Infinity Series H.100 MC3 Lite Multi-chassis Interconnect Board

TECHNICAL MANUAL

Documentation Revision 0.0: February 9, 2001

257M004

This page left intentionally blank

Contents

Infinity Series H.100 MC3 Lite Multi-Chassis Interconnect Board Technical Manual

American Tel-A-System, Inc.
800-356-9148
• 4800 Curtin Drive • McFarland, WI 53558 •
• 4145 North Service Road, Suite 200 • Burlington, Ontario L7L 6A3 •
• 257M004 •

FCC Part 15 Requirements

WARNING: This equipment generates, uses, and can radiate radio frequency energy and if not installed and used in accordance with the instruction manual, may cause interference to radio communications. Operation of this equipment in a residential area is likely to cause interference in which case the user at his own expense will be required to take whatever measures may be required to correct the interference.

FCC Part 68 Registration

This equipment is registered with the FCC under Part 68 as a component device for use with any generic PC Type computer or compatible. In order for FCC registration of this product to be retained, all other products used in conjunction with this product to provide your telephony function must also be FCC Part 68 registered for use with these hosts. If any of these components are not registered, then you are required to seek FCC Part 68 registration of the assembled equipment prior to connection to the telephone network. Part 68 registration specifies that you are required to maintain the approval and as such become responsible for the following:

- any component device added to your equipment, whether it bears component registration or not, will require that a Part 68 compliance evaluation is done and possibly that you have testing performed and make a modification filing to the FCC before that new component can be used;

- any modification/update made by a manufacturer to any component device within your equipment, will require that a Part 68 compliance evaluation is done and possibly that you have testing performed and make a modification filing to the FCC before the new component can be used;

- if you continue to assemble additional quantities of this compound equipment, you are required to comply with the FCC's Continuing Compliance requirements.

The telephone company has the right to request the registration information.

The telephone company has the right to temporarily discontinue service. They are required to provide notification and advise of the right to file a complaint.

In case of trouble, you may be required to disconnect the board from the telephone lines until the problem is resolved.

The authorized repair center is:

American Tel-A-System, Inc.
800-356-9148
4800 Curtin Drive
McFarland, WI 53558

There are no user serviceable components on the board. All repairs should be accomplished by returning the board to Amtelco with a description of the problem.

WARNING: This device contains Electrostatic Sensitive Devices. Proper care should be taken when handling this device to avoid damage from static discharges.

Canadian Customers
CP-01, Issue 8, Part 1
Section 14.1

Notice: "The industry Canada label identifies certified equipment. This certification means that the equipment meets certain telecommunications network protective, operational and safety requirements as prescribed in the appropriate Terminal Equipment Technical Requirements document(s). The Department does not guarantee the equipment will operate to the user's satisfaction.

Before installing this equipment, users should ensure that it is permissible to be connected to the facilities of the local telecommunications company. The equipment must also be installed using an acceptable method of connection. The customer should be aware that compliance with the above conditions may not prevent degradation of service in some situations.

Repairs of certified equipment should be coordinated by a representative designated by the supplier. Any repairs or alterations made by the user to this equipment, or equipment malfunctions, may give the telecommunications company cause to request the user to disconnect the equipment.

Users should ensure for their own protection that the electrical ground connections of the power utility, telephone lines and internal metallic water pipe system, ir present, are connected together. This precaution may be particularly important in rural areas.

CAUTION: Users should not attempt to make such connections themselves, but should contact the appropriate electric inspection authority, or electrician, as appropriate.

European Approvals

CE Approval

CE

EN55022 EMC declaration

This is a class B product. In a domestic environment, this product may cause radio interference in which case the user may be required to take adequate measures.

No changes or modifications to the MC3 board are allowed without explicit written permission from American Tel-A-Systems, Inc., as these could void the end user's authority to operate the device.

1.0 Introduction

The Infinity Series H.100 MC3 Lite Multi-chassis Interconnect Board is designed to provide a high capacity interconnect path between multiple computers using the H.100 bus to connect computer telephony boards within the chassis. This path is provided by fiber-optic links conforming to the MC3 standard and operating at the OC3 bit rate of 155 Mbps. Provisions are included for supporting dual counter rotating rings for redundancy or higher capacity. This board is similar to the H.100 MC3 Multi-Chassis Interconnect & Conferencing Board, but has had the additional hardware required for conferencing removed to provide a lower cost solution in those situations where only the MC3 capabilities are required.

The H.100 bus was devised by the Enterprise Computer Telephony Forum (ECTF) to provide a single telecom bus for the entire industry. It is intended for add-in boards using the PCI form factor. A wide variety of boards are available from a number of different vendors. The MC3 bus is a chassis interconnect standard promulgated by the GO-MVIP standards body.

The board is equipped with a processor that can be used to control the lower level functions of the board. The host PC controls the board using messages passed through dual-ported RAM. The board shares a common message passing and control scheme with other Infinity Series H.100 boards. This scheme is also compatible with legacy XDS boards for the MVIP-90 bus and SCbus.

The H.100 MC3 Lite Board

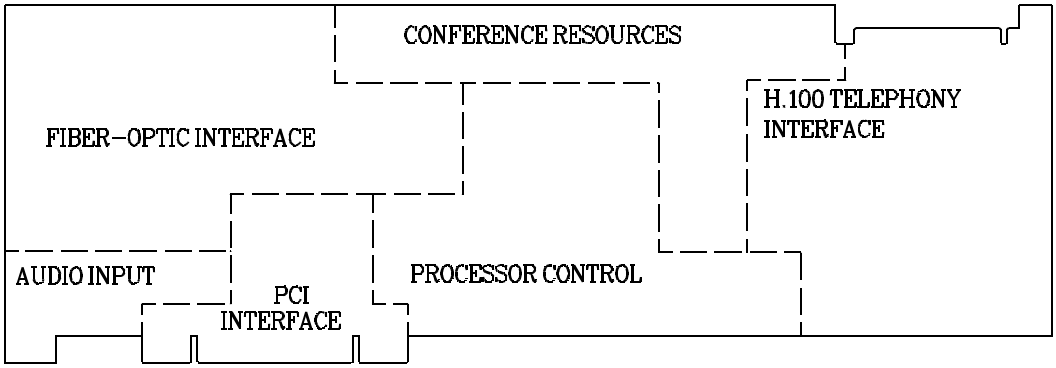

Figure 1: The H.100 MC3/Conference Board Functional Areas

1.1 Features and Capabilities

This section presents an overview of the features and capabilities of the Infinity Series MC3 Lite Multi-chassis Interconnect Board.

1.1.1 H.100 Bus

The H.100 bus is a digital bus for transporting PCM (Pulse Code Modulation) signals between telephony boards. It was created by the ECTF to provide a common bus structure for future development that would end the "bus wars" between the various legacy busses such as the SCbus and the MVIP bus.

PCM is a standard method of digitizing phone signals. It involves encoding each channel at an 8 kHz rate using eight bits. The signals from multiple channels are then combined into a frame. On the H.100 bus, each frame consists of 128 channels or timeslots. The bit rate of the H.100 bus is 8.192 MHZ. Thirty-two wires, also called streams, each carrying 128 timeslots, are combined to form the bus, and provide a total of 4096 timeslots. Two timeslots are required for a full conversation, one for each talker. For compatibility purposes with legacy busses, the first sixteen streams can also run at either 2.048 or 4.096 MHZ. with 32 or 64 timeslots respectively.

In addition to the streams, a number of other signals necessary to maintain synchronization between all the boards in the system are carried on the bus. These signals provide the clocking and framing information. Redundant clocks are provided to aid in recovery if the primary clock should fail. For interoperation with the SCbus, MVIP bus, or H.MVIP bus a number of compatibility clock signals are also defined.

The H.100 bus consists of a 68 conductor ribbon cable that is used to interconnect the boards in the system. This cable connects to a header at the upper right hand edge on each board.

1.1.2 The MC3 Bus Interface

The MC3 bus was devised by GO-MVIP as a means of providing a large number of 64 kbps channels between PC chassis using the MVIP bus for intra-chassis connections. In the interest of minimizing cost and taking advantage of existing hardware, the physical interface uses the same architecture as that used by the SONET standard operating at the OC3 bit rate of 155 Mbps.

Each link consists of a full duplex fiber-optic cable that can support 2430 channels. Seven of these channels are dedicated to framing purposes. The MC3 standard arranges two of these fiber links in dual counter rotating rings. The two rings can be used to provide redundancy against ring or chassis failure or they can be used to double the capacity. Each node of the MC3 structure provides bypass, drop and insert capabilities for each of the 64 kbps channels.

The Infinity Series H.100 MC3 Lite Board provides up to 1024 connections in each direction between the MC3 and H.100 busses.

The H.100 MC3 Lite Board

1.1.3 Clock Modes

The H.100 MC3 Lite Board can operate in a variety of clock modes. Modes are available so that the master clock can either be derived from the H.100 bus, one of the MC3 rings, or be generated internally on the MC3 Lite Board.

1.1.4 Message Passing

The board occupies 8K of memory space on the host PC. This 8K may reside anywhere within the PC's address space. As a PCI board, the address and interrupt of the board is assigned at boot time. The message passing scheme used by the Infinity Series H.100 MC3 Board is identical to that of the other Infinity Series H.100 boards, allowing for the easy combination of a variety of Infinity Series H.100 boards in a single system.

The message passing scheme and message syntax of Infinity Series H.100 boards is similar to that of the older XDS series of MVIP and SCbus boards. At the driver and API level, support is provided for both series of boards so that the H.100 boards may interoperate with legacy boards using a common interface.

1.1.5 Flash EAROM for Firmware

The firmware for the main processor is contained in Flash EAROM. This allows for easy upgrades of the firmware on the board in the field without requiring time consuming downloads every time a system boots. Once reprogrammed, the new firmware is retained even when the power is removed. The original, factory programmed firmware is also retained on board and can be accessed by installing a jumper.

1.2 How to Use This Manual

The first five sections in this manual are organized in the order you should read and use them to get started with your H.100 MC3 Lite Board. We recommend that you begin with these three steps.

1. Follow the instructions in section 2.0 (Quick Start) and 3.0 (Installation). These sections will tell you if your board is operating correctly within your system. You don't need to be familiar with the board's command set to complete this step.

2. Read section 4.0 (Initialization) to initialize the board within your system. Your application must perform these initialization procedures whenever you power-up your PC in order for the board to communicate with the PC.

3. Read section 5.0 (Communications with the PC) for an overview of how to communicate with the H.100 MC3 Lite Board. Section 5.0 includes a summary of the commands for constructing your application and details concerning system interrupts.

Before you can actually build your application, read section 6.0 (The H.100 bus, MC3 Bus and Clock Modes), 7.0 (Using the MC3 bus), and 8.0 (H.100 Bus Switching.) These sections explain, with practical examples, how the H.100 MC3 Lite Board operates and how to use the command set to achieve the desired results.

Section 9.0 explains diagnostic and error messages that may occur.

The Appendix contains diagnostic information that will be helpful if you have problems installing your H.100 MC3 Lite Board.

this page intentionally left blank

2.0 Quick Start

This section describes the first steps you should perform to determine if your Infinity Series MC3 Lite Multi-Chassis Interconnect Board is communicating correctly with your PC system. You can perform this quick check without securing the board to the PC chassis or connecting any cables.

The exact procedure will vary depending on which operating system you are running. For each operating systems, drivers are required to interface to the boards. The drivers supplied by Amtelco have tests built into them to verify communications with the boards. These drivers also come supplied with utility programs that allow the developer to test communications with the board. Please consult the appropriate documentation for the driver and operating system you are using.

Quick Start Procedure

1. Make sure the PC power is off, then insert the board into a PCI slot.

2. Turn on your PC.

3. If the Amtelco driver is not already installed, install it now, following the instructions supplied with the driver.

4. Most Amtelco drivers will display a list of boards that are installed (see the documentation for the particular driver that you are using). If the H.100 MC3 Lite board is listed, skip to step 6.

5. If the board is not listed, there may be a problem with the board not being seated correctly in the motherboard. There

may also be a problem with a memory or interrupt conflict. Power down the PC and check that the board is properly seated in the connector and repeat steps 1-4. If this does not remedy the problem, try removing any other computer telephony boards in the system. If your PC is unable to find the board, consult the number at the end of this section.

6. Run the program "xdsutil" supplied with the driver. Send the message "IN" to the H.100 MC3 Lite board. The board should respond with the message "IA".

7. Send the message "VC" to the board. Verify that the Receive Message reads: `VCxxxxvvvvPOL` (where `xxxxvvvv` is a variable indicating the firmware version).

8. If the Communications screen shows the correct command responses, your H.100 MC3 Lite Board is communicating with the PC. You may now power down the computer and attach the necessary cables (see section 3.4)

For technical assistance, call Amtelco at 1-608-838-4194 ext.168.

3.0 Installation

This section describes how to install your Infinity Series H.100 MC3 Lite Multi-Chassis Interconnect Board into your PC and how to set the switches, jumpers, and connectors. Before you begin the installation procedure, be sure to test the board as described in section 2.0 (Quick Start).

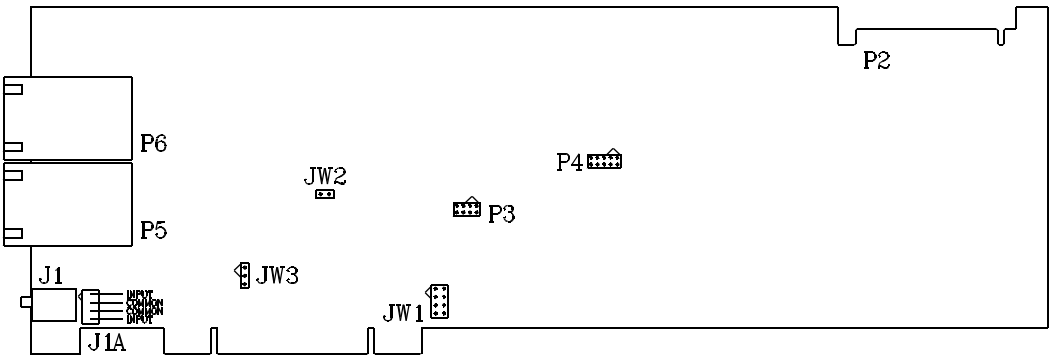

Figure 2: Location of Jumpers, Headers, and Connectors

3.1 PCI Configuration

As Infinity Series boards conform to the PCI standards, there are no switches to set to configure the H.100 MC3 Lite Board's memory address, I/O addresses, or interrupt. The PC's bios will automatically configure the board at boot time to avoid conflicts with other boards in the system.

3.2 Jumpers & Headers

The following is a complete list of all jumpers for the H.100 MC3 Lite Board:

JW1-1 Firmware Select. If firmware has been downloaded to the board, this jumper selects whether the downloaded firmware or the factory default firmware is used. When this jumper is installed, the factory default firmware is executed whenever the board is reset. When the jumper is not installed, the downloaded firmware will be executed after a reset if it is present. If no downloaded firmware is present, the factory default firmware is executed after reset.

JW1-2 DSP Firmware Select. Two separate firmware programs are included in the EAROM, one for the board processor and one for the DSP. If JW1-2 is installed and downloaded DSP firmware is present, the factory DSP firmware is executed after reset. Otherwise, the downloaded firmware is executed if present. Not used on the MC3 Lite Board. See JW1-1

JW1-3 Undefined, reserved for future use.

JW1-4 Undefined, reserved for future use..

JW2, JW3 These jumpers are used for factory testing and should not have jumpers installed.

P3 Diagnostic port. Never install jumpers here.

P4 This header is used for programming internal logic and should never be jumpered.

3.3 Connectors: P2, P5, and P6

P2 H.100 bus. The H.100 bus connector (P2) is a standard H.100 bus header. Use an H.100 bus ribbon cable to connect the H.100 MC3 Lite Board to other H.100 boards within the same PC chassis.

P5 Ring 0 RCV, Ring 1 XMT. This is one of the two fiber optic transceivers for the MC3 bus. In a normal counter rotating ring connection, the fibers from this transceiver are connected to **P6** in the next chassis in the ring. This connector is keyed to insure proper insertion.

P6 Ring 1 RCV, Ring 0 XMT. This is one of the two fiber optic transceivers for the MC3 bus. In a normal counter rotating ring connection, the fibers from this transceiver are connected to **P5** in the previous chassis in the ring. This connector is keyed to insure proper insertion.

3.4 Installation

To install the H.100 MC3 Lite Board in your system:

1. Follow the quick check procedures described in section 2.0 to verify the operation of the board.

2. If the quick check is successful, turn off the PC power and remove the board from the chassis.

3. Install any necessary board jumpers. See section 3.2 for jumper configurations.

4. Insert the board into the chassis. Seat it properly in a PCI slot in the PC chassis and tighten the screw in the back of the board to secure it.

5. Connect the H.100 bus cable to P2.

6. Connect the fiber optic cables to P5 and P6. See section 6.4 for details on the various ring configurations.

7. If the analog port is to be used, connect the music source or other compatible device.

8. Power up the computer.

4.0 Initialization

This section describes the procedures necessary to initialize the system and enable the PC to communicate with the Infinity Series H.100 MC3 Lite Multi-Chassis Interconnect Board. XDS drivers will implement some of these procedures.

4.1 PCI Initialization

The system BIOS is responsible for recognizing PCI boards and mapping them into the I/O and memory spaces as required. It is also responsible for assigning interrupts to the board. This is done through a set of on board registers which contain information specifying the memory, I/O, and interrupt needs of the board. A set of BIOS functions exist for accessing this information. A detailed description of these functions can be found in the *PCI BIOS Specification* published by the PCI SIG, the PCI Special Interest Group.

Normally, the drivers supplied by Amtelco will take care of the process of finding Infinity Series boards and establishing communications. The information in the rest of this subsection is for background only.

The configuration registers of every PCI board contain a vendor ID and device ID code. These codes are unique to each board vendor. All Infinity Series H.100 boards have the same vendor and device IDs. The vendor ID is 14E3h and the device ID is 0101h. A BIOS function exists that will find each instance of a particular vendor and device ID, and which returns with a bus and device number. The bus and device number is then used in functions to read the configuration registers.

The configuration registers contain information on the base address of the memory and I/O assigned to the board by the BIOS. A PCI board may

have up to six different base addresses. On Infinity Series H.100 boards, the first two base addresses are used by the PCI bus interface logic. The third base address which is contained in registers 18-1Bh contains the memory location of the dual-ported memory that is used to pass messages. The interrupt information is contained in register 3Ch. The information in these configuration registers can be used by a driver to address the board.

4.2 Initialization Commands

The H.100 MC3 Lite Board is initialized by sending a sequence of command messages to the board. The process of sending messages is described in detail in Section 5.0, but normally it is accomplished either with a low-level driver XMT command or the API function **xds_msg_send.** Response messages are read using the low-level driver RCV command or the API function **xds_message_receive**.

To enable communications with the H.100 MC3 Lite Board, an **IN** command message should be sent to the board. The board will respond with an **IA** message.

The board may be reset using the command message **RA**. The board will respond with an **RA** message.

Your application can now configure the H.100 MC3 Lite Board using these commands

Command	Purpose
SCmsabb(c)	Sets the clock mode for the board. The parameter m is the clock-mode. The parameter s is the clock sub-mode. The parameters a, bb, and c are used to specify additional clock control information such as compatibility modes, clock rates, local network, and CT_NETREF settings. The default mode on power-up or restart is mode 0. See section 6.0 of this manual for details of clock mode arguments.

SBabcd This command is used to define the clock rate for the lower 16 streams for compatibility with the SCbus or MVIP bus. The parameters a, b, c, d are used to set the rate for streams 0-3, 4-7, 8-11, and 12-15 respectively. The default is 8.192 MHZ. The possible settings are:

0 - 2.048 MHZ., 32 timeslots per stream
1 - 4.096 MHZ., 64 timeslots per stream
2 - 8.192 MHZ., 128 timeslots per stream

SEx Sets the encoding mode for the board. The parameter x can be either M for Mu-Law as used in North America and Japan, or A for A-Law as used in Europe and Asia. The default value is for Mu-Law.

SMx Selects the ring mode. The parameter x is used to choose between the extended mode where both rings are available and the redundant modes where one of the rings acts as the primary ring and the other as a backup. The choices for x are:

0 - extended mode, 4846 timeslots
1 - redundant mode, primary ring is 0, 2423 timeslots
2 - redundant mode, primary ring is 1, 2423 timeslots

SPstt When operating in SCbus compatibility mode, it may be desirable to restrict the number of timeslots that the board can use to transmit to the SCbus. This command restricts the total number of SCbus timeslots to s * 64 + tt, where s and tt are in hexadecimal. This command is used in conjunction with the **SX** command and should only be used when operating in the SCbus mode.

SRx Selects ring failure mode. If the board is to be operated with only a single ring due to a failure or configuration choice, this command is used to set the appropriate hardware. The choices for x are:

0 - Both rings are operational
1 - Ring 0 failure
2 - Ring 1 failure
3 - Both rings failed

STab Controls termination. Parameters a and b control termination for the H.100 and MVIP bus respectively. When set the E, termination is enable and when set to D, termination is disabled. Boards on the end of the H.100 cable should have termination enabled. When operating in MVIP compatibility mode, the MVIP termination should be enabled when the following condition exists:

For systems with five or fewer MVIP Bus connections and less than 90 pF load on the clock lines, it is adequate to place the circuit board that is the master clock source at one end of the cable and provide termination on the circuit board which is physically at the other end of the cable.

On systems with more than five MVIP connections or more than 90 pF load on the clock lines, both ends of the cable should be electrically terminated with the 1000 Ohm/ 1000 pF termination. No other boards should terminate these lines.

SXstt This command is used to set the base timeslot on the SCbus when reserving timeslots to transmit on. The parameters s and tt are hexadecimal numbers setting the lowest timeslot of the block of timeslots reserved

for the board. The number of timeslots reserved is defined by the **SP** command. This command should only be used when operating in the SCbus mode.

This page left intentionally blank

5.0 Communicating with the PC

This section describes how the PC communicates with the Infinity Series H.100 MC3 Lite Multi-chassis Interconnect Board. It includes the definitions for the H.100 MC3 Lite Board commands and responses along with a description of the mailboxes used for messaging.

The board is controlled by the host PC through a system of two mailboxes. The messages consist of short NUL-terminated ASCII strings, which are easy for the host software to compose and parse. The board is capable of buffering up to eight messages in either direction and can drive an interrupt line when it has a message for the host. Messages may not exceed 32 characters.

There are two mailboxes, one for messages to the board and one for messages from the board, and two flags associated with them. A 00h in a flag byte indicates the mailbox is free, a non-zero value indicates that the mailbox is occupied. The mailboxes and their flags are contained in an 8K block of dual-ported memory at the following offsets:

receive mailbox	1F80h
transmit mailbox	1FC0h
transmit flag	1FFCh
receive flag	1FFEh

The board's base address is determined by reading PCI Configuration Space offset 18h. The 32-bit value at this location is the base address for the dual-ported memory on the board.

To send a message, the message is placed in the mailbox and the flag is set to 01h. To read a message, the message is removed from the mailbox and the flag is cleared to 00h. This will clear the interrupt hardware.

5.1 Command and Response Protocol

This section describes the necessary step-by-step procedures for the PC to send a command to the board and to remove a response from the board.

5.1.1 Sending Commands to the Board

The basic steps to sending a command to the H.100 MC3 Lite Board are:

1. Build a command. Broadly speaking, a command is a string of ASCII characters with a null (00h) termination character.

2. Check the transmit flag. If the flag is 0, continue with the next step to put the command in memory. If the flag is not 0, wait until the flag is 0.

3. Insert the command in transmit mailbox memory beginning at the address of the transmit mailbox.

4. Write 01h to the transmit flag. This notifies the board that a message is waiting.

5.1.2 Reading Messages From the Board

1. Check the receive flag. If the flag is 0, there is no message. If it is non-zero, a message is waiting. Continue with the next step to read the message.

2. Remove the message from memory, starting at the address of the receive mailbox. The message is a NUL-terminated ASCII string.

3. Write 0h to the receive flag.

5.1.3 Reading Board Information

A range of board information is included in memory so that it can be checked without sending a message:

Type of Information	Offset Address
Board ID	1F00-1F03
Firmware Version	1F04-1F07
Number of transmit timeslots	1F10-1F11
Base timeslot	1F12-1F13
Base timeslot of conferences	1F14-1F15
Base timeslot of analog port	1F16-1F17
Clock mode settings	1F18-1F1B
Board configuration	1F1C-1F1E
Clock status bits	1F1F

Note: The number of reserved transmit timeslots, and base timeslots are used only in the SCbus compatibility mode when reserving transmit timeslots. The base timeslots for conferences and analog ports are for compatibility with the MC3/Conference Board.

The board stores its identity upon power up or a hardware restart. The phrase **Restart PCI MC3 (C) Amtelco 2001** appears in the receive mailbox. The receive flag is not set and no interrupt is generated.

5.2 Interrupts

The H.100 MC3 Lite Board can generate an interrupt to the PC indicating that a message is available. The interrupt for PCI boards is assigned by the BIOS or Operating System at boot time. The assignment is dependent on which PCI slot the board is in. The interrupt line is usually shared by more than one device. If multiple Infinity Series boards are installed they may or may not all share the same interrupt line.

In order for an Infinity Series board to send interrupts to the PC, the PCI

Interface circuit on the board must be programmed to enable interrupts. This is accomplished by setting bits 0 and 3 in the board's Interrupt Control/Status Register. This is a byte-wide register located at an offset of 69h from PCI Base Address 0. PCI Base Address 0 is contained in PCI Configuration Space register 10h. The Base address is a 32-bit value and is mapped into memory.

When an Infinity Series board sends a message, it generates a local interrupt to the PCI Interface circuit on the board. If the PCI Interface circuit has been programmed to generate interrupts to the PC, the local interrupt is passed through to the PC. When the PC receives an interrupt, its Interrupt Service Routine (ISR) should check the Infinity board's receive flag to see if a message is pending (i.e. the receive flag is non-zero). It should then process the message for the board and write a 0 to the board's receive flag.

5.2.1 Interrupt Initialization

1. Clear the board's receive flag.

2. Read the PCI Base Address 0 from PCI Configuration Space offset 10h (this must be a 32-bit access).

3. Set bits 0 and 3 of PCI Base Address 0 + 69h. Do not modify any other bits in this register. This register is a byte-wide memory mapped register.

5.2.2 Step-by-Step Interrupt Processing Summary

1. Check to see if the receive flag is non-zero.

2. Remove the message from the receive mailbox.

3. Write 0h to the receive flag.

4. Re-enable the interrupt controller on the PC.

5.3 Commands and Responses

This section gives a general overview of the H.100 MC3 Lite Board commands and responses. The commands are grouped by function and then listed in alphabetical order by two-letter command. Refer to sections 7.0 through 10.0 for examples and explanations of how to use these commands.

5.3.1 Characteristics of Command Strings

< All commands consist of null (00h) terminated ASCII strings.

< There are no spaces or other delimiters between parameters in the commands.

< All letters in command strings must be UPPERCASE unless otherwise noted.

< Lowercase monospaced letters (such as xx) in the following command references represent parameters within commands. Each letter represents one ASCII digit.

< Numeric parameters are always hexadecimal numbers.

5.3.2 Command Parameters

The following table documents the common parameters for many of the commands listed in the next sections. Other less common parameters are defined with individual commands.

Common Command Parameters		
Parameter	**Definition**	**Values**
pp	Output pattern value	00-FFh
iiii	1st & 2nd i, H.100 receive stream	00-1Fh
	2nd & 3rd i, H.100 receive timeslot	00-7Fh
oooo	1st & 2nd o, H.100 transmit stream	00-1Fh
	3rd & 4th o, H.100 transmit timeslot	00-7Fh
yyyy	1st & 2nd y, MC3 transmit stream	00-4Bh
	3rd & 4th y, MC3 transmit timeslot	00-3Fh
zzzz	1st & 2nd z, MC3 receive stream	00-4Bh
	3rd & 4th z, MC3 receive timeslot	00-3Fh
bsstt	MVIP terminus, b - bus, ss - stream, tt - timeslot	C, H, L, X 0000-4B7F

5.3.3 Commands to the H.100 MC3 Lite Board

Note that sections 7.0-9.0 provide supplemental information and examples for the commands and messages documented here.

MC3 Bus Commands

XCooooiiiiyyyyzzzz Connect H.100 timeslot iiii to MC3 transmit timeslot yyyy and MC3 receive timeslot zzzz to H.100 timeslot oooo

XDIoooo Disable output to H.100 timeslot oooo from MC3 bus

XDOyyyy Disable output to MC3 timeslot yyyy

XLIoooozzzz One-way audio from MC3 timeslot zzzz to H.100 timeslot oooo

XLOyyyyiiii One-way audio from H.100 timeslot iiii to MC3 timeslot yyyy

XPIooopp Output pattern pp on H.100 timeslot oooo
XPOyyyypp Output pattern pp on MC3 timeslot yyyy

MVIP Compatibility Commands

MObsttD Set_output disable mode, bsstt - output terminus
MObsttEbsstt Set_output enable mode, bsstt - output terminus, bsstt - input terminus
MObsttPpp Set_output pattern mode, bsstt - output terminus, pp - pattern value
MTD Disable output to the CT Bus (tristate)
MTE Enable output to the CT Bus

Interrupt Control Commands

IF Disable transmit interrupts and messages
IN Enable transmit interrupts and messages

Reset Commands

RA Reset All

Setup Commands

SBabcd Set bit rate for streams 0-3, 4-7, 8-11, and 12-15
0 - 2.048 MHZ.
1 - 4.096 MHZ.
2 - 8.192 MHZ.
SCmsabb(c) Set clock mode m submode s, arguments a, bb, & c
SEm Set encoding mode, m = M - mu-Law, A - A-Law
SLx Set Loopback Mode x = 0-F, 0 = no loopback
bit 0 - TLBB, 1 - FLBB, 2 - TLBA, 3 - FLBA
SMx Set Ring Mode, x = 0 extended mode, both rings available
1 - Redundant rings, Ring 0 primary ring
2 - Redundant rings, Ring 1 primary ring
SPstt Set maximum number of transmit timeslots to stt
SRx Select Ring Failure bits, x = 0, no failure, x = 1 ring 0, x = 2 ring 1, x = 3 both rings
STab Set bus termination, a = H.100 bus, b = MVIP bus
SXstt Set base SCbus transmit timeslot to stt

Version Requests

VA	Checksum of alternate segment request
VC	Version request

Download Commands

@xxxx	Download 1K block to address xxxx
@Es	Erase segment s
GA	Jump to Alternate Program
GM	Jump to Main Program
@Ws	Write from RAM to segment s

Diagnostics

QC	Query Clock Mode information
QHcrrrr	Query CT812, c = CT812, rrrr = register
QMRx0zzz	Query MC3 Receive, x = MT90840, zzz = internal stream & timeslot
QMTxyyyy	Query MC3 Transmit, x = MT90840, yyyy = MC3 stream & timeslot
QObsstt	Query Output for terminus bsstt
QS	Query MC3 bus status
QXxxx	Query SCbus transmit timeslot for port xxx

5.4.4 Responses from the H.100 MC3 Lite Board

Acknowledgments

A(msg)	Verify set response to a message
IA	Interrupt On acknowledge
RA	Reset all acknowledge

Error Messages

ECxx	Clock error bits xx
EFr	Ring r failure
EPoooo	Path error for H.100 output timeslot xxx
ERr	Ring r restored
ESrxx	Ring r status error, status xx

Query Responses

QCmsabrrttkkrsmxsy	Reply to Query Clock Mode m - mode, s - submode, a, b - arguments rr - stream rate byte, tt, kk - enable flags rs - reset byte, mx - mux byte, sy - SYN-155 byte
QH0rrrrdddddd	Reply to CT812 query, dddddd is register data
QMRx0zzzhhlldd	Reply to Query MC3 Receive, hh = receive path connection memory high ll = receive path connection memory low dd = transmit path data memory
QMTxyyyyhhlldd	Reply to Query MC3 Transmit hh = transmit path connection memory high ll = transmit path connection memory low dd = receive path data memory
QObssttm(bsstt)	Query_output reply, bsstt - output terminus, m - mode (bsstt) input terminus
QSsstt	MC3 bus status ss for ring 0, tt for ring 1
QXxxxstt	SCbus transmit timeslot for port xxx
QXxxxZ	No SCbus transmit timeslot set for port xxx

Diagnostic Responses

VAcccc	Alternate segment checksum, cccc - checksum
VCccccvvvviiix	Version request response, cccc - checksum vvvv - version, iii - board identifier POL - PCI MC3 Lite board, x - population level
U(msg)	undefined or unparseable message

this page intentionally left blank

6.0 The H.100 Bus, Clock Modes & MC3 Bus

The Infinity Series H.100 MC3 Lite Multi-Chassis Interconnect Board provides a means of connecting the fiber-optic MC3 interchassis bus to the H.100 bus. Through this bus, the MC3 channels can be connected to other H.100 compatible boards. It is capable of operating in a variety of clock modes compatible with H.100 and MC3 operation. In addition, the board is capable of interoperating with legacy MVIP and SCbus boards.

6.1 The H.100 Bus

The H.100 bus consists of 32 Pulse Code Modulation (PCM) streams operating at an 8.192 MHZ. clock rate. Each stream contains 128 timeslots, for a total of 4096 timeslots. In addition to the PCM data signals, there are a number of bit, frame, and network reference signals that are used to synchronize the operation of multiple boards. For interoperation with the legacy SCbus, MVIP-90 bus and the H-MVIP bus there are some additional clock signals that are included on the bus.

For the purposes of commands, a particular H.100 timeslot is referred to by a four digit hexadecimal number. The first two digits are the stream number, while the last two digits are the timeslot within the stream. Streams range from 00h to 1Fh, and timeslots from 00-7Fh.

The physical H.100 bus is a 68 conductor ribbon cable that connects the various boards in the system. As in any such bus, termination is important for its proper operation. The board at each end of the H.100 cable must have the proper termination installed or enabled, while any board between the ends must not terminate the bus. For the H.100 MC3 Lite Board, termination is enabled using a command of the form **STab** where a controls the H.100 termination and b the MVIP bus termination. Termination is

enabled if a is “E” and disabled if a is “D”.

6.1.1 Legacy Bus Compatibility

The H.100 specification provides for inter-operability with several common legacy PCM busses. These include the SCbus, the MVIP-90 bus and the H-MVIP bus. Because these busses run at different bit rates than the 8.192 MHZ. of the H.100 bus, provisions exist in the specification to run the first 16 streams at either 4.096 MHZ. or 2.048 MHZ. For inter-operability with the SCbus, these streams typically should be run at 4.096 MHZ. (2.048 MHZ. and 8.192 MHZ are also possible choices) and with the MVIP-90 bus they should be run at 2.048 MHZ. H-MVIP runs these streams at either 2.048 or 8.192 MHZ. depending on whether MVIP-90 compatibility is desired.

On the H.100 MC3 Lite Board, the bit rate of the first 16 streams is set using the “SB” command. This command takes the form **SBabcd** where the parameters a, b, c, and d select the bit rate for streams 0-3, 4-7, 8-11, and 12-15 respectively. The choices for these parameters are:

0 - 2.048 MHZ.
1 - 4.096 MHZ
2 - 8.192 MHZ.

Thus to operate with the SCbus at 4.096 MHZ. the command would be **SB1111** and to operate with the MVIP-90 bus **SB0000**. The default selection for these streams is the H.100 rate of 8.192 MHZ.

When operating in a compatibility mode, the timeslot in board commands range from 00 to the maximum number of timeslots allowed by the bit rate. At 2.048 MHZ. timeslots within a stream are numbered 00-1Fh and at 4.096 the timeslots are 00-3Fh. MVIP-90 bus streams are numbered 00-0Fh. This numbering corresponds to the DSo/DSi convention according to the following table:

H.100 stream	MVIP-90 stream	MVIP-95 stream	H.100 stream	MVIP-90 stream	MVIP-95 stream
00h	DSo0	HDS0	08h	DSo4	HDS8
01h	DSi0	HDS1	09h	DSi4	HDS9
02h	DSo1	HDS2	0Ah	DSo5	HDS10
03h	DSi1	HDS3	0Bh	DSi5	HDS11
04h	DSo2	HDS4	0Ch	DSo6	HDS12
05h	DSi2	HDS5	0Dh	DSi6	HDS13
06h	DSo3	HDS6	0Eh	DSo7	HDS14
07h	DSi3	HDS7	0Fh	DSi7	HDS15

6.2 Clock Modes

The H.100 bus specification defines a variety of clock signals. Two clock signals CT bus A and CT bus B are provided for redundancy. In addition, a signal called CT_NETREF is defined which may be referenced to an external clock source such as a T1 or E1 span. This signal exists to aid in recovery if the primary clock source should fail. The specification also includes clock signals for compatibility with both the MVIP90 and SCbus.

The clock mode must be set before any connections can be made with other boards. The clock mode is set using the Set Clock command “SCmsabbc”, where m is the clock mode, s is the sub-mode, and a, bb, and c are additional arguments used to select clock sources and specify compatibility modes. The default clock mode on a power up is to provide a local clock, but to neither source clock signals to the bus or derive the clock from the bus. The possible clock modes are:

0 no clocks to or from the bus
1 clocks slaved to the CT bus
2 the board is master CT bus clock A
3 the board is master CT bus clock B
4 the board is secondary master for CT bus clock A
5 the board is secondary master for CT bus clock B

Connections are possible only when all boards within a system are synchronized to the same clock. Only one board in a system can provide the H.100 bus clock. The other boards in the system must slave their internal clocks to the master. If the H.100 MC3 Lite Board is to use the H.100 bus clock, this clock must be provided by another board before switching can be accomplished.

6.2.1 Slave Mode

In the Slave Mode, the H.100 MC3 Lite board derives its clocks from one of the clock signals on the CT bus. The clock signal is selected with the submode argument in the **SC** command. The possible clock signals are:

0 - CT bus clock A
1 - CT bus clock B
2 - SCbus clock at 2 MHZ.
3 - SCbus clock at 4 MHZ.
4 - SCbus clock at 8 MHZ.
5 - MVIP90 clock signal
6 - CT bus clock A, auto-fallback mode
7 - CT bus clock B, auto-fallback mode

Argument a is used to set the CT_NETREF mode, while argument bb is used to select the source of CT_NETREF. The choices for argument a are:

0 - No CT_NETREF output
1 - CT_NETREF_1 output is enabled
2 - CT_NETREF_2 output is enabled

It should be noted that CT_NETREF_2 is defined only for the H.110 bus and not the H.100 bus. It is included for upward compatibility. The CT_NETREF source can be either MC3 Ring 0 if argument bb is 00 or MC3 Ring 1 if argument bb is 01.

6.2.2 Primary Master Mode

In modes 2 or 3, the board supplies the CT master clocks A or B respectively. Other boards on the H.100 bus will synchronize to one of these clocks. The source of the clock is selected by the submode argument s. The choices are:

0 - freerun, the board's internal clock
1 - CT_NETREF_1
2 - CT_NETREF_2 (not present on the H.100 bus)
3 - a local network, either Ring 0 or Ring 1
4 - a local network, either Ring 0 or Ring 1 with auto-fallback to CT_NETREF

For submode 1 and 2, argument bb will select the frequency of the CT_NETREF signal. The choices are:

00 - 8 kHz. (frame rate)
01 - 1.536 MHZ. (T1 rate)
02 - 1.544 MHZ. (T1 extended superframe rate)
03 - 2.048 MHZ. (E1 rate)

For submodes 3 and 4, argument bb will select either the MC3 Ring 0 if 00 or Ring 1 if 01. For submode 4, the optional argument c will specify the frequency of CT_NETREF.

For all modes, argument a will select the legacy bus compatible clocks that the board will supply. The options are:

0 - no compatibility clocks
1 - SCbus clocks at 2 MHZ.

2 - SCbus clocks at 4 MHZ.
3 - SCbus clocks at 8 MHZ.
4 - MVIP 90 clocks
5 - H-MVIP clocks

6.2.3 Secondary Master Modes

When a board is operating as a secondary master, it uses the other clock signal as a source, i.e. if a board is the secondary master for CT clock B, it uses CT clock A as a source and provides CT clock B. If the primary clock fails, the secondary master then becomes the clock master. Typically, one board will be set as the master for clock A and another board as the secondary master for clock B, or vice versa. If the clock source specified by the submode is either of the CT_NETREF signals or a local network, the board will automatically fall back on that source if the primary clock should fail. If set to free-run, it will fall back to a PLL that was locked to the primary master clock.

In all secondary master modes, if the primary master fails, the board will automatically become the new primary master. If the original primary master is restored, the clock mode for the original secondary master must be reset.

When operating in secondary master mode, the arguments s, a, and bb are the same as when operating as a primary master.

6.2.4 Clock Fallback

The H.100 Specification details a scheme for automatically recovering from a clock failure. One of the CT bus clocks, either A or B is designated the master clock. The other clock is the secondary master and is generated by a different board than the primary clock. While the primary clock is valid, the secondary clock is locked to it. If the primary clock should fail, the secondary clock takes over using a local oscillator, CT_NETREF or a local network as the source. Boards that are slaves should automatically fall back to the secondary clock. After a failure of the master clock, system

software should designate new primary and secondary clocks. The new primary may be the previous secondary clock master. For Infinity Series H.100 boards, this will involve sending a set clock command with the new primary clock information.

When an Infinity Series board is set for automatic fallback, the board will automatically switch to the secondary clock if the primary clock fails. When this occurs, the board will send an "EC" message indicating the failure. When the application designates a new primary master, it should send a new clock mode command to the board even though auto-fallback may have occurred.

6.2.5 Clock Errors

If the board detects a problem with the clocks, it will generate a clock error, which notifies the application that it should take appropriate action. Clock errors are reported in the Clock Error Bit message, **ECxx** where the **xx** is a hexadecimal value in which each bit identifies the specific error. A value of 1 indicates an error condition. The bits are as follows:

bit	Error Description
0	CT bus clock A
1	CT bus clock B
2	SCbus clocks
3	MVIP bus clocks
4	Master PLL error
5	Frame Boundary

6.3 The MC3 Bus

The MC3 bus is an interchassis connection mechanism that was defined by the GO-MVIP Technical Committee. It uses many of the same concepts and physical standards as the OC3 SONET specification, but is unique and not intended for direct interconnection with SONET equipment.

The MC3 bus consists of two full-duplex fiber-optic connections operating at a 155 Mbps bit rate. This gives 2430 64 kbps channels on each ring, of which seven are required for framing. The remaining 2423 channels are available for use in transporting 64 kbps information between chassis. The separation between chassis can be as great as 2000 meters. The MC3 bus can be operated in two modes. In the first mode, the two rings can be combined to give a total of 4846 channels or timeslots. In the second mode, the two rings can be arranged as redundant counter rotating rings. In this mode signals can be routed around any one break between chassis.

For the purposes of compatibility with the XDS MC1 Multi-Chassis board, the MC3 bus is divided up into logical streams of 64 timeslots each. Each ring has 38 of these logical streams, though the last stream does not have a full 64 timeslots. Timeslots are referenced using a four digit hexadecimal number where the first two digits indicate the stream and the last two the timeslot. It should be noted that this arrangement is merely a logical convention and each frame on a ring actually consists of 2430 timeslots.

6.3.1 MC3 Ring Errors

The MC3 rings are also a possible source of errors. If a ring failure is detected, it will be reported with a message of the form **EFr** where r is the ring. Restoration will be reported with a message of the form **ERr.** If the failed ring was being used to provide a reference for the chassis, it may be necessary to change clock modes. Other error conditions may also be detected. These are reported with a message of the form **ESrxx** where r is the ring and xx is the value of the ring status register. The bits in this register are:

bit	description
0	LOS - loss of signal
1	LOF - loss of frame
2	OOF - out of frame
3	RFE - receive frame error
4	B1ERR - bit error

6.4 MC3 Ring Configurations

The standard MC3 configuration takes the form of two counter rotating rings. That is, signals in one ring move between the chassis in one direction and the signals in the other ring move in the opposite direction. To make this configuration, the fiber plugged into P5 in one chassis is plugged into P6 of the next chassis in the system. This pattern is repeated until the fiber finally wraps around itself and is plugged into P6 of the first chassis in the circle. As each connector consists of the receive fiber for one ring and the transmit fiber for the other ring, this configuration completes the two rings rotating in opposite directions.

If the fiber should be interrupted between two chassis, or if there should be a chassis failure, the boards on either side of the break can be set so that signals loop back on themselves. Received signals, instead of being transmitted on the same ring, are transmitted on the other ring in the opposite direction. This forms a completed loop avoiding the broken segment or off-line chassis. The ability to avoid a segment or chassis allows maintenance to be performed while keeping the rest of the system running.

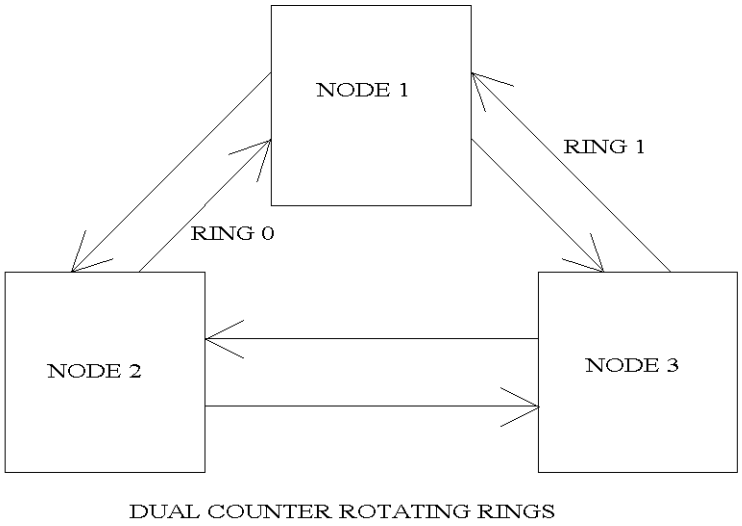

Figure 3: Three Node MC3 Ring

The H.100 MC3 Lite Board

It is also possible with a two chassis system to connect the chassis using only one fiber pair. This may be desirable in some small configurations for the sake of simplicity or to keep the cost down.

6.5 Configuration Information

Information on the clock mode setting, stream rates, and other configuration settings is available in the dual-ported memory in an eight byte block beginning at an offset of 1F18h. The first four bytes are the clock mode, the submode, and the a and bb arguments from the set clock command **SC**. The next byte contains the stream rate information from the **SB** command with bits 0-1 containing the value for streams 0-3, bits 2-3 for streams 4-7, and so on. Bits 0 and 1 of the sixth byte indicate the state of the H.100 and MVIP termination, respectively, with a value of 1 being the enabled state. Bits 0 and 1 of the seventh byte indicate whether conferencing and the audio port have been enable. The eighth byte contains the clock error status bits. These are in the same order as in the **EC** clock error message (Sec. 6.2.5)

7.0 Using the MC3 Bus

This section describes the operation and use of the MC3 Bus. It will explain the steps required for initialization, how to make and break connections, how to take advantage of the redundancy the MC3 architecture provides, how to handle errors, and some of the diagnostic modes.

7.1 Initialization

The most important consideration in properly initializing a system is the configuration of the clocks. Before connections can be established, a uniform set of clocks must be set up and synchronized. This is necessary so that a particular timeslot in a frame can be identified.

One and only one board in the system can serve as the master clock. This clock can be derived from an external digital circuit such as a T1 or ISDN interface or it can be generated by a local on board oscillator. If there is one or more T1 or E1 circuits coming from the public switched network, one of these **must** be the master clock source. This clock is then placed on the H.100 bus for the other boards in that chassis to use as a reference. In the case of an MC3 board, this clock is used to generate the framing on the MC3 rings. Other MC3 boards in the system then reference their clocks from the MC3 framing and use these clocks to drive the CT bus clocks in all the other chassis in the system. If there are no external digital circuits to serve as the ultimate clock reference, one of the MC3 boards may be selected to derive its clock from the local oscillator on that board. This board then drives the clock for both the H.100 and MC3 busses.

Before connections can be made, the clock configuration of all boards in the system must be set. On the H.100 MC3 Lite Board, this is done using the Set Clock command. This command takes the form **SCmsabb(c),**

where m is the clock mode, s is the clock sub-mode, and a, bb, and c specify reference frequency and which local network (MC3 Ring) is the clock source if the board is a clock master. Note that all but one of the MC3 Boards must be a clock master deriving it's clock from one of the two rings. This means that for all but one of the H.100 MC3 boards the clock mode must be "2" or "3", the clock submode must be "3" or "4" (clock source is the local network), and the local network bb must be either Ring 0 or Ring 1. The remaining board must either be running in slave mode or be running as a bus master in clock mode "2" or "3" with the clock submode set to "0" (freerun). If there are digital network connections in one of the chassis, that is a T1, E1, Primary Rate ISDN or Basic Rate ISDN board, that board must be the clock master in that chassis and the MC3 board in that chassis must be slaved to the bus and providing the clocking to the MC3 bus. If more than one chassis has such a board, then one chassis should be picked as the master (a T1, E1, or Primary Rate ISDN is preferred over a Basic Rate ISDN board).

For interoperability with chassis using the XDS MVIP MC3 Board, the XDS MVIP MC3 boards should be running in clock mode "3" unless that chassis contains the ultimate clock source, in which case the XDS MC3/Conferencing Board clock mode will be "0" or "1" if the MVIP bus in that chassis is the source of the clock reference, or "2" if the clock is being derived from the local oscillator on the board.

With the clock mode selected, it may be necessary to select the ring mode. The default is the extended mode where timeslots on both rings are available for use. To select the redundant mode, the Set Ring Mode command must be used. This command takes the form **SMx** where x is the mode. If x is "1", Ring 0 is the primary ring with ring 1 reserved for fallback in case of a failure. If x is "2", Ring 1 is the primary ring. An x value of "0" selects the extended mode.

If the MC3 configuration is anything other than an extended counter rotating ring, it will also be necessary to set the ring failure bits. This is done using the "SR" or Set Ring Failure command. This command takes the form **SRx** where x selects the ring failure mode. If x is "0", there is no failure, if x is "1", ring 0 has failed, and if x is "2", ring 1 has failed. The

default mode is 0 where full counter rotating rings are assumed to be functioning. If only one fiber pair is functioning, or one of the redundant modes is selected, then one of the failure modes must be selected. If the working fiber is connected to P5, or Ring 0 is the primary ring, then the failure mode should be "2", if the working fiber is connected to P6, or Ring 1 is the primary ring, the failure mode selected should be "1".

7.2 MC3 Switching Commands

Connections between the MC3 and H.100 busses are controlled using the "X" commands. These commands consist of the letter "X" followed by one or two additional letters specifying the command. There will also be one to four arguments indicating the MC3 and H.100 timeslots involved in the command. The timeslot arguments are four hexadecimal digits for the H.100 bus and four hexadecimal digits for the MC3 bus. The first two digits indicate the stream and the last two digits indicate the timeslot within the stream. H.100 streams have 128 timeslots while the MC3 streams have 64 (See section 6.1 and 6.3). In the extended mode, MC3 streams have a range of 00-4Bh and in the redundant mode the range is 00-25h. Note that because of framing signals, streams 25h and 4Bh have only 55 timeslots, or timeslots in the range 00-36h. "Inward" connections are from the MC3 bus to the H.100 bus and "outward" connections are from the H.100 to MC3 bus.

A full-duplex connection between the H.100 bus and the MC3 bus is made using a command of the form **XCooooiiiiyyyyzzzz** where oooo specifies the H.100 transmit timeslot, iiii specifies the H.100 receive timeslot, yyyy specifies the MC3 transmit timeslot and zzzz specifies the MC3 receive timeslot.

Half duplex connections can be made using the "XLI" and "XLO" commands which make connections to and from the H.100 bus respectively. The "XLI" command takes the form **XLIoooozzzz** where oooo is the H.100 transmit timeslot and zzzz is the MC3 receive timeslot. The "XLO" command takes the form **XLOyyyyiiii** where yyyy is the MC3 transmit timeslot and iiii is the H.100 receive timeslot. An "XLI" and an

"XLO" command can be combined to form a full duplex connection.

To disable a connection, the "XDI" and "XDO" commands are used to disable output to the H.100 and MC3 bus respectively. These commands take the form **XDIoooo** and **XDOyyyy** where oooo is an H.100 timeslot and yyyy is an MC3 timeslot.

As an example, to make a connection from H.100 timeslot 1 stream 0 to the 131st timeslot on the MC3 bus, and from the 196th timeslot on the MC3 bus to H.100 timeslot 35 on stream 1, the command would be **XC0123000102020303** where 0123 is timeslot 35 (23h) on H.100 stream 1, 0001 is timeslot 1 on H.100 stream 0, 0202 is timeslot 131, and 0303 is timeslot 196 on the MC3 bus. Note that timeslot numbers begin with 0. The same connection could be made with the commands **XLI01230303** and **XLO02020001**. To disable the output to the H.100 bus the command would be **XDI0123,** while the output to the MC3 bus would be disabled by the command **XDO0202.**

There may be occasions when it is necessary to output a fixed pattern on the H.100 or MC3 bus. This can be for diagnostic purposes or for outputting a "silence" pattern. This can be done with a command of the form **XPIooooopp** where oooo is the H.100 timeslot and pp is the pattern value or **XPOyyyypp** where yyyy is the MC3 timeslot and pp is the hexadecimal value of the byte to be output.

7.3 Ring Errors

A number of factors can cause errors. These are signaled in a message of the form **ESrxx** where r is the ring and xx is a status byte indicating the error type. The bits in this byte are as follows:

bit	description
0	LOS - loss of signal
1	LOF - loss of frame
2	OOF - out of frame
3	RFE - receive frame error
4	B1ERR - a parity error detected
6	FSA - frame slip alarm
7	TXPPA - transmit phase alignment alarm

Any of these errors may indicate a problem with the fiber connection or clocking.

If the errors on bits 0-3 persist for more than 150 msec. then a ring failure has occurred. This will be signaled by a message of the form **EFr** where r is the ring number. A recovery occurs when none of these bits indicates an error condition for 1.5 seconds. This will be signaled by a message of the form **ERr** where r is the ring.

7.4 Ring Redundancy & Fallback

The dual counter rotating ring architecture of the MC3 bus allows for redundancy and dynamic fallback in case of a failure in a ring. To take advantage of the redundancy, the MC3 bus must be operated in one of the two redundant modes, with either Ring 0 or Ring 1 as the primary ring. When in the redundant mode, signals are transmitted on the primary ring, and the corresponding timeslot on the secondary ring is set to bypass. Signals from the primary ring are connected to the H.100 bus while the secondary ring remains in reserve.

To set up an H.100 MC3 Lite Board to operate in the redundant mode, it is necessary to set up both the ring mode with an “SM” command, and the ring failure mode with an “SR” command. For example, if the primary ring is Ring 0, then the commands would be **SM1** and **SR0**. If the primary ring is Ring 1, then the commands would be **SM2** and **SR0**. If the board is a clock master, and the clock submode of the board is “3” or “4", then it will be necessary to select the primary ring as the clock source.

A ring failure is signaled by an error message of the form **EFr** where r is the failed ring. When a ring interruption occurs, the boards on either side of the break will indicate errors. If the error is in the primary ring, it will be necessary to change both the ring mode and the ring failure mode. For example, if Ring 0 is the primary ring, and a failure in the primary ring is detected, then the ring mode must be changed so that Ring 1 is the primary ring by issuing an **SM2** command. It will also be necessary to change the failure mode by issuing an **SR1** command. If the board is a clock master and the source is the local network, it must be changed so that the clock source is Ring 1. This is done by issuing an **SC23001** command if the board is clock master A or **SC33001** if clock master B. If the failure is in the secondary ring, only the ring failure mode must be changed. For example, if Ring 1 fails, an **SR2** should be issued.

Note that the "SR" command is used to set any required loopbacks for redundancy. The "SL" command is not used for this purpose and is only used for testing purposes.

A ring recovery is signaled by an **ERr** message where r is the recovered ring. If the recovery is for the ring that had originally been the primary ring, the **"SM"**, **"SR"**, and **"SC"** commands to reestablish the primary ring should be issued. For example, if Ring 0 is restored, the commands would be **SM1, SR0,** and **SC23000** or **SC33000.** If the recovery is for the secondary ring, only an **SR0** command should be issued.

Under some circumstances a ring failure may also require changes to the clock modes of other boards on the affected ring. For a more complete discussion, see Appendix B.

7.5 Loopback Modes

The rings can be placed into several loopback modes for diagnostic purposes. This is done by setting one or more of the loopback bits using the Set Loopback command. This command has the form **SLx** where x is the value of the loopback nibble. The bits in this nibble are

bit	description
0	TLBB - Terminal Loopback B
1	FLBB - Facilities Loopback B
2	TLBA - Terminal Loopback A
3	FLBA - Facilities Loopback A

A Facilities Loopback will send data coming in on one ring out the other ring. FLBA will cause data from Ring 0 to be output on Ring 1, and FLBB will cause data from Ring 1 to be output on Ring 0.

A Terminal Loopback will take data to be output on one ring and loop it back as an input on the corresponding timeslot on the other ring. TLBA will take data to be output on Ring 1 and loop it back as incoming data on Ring 0. TLBB will take data to be output on Ring 0 and loop it back as incoming data on Ring 1.

this page intentionally left blank

8.0 CT Bus Switching

This section describes the operation of the Computer Telephony bus switching capabilities of the Infinity Series MC3 Lite Multi-Chassis Interconnect Board. Topics include the H.100 bus, the operation of the board when operating in SCbus or MVIP bus compatibility mode and the operation of the analog port. In this chapter, the term "Computer Telephony" or CT bus shall include not only the H.100 bus, but also the MVIP-90, H-MVIP, and SCbus

8.1 The H.100 Switching Hardware

The H.100 MC3 Lite Board consists of a single switch blocks dedicated to switching between the MC3 bus and the CT bus. The MC3/Conference board includes another switch block for switching between the computer telephony bus and the conferencing facilities. A total of 1024 inputs from and 1024 outputs to the CT bus are available on the board.

8.2 SCbus Compatibility

The SCbus is a 16 stream bus. Each stream on the bus normally operates at 4.096 MHZ. and has 64 timeslots per stream for a total of 1024 timeslots. However, the SCbus may optionally run at either 2.048 MHZ. with 32 timeslots per stream for a total of 512 timeslots or 8.192 MHZ. with 128 timeslots per stream and a total of 2048 timeslots on the 16 streams. When inter-operating H.100 MC3 Lite board with the SCbus at either 2.048 or 4.096 MHZ. the bit rate on the lower 16 H.100 streams must be set appropriately using the "SB" command.

The SCbus uses a 26 conductor ribbon cable. Because of this, an adapter must be used between the P2 connector on the H.100 board and the

SCbus.

8.2.1 Timeslot Assignment

Typically, drivers and libraries conforming to the SCbus specification use a scheme called “timeslot assignment” to insure that no two devices are transmitting on the same timeslot. Not only will having two transmitters on the same timeslot degrade audio signals, but they also may damage some SCbus boards. To prevent this, each device or “port” is assigned a unique timeslot to transmit on during boot-up. For purposes of timeslot assignment, each possible connection from the MC3 to SCbus can be considered a port.

Transmit timeslots are reserved on the H.100 MC3 Lite board using the “SX” command. The same command is used to reserve timeslots on XDS SCbus boards. The command takes the form **SXstt** where s is the stream and tt is the timeslot on that stream of the timeslot reserved for the first port on the board. As an example, if the first timeslot reserved for the board is 100 in decimal, then the message **SX124** would be sent (timeslot 100 corresponds to stream 1, timeslot 24h).

As the H.100 MC3 Lite Board can transmit on a maximum of 1024 timeslots, it would be possible for the board to require all the timeslots available on the SCbus running at 4.096 MHZ. To prevent this, the “SP” command is used to restrict the number of “ports” that the board will present to the SCbus. This command takes the form **SPstt** where stt is the number of streams and timeslots to be reserved for the board. This command assumes that there are 64 timeslots per stream.

As an example, 256 timeslots can be reserved for the board by sending a command of the form **SP400** (4 streams of 64 timeslots each). If the base timeslot of the board is set at 96, a command of the form **SX120** should be sent to the board. 256 timeslots from 96 to 351 (120-51Fh) would be assigned to the MC3 ports.

When timeslots have been assigned, the actual output timeslot value is no

longer used in commands that control outputs to the CT bus. Instead, a port number is used. This port number will be from 0 to the maximum number of timeslots reserved and will be represented in stream timeslot notation (0stt) in the commands. In the example above, the MC3 function would have 256 ports and these would run from 0000-033F. As an example, to connect MC3 timeslot 0000 to the 96th timeslot reserved for the board, the command would be **XLI011F0000**.

Normally, the timeslot assignment process is carried out as part of the initialization and loading of the driver. A configuration file is used to specify the number of timeslots to be reserved for the board. The function **xds_xmt_timeslot** is used for finding the transmit timeslot of a port. To aid this process, information on the number of reserved timeslots and the base timeslot is presented in the dual-ported memory. This information is available at the following locations:

1F10h	total number of timeslots for the board
1F12h	1st timeslot assigned to MC3 connections
1F14h	1st timeslot assigned to conferencing
1F16h	timeslot assigned to the analog port

The base timeslots reserved for conferencing and the analog port are there for compatibility with the MC3/Conference board and will always be set to 0FFFFh indicating no timeslots are reserved for these functions.

The command **QXstt** can also be used to inquire as to which timeslot is reserved for a port stt. The reply takes the form **QXsttabb** where a is the stream and bb is the timeslot on the stream that is reserved for stt. If no timeslots have been reserved on the board, the response will take the form **QXsttZ.**

8.3 MVIP Compatibility

The MVIP-90 bus has 16 streams with 32 timeslots each. The streams run with a bit rate of 2.048 MHZ. When the H.100 MC3 Lite Board is inter-operating with the MVIP bus, the bus rate on the lower 16 streams must be

set using the “SB” command.

The MVIP-90 bus uses a 40 pin ribbon cable. To connect the H.100 MC3 Lite Board to the MVIP-90 bus an adapter must be used. The MVIP rules for termination must also be followed (see Section 4.2).

Timeslots on the MVIP bus are normally paired, that is timeslot x on DSoy is paired with timeslot x on DSiy. One timeslot of a pair must be defined as an input and the other as an output. With most MVIP boards, an attempt to use both timeslots of a pair as inputs or as outputs will result in a conflict. The table in Section 6.1.1 gives the association between H.100 stream numbers and the DSi and DSo streams.

MVIP-90 applications normally assign timeslots dynamically, so the fact that the H.100 MC3 Lite Board can transmit on more timeslots than are available on the MVIP bus is not a concern. Connections that are not enabled are tri-stated.

The H-MVIP bus has 24 streams. The H-MVIP specification has several modes. One mode is compatible the MVIP-90 specification, that is the lower 16 streams run at a 2.048 MHZ. rate while streams 16-23 run at 8.192 MHZ. This mode can be treated as the MVIP-90 case. In another mode, all 24 streams run at the same 8.192 MHZ. rate as the H.100 bus. In this case, only the physical cabling between the H-MVIP bus and the H.100 bus and clock issues need to be addressed.

8.3.1 MVIP Compatibility Commands

Several commands exist for compatibility with the MVIP-95 driver specification. This specification uses the concept of a “terminus” to define an input or output timeslot. The terminus argument consists of three parts, a bus, a stream within the bus, and a timeslot on that stream. In MVIP compatibility messages, a terminus is represented by a five character string. The first character indicates the bus. Valid bus selections are “H” for the H.100 CT bus and “X” for the MC3 bus.

In the MVIP compatibility mode, connections are controlled using the Set Output command **MO**. This command takes the form **MObssttm**, where "bsstt" is the output terminus being controlled, and m is the mode. Valid modes are "D" for disable, "E" for enable, and "P" for pattern output. In the enabled mode, the input terminus follows the mode character, and in the pattern mode, a two digit hexadecimal number representing the value of the byte to be output follows the mode. As an example, the message "MOH0123EX0000" would enable a connection from the MC3 bus timeslot 0, stream 0 to the H.100 timeslot 23h, stream 1.

A query command **QObsstt** is also available to query the state of the output terminus "bsstt". This command corresponds to the Query_Output command in the MVIP-95 specification. The response takes the form **QObssttm(bsstt)** where "bsstt" is the output terminus, "m" is the mode, and if the mode is enable, the second "bsstt" is the input terminus.

this page intentionally left blank

9.0 Diagnostics & Error Messages

9.1 Diagnostic Commands

Several diagnostic commands are available:

VA Used to request the checksum of the firmware in the alternate segment of the board. This is returned in a message of the form `VAxxxx` where `xxxx` is the checksum of the firmware in the alternate segment of ROM.

VC Used to request the version of the firmware on the board. The version information is returned in a message of the form `VCxxxxyyyyPOL`, where `xxxx` is the checksum of the firmware stored in the main segment of ROM, `yyyy` is a four-digit version number and `POL` indicates the board type. This message takes the same form with all Infinity Series boards, and can be used to determine the configuration of the system.

QHcrrrr Queries the contents of the CT812 chip c, for register `rrrr`. The contents are returned as the `24 bit value dddddd` values in the message `QHcrrrrdddddd`. This command refers to the details of the internal switching circuitry, and is ordinarily of limited use to the application.

QMRx0zzz This command queries the MC3 90840 interface chip `x` for the contents of the receive path connection memory and transmit path data memory for the internal timeslot `zzz`. The results are returned in a message of the form `QMRx0zzzhhlldd` where hh and ll are the high and low

byte contents of the receive path connection memory and dd is the contents of the transmit path data memory. This command refers to the details of the internal switching circuitry, and is ordinarily of limited use to the application.

QMTxyyyy This command queries the MC3 90840 interface chip `x` for the contents of the transmit path connection memory and receive path data memory for the internal timeslot zzz. The results are returned in a message of the form `QMTxyyyyhhlldd` where hh and ll are the high and low byte contents of the transmit path connection memory and dd is the contents of the receive path data memory. This command refers to the details of the internal switching circuitry, and is ordinarily of limited use to the application, though it may be used to look for a pattern byte on the MC3 bus.

QXxxx Queries the transmit timeslot reserved for port `xxx`. This command is only valid if the board is operating in SCbus mode. The reply takes the form `QXxxxstt` where `stt` is the stream and timeslot for the port. If no timeslot is assigned, the reply will be `QSxxxZ`. This command is common to all Infinity Series boards operating in the SCbus mode and all XDS SCSA boards. The H.100 MC3 Lite Board is capable of transmitting on 1024 timeslots, but is typically restricted to fewer timeslots (See section 8.2.1 for details of SCbus timeslot assignment).

10.2 Error Messages

The board will detect a number of error conditions and respond with appropriate error messages. These messages are:

ECxx A clock error bit event `xx` has occurred. The value xx is a hexidecimal number where the bits are (a bit value of 1 is an error)

bit	description
0	CT bus clock A
1	CT bus clock B
2	SCbus clocks
3	MVIP bus clocks
4	Master PLL error
5	Frame Boundary

EFr A failure of ring `r` has been detected.

EPxxxx An attempt at switching has failed because all connections between the MC3 and H.100 bus are used. The command was for stream and timeslot `xxxx`.

ERr Restoration of ring `r` has been detected.

ESrxx A change in the ring status error bits has been detected for ring r. The status bits are reported in xx. The bit values are

bit	description
0	LOS - loss of signal
1	LOF - loss of frame
2	OOF - out of frame
3	RFE - receive frame error
4	B1ERR - bit error
6	FSA - frame slip alarm
7	TXPPA - transmit phase alignment alarm

U*[cmnd]* If the board does not recognize a command message, or if it does not have the appropriate number of arguments, the same message will be returned by the board preceded by a U to indicate an undefined message.

10.3 QM Queries

The QMT and QMR commands can be used to query the contents of the connection and data memories of the chips used to interface to the MC3 rings. Four MT90840 chips control the switching between the MC3 rings and an internal bus, with two chips used for each ring. This internal bus has 1024 timeslots and is used to connect the MT90840 chips with the CT812 chips used to interface to the H.100 bus. These internal timeslots are assigned dynamically with the first connection established using the first timeslot, the second connection the next timeslot and so on. When a connection is disabled, the internal timeslot is released for use.

The QMT command can be used to read the connection memory and data associated with the MC3 bus. The command takes the form QMTxyyyy where x specifies the chip and yyyy specify the MC3 stream and timeslot. Each "stream" has 64 timeslots and the range runs from 0000 to 2536. The chips 0 and 1 are associated with Ring 0 and chips 2 and 3 are associated with Ring 1.

The response takes the form QMTxyyyyhhlldd where xyyyy are as in the query, hh is the high byte of the connection memory, ll the low byte, and dd the contents of the data memory. The top three bits of the high byte are control bits. The value of the high byte will be C0h or 00h when output is disabled, 8xh when enabled, and A0h when outputting a pattern. The low byte and bit 0 of the high byte contain the internal bus timeslot address. When outputting a pattern, the low byte will be the value of the pattern. The data byte will contain the value on the MC3 bus timeslot.

The QMR command can be used to read the connection memory and data associated with the internal bus. The command takes the form QMRx0zzz where x specifies the chip and 0zzz specifies the internal stream and timeslot. Each "stream" has 64 timeslots and the range runs from 0000 to 073F. Chips 0 and 1 are associated with Ring 0 and chips 2 and 3 are associated with Ring 1.

The response takes the form QMRx0zzzyhhlldd where x0zzz are as in the

query, hh is the high byte of the connection memory, ll the low byte, and dd the contents of the data memory. The top four bits of the high byte are control bits. The value of the high byte will be 0xh when output is disabled and 3xh when enabled. The remaining bits of the connection memory are the MC3 bus timeslot address. This command reads the actual contents of the MT90840. Because several timeslots are used for framing purposes, the address bits do not contain the MC3 stream and timeslot value as used in commands. For timeslots 0000-0407, the address bits are incremented by 1 over the command timeslot and for timeslots 0408-2536 the address bits are incremented by 2. For example, if the address bits read 011h, this indicates stream 00, timeslot 10, and if the address bits read 142, this means stream 5, timeslot 0. The data bits are the value present on the local timeslot.

this page intentionally left blank

Appendix A: Environmental Specifications

The Infinity Series H.100 MC3 Lite Multi-Chassis Interconnect Board meets the following environmental specifications:

TEMPERATURE EXTREMES:

Operating: 0EC (+32EF) to +50EC (+122EF).
Storage: –40EC (–40EF) to +70EC (+158EF).

AMBIENT HUMIDITY:

All boards will withstand ambient relative humidity from 0% to 95% non-condensing in both operating and storage conditions.

MECHANICAL:

All Infinity Series H.100 boards conform to PCI-SIG mechanical specifications for full-length PCI cards.

MTBF:

50,000 hours.

ELECTRICAL REQUIREMENTS:

+5 volts ±5% @ 4.0 amps maximum.
-12 volts @ 15mA. maximum
+3.3 volts, -5 volts, & +12 volts not required

this page intentionally left blank

Appendix B: Notes on H.100 MC3 Redundancy

B.1 Introduction

MC3 is designed to provide for fault tolerance in systems by allowing users to implement redundant operation. This document describes the procedures that may be used to implement the fault tolerant design and how to recover from a ring failure.

B.2 Term Definition

First, some terms need to be defined. In some failure modes, it is necessary to know the arrangement of chassis in a ring. For purposes of this document, the terms upstream and downstream will be used when discussing the topology of the ring. Let's say there are three chassis in a system, labeled Chassis A, Chassis B, and Chassis C. Chassis A transmits to Chassis B on ring 0, Chassis B transmits to Chassis C on ring 0, and Chassis C transmits to Chassis A on ring 0. Since the rings counter-rotate, Chassis A therefore transmits to Chassis C on ring 1, Chassis C transmits to B on ring 1, and B to A on ring 1. Given this topology, we can consider the locations of Chassis A and C with respect to B. By definition, Chassis A is upstream of Chassis B on ring 0 and downstream of Chassis B on ring 1 (because Chassis B is receiving data from Chassis A on ring 0 and transmitting to Chassis A on ring 1). Likewise, Chassis C is downstream of Chassis B on ring 0 and upstream of Chassis B on ring 1.

This terminology works well when describing adjacent chassis. However, it is sometimes necessary to consider all chassis in the system. Since a ring topology is used, it could be argued that a given chassis is downstream of all other chassis. When a ring break occurs, it will be necessary to consider what happens to the clocking for the system. In order to do this rationally, the chassis that is providing the master clock for the system will be considered the origin of the ring.

If we consider the three chassis example above, with chassis A providing the master clock for the system, and assume a break occurs between chassis B and C, we can consider chassis A and B to be upstream of the break on ring 0. Chassis C and A are upstream of the break on ring 1.

B.3 Setting up Redundant Operation

Before setting up redundant operation, the clock configuration for the system must be established and a primary and secondary ring must be defined. The clock configuration requires that one chassis provide the master clock for the system. Generally, this clock is derived from a digital network trunk, although the MC3 board is capable of providing a free-running clock. All other chassis will be set up to derive their clock from the designated primary ring.

Once the clocks have been configured for all chassis, redundant mode is enabled by issuing an "SMx" command, where "x" is '1' if ring 0 is the primary ring or '2' if ring 1 is the primary ring. At this point, redundant operation is enabled. Note that in commands dealing with connections, all commands must use streams on ring 0. It is not valid to use a stream on ring 1. Thus, the command "XC0000000126002601" is not valid. Instead, "XC0000000100000001" would establish a connection.

B.4 Responding to a Ring Break

The main actions necessary to respond to a ring break are to loop back data on each side of the break so that all chassis can still access all data and to reconfigure clock modes as appropriate. The required actions for a given chassis will depend on its location in the ring with respect to the break. The following paragraphs describe the specifics.

1. A chassis detects a failure in the secondary ring. This chassis simply needs to issue an "SRx" command to appropriately reroute the data. If the primary ring was 0, the chassis will issue an "SR2". If the primary ring was 1, the chassis will issue an "SR1".

2. A chassis detects a failure in the primary ring. This chassis must implement several procedures. First, it will need to issue an "SMx" command to make the secondary ring primary for as long as the break exists. If ring 0 was primary, an "SM2" command will be issued. If ring 1 was primary, an "SM1" will be issued.

Next, an "SRx" command will be issued to appropriately reroute the data. If ring 0 was primary, an "SR1" will be issued. If ring 1 was primary, an "SR2" will be issued.

Third, the clocks for the system will need to be reconfigured. This is where the application must know the physical ring topology. Basically, all chassis from the failed chassis downstream to the clock master on the primary ring must switch clock modes to derive their clock from the secondary ring.

3. A chassis detects failures in both rings. This would generally occur in a maintenance situation and will require recovery procedures that combine the two failures. Usually, the chassis next to it in one direction would get a failure in one ring and the chassis next to it in the other direction would get a failure in the other ring. A brief example is given below.

If one of the rings is restored, the appropriate "SRx" command should be issued for the remaining failed ring. Thus, if ring 1 is restored and ring 0 is still failed (with ring 0 the primary ring), an SR1 should be issued. When the second ring is restored, an "SR0" should be issued.

4. Multiple failures are detected in a single ring. Generally, these can be dealt with as the failures are received as though no other failures occurred. However, when recovering from a multiple failure condition, the appropriate clock modes must be retained. For example, if ring 0 is primary and two chassis detect a ring 0 failure, both will issue "SR1" commands and

"SM2" commands. Both will also reconfigure clocks on all ring 0 downstream chassis to derive their clocks from ring 1. However, if one of the rings recovers, all chassis downstream of the other failed chassis will still need to continue deriving their clocks from ring 1. Thus, if the chassis that recovers first is downstream of the other, the clock mode should not be changed. If the chassis that recovers first is upstream of the other, only those chassis from the recovered one to the failed one will have their clock reconfigured.

B.5 Examples of Some Specific Cases

A. Assume the three chassis configuration described in the Term Definition section. Chassis A is providing the master clock, chassis B and C are deriving their clock from ring 0. Ring 0 has been designated as primary by issuing "SM1" commands to all three chassis.

Chassis A detects a failure on ring 0. It will issue an "SM2" and an "SR1" command. Since it is the master clock, it does not need to change clock modes. Furthermore, since chassis B and C are downstream on ring 0, their clock modes do not need to change.

Typically, if chassis A detects a failure on ring 0, chassis B will detect a failure on ring 1. It will therefore issue an "SR2" command.

B. The system is setup as described in Example A. Chassis B detects a failure on ring 0. It will issue an "SM2" command and an "SR1" command. Neither chassis B nor chassis C can derive their clock from ring 0 any more because both are downstream of the break on ring 0. Thus, both Chassis B and Chassis C will change their clock modes so that they are deriving their clock from ring 1.

C. The system is setup as described in Example A. Chassis C detects a failure on ring 0. It will issue an "SM2" command and an "SR1" command. Since chassis B is upstream of the break on ring 0, it does not need to change its clock mode. However, chassis C will need to change clock modes so that it will begin deriving its clock from ring 1.

D. The system is setup as described in Example A. Chassis B detects failures on both rings. Generally, this means that chassis A will detect a ring failure on ring 1 and chassis C will detect a failure on ring 0. Assuming this to be the case, chassis B is simply no longer in the loop and cannot recover.

If, for some reason, chassis A and C do not get ring failures, chassis B should notify chassis A that there was a failure on ring 1 (as though chassis A did get a ring failure on ring 1). Chassis A will then issue an "SR2" command. Likewise, chassis B should notify chassis C that there was a failure on ring 0 (as though Chassis C did get a ring failure on ring 0). Chassis C will then issue an "SM2" command and an "SR1" command, as well as changing clock modes so that it is deriving its clock from ring 1.

B.6 Recovering from a Ring Break

When a ring is restored after a break, recovery is basically reversing what was done when the break occurred. If a chassis detects a recovery on the secondary ring, it will issue an "SR0" command. If a chassis detects recovery on the primary ring, it will issue an "SR0" command, an "SMx" command, where "x" is '1' if ring 0 is primary and '2' if ring 1 is primary, and will set all clocks back to their original configuration. (Clock reconfiguration may be complicated if multiple chassis detected a failure on the primary ring. This situation was described in failure mode 4 above.)

If errors had occurred on both rings, the application must deal with the recovery of one ring by acting as though the remaining ring had just failed. Thus, the application will not issue an "SR0" command if only one ring recovers. Instead, it will issue an "SR" command appropriate for the ring that is still failed. Clock modes will have to be dealt with similarly.

B.7 Combining Redundant and Extended Operation Modes

In some cases, customers may desire to run with the total MC3 bandwidth available. However, if there is a break, they would like to recover as gracefully as possible, although some connections may be lost. This can be

done using the XDS MC3 board if some basic guidelines are followed.

First, the application should select a primary ring. This ring should be used for all connections until its bandwidth is used up. The secondary ring may then be used for overflow connections.

If a ring failure is detected, the application should issue "SMx" commands to all chassis, where "x" is '1' if ring 0 was primary, or '2' if ring 1 was primary. This will automatically disconnect all connections on both rings, and will enter the redundant mode. The application will then need to reestablish all connections, as appropriate, given the remaining bandwidth. After the connections have been restored, the system is treated as a standard redundant configuration.

The application may return to extended mode operation at any time by issuing "SM0" and "SR0" commands. ("SR0" commands are only needed if a chassis was in mode "SR1" or "SR2".) When it does so, all connections will again be terminated by the board and will need to be re-established via software control.

B.8 Determining System Topology

In order to implement appropriate fallback techniques, the location of each chassis in the ring must be known. This could be handled manually by having the user enter configuration data. However, it is also possible to use patterns on the MC3 bus to determine where chassis are with respect to each other.

The manual method requires keeping track of where each fiber optic cable is connected. On a given board, there are two duplex fiber connectors. With the board mounted vertically in a standard PC, the bottom position of the top connector is the Ring 0 input. The top position of this connector is the Ring 1 output. Similarly, the bottom position in the bottom connector is the Ring 1 input and the top position of the bottom connector is the Ring 0 output. When cabling multiple chassis, the ring 0 output of one chassis goes to the Ring 0 input of the next chassis. The Ring 1 output of this second chassis goes back to the Ring 1 input of the first chassis. This

configuration is continued around the ring until all chassis are in the ring.

The application software can also determine the ring topology by applying patterns appropriately to the MC3 bus. The procedure is essentially the same whether extended mode or redundant mode is being used. However, in redundant mode, the procedure will determine the ordering of chassis on the primary ring, whereas in extended mode, the ordering of chassis can be determined on either ring based on the commands that are issued.

To enable chassis identification, every chassis should output a unique ID code in the range 00h - 0FFh on a timeslot on one of the rings. All chassis should use the same timeslot and the same ring. This may be accomplished with the "XPO" command. For example, if timeslot 0 on ring 0 is to have the ID code, the application would issue "XPO0000pp", where 'pp' is the pattern to output.

Next, each chassis must query the selected ID timeslot to determine which chassis is transmitting to it. This can be done with the "QMT" command. For the given example, the command issued would be "QMT00000". The board would respond with "QMT00000xxxxpp", where the 'x' arguments may be ignored and 'pp' is the value of the pattern on ring 0 timeslot 0. This pattern is the ID of the board immediately upstream of the chassis on ring 0. It also indicates the downstream chassis on ring 1. Generally, it is probably best to configure all patterns in all chassis first. Then, after a nominal delay to ensure that the pattern is available in all chassis, the "QMT" commands may be issued.

If ring 1 is used for the pattern, the "QMT" command will have the format "QMT2sstt", where 'ss' will be the stream number minus 26h and 'tt' is the timeslot used.

Once the IDs have been determined, the patterns may be disconnected. For the example, the command "XDO0000" could be used in all chassis.

The topology derived from the ID codes may be saved and used to determine how to configure clock modes when a ring failure occurs.

B.9 Maximum Timeslot Utilization

When the multi-chassis system is used for simple full-duplex connections, each party in the connection can use the same MC3 timeslot for maximum capacity. Thus, if two parties are involved in a full duplex conversation and one is transmitting on ring 0 timeslot 0, the other can also transmit on ring 0 timeslot 0. Each would therefore listen to ring 0 timeslot 0. This mode of operation provides for a total of 4846 full-duplex conversations in extended mode and 2423 full-duplex conversations in redundant mode.

In broadcast configurations, where multiple devices listen to the same transmitter, the transmitting device must use a unique timeslot. No other devices can transmit on this timeslot, but as many devices as necessary (depending on the switching capacity of the system), can listen to this timeslot.