PowerPC 970MP 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 the specified frequencies are detailed in the *IBM PowerPC 970MP 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 970MP Processor

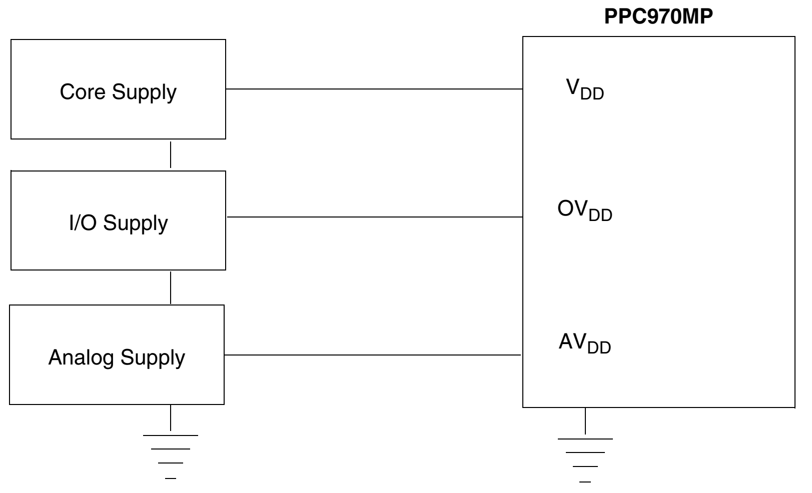

8.2.1 Board Layout Considerations

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 28 to 125 A, careful consideration must be paid to power supply topology and layout. As an example, a 2.5 GHz part operating with a V_{DD} of 1.2 V and a maximum power of 115 W, draws 95 A.

Ninety-five 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 PPC970MP/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 described in the *IBM PowerPC 970MP RISC Microprocessor Datasheet*.

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 OV_{DD}, 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 PPC970MP/CPC945 Evaluation Board Example

The PPC970MP/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 discussed in the following paragraphs.

8.3.1 Power Delivery to the PowerPC 970MP 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 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. The PowerPC 970MP processor features three pairs of Kelvin voltage and ground pins to assist in analyzing on-chip noise and voltage compression. Pins CP0_KELV_V0 and CP0_KELV_GND0 are the V0-GND pair. Pins CP1_KELV_V1 and CP1_KELV_GND1 are the V1-GND pair. Pins KELV_OVDD and KELV_GND2 are the OV_{DD}-GND pair. Never tie these pins into 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 *Using the Kelvin Voltage Sense Pins in the PPC970MP 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 PPC970MP/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 970MP RISC Microprocessor Datasheet*.

Figure 8-2. PLL Power Supply Filter Circuit

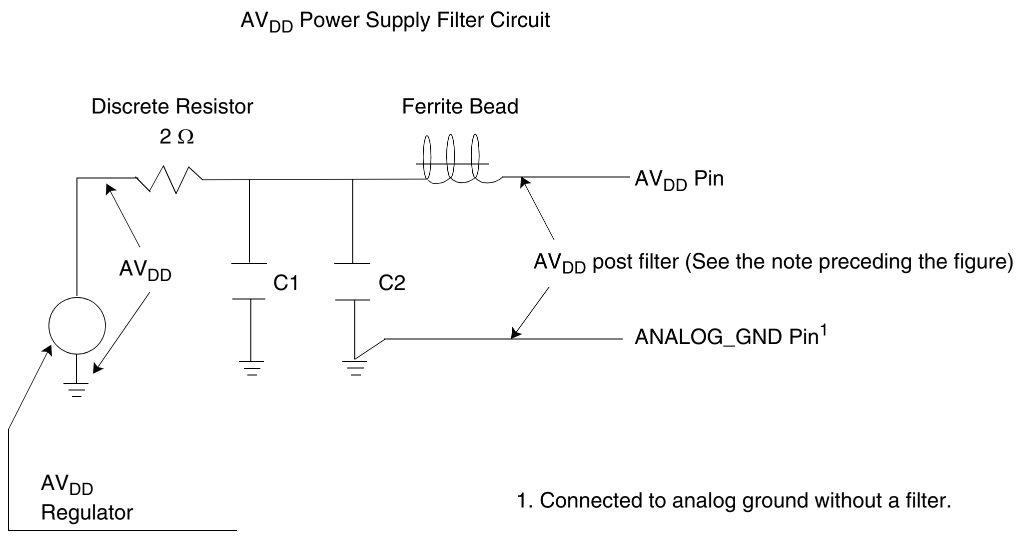

8.3.2 Using VRM Modules

Standard voltage regulation modules (VRMs) can be used for the power supply modules. Although, the PowerPC 970MP 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 970MP processor does not have voltage identification (VID) bits output from the processor, but the service processor can read the fuse ring to obtain the 3-bit VFC data and translate the VFC data 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 970MP processor because of the assortment of interfaces supported (see *Figure 8-3 PowerPC 970MP 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 970MP processors, HyperTransport, and PCI Express (PCI-E). The CPC945 bridge requires core, PLL, and 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 that are not be 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 users 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 970MP 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 970MP 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 systems architecture it might not be practical to share supplies.

In the PPC970MP/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 970MP Block Diagram

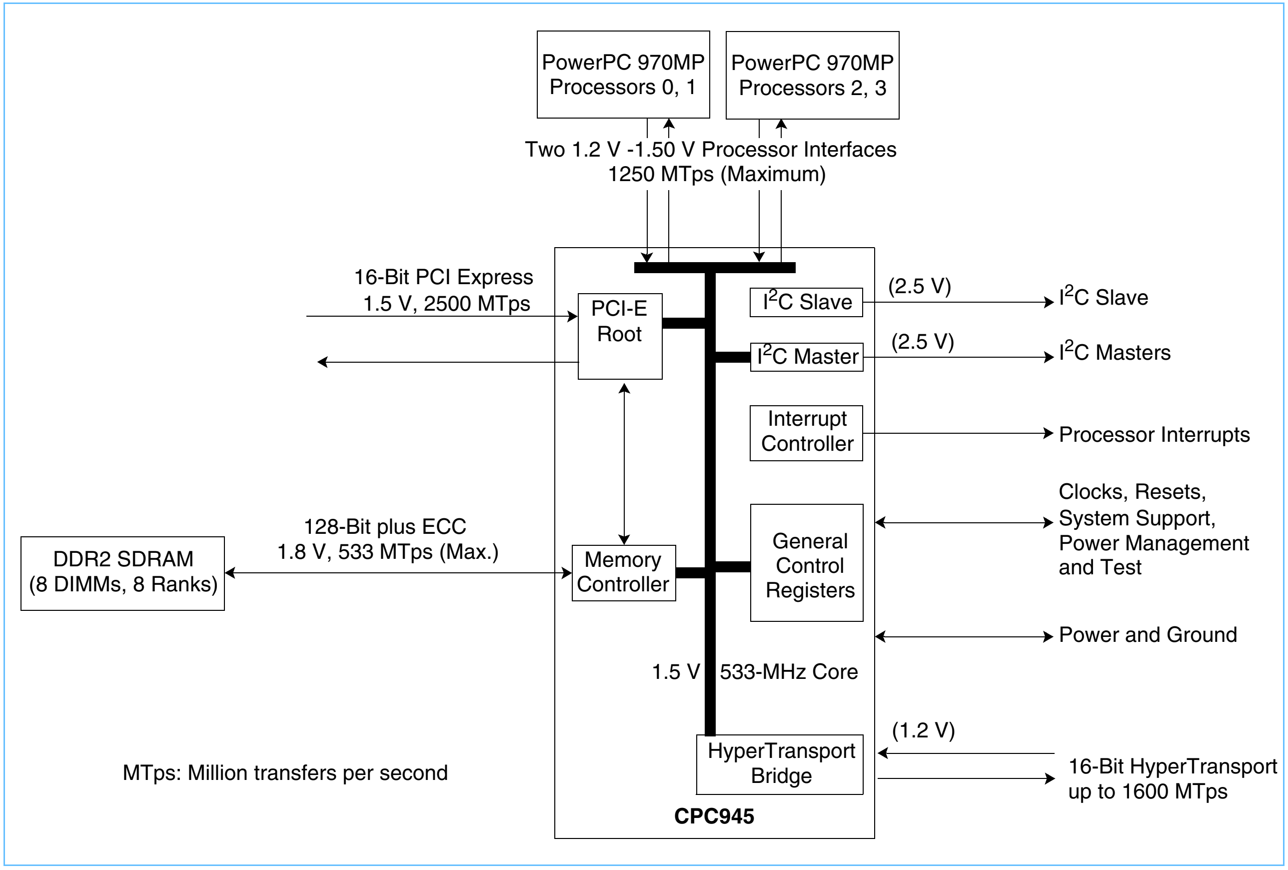

8.3.4 Power Delivery to the HyperTransport Tunnel

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 970MP/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 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, of 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 the 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 970MP/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

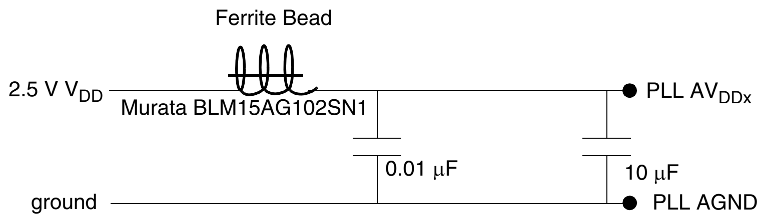

Figure 8-5. Analog V_{DD} Filtering for the Processor Interface and DDR2 Interface Phase-Locked Loops

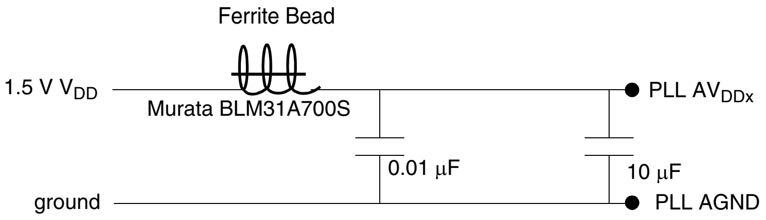

8.3.8 Processor Voltage Sequencing

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 sequence until the V_{DD}, OV_{DD} and AV_{DD} voltages are within their specified tolerances. The following considerations apply, according to the *IBM PowerPC 970MP RISC Microprocessor Datasheet*:

- The power supply ramping order does not matter as long as the supplies reach their final destination in 2000 ms.
- V0 cannot exceed OV_{DD} by more than 0.8 V, except for 2000 ms during power up or down. (During those intervals, it is allowed to be approximately 1.55 V.)
- OV_{DD} cannot exceed V0 by more than 0.8 V except for 2000 ms during power up or down. (During those intervals, it is allowed to be approximately 1.5 V.)
- AV_{DD} cannot exceed V0 by more than 2.5 V except for 2000 ms during power up or down. (During those intervals, it is allowed to be approximately 2.75 V.)
- The V1 circuitry is independent of V0, AV_{DD}, and OV_{DD}.

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.