

FEATURES

Please contact Galileo Technology for possible updates before finalizing a design.

- Single-chip 4-port Switched Fast Ethernet Controller
 - Provides packet switching functions between four on-chip Fast Ethernet ports and the PCI bus
 - Switch expansion via 66MHz PCI bus (2Gbps)
 - Designed for Fast Ethernet Switches from 4 ports to 24 ports
- GalNet® Architecture Family Member
 - Connects to other GalNet Family Devices through 33MHz-to-66MHz PCI bridges
 - 100% software compatible with GT-48002A
- Incorporates four 802.3 compliant 10/100Mbps Media Access Controllers
 - Direct Interface to MII (Media Independent Interface)
 - Half/Full Duplex Support (up to 200 Mbps/port)
 - IEEE 802.3 100Base-TX, T4, and FX compatible
- Full MII Management Support (MDC/MDIO) via CPU
- Auto-negotiation supported through MII Interface
- Direct support for packet buffering
 - Glueless interface to 2 or 4 Mbyte of 35ns EDO DRAM
 - Up to 2K buffers, 1536-bytes each, dynamically
- High-Performance Distributed Switching Engine
 - Performs forwarding and filtering at full wire speed
 - 148,800 packets/s on each Ethernet port
 - Flexible software or hardware intervention in packet routing decisions
 - allocated to the receive and PCI ports
- Virtual LAN Support
 - Port based virtual LANs
 - Ability to define "super-VLANs" that span multiple VLANs
- Quality-of-Service Queuing
 - Priority queuing based on port number or MAC address
- Advanced address recognition
 - Intelligent address recognition mechanism enables forwarding rate at full wire speed
 - Self-learning mechanism
 - Supports up to 8K Unicast addresses and unlimited Multicast/Broadcast addresses
 - Broadcast storm filtering
- 66MHz Fast PCI interface for switch expansion and management CPU connection
 - Up to 6 GT-48004A devices per PCI bus segment without PCI-to-PCI bridging (24 ports)
 - Up to 32 GalNet devices in a single switch
 - Standard CPU connection for management (66 MHz PCI connection through GT-641xx chips)
- Extensive network management support
 - Repeater MIB and PCI counters
 - Address aging support
 - Hardware assist for Spanning Tree algorithm
 - RMON Station-to-Station connectivity matrix
 - CPU access to Address Table
 - Ability to define static addresses
 - Monitoring (sniffer) mode
- Packet sampling management technology
 - Takes "snapshots" of packets at programmable intervals
 - Allows for the implementation of RMON with low-cost CPUs
- High observability LED interface
 - Dual 3 pin serial LED interfaces give access to over 80 internal status signals
- 329 pin BGA package
 - Advanced 0.35 micron CMOS process
 - 3.3V supply, 5V tolerant I/O

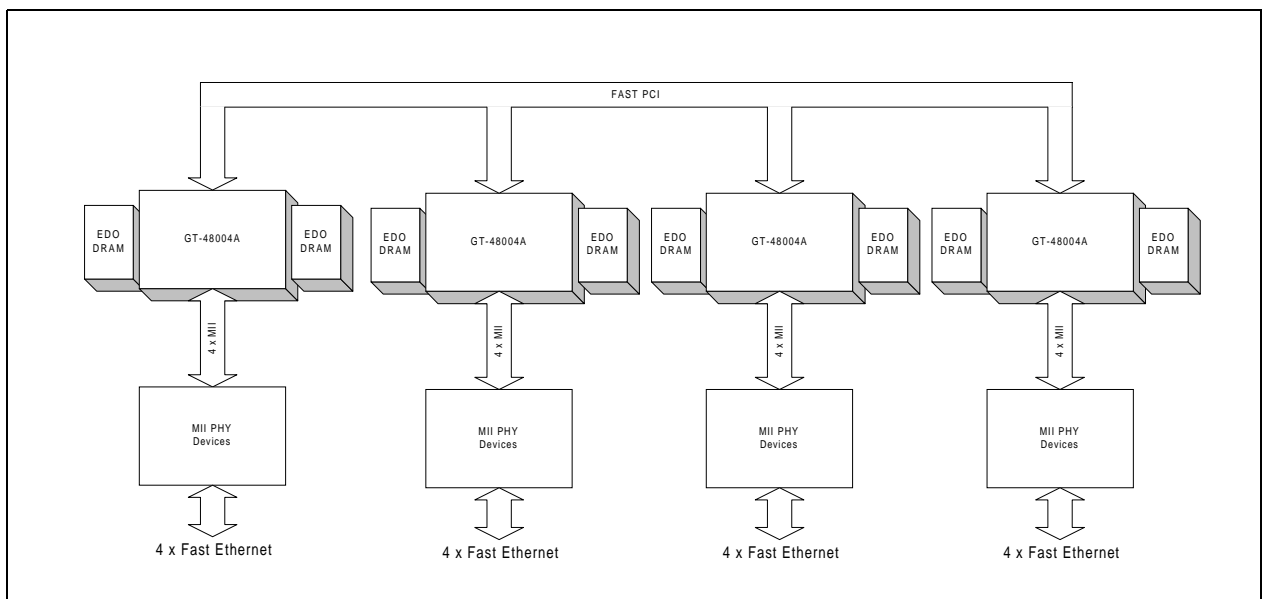


TABLE OF CONTENTS

1. Functional Overview	7
1.1 The GalNet Switching Architecture.....	7
1.2 Fast Ethernet Ports.....	7
1.3 Address Recognition	8
1.4 CPU Packet Routing.....	8
1.5 Intervention Mode.....	8
1.6 VLAN Support.....	8
1.7 Quality-of-Service/Priority Support	9
1.8 Network Management Features	9
1.9 DRAM Interface	9
1.10 Fast PCI Interface.....	9
2. Pin Information	10
2.1 Logic Symbol.....	10
2.2 Pin Functions and Assignment	11
3. Internal Architectural Overview	16
3.1 Internal Block Diagram	16
3.2 Fast Ethernet Unit Block Diagram	17
3.3 Packet Forwarding in the GT-48004A	17
4. Operational Overview	18
4.1 Enabling/Disabling the GT-48004A	18
4.2 Basic Operation	18
4.3 Address Learning	19
4.4 Packet Buffering	19
4.5 Packet Forwarding.....	19
4.6 The GalNet Protocol	19
4.7 Terminology.....	19
5. MAC Address Learning Process	20
5.1 Address Recognition	20
5.2 Recovery Process	20
5.3 Address Aging	21
5.4 Static Addresses.....	21
5.5 Address Recognition Failure	21
6. GT-48004A Buffers and Queues	22
6.1 Rx Buffer Threshold Programming	23
7. Packet Forwarding	24
7.1 Forwarding a Unicast Packet to a Local Port	24
7.2 Forwarding a Unicast Packet to a Port in a Different GalNet Device or FEU	24
7.3 Forwarding a Multicast Packet	25
7.3.1 Local Ports.....	25
7.3.2 Between GalNet Devices or FEUs.....	25
7.3.2.1 CPU Disabled.....	25

7.3.2.2 CPU Enabled.....	25
7.4 Forwarding a Packet to the CPU Directly	26
7.5 Forwarding a Packet from the CPU to a GalNet Device	28
7.6 CRC Generation	28
7.7 Tx Watchdog Timer.....	29
8. Device Table Operation	30
8.1 Automatic Device Table Initialization	30
8.2 Manual Device Table Initialization	30
8.3 Programming Device Numbers	30
9. Unicast Intervention Mode	31
9.1 Unicast Intervention Mode Address Space.....	32
10. Address Table	33
11. GalNet Messaging Protocol	35
11.1 GalNet Protocol Region	35
11.2 GalNet Messages Between Devices.....	37
11.2.1 NEW_ADDRESS Message between GalNet devices	37
11.2.2 BUFFER_REQUEST Message between GalNet devices	38
11.2.3 START_OF_PACKET Message between GalNet devices.....	38
11.2.4 PACKET_TRANSFER Message between GalNet devices	39
11.2.5 END_OF_PACKET Message between GalNet devices.....	39
11.3 GalNet Messages Between a GalNet Device and a CPU.....	40
11.3.1 NEW_ADDRESS Message (GalNet to CPU).....	40
11.3.2 NEW_ADDRESS Message (CPU to GalNet).....	41
11.3.3 BUFFER_REQUEST Message (GalNet to CPU).....	42
11.3.4 BUFFER_REQUEST Message (CPU to GalNet).....	42
11.3.5 START_OF_PACKET Message (GalNet to CPU)	43
11.3.6 START_OF_PACKET Message (CPU to GalNet)	43
11.3.7 PACKET_TRANSFER Message (GalNet to CPU 16 Block Buffer)	44
11.3.8 PACKET_TRANSFER Message (GalNet to CPU in Unicast Intervention Mode)	44
11.3.9 PACKET_TRANSFER Message (CPU to GalNet).....	45
11.3.10 END_OF_PACKET Message (GalNet to CPU 16 Block Buffer)	46
11.3.11 END_OF_PACKET Message (GalNet to CPU in Unicast Intervention Mode)	46
11.3.12 END_OF_PACKET Message (CPU to GalNet).....	47
12. Fast PCI Bus Operation	48
12.1 Separate Logical PCI Interfaces for Each FEU.....	48
12.2 Interfacing Management Processors to Fast PCI	48
12.3 PCI Configuration Header Registers.....	48
12.4 Accessing DRAM and Internal Registers through the PCI Interface.....	48
12.5 Fast PCI Bandwidth/Performance Issues	48
12.6 Plug-and-Play Considerations In PCI Systems.....	49
12.7 PCI Bus in Stand-Alone Systems	49
12.8 PCI Bus Arbiter	50
13. Fast Ethernet Interfaces	51
13.1 10/100 MII Compatible Interface.....	51

13.2 Media Access Control (MAC)	51
13.3 Auto-negotiation	51
13.3.1 Disabled	51
13.3.2 Enabled	51
13.3.3 Auto-negotiation Control Per Port	52
13.3.4 Auto-Negotiation, Software Detection	53
13.4 Backoff Algorithm Options	53
13.5 Data Blinder	53
13.6 Inter-Packet Gap (IPG)	53
13.7 10/100 Mbps MII Transmission	53
13.8 10/100 Mbps MII Reception	54
13.9 10/100 Mbps Full-Duplex Operation	55
13.10 Illegal Frames	55
13.11 Partition Mode	55
13.11.1 Enabling Partition Mode	55
13.11.2 Entering Partition State	55
13.11.3 Exiting from Partition State	55
13.12 Back-pressure	55
13.13 VLAN Tagging Support	56
14. MII Management Interface (SMI)	57
14.1 SMI Cycles	57
14.1.1 SMI Timing Requirements	57
14.2 Link Detection and Link Detection Bypass (ForceLinkPass*)	59
15. Network Management Support	60
15.1 Repeater MIB and PCI Counters	60
15.2 Monitoring (Sniffer) Mode	60
15.3 Spanning Tree Support	60
15.4 Broadcast Storm Filtering	61
16. HP-EASE Packet Sampling Technology	62
16.1 HP EASE Technology Overview	62
16.2 EASE Functionality on the GT-48004A	63
16.3 Ease_Register	63
16.4 EASE Interrupts	63
16.5 Sampled Packet Indication	64
16.6 Error Source Indications	65
16.7 Enabling/Disabling EASE Functionality	66
16.8 Interaction With Other GT-48004A Features	66
17. DRAM Interface and Usage	67
18. LED Support	68
18.1 LED Indications Interface Description	68
18.2 Detailed LED Signal Description	68
18.2.1 Primary Port Status LED	68
18.2.1.1 Primary Port Status LED in Mode 0: (LEDMode input is LOW)	68
18.2.1.2 Status LED blink timing (Mode 0)	68
18.2.1.3 Primary Port Status LED (Mode 1): (LEDMode input is HIGH)	69
18.2.2 Transmit data in progress	70

18.2.3 Receive data in progress.....	70
18.2.4 Collision active	70
18.2.5 Full/Half duplex.....	70
18.2.6 Receive Buffer Full	70
18.2.7 Forwarding of unknown packets enabled.....	70
18.2.8 The port is configured as Sniffer	70
18.2.9 Link Fail State.....	70
18.2.10 Partition State.....	70
18.2.11 Secondary Port Status LED	70
18.2.12 Pure Port Status LED	70
18.3 LED Signals Timing Type	71
18.3.1 Static LED Signals.....	71
18.3.2 Dynamic Internal Signals:.....	71
18.4 Serial LED Interface Description.....	71
18.4.1 Table of Internal Activities/Status Driven via the Serial LED Interface.....	72
19. Interrupts	73
<hr/>	
20. RESET Configuration	73
<hr/>	
20.1 Configuration Pins.....	73
20.2 Configuration Input Timings	73
21. Switch Expansion	75
<hr/>	
22. Development Tools	76
<hr/>	
22.1 Evaluation Platforms/Reference Designs	76
22.2 Verilog Models	76
22.3 Complimentary Products.....	76
23. Register Tables	77
<hr/>	
23.1 Register Map.....	77
23.2 Internal Control Registers	78
23.3 Port MIB Counters (2 Blocks), Offset: 0x040000 - 0x0400ac	87
23.4 PCI Global Counters, Offset: 0x140040 - 0x140044.....	91
23.5 SMI Register, Offset: 0x14004c.....	91
23.6 PCI Configuration Registers	91
24. GT-48004A PINOUT TABLE, 329 pin BGA (sorted by ball number)	95
<hr/>	
25. DC Characteristics - PRELIMINARY/SUBJECT TO CHANGE	98
<hr/>	
25.1 Absolute Maximum Ratings	98
25.2 Recommended Operating Conditions.....	98
25.3 DC Electrical Characteristics Over Operating Range	98
25.4 Thermal Data	99
26. AC Timing - PRELIMINARY/SUBJECT TO CHANGE	100
<hr/>	
27. Functional Waveforms	103
<hr/>	

27.1 PCI Read/Write Cycle.....	104
--------------------------------	-----

28. Packaging	105
----------------------	------------

1. Functional Overview

The GT-48004A is a high-performance, low-cost, Switched Fast Ethernet Controller that provides packet switching functions between four, on-chip, 10/100Mbps, auto-negotiated ports and the 2 Gbps Fast PCI backplane. The GT-48004A uses the innovative GalNet® Switching Architecture to allow expansion to additional Ethernet and Fast Ethernet ports.

The GT-48004A is a higher-speed derivative of the two-port GT-48002A. Two GT-48002A functional blocks are integrated into the GT-48004A and the following enhancements are added:

- The DRAM bus bandwidth is doubled from the GT-48002A (35 ns EDO DRAMs are used versus 70 ns on the GT-48002A)
- A 66MHz Fast PCI bus is used for interconnect on the GT-48004A versus a standard 33MHz PCI bus on the GT-48002A. Galileo will be offering the GT-64120 device which will bridge between 64-bit MIPS processors and Fast PCI.

From a software standpoint, the GT-48004A looks like two independent GT-48002A devices, thus making software migration extremely simple for customers already using GalNet® Architecture Devices.

The GT-48004A integrates two GT-48002A Fast Ethernet Switch units, and each unit is given separate access to the PCI bus. This means that from a hardware point of view, the GT-48004A looks like two GT-48002A devices sharing the “bussed” PCI signals. Non-bussed PCI signals (like IdSel*) are brought to separate package pins. See below for more details.

Some features were deleted from the GT-48002A functionality in the GT-48004A. The parallel LED interface was removed due to pinout constraints. The RMON FIFO support was also removed; primarily because it is almost never used and it takes 2 pins. All RMON/SNMP counter functionality remains in the GT-48004A.

1.1 The GalNet Switching Architecture

The GalNet Switching Architecture is based on a proprietary messaging protocol using the industry standard PCI bus as a medium. GalNet devices are designed to connect seamlessly allowing packets to be switched between devices without processor intervention (see Figure 1). Each GalNet device acts as an intelligent agent, sharing information between all other devices in the system. For example, when one GalNet device learns a new address, it automatically updates all other GalNet devices via the NEW_ADDRESS message. GalNet messages are defined as *write-only* (request/response) in order to achieve the maximum bandwidth from the PCI bus.

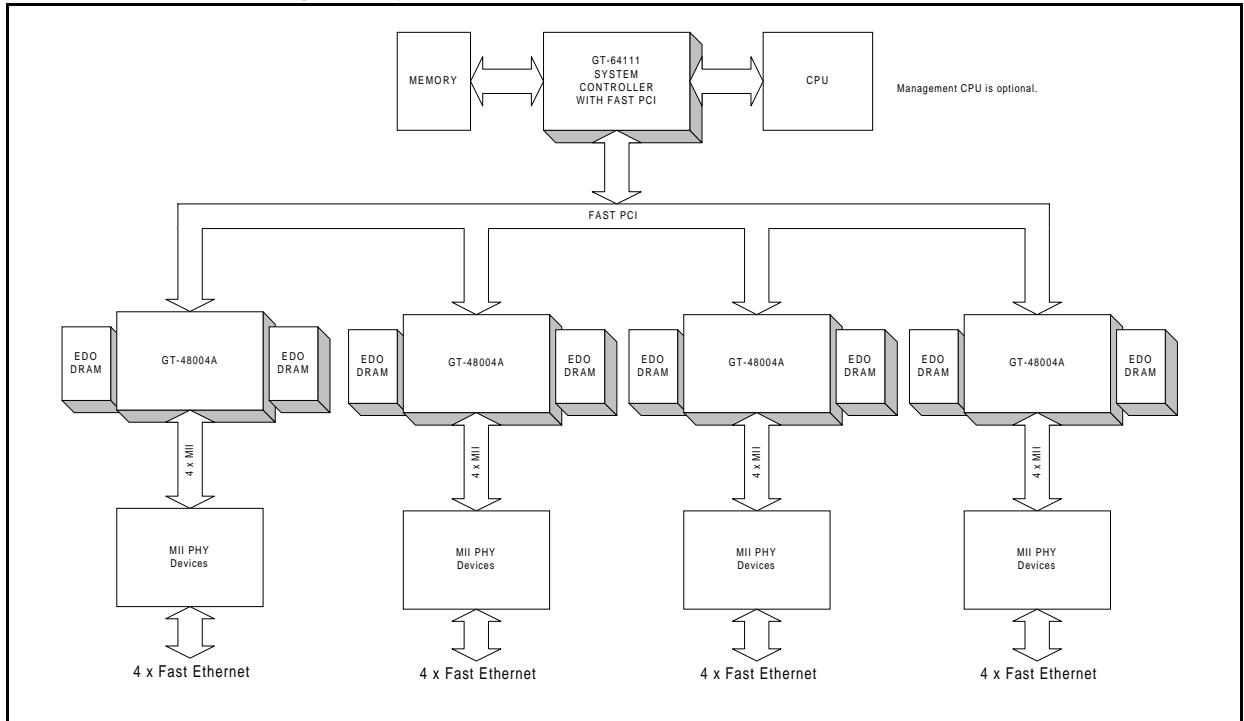
The GalNet Architecture Family currently consists of three products: the GT-48001A (eight ports of 10BaseX), the GT-48002A (two ports of 100BaseX), and the GT-48003¹ (two ports of 100VG-AnyLAN). In addition, Galileo Technology provides a number of other complementary PCI interface products for popular microprocessors.

1.2 Fast Ethernet Ports

The GT-48004A integrates four Fast Ethernet ports. Each port works at 10/100Mbps (half-duplex) or 20/200Mbps (full-duplex). Four Media Independent Interfaces (MII) are provided for glueless connection to off-the-shelf PHY chips. The GT-48004A supports full auto-negotiation for capable PHYs. Therefore, the speed (10 or 100 Mbps) and duplex (half or full) which the PHY resolves to operate is automatically reported to the GT-48004A in both managed and unmanaged systems. The port can also be forced to operate in a certain mode, if so desired. Each port includes the Media Access Control function (MAC) and six LEDs for Link Status, Collision, Receive Transmit, Half/Full Duplex and Receive Buffer Full indications. The GT-48004A incorporates full MII management support. The MDC/MDIO pins are directly controlled by the CPU.

1. The GT-48003 is now obsolete.

Figure 1: Typical 16-port 10/100-Mbps Switch Implementation



1.3 Address Recognition

Each GT-48004A in a system can recognize up to 8,000 different Unicast MAC addresses and unlimited Multicast/Broadcast MAC addresses. An intelligent address recognition mechanism enables filtering and forwarding packets at full Fast Ethernet wire speed. Hardware assist for address aging and static addresses is also included.

The GT-48004A provides an address self-learning mechanism. Each device has a private address table located in its DRAM array. As the GT-48004A learns new addresses, it updates all address tables in the system via the GalNet messaging protocol.

1.4 CPU Packet Routing

The GT-48004A has the capability to automatically forward certain packets to the CPU for routing including:

- Unicast packets with a destination address tagged for the CPU
- Multicast packets
- Unknown packets (packets with MAC addresses that have not been recognized)

This gives the system designer the flexibility to decide how to handle such packets in a managed system.

1.5 Intervention Mode

The GT-48004A incorporates an enhanced feature called *intervention mode*. This feature permits software or hardware intervention in the packet routing decision mechanisms for unicast packets. When intervention mode is selected for a MAC address, the GT-48004A sends the packet to the system CPU instead of switching it as usual. This capability can be used for many functions including: layer 3 routing, security, virtual LAN support, filtering and management.

1.6 VLAN Support

The GT-48004A supports port based VLANs that allow for efficient filtering of broadcast/multicast domains as well as isolation of unicast traffic for security. Each port can be defined to be a member of a single VLAN, or as a member of a "Super-VLAN" that spans one or more simple VLANs. VLAN support in the GT-48004A is described in detail in a sep-

erate document, available on our website.

1.7 Quality-of-Service/Priority Support

The GT-48004A includes the capability to prioritize traffic based on source or destination MAC address, or input port. Two output priorities are supported. Priority support in the GT-48004A is described in detail in a separate document, available on our website.

1.8 Network Management Features

The GT-48004A provides comprehensive management capabilities enabling the switch OEM to implement a wide range of network management features.

For OEMs offering RMON capability, the GT-48004A provides per-port statistics counters and PCI traffic counters. Also included in the GT-48004A is a packet sampling capability that can be used to implement RMON using a low cost CPU. Each port has the ability to take “snapshots” of packet data and counters at programmable intervals. These samples are forwarded to the management CPU for processing. Source addresses of every errored packet are also sent to the CPU allowing switch OEM to support error counters in RMON host and matrix groups.

Also included in the GT-48004A is hardware assistance for address aging and bridge spanning tree algorithms.

1.9 DRAM Interface

Each GT-48004A device in the system requires two arrays of fast EDO DRAM (35ns). The DRAMs are used to store the incoming/outgoing packets as well as the address table and other device data structures. The interface to EDO DRAM is glueless; all signals needed to control EDO devices are included. For more information, see Section 17. The 35ns EDO DRAMs required for 66MHz operation are available from several vendors including Silicon Magic (www.simagic.com) and Mosel-Vitellic (www.mosel-vitellic.com). These high-speed DRAMs are priced at almost the same cost as standard 60ns EDOs.

NOTE: Each bank of EDO is used to service 2 Fast Ethernet Ports. Bank A supports ports 0 and 1; Bank B supports ports 2 and 3. (The GT-48002A device required only 1 bank for 2 Fast Ethernet ports.)

1.10 Fast PCI Interface

Switch expansion and access to internal management features is possible with the GT-48004A's PCI interface. The GT-48004A can be either a master initiating a PCI bus operation or a target responding to a PCI bus operation. Up to six GT-48004A devices can reside on the same PCI bus segment, forwarding packets from one port to the other without CPU intervention. By using PCI-to-PCI bridge devices¹, the switch can be expanded to up to 32 GalNet devices.

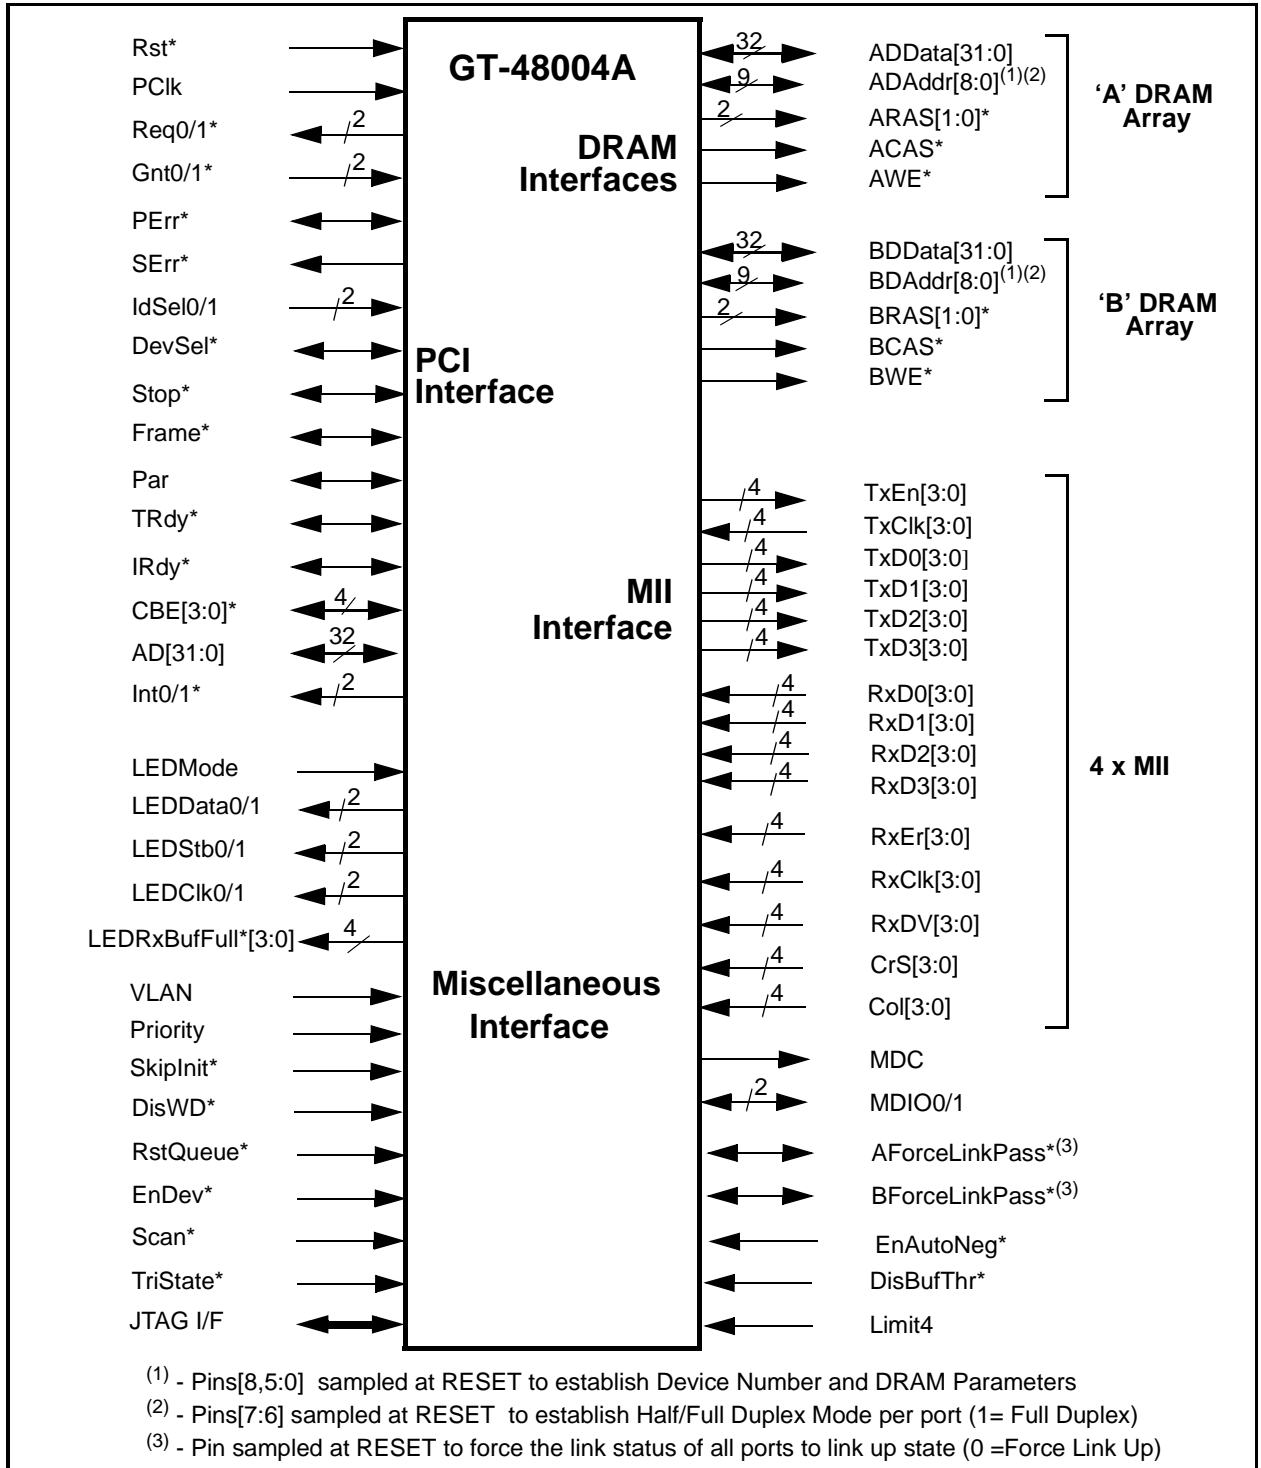
The PCI bus may also be used to connect to an optional CPU for management, or to connect other high speed LAN adapters such as ATM and FDDI.

NOTE: The GT-48004A uses a Fast PCI bus which is clocked between 33MHz and 66MHz.

1. 32 to 64 bit PCI bridges are recommended when bridging several devices for additional bandwidth. Contact Galileo for more information.

2. Pin Information

2.1 Logic Symbol



2.2 Pin Functions and Assignment

Symbol	Type	Description
PCI Bus Interface		
Rst*	I	RESET: Active LOW. Rst* must be asserted for at least 10 PCI clock cycles. When in the reset state, all PCI output pins are tristated and all open drain signals are floated. Following Rst* deassertion, the GT-48004A clears the internal buffers and initializes the address table in the DRAM. The address table initialization takes 165,000 CLK cycles to complete. Any incoming packets during the address table initialization, are ignored.
PClk	I	Clock: Provides the timing for the GT-48004A internal units. All functional units except for the serial interfaces use this clock. Clk also provides timing for PCI bus transactions. The clock frequency is 50MHz (target of 66MHz).
Req0/1*	O	Bus Request0/1: Asserted by the individual 2 port Fast Ethernet units within GT-48004A to indicate to the PCI bus arbiter that this unit requires mastership of the bus. <i>Note that even a 4 port single chip system will need to provide external PCI arbitration between the two on-chip units.</i>
Gnt0/1*	I	Bus Grant 0/1: Indicates to the individual Fast Ethernet units GT-48004A that access to the PCI bus is granted.
PErr*	I/O	Parity Error: Asserted when a data parity error is detected on the PCI bus.
SErr*	O	System Error: Asserted by the GT-48004A when an address parity error is detected on the PCI bus. The GT-48004A asserts SErr* two cycles after the failing address. This output features an open-collector driver.
IDSel0/1	I	Initialization Device Select 0/1: Asserted by a PCI bus master to gain access to the GT-48004A's configuration headers for each Fast Ethernet unit during configuration read/write transactions.
DevSel*	I/O	Device Select: Asserted by the target of the current PCI access. When the GT-48004A is a bus master, it expects the target to assert DevSel* within 5 bus cycles, confirming the access. If the target does not assert DevSel* within the required bus cycles, the GT-48004A aborts the cycle. As a target, the GT-48004A asserts DevSel* as a "medium speed" PCI device (two cycles after the assertion of Frame*).
Stop*	I/O	Stop: Indicates that the current target is requesting the bus master to stop the current transaction. As a master, the GT-48004A responds to the assertion of Stop* by either disconnecting, retrying, or aborting. As a target, the GT-48004A asserts Stop* to force a retry.
Frame*	I/O	Cycle Frame: Asserted by the GT-48004A to indicate the beginning and duration of a master transaction. Frame* is asserted to indicate the beginning of the cycle. While Frame* is asserted, data transfer continues. Frame* is deasserted to indicate that the next data phase is the final data phase transaction. Frame* is monitored when the GT-48004A acts as a target, to detect a configuration or memory transaction.
Par	I/O	Parity: Calculated by the GT-48004A as an even parity bit for the AD[31:0] and CBE[3:0]* lines.
TRdy*	I/O	Target Ready: Indicates the target agent's ability to complete the current data phase of the transaction. A data phase is completed on any clock when both TRdy* and IRdy* are asserted. Wait cycles are inserted until both IRdy* and TRdy* are asserted together.

Symbol	Type	Description
IRdy*	I/O	Initiator Ready: Indicates the bus master's ability to complete the current data phase of the transaction. A data phase is completed on any clock when both TRdy* and IRdy* are asserted. Wait cycles are inserted until both IRdy* and TRdy* are asserted together.
AD[31:0]	I/O	Address/Data: 32-bit multiplexed PCI address and data lines. During the first clock of the transaction, AD[31:0] contains a physical byte address (32 bits). During subsequent clock cycles, AD[31:0] contains data.
CBE[3:0]*	I/O	Bus Command/Byte Enable: During the address phase of the PCI transaction, CBE[3:0]* provide the Bus Command. During the data phase, CBE[3:0]* provide Byte Enables, which determine which bytes carry valid data.
Int0/1*	O	Interrupt Request Line0/1 : Int* is asserted byand FEU in the GT-48004A when one (or more) of the bits in the Interrupt Cause register(s) are set. These outputs use an open-collector driver and can be wired-ORed in most applications.
DRAM Interfaces: 'A' interface is for Fast Ethernet Unit 0; 'B' is for Fast Ethernet Unit 1		
ADData[31:0]	I/O	DRAM Data for Fast Ethernet Unit 0: 32-bit EDO DRAM data bus. These signals connect directly to the data input/output pins of the DRAM devices.
BDData[31:0]	I/O	DRAM Data for Fast Ethernet Unit 1: same as above.
ADAddr[8:0]	I/O	DRAM Multiplexed Address Bus for Fast Ethernet Unit 0: In normal operation, ADAddr[8:0] contain the DRAM multiplexed row/column address. During RESET, these multiplexed pins are sampled by the GT-48004A to indicate the Device Number and the DRAM Parameters (see RESET Configuration Section 20.) Values are determined by connecting pull-up/pull-down resistors. The Device Number and the DRAM Size are read by the CPU from the Status register.
BDAAddr[8:0]	I/O	DRAM Multiplexed Address Bus for Fast Ethernet Unit 1: same as above.
ARAS[1:0]*	O	Row Address Strobes for Fast Ethernet Unit 0: DRAM row address strobes. ARAS[0]* is used for Bank 0. ARAS[1]* is used for Bank 1.
BRAS[1:0]*	O	Row Address Strobes for Fast Ethernet Unit 1: same as above.
ACAS*	O	Column Address Strobe for Fast Ethernet Unit 0: DRAM column address strobe. The GT-48004A always accesses 32-bit values and does not require a separate ACAS* for each byte.
BCAS*	O	Column Address Strobe for Fast Ethernet Unit 1: same as above.
AWE*	O	Write Enable for Fast Ethernet Unit 0: DRAM write enable.
BWE*	O	Write Enable for Fast Ethernet Unit 1: same as above.
Media Independent Interface (MII)		
TxEn[3:0]	O	Transmit Enable: Active HIGH. This output indicates that the packet is being transmitted. TxEn is synchronous to TxClk.
TxClk[3:0]	I	Transmit Clock: Provides the timing reference for the transfer of TxEn, TxD signals. TxClk frequency is one fourth of the data rate (25 MHz for 100Mbps, 2.5 MHz for 10Mbps). TxClk nominal frequency should match the nominal frequency of RxClk for the same port.
TxD0[3:0]	O	Transmit Data 0: Outputs the Port0 Transmit Data. Synchronous to TxClk[0].
TxD1[3:0]	O	Transmit Data 1: Outputs the Port1 Transmit Data. Synchronous to TxClk[1].
TxD2[3:0]	O	Transmit Data 2: Outputs the Port2 Transmit Data. Synchronous to TxClk[2].

Symbol	Type	Description
TxD3[3:0]	O	Transmit Data 3: Outputs the Port3 Transmit Data. Synchronous to TxClk[3].
Col[3:0]	I	Collision Detect: Active HIGH. Indicates a collision has been detected on the wire. This input is ignored in full-duplex mode, and in half-duplex mode when TxEn of the same port is LOW. Col is not synchronous to any clock.
RxD0[3:0]	I	Receive Data 0: Port 0 Receive Data. Synchronous to RxClk[0].
RxD1[3:0]	I	Receive Data 1: Port 1 Receive Data. Synchronous to RxClk[1].
RxD2[3:0]	I	Receive Data 2: Port 0 Receive Data. Synchronous to RxClk[2].
RxD3[3:0]	I	Receive Data 3: Port 1 Receive Data. Synchronous to RxClk[3].
RxEr[3:0]	I	Receive Error. Active HIGH. Indicates that an error was detected in the received frame. This input is ignored when RxDV for the same port is inactive.
RxCk[3:0]	I	Receive Clock. Provides the timing reference for the transfer of the RxDV, RxD, RxEr signals (per port). Operates at either 25 MHz (100Mbps) or 2.5 MHz (10Mbps). The nominal frequency of RxClk (per port) should match the nominal frequency of that port's TxClk.
RxDV[3:0]	I	Receive Data Valid: Active HIGH. Indicates that valid data is present on the RxD lines. Synchronous to RxClk. This input is ignored when it represents loopback of the transmitted packet in 10BaseT mode half-duplex.
CrS[3:0]	I	Carrier Sense: Active HIGH. Indicates that either the transmit or receive medium is non-idle. CrS is not synchronous to any clock.
MDC	O	Management Data Clock: 1 MHz clock. Provides the timing reference for the transfer of the MDIO 0/1 signal. This output may be connected to the PHY devices of all ports.
MDIO0/1	I/O	Management Data Input/Output 0/1: This bidirectional line is used to transfer control information and status between the PHY and the GT-48004A. It conforms with IEEE Std 802.3. This signal may be connected to the PHY devices of both ports. When not in use, this pin must be connected to a pull-down resistor. Each 2 port Fast Ethernet unit has its own MDIO line.
Miscellaneous Interface Pins		
EnAutoNeg*	I	Enable Auto-negotiation: Active LOW. The GT-48004A controls the auto-negotiation process and configures both ports to the correct speed and duplex as resolved by each port's PHY. When HIGH, auto-negotiation is disabled and the duplex setting of both ports is set based on the RESET configuration. See Section 13.3.3 for more information on Auto-negotiation Control Per Port.
AForceLinkPass*	I/O	Fast Ethernet Unit 0 ("A") Force Link Pass: Active LOW. This pin is sampled on Rst*. When connected HIGH, the link status of ports 0 and 1 is read through the SMI (MDC/MDIO interface) from the PHY devices (register#1, bit#2). When connected LOW, the link status of ports 0 and 1 remains in the "link is up" state regardless of the PHY's link bit value. This pin should be connected to either a pull-up (normally) or a pull-down resistor (to force the link pass). Following Rst* deassertion, this pin becomes an output (unused - value is undefined).

Symbol	Type	Description															
BForceLinkPass*	I/O	Fast Ethernet Unit 1 (“B”) Force Link Pass: Active LOW. This pin is sampled on Rst*. When connected HIGH, the link status of the ports 2 and 3 is read through the SMI (MDC/MDIO interface) from the PHY devices (register#1, bit#2). When connected LOW, the link status of ports 2 and 3 remains in the “link is up” state regardless of the PHY’s link bit value. This pin should be connected to either a pull-up (normally) or a pull-down resistor (to force the link pass). Following Rst* deassertion, this pin becomes an output (unused - value is undefined).															
LEDMode	I	LED Mode Select: Affects Port status LED, LEDClk frequency and LED ON time values. For more information, see Section 18. 0 - select LEDMode 0 1 - select LEDMode 1															
LEDData 0/1	O	LED Data: LED indicators (Link Status, Receive, Transmit, Collision, Unknown, Port Sniffer, and Half/Full Duplex) of each port. The data is shifted out in 128 bit long frames using the LEDClk and LEDStb pins. Each 2 port Fast Ethernet unit has its own LEDData pin.															
LEDStb 0/1	O	LED Strobe: Indicates the beginning of valid data frame on the LEDData pin. Each 2 port Fast Ethernet unit has its own LEDStb pin.															
LEDClk0/1	O	LED Clock 0/1: 1 MHz clock (at LEDMode 0), 202 KHz clock (at LEDMode 1). This output is used to clock the LEDStb and LEDData outputs. During RESET, LEDClk is tristated. Each 2 port Fast Ethernet unit has its own LED clock.															
LEDRxBufFull*[3:0]	O	LED Receive Buffer Full: Active LOW LED output Indicates (per port) receive buffer full condition (programmable limit exceeded). An external driver is required to drive LEDs.															
RstQueue*	I	Reset Transmit Queues: When asserted, all internal transmit and receive queues are cleared. All GT-48004A state machines are moved to their initial state.															
EnDev*	I	Enable Device: Enables serial and PCI ports. When asserted LOW, all serial ports and the PCI port are active. When deasserted, the ports and the PCI are disabled (see Section 4.1).															
Scan*	I	<p>Scan: This pin together with TriState* indicate the GT-48004A mode of operation as follows:</p> <table border="1" data-bbox="618 1290 1274 1551"> <thead> <tr> <th>Scan*</th> <th>TriState*</th> <th>Mode</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>1</td> <td>Normal operation</td> </tr> <tr> <td>0</td> <td>1</td> <td>Factory test mode (reserved)</td> </tr> <tr> <td>1</td> <td>0</td> <td>The GT-48004A drives all outputs and I/O pins to high impedance.</td> </tr> <tr> <td>0</td> <td>0</td> <td>Factory test mode (reserved)</td> </tr> </tbody> </table> <p>Factory test modes are reserved and are not to be used in-system. Failure to observe this restriction could result in damage to the device.</p>	Scan*	TriState*	Mode	1	1	Normal operation	0	1	Factory test mode (reserved)	1	0	The GT-48004A drives all outputs and I/O pins to high impedance.	0	0	Factory test mode (reserved)
Scan*	TriState*	Mode															
1	1	Normal operation															
0	1	Factory test mode (reserved)															
1	0	The GT-48004A drives all outputs and I/O pins to high impedance.															
0	0	Factory test mode (reserved)															
TriState*	I	Tri-State: This pin together with Scan* indicate the GT-48004A mode of operation as described above.															

Symbol	Type	Description
DisWD*	I	Disable Watchdog Timer: Active LOW. DisWD* controls the enabling (HIGH) or disabling (LOW) of the Tx Watchdog Timer on all ports.
DisBufThr*	I	Buffer Threshold: Active LOW. This pin externally enables or disables the buffer threshold. When HIGH, the buffers allocated to the ports and the PCI are limited to the number written in the Rx Buffer Threshold Register. When LOW, the buffers are dynamically allocated to the ports and the PCI bus (i.e. there is no limitation on the buffers' allocation.)
Limit4	I	Backoff Algorithm: This pin selects the number of retransmit attempts after a collision will occur before the back-off algorithm is restarted. When LOW, 16 retransmit attempts after a collision must occur before the back-off algorithm is restarted (802.3 standard). When HIGH, 4 retransmit attempts after a collision must occur before the back-off algorithm is restarted (more aggressive.)
SkipInIt*	I	Skip Initialization Stage: Active LOW. This pin controls the initialization stage of the GT-48004A. When asserted, the GT-48004A skips the initialization stage (clearing the address table which is stored in the DRAM), upon the deassertion of Rst*. This pin is typically used for testing and the default state is to pull this pin HIGH.
VLAN	I	VLAN Enable: Active HIGH. This pin enables VLAN support in the GT-48004A when connected to VCC. Connect to GND to disable VLAN mode.
Priority	I	Priority Enable: Active HIGH. This pin enables priority support in the GT-48004A when connected to VCC. Connect to GND to disable priority mode.
JTAG		
JTRST*	I	JTAG Reset: Asynchronous reset to test logic.
JTCLK	I	JTAG Clock: Clock for test logic. JTMS and JTDI are received on the rising edge, JTDO is driven from the falling edge. This signal determines the shifting rate.
JTMS	I	JTAG Mode Select: A broadcast signal which controls test logic operation.
JTDI	I	JTAG Data In: Serial data input.
JTDO	O	JTAG Data Out: Serial data output. Tri-state changes on negative change of JTCLK.

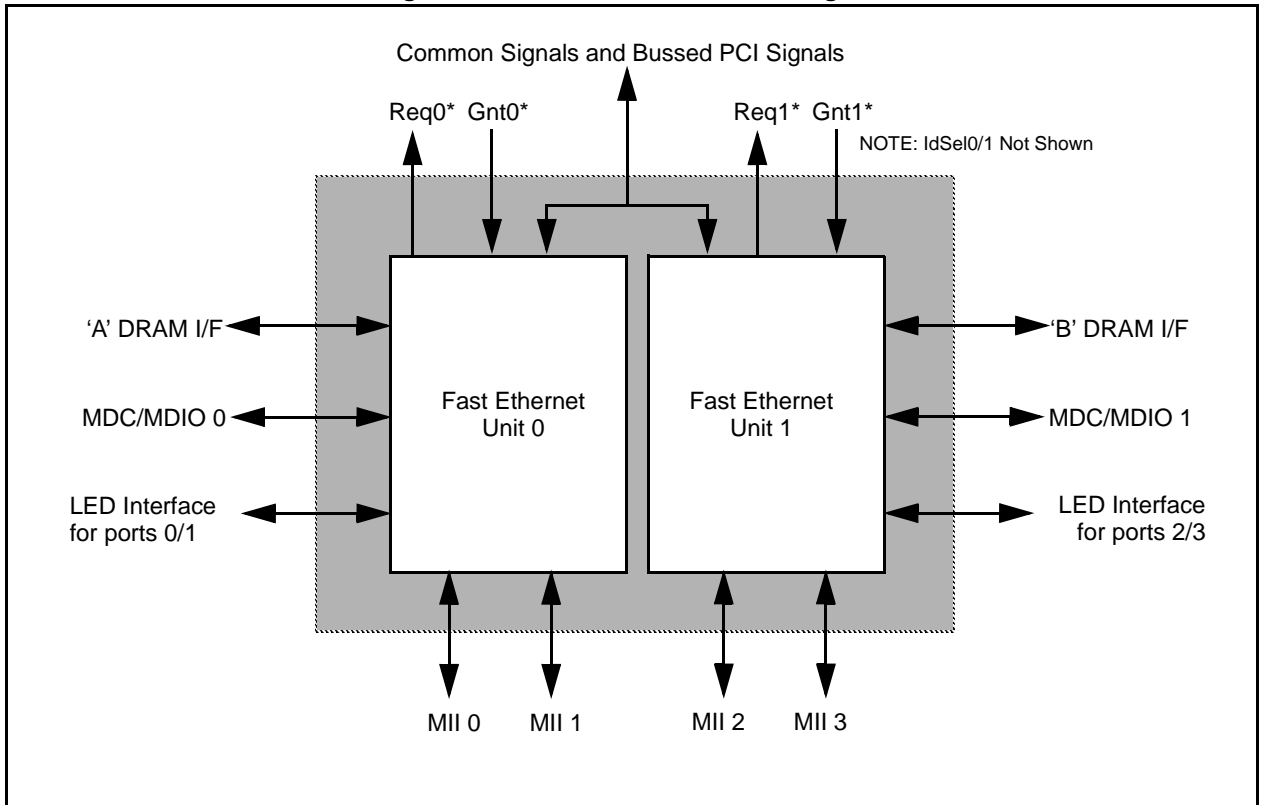
3. Internal Architectural Overview

This section describes the internal architecture of the GT-48004A. Knowledge of Galileo's other GalNet devices is helpful (but not required) to understand this section.

3.1 Internal Block Diagram

The GT-48004A is internally divided into two separate Fast Ethernet Units (FEUs). Each FEU has its own DRAM, switching engine, MACs, etc. The two FEUs are tied together on-chip only via the PCI bus. (Customers familiar with the GT-48002A will recognize that the GT-48004A is essentially two GT-48002A's integrated into a single chip and "shrunk" onto a faster 0.35 micron process.) Figure 2 shows a simplified internal block diagram.


Figure 2: GT-48004A Internal Block Diagram



3.2 Fast Ethernet Unit Block Diagram

Figure 3 shows a block diagram of each FEU.


Figure 3: FEU Block Diagram



3.3 Packet Forwarding in the GT-48004A

Packets within the same FEU are forwarded internally in the GT-48004A. Packets that must cross from one FEU to another are forwarded across the Fast PCI bus. This means that designs using a single GT-48004A for four Fast Ethernet ports must still supply PCI clock and must implement a simple external arbiter and must properly tie-off all PCI signals (see Figure 4 and see Section 12.7 on page 49.)

Figure 4: External PCI Arbiter in 4 Port System



4. Operational Overview

The GalNet Architecture Family of switching devices has been defined as an extensible, scalable architecture for the switching of packetized data. The GalNet Architecture Family currently supports Ethernet (GT-48001A), Fast Ethernet (GT-48002A and GT-48004A), and 100VG-AnyLAN (GT-48003.)

All GalNet Architecture Family devices act as distributed intelligent agents within a switching system. Each GalNet device makes switching decisions independent of other devices in the system, and can communicate information regarding the network to all other agents. This distributed processing approach is a significant performance improvement over switching architectures that rely on centralized switching “engines” or single-point address recognition devices. Unlike centralized resource approaches, GalNet designs can actually add packet processing capability as additional ports are added.

The GalNet Architecture Family uses a “store-and-forward” switching approach. Store-and-forward was chosen for the following reasons:

- Store-and-forward switches allow switching between differing speed media (e.g. 10BaseX and 100BaseX.) Such switches require the large elastic buffers that are provided by the EDO DRAM arrays.
- Store-and-forward switches improve overall network performance by acting as a “network cache”, effectively buffering packets during periods of heavy congestion.
- Store-and-forward switches prevent the forwarding of corrupted packets by analyzing the frame check sequence (FCS) before forwarding to the destination port.

A typical unmanaged GalNet Architecture system is extremely simple to implement as shown in Figure 1 on page 8. No CPU is needed as each GalNet device is intelligent and capable of sharing network information and packet data autonomously. A CPU may be added to provide network management capability.

4.1 Enabling/Disabling the GT-48004A

Ports of the GT-48004A can be enabled and disabled depending on the combination of:

- An external hardware pin, EnDev* (LOW - device enabled, HIGH - device disabled)
- EnableDevice, bit 27 of the Global Control Register ('0' - device status based on EnDev*, '1' - device enabled)
- PortEn, bit 0 of each Port Control Register, ('0' - port disabled, '1' - port enabled)

When a port is disabled, no packets will be received or transmitted from the serial ports or the PCI bus. Even though ports are disabled, another PCI master can read from and write to the GT-48004A's registers. See Table 1 for enabling or disabling ports of the GT-48004A.

Table 1: Enabling/Disabling Ports of the GT-48004A

EnDev* Pin	EnableDevice Bit	PortEn Bit	Port Status
LOW	0	0	Disabled
LOW	0	1	Enabled
LOW	1	0	Disabled
LOW	1	1	Enabled
HIGH	0	0	Disabled
HIGH	0	1	Disabled
HIGH	1	0	Disabled
HIGH	1	1	Enabled

4.2 Basic Operation

The basic operation of the GT-48004A is quite simple. The GT-48004A receives incoming packets from the Fast Ethernet wire, searches in the Address Table for the Destination MAC Address and then forwards the packet to the appropriate port. The destination port can be either be local (one of the GT-48004A's ports) or in a different GT-48004A device that resides on the same PCI bus. If the destination address is not found, the GT-48004A treats the packet as a multi-

cast packet and forward the packet to all ports of all devices in the system specified to forward unknown packets.

The GT-48004A automatically learns the port number of attached network devices by examining the Source MAC Address of all incoming packets. If the Source Address is not found in the GT-48004A's Address Table, the device adds it to the table (with an indication of on which port the address resides). The GT-48004A then notifies other GalNet devices in the system of the new address via a NEW_ADDRESS message.

4.3 Address Learning

The GT-48004A can learn up to 8K unique MAC addresses. Addresses are stored in the Address Table located in DRAM. The Address Table is managed automatically by the GT-48004A (i.e. a new address is automatically added to the Address Table). The GT-48004A's address learning process is outlined in Section 5.

The Address Table includes information regarding target port, aging status, static/dynamic status, and flags to force processor intervention. The management CPU has the ability to insert, remove or modify the entries.

4.4 Packet Buffering

Incoming packets are buffered in the DRAM array. These buffers provide elastic storage for transferring data between low-speed and high-speed segments. The packet buffers are managed automatically by the GT-48004A.

4.5 Packet Forwarding

Once an address has been learned, and the packet is buffered, it must be forwarded. The packet forwarding mechanism for the GT-48004A is handled automatically based on the destination address. Optionally, the CPU can be involved in unicast packet forwarding decisions by using *intervention mode*. If a CPU is utilized for system management functions, multicast packets will be forwarded to the CPU for forwarding decisions.

4.6 The GalNet Protocol

The GT-48004A uses a proprietary inter-chip communication protocol on the PCI bus known as the GalNet Protocol messages. The protocol consists of five groups of messages: NEW_ADDRESS, BUFFER_REQUEST, START_OF_PACKET, PACKET_TRANSFER, and END_OF_PACKET.

All GalNet messages are *write-only*. For example, a GT-48004A may request a buffer location in another GalNet device by writing a BUFFER_REQUEST message to the target device. The target device responds by writing a START_OF_PACKET message to the requesting GT-48004A. Read transactions are strictly avoided since they tend to stall the PCI bus, thereby wasting precious bandwidth.

4.7 Terminology

It is important to understand the basic terminology used to describe the GalNet Architecture Family before getting into the detailed description. Table 2 explains the terms used throughout this document.

Table 2: Terminology

Term	Definition
Address Table	The Address Table is a data structure in the GT-48004A's DRAM that contains all learned MAC addresses, and routing information associated with those addresses.
Source Address	The Source Address (SA) is the MAC address from which a received packet was sent.
Destination Address	The Destination Address (DA) is the MAC address to which a received packet was sent.
Device Number	Each GalNet device in a system (including the CPU, if any) has a specific Device Number. There are 32 possible Device Numbers.
Port Number	Each Ethernet port on a GalNet device has an associated port number. The GalNet device associates Port Numbers with the MAC addresses located on those ports.

5. MAC Address Learning Process

The GT-48004A has a self-learning mechanism for learning the MAC addresses of attached Fast Ethernet devices in real time. The GT-48004A searches for the Source Address (SA) of an incoming packet in the Address Table and acts as follows:

If the SA was not found in the Address Table (a new address), the GT-48004A waits until the end of the packet (Non errored packet) and updates the Address Table. It also notifies the other GT-48004A devices and the CPU by sending a separate NEW_ADDRESS message to each GalNet device on the PCI bus. If a CPU is enabled in the system (bit 10 of the Global Control Register, 0x140028), the NEW_ADDRESS message can optionally be forwarded to the CPU (bit 7 in the Global Control Register, 0x140028). This message contains the new MAC address, the Device Number and the Port Number.

29. If the SA was found, the GT-48004A compares all fields of the NEW_ADDRESS message to the entry in the Address Table. If any fields differ (and the 'Static' bit in the Address Table is '0'¹), the GT-48004A updates the entry with the new information in the NEW_ADDRESS message and notifies the other GT-48004A devices and the CPU. If the device and port numbers are equal, the packet is not switched (i.e. this packet was destined for a device on the same network segment and does not require forwarding.)

30. If the SA was found in the Address Table, the Aging bit is set. This is done to indicate to the aging software that this address was accessed recently.

The CPU can access the Address Table to modify, remove or to add a MAC address. This is accomplished by performing a NEW_ADDRESS message (Section 11.3.2) on the PCI bus.

5.1 Address Recognition

The GT-48004A forwards the incoming packets to the appropriate port(s) according to Destination Address (DA) as follows:

1. If the DA is a Unicast address and the address was found in the Address Table, the GT-48004A acts as follows:
 - If the Port Number and the Device Number are equal to the Port/Device on which the packet was received, the packet is discarded.
 - If the Port Number is different, but the Device Number is equal, the packet is forwarded to the appropriate local port.
 - If the Device Number is different, the packet is forwarded to the appropriate GT-48004A device via the PCI bus.
2. If the DA is a Unicast address and the address was not found (Unknown), the GT-48004A acts as if the unknown packet is a Multicast packet and forwards it to ports and the devices that have been programmed to receive unknown packets (bit 4 in the Command register).
3. If the DA is a Multicast address, the packet is forwarded to all the local ports (except for the port in which the packet was received). It is also forwarded to all the other devices via the PCI bus. This procedure is outlined in Section 7.3.

5.2 Recovery Process

The purpose of the Recovery Process is to guarantee that Address Tables entries in all the devices correlate. When the packet is Unknown, the source GT-48004A sends a NEW_ADDRESS message to all the devices. Each device searches its own Address Table for the new address. More than one device can find the address, but only one device "owns" this address (the Device Number written in the Address Table is equal to its own Device Number). This particular device updates the source GT-48004A's Address Table with the new address (by in turn sending it a NEW_ADDRESS message).

1. Static Address Table entries cannot be modified, as described in Section 5.4.

5.3 Address Aging

The GT-48004A includes hardware support for address aging, which requires the use of a CPU to be implemented. The GT-48004A gives an indication to the CPU of the relative “age” of an address by setting the Aging bit in the address table when it receives a packet. The CPU reads the Aging bit each aging period (the default from the 802.12d specification is 300 seconds) and clears the bit. If in the next period, the bit remains clear, the CPU knows that this station didn't transmit any packet in this period of time and the address can be removed from the table.

Note that the aging bit is set only in the GT-48004A device that received the packet; other GT-48004A's in the system are not notified since, by definition, the receiving device “owns” the address.

The only time the GT-48004A will delete an entry by itself is when a user changes the location of a station. The GT-48004A will then automatically learn the new location of the station the next time the station sends a packet.

5.4 Static Addresses

The GT-48004A includes support for “static” MAC addresses. IEEE 802.1d chapter 3.9.1: “Static entries may be added to and removed from the filtering database under explicit management control. They are not automatically removed by any time-out mechanism”. This means that when an address is selected to be static, it will not be removed from the address table during aging.

During normal address recognition, if an address is static, the GT-48004A will not update its address table parameters and will not send a NEW_ADDRESS message over the PCI (if the address has changed ports.)

5.5 Address Recognition Failure

It is possible that an address recognition cycle will fail when more than 8K addresses have already been entered into the address table. In the case of an address recognition failure the packet will be treated as unknown and forwarded to all ports. An interrupt is also generated to the CPU (if any.)

Address recognition failures are not fatal and do not need to be handled (e.g. designers of unmanaged systems need not worry about them.) Managed systems may want to “clean” the address table of old addresses when such a failure occurs (see “Address Aging” on page 21).

6. GT-48004A Buffers and Queues

The GT-48004A incorporates three transmit queues for the 2 Ethernet ports and the PCI bus port, and one common receive buffer area in each FEU. The receive buffers as well as the transmit queues are located in the DRAM along with the Address Table. DRAM address mapping for each FEU in the GT-48004A is shown in Table 3. The GT-48004A data structure components are the following:

- **Receive Buffer** - A common Receive Buffer area for all ports. The buffer is divided into 320 or 1008 blocks (depending on the DRAM size) of 1.5KBytes (1536 bytes) each. Each block contains the entire packet.
- **Rx Empty List** - A list of 320 or 1008 bits. Each bit contains the status of its appropriate receive block in the DRAM (empty or occupied).
- **Tx Descriptors** - A set of 3 transmit descriptor rings. Each ring contains 1024 descriptors. The descriptor size is one 32-bit word and contains the Block Address divided by 0x600 (1.5K), the byte count and the packet type (Multicast or Unicast).
- **Read/Write Pointers** - 3 pairs of pointers to the transmit descriptors.

Figure 5: GT-48004A Buffers and Queues

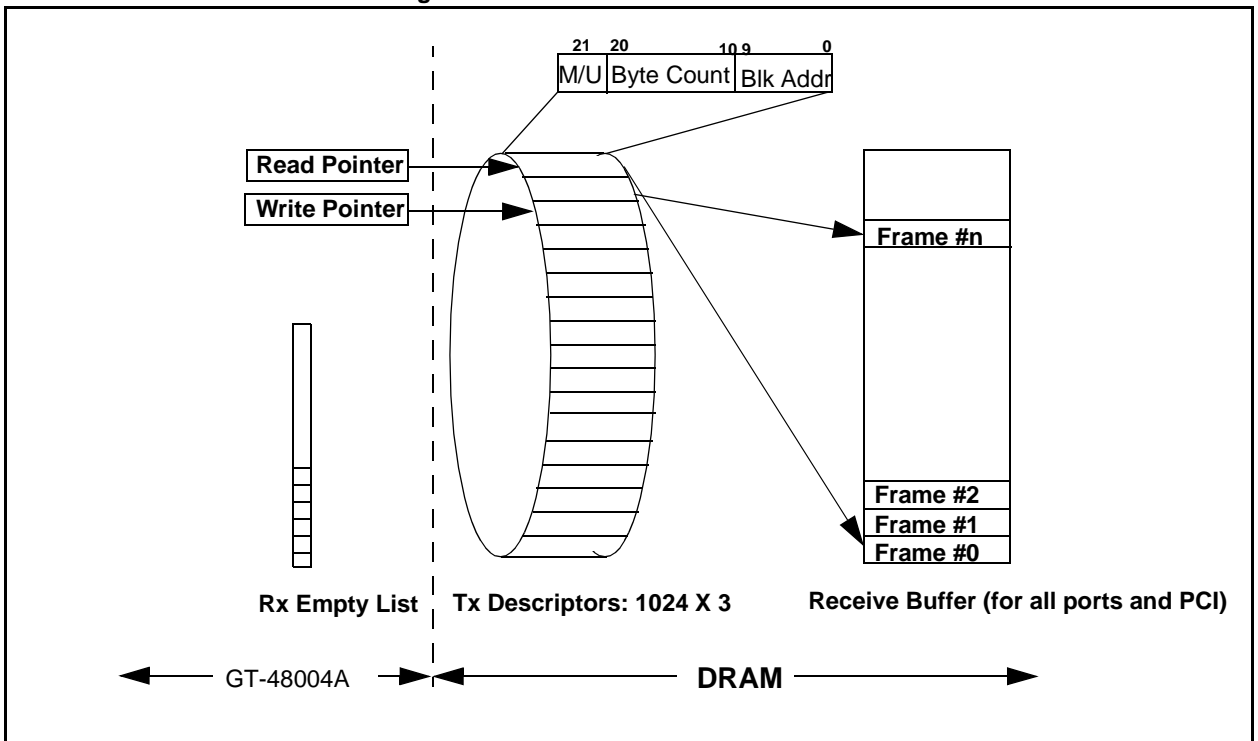


Table 3: GT-48004A DRAM Address Mapping per FEU

Memory	Description	1Mbyte	2Mbyte
Rx Buffers	320 Buffers for 1Mbyte 1008 Buffers for 2 Mbyte	0x00000 - 0x7cfff	0x00000 - 0x17cfff
PCI Tx Descriptor		0x7d000 - 0x08efff	0x17d000 - 0x18efff
PCI Rx Descriptor		0x8f000 - 0x9afff	0x18f000 - 0x19afff
Reserved	4KBytes	0x9b000 - 0x09bfff	0x19b000 - 0x19bfff
Tx Descriptor		0x09c000 - 0x09ffff	0x19c000 - 0x19ffff
Address Table		0xa0000 - 0xfffff	0x1a0000 - 0x1fffff

6.1 Rx Buffer Threshold Programming

The number of receive buffers allocated to each port of the GT-48004A is controlled by the RxBufThr field in the Rx Buffer Threshold register and the DisBufThr* pin. The default value is 80 buffers per port for 1MB DRAM and 140 buffers per port for 2MB DRAM. If the buffer threshold is disabled (by clearing the BufThrEn bit in the Global Control register or if the DisBufThr* pin is held LOW) the GT-48004A dynamically allocates the buffers to the 2 ethernet ports in each FEU and the PCI bus port. In other words, there are no limits on each buffers' allocation. The Rx Buffer Threshold value can be used to tune performance during development. See Table 4 for Rx Buffer Threshold settings.

If a received packet overflows the Rx buffer allowance, then that packet will be discarded and the Dropped Packets counter will be incremented. The overflow of the Rx buffer allowance is also indicated by the "Receive Buffer Full" LED in the Serial and Parallel LED Interfaces.

Table 4: Setting the Rx Buffer Threshold

DisBufThr* Pin	BufThrEn Bit	Result
HIGH	1	Receive Buffer Size is limited to the value of RxBufThr
HIGH	0	Dynamic Receive Buffer Allocation
LOW	1	Dynamic Receive Buffer Allocation
LOW	0	Dynamic Receive Buffer Allocation

7. Packet Forwarding

The following sections describe the procedures for forwarding packets in the following situations:

- A unicast packet to a local port in the same GalNet device (Section 7.1)
- A unicast packet between GalNet devices (Section 7.2)
- A multicast packet to local ports in the same GalNet device (Section 7.3.1)
- A multicast packet between GalNet devices in a system without a CPU (Section 7.3.2.1)
- A multicast packet between GalNet devices in a system with a CPU (Section 7.3.2.2)
- A packet destined for the CPU, multicast, or unicast packet from a GalNet device to the CPU (Section 7.4)
- A unicast/multicast packet from the CPU to a GalNet device (Section 7.5)

7.1 Forwarding a Unicast Packet to a Local Port

The sequence for forwarding a unicast packet to a local port (port in the same GT-48004A) is as follows:

1. The incoming packet is fed to the Rx FIFO (there is an 20x32-bit Rx FIFO per port) and is transferred to an empty block in the Receive Buffer area of DRAM.
2. In parallel, an address recognition cycle is performed for both the DA and the SA. The GT-48004A uses the DA's corresponding Port Number to queue the packet to the appropriate local port.
3. At the end of an error-free packet transfer, packet information is written to the appropriate port's transmit descriptor. This information includes the Byte Count and the Receive Buffer Block Address which is pointed to by the Write Pointer.
4. The Write Pointer of the outgoing port's transmit descriptor is incremented. The target GalNet device transmits whenever the Write Pointer is not equal to the Read Pointer.
5. At the end of the packet transmit process, the target GalNet device increments the Read Pointer and clears the appropriate bit in the Empty List.

7.2 Forwarding a Unicast Packet to a Port in a Different GalNet Device or FEU

The sequence for forwarding a Unicast packet to a port in a different GalNet device located on the PCI bus or in a different Fast Ethernet Unit (FEU) on the same GT-48004A, is as follows:

1. The incoming packet is fed to the Rx FIFO and is transferred to an empty block in the Receive Buffer area of DRAM.
2. In parallel, an address recognition cycle is performed for both the DA and the SA. The GT-48004A uses the DA's corresponding Port Number and Device Number to queue the packet to the appropriate GalNet device and port.
3. At the end of an error-free packet transfer, packet information is written to the PCI's transmit descriptor. This information includes the byte count and Receive Buffer Block Address which is pointed to by the Write Pointer. When the PCI's transmit descriptor's Write Pointer is not equal to the Read Pointer, the source GalNet device sends a BUFFER_REQUEST message (Section 11.2.2) to the appropriate target GalNet device indicating that there is a packet for transmission across the PCI bus.
4. The target GalNet device receives this message and allocates a buffer in its DRAM. This target device then sends a START_OF_PACKET message (Section 11.2.3) back to the source GT-48004A indicating it is ready to receive the packet.
5. The source GT-48004A transfers the packet with the PACKET_TRANSFER message (Section 11.2.4) using PCI master operations in multiple eight 32-bit bursts. The packet is buffered in the Receive Buffer area of the target device's DRAM. After the entire packet has been transmitted, the source GT-48004A performs an additional write transaction by sending the END_OF_PACKET message (Section) indicating completion of the packet transfer. This message contains the Byte Count, the target Port Number, the Rx Block address, and the Packet Type. The

source GT-48004A also clears the appropriate bit in its Empty List.

6. Some packet information included in the END_OF_PACKET message is written to the appropriate transmit descriptor in the target device. This information includes the Byte Count and the Receive Buffer address which is pointed to by the Write Pointer.
7. The Write Pointer of the outgoing port's transmit descriptor is incremented. The target GalNet device transmits whenever the Write Pointer is not equal to the Read Pointer.
8. At the end of the packet transmit process, the target GalNet device increments the Read Pointer and clears the appropriate bit in the Empty List.

7.3 Forwarding a Multicast Packet

The GT-48004A forwards Multicast packets to all local ports and devices using the same mechanism as described for Unicast packets. The GT-48004A has the ability to forward multicast packets to a management CPU for intervention routing, if desired.

7.3.1 Local Ports

For local ports in the same device, the packet is queued to all transmit ports except for the port which the packet arrived and the packet is transferred with the same procedure outlined in Section 7.1 to each port.

7.3.2 Between GalNet Devices or FEUs

Forwarding multicast packets to other GalNet devices is handled differently depending if a CPU is disabled or enabled in the system or not.

7.3.2.1 CPU Disabled

Systems which do not utilize a CPU (bit 10 of the Global Control Register, 0x140028, is not set) will automatically forward multicast packets to all of the local ports and devices with the following procedure:

1. The incoming packet is fed to the Rx FIFO and is transferred to an empty block in the Receive Buffer area of DRAM.
2. In parallel, an address recognition cycle is performed for the SA. The DA marks this packet as a multicast packet. At the end of a good packet transfer, packet is forwarded to all of the local ports according to Section 7.3.1. This multicast packet is also forwarded to the other GalNet devices in the system with the same procedure as forwarding a unicast packet from one device to another. This procedure is outlined in Section 7.2. There is a single, separate, multicast packet transfer from the source GT-48004A to each of the GalNet devices in the system. Bit 21 of Data 0 of the BUFFER_REQUEST message (Section 11.2.2) will be set to indicate this is a multicast packet.

7.3.2.2 CPU Enabled

Systems which utilize a CPU (bit 10 of the Global Control Register, 0x140028, is set), will **always** have multicast packets forwarded to the CPU. This allows the CPU to intervene, if necessary, and redirect or update the multicast packet before forwarding. Control of forwarding multicast packets to all of the ports is set by bit 22 of the Global Control Register. The default setting is to forward all multicast packets to all of the ports in the system as well as the CPU. If this bit is set, multicast packets will go *only* to the CPU.

Assuming bit 22 is not set and all multicast packets are forwarded to the CPU as well as all of the ports, the procedure for handling multicast packets is as follows:

1. The incoming packet is fed to the Rx FIFO and is transferred to an empty block in the Receive Buffer area of DRAM.

2. In parallel, an address recognition cycle is performed for both the DA and the SA. The DA marks this packet as a multicast packet. At the end of a good packet transfer, packet is forwarded to all of the local ports according to Section 7.3.1. This multicast packet is also forwarded to the other GalNet devices in the system with the same procedure as forwarding a unicast packet from one device to another. This procedure is outlined in Section 7.2. There is a single, separate, multicast packet transfer from the source GT-48004A to each of the GalNet devices in the system. Bit 21 of Data 0 of the BUFFER_REQUEST message (Section 11.2.2) will be set to indicate this is a multicast packet.
3. Packet information is also written to the PCI's transmit descriptor which instructs the GT-48004A to send this multicast packet to the CPU. This multicast packet is then forwarded directly to the CPU with the procedure outlined in Section 7.4.

Again, if bit 22 is set, all multicast packets will *only* be forwarded to the CPU and not to the local ports, nor other ports of other GalNet devices. The CPU can then decide to what ports the multicast packet should be sent to. Only one packet needs to be sent to each GalNet device and each device will automatically forward the packet only to the ports that the CPU tagged for that specific Multicast packet. These ports are tagged in bits [29:22] of the END_OF_PACKET message.

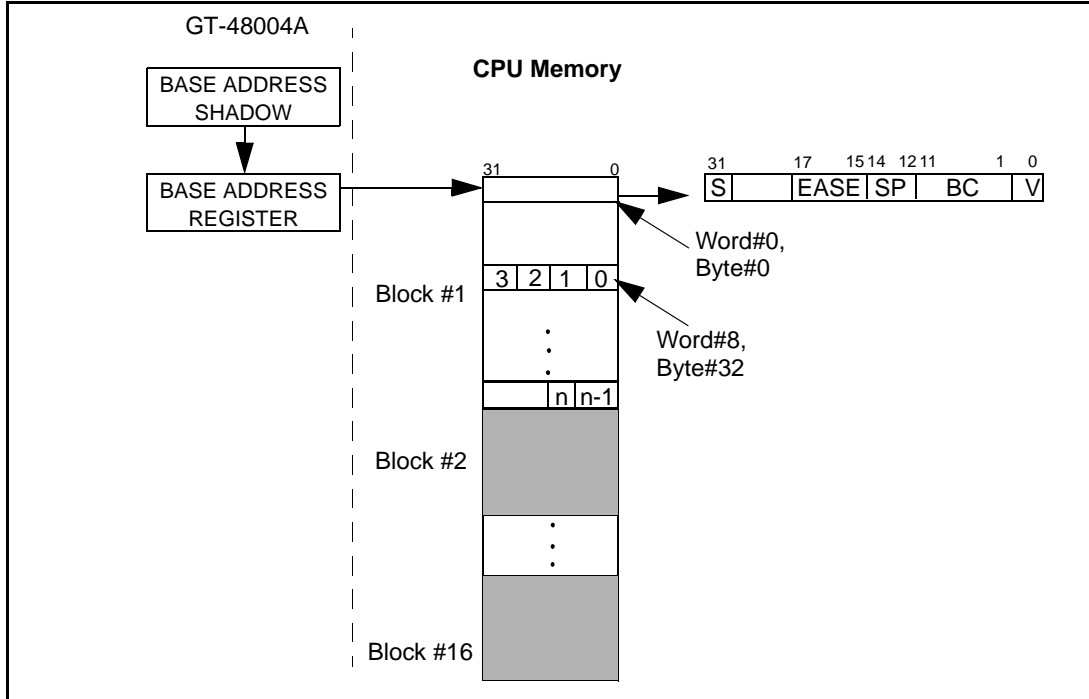
7.4 Forwarding a Packet to the CPU Directly

Systems which utilize a CPU (bit 10 of the Global Control Register, 0x140028, is set) will forward certain packets directly to the management CPU's memory. These packets include:

- Unicast packets destined for the CPU (Device Number in the address table is equal to the CPU number)
- Multicast packets
- Unknown packets (if set by bit 6 in the Global Control Register, 0x140028)
- BPDU messages
- Sniffer packets when the CPU is the target sniffer
- EASE packets

The GT-48004A contains two pointers to a sixteen block buffer area in the CPU's memory space. The registers are called the CPU Base Address (CBA) and CPU Base Address Shadow (CBAS.) These registers are physically located at the same address in the GT-48004A (0x140034). The first write to this register updates the CBA register, the second write updates the shadow register (CBAS). Figure 6 shows the data structure in the CPU memory.

Figure 6: CPU Memory Data Structure



The data structure components are the following:

- **CPU Base Address Register (CBA)**¹ - A register that points to the beginning of a sixteen block area in CPU memory.
- **CPU Base Address Shadow Register (CBAS)** - A second register that holds a pointer to a second sixteen block area in the CPU main memory. The value in the Shadow register propagates into the Base Address register after sixteen packets are transferred to the main memory.
- **Buffer Area** - The CPU buffer area consists of 16 blocks of 2Kbytes each. The first word (Word #0) of each block contains the Sniffer Indication [31], EASE indication [17:15], Source Port numbers [14:12], Byte Count (bits [11:1]), and the Valid bit (bit 0). These bits are written as the END_OF_PACKET (Section 11.3.10) message transferred from the source GT-48004A to the CPU at the end of a valid packet transfer. Words 1 to 7 are left empty for user purposes.

The communication between the GT-48004A and the CPU follows this sequence:

1. CPU updates the CBA (1st write to 0x140034).
2. CPU updates the CBAS (2nd write to 0x140034).
3. GT-48004A transfers 16 packets to the CPU main memory and asserts the Int* at the end of each packet transfer.
4. The CPU must count sixteen interrupts and then update the Shadow register in the GT-48004A (write to 0x140034). Also, the BufWrap interrupt can be checked instead of counting sixteen interrupts.

Steps 3-4 are repeated. The packet transfer to the CPU is done as follows:

1. The incoming packet is fed to the Rx FIFO and is transferred to an empty block in the Receive Buffer area of

1. Assuming CBA and CBAS have already been written to, any writes to 0x140034 will update the CBAS register ONLY. In other words, CBA can only be updated on the first write, and after 16 packets have been written to the CPU where the CBA will take the value of the CBAS.

DRAM.

2. In parallel, an address recognition cycle is performed for both the DA and the SA. If the packet is resolved to be a packet of the type listed at the beginning of this section, packet information is written to the PCI's transmit descriptor. This information includes the byte count and Receive Buffer Block Address which is pointed to by the Write Pointer. When the PCI's transmit descriptor's Write Pointer is not equal to the Read Pointer, the source GT-48004A sends the packet DIRECTLY to the appropriate block in the CPU main memory with the PACKET_TRANSFER (Section 11.3.7) message using PCI master operations in multiple eight 32-bit bursts. The data is entered starting at the 8th word (33rd byte) of the next free block. Words 1 to 7 are left empty for user purposes. The address of this packet transfer is based upon the CBA.
3. At the end of the packet transfer, the source GT-48004A sends the END_OF_PACKET message (Section 11.3.10) to the first word of the block (Word #0). It also sends an interrupt via Int* to the CPU, increments the Read Pointer, and clears the appropriate bit in its Empty List.

The CPU now has the packet buffered in its buffer area. This gives the CPU the ability to intervene in the packet's routing or to modify the contents of the packet.

7.5 Forwarding a Packet from the CPU to a GalNet Device

The sequence for forwarding a packet from the CPU to a port in a GalNet device is the same for unicast and multicast packets. The procedure is as follows:

1. The CPU sends a BUFFER_REQUEST message (Section) to the appropriate target GalNet device indicating that there is a packet ready for transmission across the PCI bus.
2. The target GalNet device receives this message and allocates a buffer in its DRAM. This target device then sends a START_OF_PACKET message (Section 11.3.5) back to the CPU indicating it is ready to receive the packet. This START_OF_PACKET message is sent to a CPU buffer location indicated by the Start Packet Base Address register (0x140038). This address points to a CPU buffer area that can hold up to 32 START_OF_PACKET messages. The CPU must poll this structure to know when a target GT-48004A is ready to receive a packet, and to what address the packet should be sent. If the Byte Count field in the START_OF_PACKET message is 0, then the CPU should not write the PACKET_TRANSFER message nor the END_OF_PACKET message to the device. This indicates that either the buffers are full in the target GalNet device, or that the link is down on the target port.
3. The CPU transfers the packet with the PACKET_TRANSFER message (Section 11.3.9) using PCI master operations in multiple eight 32-bit bursts. The packet is buffered in the Receive Buffer area of the target device's DRAM. After the entire packet has been transmitted, the CPU performs an additional write transaction by sending the END_OF_PACKET message (Section 11.3.12) indicating completion of the packet transfer. This message contains the Byte Count, the target Port Number, the Rx Block address, and the Packet Type. It also includes a bit which commands the GT-48004A to generate CRC for this out-going packet or not. If the packet is a multicast packet, the packet will only be forwarded to the ports tagged in bits [29:22] of the END_OF_PACKET message (The CPU should NEVER write '1' to bits 29:22 that had '0' in the same bits of the START_OF_PACKET message.)
4. Some packet information included in the END_OF_PACKET message is written to the appropriate transmit descriptor in the target device. This information includes the Byte Count and the Receive Buffer address which is pointed to by the Write Pointer.
5. The Write Pointer of the outgoing port's transmit descriptor is incremented. The target GalNet device transmits whenever the Write Pointer is not equal to the Read Pointer.
6. At the end of the packet transmit process, the target GalNet device increments the Read Pointer and clears the appropriate bit in the Empty List.

7.6 CRC Generation

As mentioned in Section 7.5, the GT-48004A also includes a CRC generator for packets sent out by the CPU. The

CRC generator enhances system performance by implementing the CPU intensive packet CRC calculation in hardware. Note that CRC generation is not required for packets transferred between GalNet devices, since the CRC is already appended to these packets.

CRC generation for CPU generated packets is enabled through the EnCRC bit in the Global Control register (bit 29). In addition, the CPU must set the GenCRC bit in the END_OF_PACKET (Section 11.3.12) message sent to the GT-48004A at the completion of forwarding the packet.

7.7 Tx Watchdog Timer

The GT-48004A includes a transmit watchdog timer for each transmit queue. For 100Mbps operation, the default value of the timer is 63msec; the range is between 10.5 to 168msec. For 10Mbps operation, the default value of the timer is 630msec and the range is between 105msec to 1680msec. The timer measures the time between the transmission of two consecutive outgoing packets. When the timer expires, the GT-48004A clears the appropriate used blocks and sends an interrupt to the CPU via Int*.

The transmit watchdog timer prevents transmission problems on one port from blocking traffic to other ports. In managed systems, the timers also provide a mechanism to notify the CPU of a possible system problem.

The Tx Watchdog Timer can be externally disabled by holding the DisWD* pin LOW. The default is to leave the Tx Watchdog Timer enabled (DisWD* HIGH.)

8. Device Table Operation

The GalNet Architecture supports a maximum of 32 devices per system. Each device needs to know of the existence of all other GalNet devices in the system in order to communicate information and packet data. The Device Table is a 32-bit register that uses a single bit for each possible Device Number to indicate its presence/absence in the system.

8.1 Automatic Device Table Initialization

Upon RESET, each FEU in each GT-48004A in the system sets all of its Device Table bits to '1'. This is essentially a "guess" by the FEU that *all* Device Numbers are in use. In a system with any PCI device that is NOT a GalNet device (CPU, graphics card, etc.), it is recommended that the CPU initialize the device table before the GalNet device is enabled (or before packet transmission starts).

Specific Device Table bits are cleared, and the corresponding Device Numbers are logically eliminated from the system, by:

- The receipt of a PCI Master Abort when the GT-48004A attempts to access the corresponding Device Number
- Management CPU programming

Device Table "cleaning" is handled automatically by each device upon receipt of the first unknown packet following a RESET. This process discovers all other GalNet devices in the system without processor intervention:

1. The GT-48004A receives the first good packet from any Ethernet port and places it in the buffer area in DRAM.
2. Since the Address Table is cleared (immediately following RESET) the packet must be forwarded to all ports, including those located on other GalNet devices.
3. The GT-48004A attempts to allocate a buffer in all 31 possible additional GalNet devices in the system. This is done by issuing 31 BUFFER_REQUEST messages- one to every possible device in the system.
4. Each BUFFER_REQUEST message that fails due to a PCI Master Abort results in the corresponding Device Number bit being cleared in the Device Table. A PCI Master Abort only occurs when there is no target device at the requested address (i.e. the target device doesn't exist.)

NOTE: Each FEU in a GT-48004A performs the above process separately.

8.2 Manual Device Table Initialization

Alternatively, the CPU can notify each FEU within the GT-48004A of the existence of the GalNet devices. This is done by setting bit 2 in the global control register (DevTabMod.) In this mode, the CPU is responsible for notifying the GT-48004A of the existence of all GT-4800x devices in the system. The CPU does this by writing the appropriate values to each device table.

"Manual" mode must be used when there are multiple PCI buses in the system separated via PCI-to-PCI bridges.

8.3 Programming Device Numbers

The Device Number for each FEU in the GT-48004A is set after RESET by the values on the xDAddr[4:0] pins. The device number can be changed by writing the new device number to [26:22] of the DRAM/Internal Register Base Address Register at 0x010 (each FEU is addressed separately.) Typically, once the number of a device has been changed, the device table of all devices in the system will also need to be updated.

NOTE: Each FEU in a GT-48004A must have a different device number.

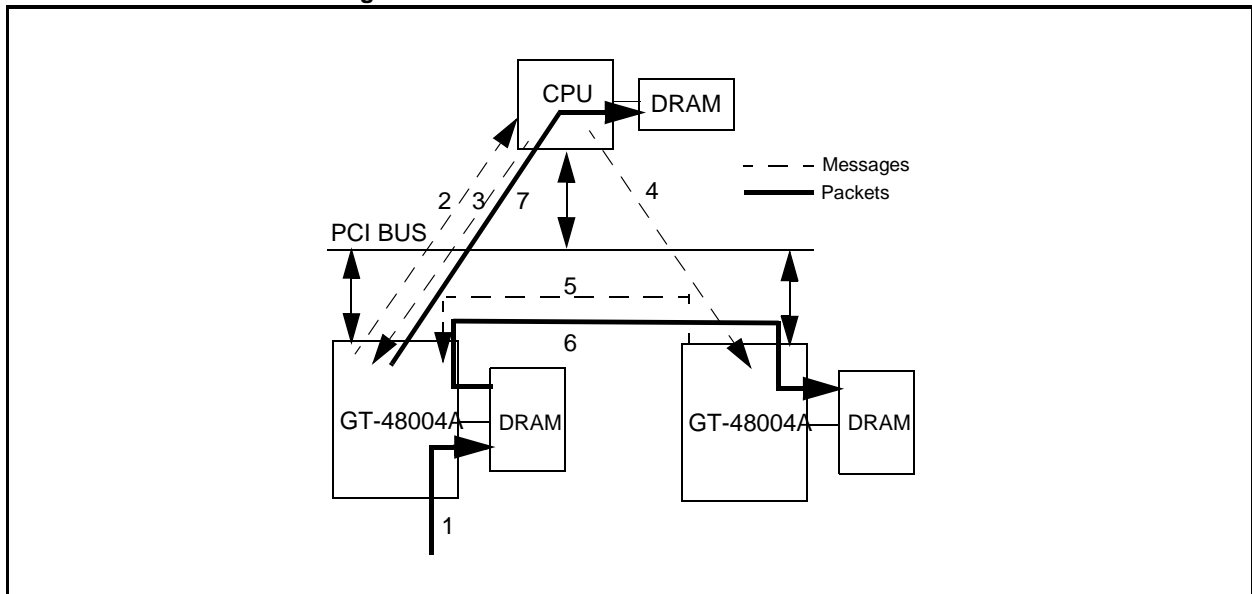
9. Unicast Intervention Mode

The GT-48004A supports a powerful feature called Intervention Mode for unicast packets. Intervention in Unicast traffic is optional per MAC address. Each entry in the Address Table includes an Intervention bit for the destination address (bit 62) and the source address. When an Intervention mode bit is set, the GT-48004A will *not* forward the packet automatically to the destination device. Instead, the GalNet device will send a BUFFER_REQUEST message (Section 11.3.3) to CPU memory. The BUFFER_REQUEST messages will be sent to the buffer area in the CPU main memory which contains 256 entries of dual 32-bit words. The location of this CPU memory is specified by the CPU Intervention Base Address at 0x140048. The buffer area is pointed to by this Base Address register and a Shadow register at this same offset. The BUFFER_REQUEST message includes information about the routing of the packet (Source and Target port/device numbers). The CPU then has the option to:

- discard the packet
- forward the packet to a specific device
- request the packet be transferred to CPU buffer memory

Figure 7 shows a Unicast packet transfer using Intervention mode.

Figure 7: Unicast Packet Transfer With Intervention



The sequence illustrated is as follows:

1. The incoming packet is fed to the Rx FIFO and transferred to an empty block in the Receive Buffer area of DRAM.
2. If a corresponding Intervention bits is set, the device sends a BUFFER_REQUEST message (Section 11.3.3) to the CPU's Intervention Base Address. The BUFFER_REQUEST includes the source port and the source buffer address

The CPU then has the following options:

1. Discard the packet.
 - The CPU sends a START_OF_PACKET message (Section 11.3.6) to the source device with the Byte Count field cleared with all 0s (arrow #3).
2. Forward the Packet to a port in the same device

- The CPU sends a BUFFER_REQUEST (Section 11.2.2) to back to the same device where the target device number is equal to the source device number. The correct port number is also specified. This BUFFER_REQUEST is of the same format that is used between GalNet devices (arrow #2).
 - The packet is then sent out on a port of same source device as specified by the CPU.
3. Forward the packet to a different device.
- The CPU sends a BUFFER_REQUEST (Section 11.2.2) to the destination device. This BUFFER_REQUEST is of the same format that is used between GalNet devices and uses the source device number as the Source Device instead of the CPU device number (arrow #4).
 - The target device allocates a buffer and sends a START_OF_PACKET message (Section 11.2.3) to the source device (arrow #5).
 - The source device transfers the packet with the PACKET_TRANSFER message (Section 11.2.4) followed by an END_OF_PACKET message (Section 11.2.5) to the target device (arrow #6).
4. Take the packet.
- The CPU sends a START_OF_PACKET message (Section 11.3.6) back to the source GalNet device. The target device in this message is the CPU number. The source GT-48004A device transfers the packet with the PACKET_TRANSFER message (Section 11.3.8) using PCI master operations in multiple eight 32-bit bursts. The first data word of the packet will be written to the second word of the buffer. The first word is left empty for the END_OF_PACKET message. NOTE: Unicast packets tagged for intervention mode are **not** forwarded to the same buffer space that is specified by the CPU Buffer Base Address (CBA) at offset 0x14140034).
 - At the end of the packet transfer, the source GT-48004A sends the END_OF_PACKET message (Section 11.3.11) to the first word of the buffer. It also sends an interrupt via Int* to the CPU, and clears the appropriate bit in its Empty List.

The CPU buffer now contains the packet whose destination address was marked for intervention. The CPU can perform a number of functions with the packet including layer 3 routing, security, virtual LAN support, filtering and management. After the CPU has modified the packet, the packet can be transferred to the appropriate target GalNet device with the procedure outlined in Section 7.5.

9.1 Unicast Intervention Mode Address Space

Prior to the CPU receiving unicast packets tagged for intervention mode, an address space for GalNet devices to forward packets to must be allocated. This address space should reside in the GalNet protocol region (see Section 11.1). Therefore, this buffer region will have a similar address to a GalNet device. The buffer address format is shown in Table 5.

Table 5: Buffer Address for Unicast Intervention Mode Packets


PCI Bits	Description
Address	
[31:27]	GalNet Protocol Region
[26:22]	CPU Device Number (bits [12:8] in 0x140030)
[21]	'1' (DRAM)
[20:0]	DRAM location

10. Address Table

The GT-48004A's Address Table resides in the local DRAM array. Figure 8 shows the Address Table structure. The Address Table structure occupies approximately 320Kbytes and is entirely controlled and initialized by the GT-48004A. Following RESET, the GT-48004A initializes the Address Table by setting all Valid Bits in the table to '0'. Every new address that is learned has its entry's Valid Bit set to '1'. In order to remove an address entry from the table, ONLY the Skip Bit should be set to '1' (the Valid Bit should not be modified).

Modifications to the Address Table are normally made through GalNet protocol message requests by the CPU. It is possible to access the Address Table directly, however, this mode is not recommended. Please contact Galileo if you feel your application requires direct Address Table access.

Figure 8: Address Table Format



Address Table Bits Description

Bit	Description
V	Valid: Indicates this entry is taken or previously was taken 0 - Not valid 1 - Valid
Sk	Skip - Indicates if this entry is currently being used 0 - Active entry, do not skip 1 - Inactive entry, overwriting this entry is allowed
A	Aging - This bit is used for the Aging process 0 - Cleared by the CPU 1 - Set by the GT-48001A upon receiving a packet from the station corresponding to this entry
Addr	48 bit MAC address
Dev#	Device Number - 5 bit GalNet device number indicating which of a maximum of 32 devices in the system is associated with this address.
Port#	Port Number - 3 bit Port number (for compatibility with the GT-48001A) indicating which of the eight ports of a GT-48001A or which of the 2 ports of a GT-48004A is associated with this address.
R	Reserved
St	Static - Indicates whether an entry can be modified or not 0 - The entry can be modified 1 - The entry is static, the Dev# and Port# cannot be modified
M	Multiple - Meaningful when bit St is set 0 - Forward this packet only to the destination port 1 - Forward this packet to all ports (as Unknown)
Id	Intervention for Destination Addresses 0 - Don't activate Intervention mode for this address as a destination 1 - Activate Intervention mode
Is	Intervention for Source Addresses 0 - Don't activate Intervention mode for this address as a source 1 - Activate Intervention mode

11. GalNet Messaging Protocol

The GalNet Messaging Protocol is comprised of the messages passed over the PCI bus between

- GalNet family members
- a GalNet family member and a CPU

The messages are encoded in both the address and the data phases of the PCI transfers. Five groups of messages that are currently defined: NEW_ADDRESS, BUFFER_REQUEST, START_OF_PACKET, PACKET_TRANSFER, and END_OF_PACKET.

GalNet messages are write-only (request/response). For example, a GalNet device that needs to transfer data to another GalNet chip starts the transfer by requesting a buffer (BUFFER_REQUEST) from the target. The target responds with an address to which the packet should be transferred to (START_OF_PACKET).

The GalNet messaging protocol allows messages to be interleaved. For example, a device may send out multiple BUFFER_REQUEST messages to other devices before receiving a START_OF_PACKET message reply.

Write only messaging was used since it better utilizes PCI bus bandwidth. Reads on the PCI bus tend to stall the bus, and the reading device thereby degrading overall system performance.

Please note that there are subtle differences in the GalNet message format when transferring messages between devices (Section 11.2) and between a device and a CPU (Section 11.3).

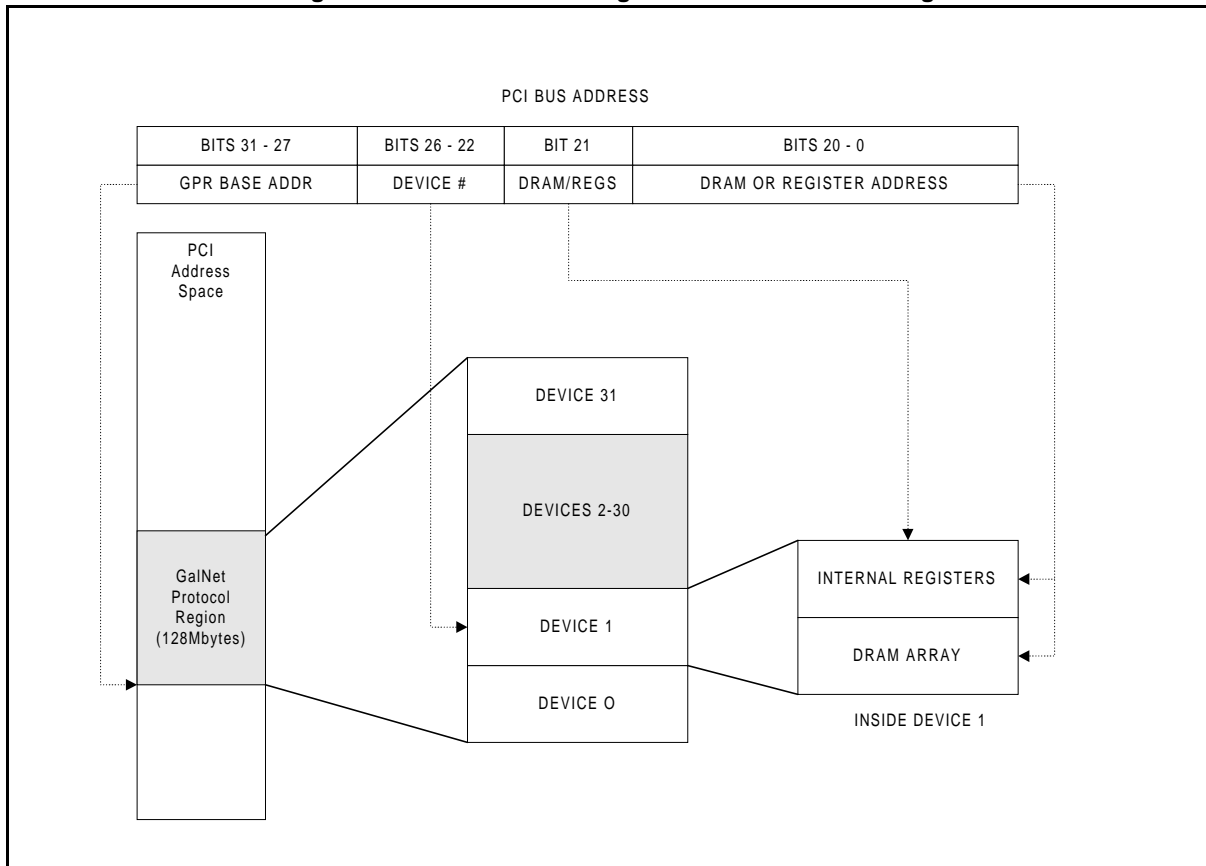
11.1 GalNet Protocol Region

All GalNet devices in a system must reside within a single 128Mbyte region in the PCI memory address space known as the GalNet Protocol Region (GPR). The base address of the GPR is set in bits 31:27 of the DRAM/Internal Base Address Register at offset 0x10 in the GT-48004A's PCI configuration header.¹ All GalNet devices in a system MUST have the same value in this field. GalNet devices default to a value of '00001' in the GPR bits following RESET.

Each GalNet device in the system occupies a separate 4Mbyte "slice" of the GalNet Protocol Region. This "slice" is decoded by the Device Number field in the DRAM/Internal Base Address Register (bits 26:22). The individual device's address space is further divided by the DRAM/Register bit (bit 21). This bit determines whether an access to a target device is directed to the DRAM array or to the registers.

1. Offset 0x10 in the PCI configuration register is known as Base Address Register 0 for standard PCI devices.

Figure 9: GalNet Protocol Region and Address Decoding



11.2 GalNet Messages Between Devices

The five groups of GalNet messages as described in the sections below are sent from one GalNet device to another. The data format for all GalNet messages is *little-endian*, which is the PCI standard.

11.2.1 NEW_ADDRESS Message between GalNet devices

A message sent from one GalNet device to another to alert of a new address.

Table 6: NEW_ADDRESS Message between GalNet devices

PCI Bits	Description
Address	
[31:27]	GalNet Protocol Region
[26:22]	Target Device Number
[21]	'0' (Internal registers)
[20:18]	'100' (NEW_ADDRESS message)
[17:0]	'0'
Data 0	
[31:3]	MAC address [19:47]
[2]	Aging
[1]	Skip
[0]	Valid
Data 1	
[31]	'1'
[30]	Static
[29]	Address Unknown/New address
[28]	Multiple
[27]	'0'
[26:24]	Port Number
[23:19]	Device Number
[18:0]	MAC address [0:18]
Data 2	
[24]	Intervention mode for Destination Address
[25]	Intervention mode for Source Address

The NEW_ADDRESS message may also be posted as a 'question' to other GalNet devices during the recovery process (see Section 5.2 on page 20). The indication is via the Address Unknown/New Address bit (bit 29). If bit 29 is clear, then the NEW_ADDRESS message is a 'question'. When bit 29 is set, the message indicates a 'real' address to put into the target devices' address tables.

The NEW_ADDRESS message must be transferred in a burst of three words over the PCI.

11.2.2 BUFFER_REQUEST Message between GalNet devices

A message sent from a source GalNet device to a target GalNet device to request a buffer.

Table 7: BUFFER_REQUEST Message between GalNet devices

PCI Bits	Description
Address	
[31:27]	GalNet Protocol Region
[26:22]	Target Device Number
[21]	'0' (Internal registers)
[20:18]	'101' (BUFFER_REQUEST message)
[17:13]	Source Device Number
[12:0]	'0'
Data 0	
[31]	Sniffer (0-Sniffer type message)
[30]	Unknown (0-Unknown message)
[29:28]	'0'
[27:25]	Source Port Number
[24:22]	Target Port Number
[21]	Multicast/Unicast (0-Unicast)
[20:10]	Byte Count
[9:0]	Source Buffer Address (divided by 0x600)

11.2.3 START_OF_PACKET Message between GalNet devices

A message from a target GalNet device to the source GalNet device which contains the Empty Receive Buffer address.

Table 8: START_OF_PACKET Message between GalNet devices

PCI Line	Description
Address	
[31:27]	GalNet Protocol Region
[26:22]	Source Device Number
[21]	'0' (Internal registers)
[20:18]	'110' (START_OF_PACKET message)
[17:0]	'0'
Data 0	
[31]	Sniffer (0 - Sniffer type message)
[30]	'0'
[29:22]	Target Port Number (1bit per each port; bit 22 - Port 0, bit 23 - Port 1 etc.)
[21]	Multicast/Unicast (0-Unicast)
[20:10]	Byte Count
[9:0]	Target Buffer Address (divided by 0x600)
Data 1	
[31:18]	'0'
[17:15]	Source Port Number
[14:5]	Source Buffer Address (divided by 0x600)
[4:0]	Target Device Number

The START_OF_PACKET message must be transferred in a burst of two words over the PCI.

11.2.4 PACKET_TRANSFER Message between GalNet devices

A burst of up to 8 32-bit words from the source GalNet device to the target GalNet device which contains the packet.

Table 9: PACKET_TRANSFER Message between GalNet devices

PCI Line	Description
Address	
[31:27]	GalNet Protocol Region
[26:22]	Target Device Number
[21]	'1' (DRAM)
[20:0]	DRAM location
Data 0	
[31:0]	Data 0
----	----
----	----
Data 7	
[31:0]	Data 7

The PACKET_TRANSFER message can be transferred using any burst size over the PCI. The GalNet device uses a burst of 8 words.

11.2.5 END_OF_PACKET Message between GalNet devices

A message from the source GalNet device to the target GalNet device which indicates the end of the packet.

Table 10: END_OF_PACKET Message between GalNet devices

PCI Line	Description
Address	
[31:27]	GalNet Protocol Region
[26:22]	Target Device Number
[21]	'0' (Internal registers)
[20:18]	'111' (END_OF_PACKET message)
[17:0]	'0'
Data 0	
[31]	CRC (0 - Do not append, 1 - Append) Always '0', GalNet never appends CRC
[30]	'0'
[29:22]	Target Port Number (1bit per each port; bit 22 - Port 0, bit 23 - Port 1 etc.)
[21]	Multicast/Unicast (0-Unicast)
[20:10]	Byte Count
[9:0]	Target Buffer Address (divided by 0x600)

11.3 GalNet Messages Between a GalNet Device and a CPU

When a management CPU is used in a GalNet switch system, bit 10 (CPUEn) of the Global Control Register, 0x140028, must be set.

The five groups of GalNet messages described in the sections below are sent from one GalNet device to a CPU and from a CPU to a GalNet device.

11.3.1 NEW_ADDRESS Message (GalNet to CPU)

A message from a GalNet device to the CPU that contains information about a new MAC Address. Bit 7 (ForwNewAdd) of the Global Control Register (0x140028) must be set.

Table 11: NEW_ADDRESS Message (GalNet to CPU)

PCI Bits	Description
Address	
[31:8]	NEW_ADDRESS Base Address (Offset: 0x14003c)
[7:3]	Offset Counter (32 total entries)
[2:0]	'000'
Data 0	
[31:3]	MAC address [19:47]
[2]	'1'
[1]	'0'
[0]	'0'
Data 1	
[31]	'1'
[30]	Reserved
[29]	Address Unknown/New address (0 - Unknown, 1 - New Address Message)
[28]	Reserved
[27]	Reserved
[26:24]	Port Number
[23:19]	Device Number
[18:0]	MAC address [0:18]

The NEW_ADDRESS message must be transferred in a burst of two words over the PCI from a GalNet device to the CPU.

11.3.2 NEW_ADDRESS Message (CPU to GalNet)

A message from the CPU to a GalNet device that contains information about a new MAC Address.

Table 12: NEW_ADDRESS Message (CPU to GalNet)

PCI Bits	Description
Address	
[31:27]	GalNet Protocol Region
[26:22]	Target Device Number
[21]	'0' (Internal registers)
[20:18]	'100' (NEW_ADDRESS message)
[17:0]	'0'
Data 0	
[31:3]	MAC address [19:47]
[2]	Aging
[1]	Skip
[0]	Valid (should always be '1')
Data 1	
[31]	'1'
[30]	Static
[29]	Address Unknown/New address (0 - Unknown, 1 - New Address Message)
[28]	Multiple
[27]	'0'
[26:24]	Port Number
[23:19]	Device Number
[18:0]	MAC address [0:18]
Data 2	
[24]	Intervention Mode for Destination Address
[25]	Intervention Mode for Source Address

The NEW_ADDRESS message must be transferred in a burst of three words over the PCI.

Data 2 must be sent when sending the NEW_ADDRESS message. After every NEW_ADDRESS message sent by the CPU, a dummy NEW_ADDRESS message (with a non existing MAC address) should be sent with Data2 cleared (all 0s). The dummy MAC address can be all zeros or another MAC address that doesn't exist in the network.

11.3.3 BUFFER_REQUEST Message (GalNet to CPU)

A message from the source GalNet device to the target CPU to request a buffer. This message is only used when a unicast packet's destination address is marked for Intervention Mode and will be forwarded to the CPU buffer.

Table 13: BUFFER_REQUEST Message (GalNet to CPU)

PCI Bits	Description
Address	
[31:11] [10:3] [2:0]	CPU Intervention Base Address (Offset: 0x140048) Offset Counter (256 total entries) '000'
Data 0	
[31] [30] [29:28] [27:25] [24:22] [21] [20:10] [9:0]	Sniffer (0-Sniffer type message) Unknown (0-Unknown message) '0' Source Port Number Target Port Number Multicast/Unicast (0-Unicast) Byte Count Source Buffer Address (divided by 0x600)
Data 1	
[31:29] [28:24] [23:0]	'0' Target Device Number '0'

The BUFFER_REQUEST message must be transferred in a burst of two words over the PCI.

11.3.4 BUFFER_REQUEST Message (CPU to GalNet)

A message from the source CPU to the target GalNet device to request a buffer.

Table 14: BUFFER_REQUEST Message (CPU to GalNet)

PCI Bits	Description
Address	
[31:27] [26:22] [21] [20:18] [17:13] [12:0]	GalNet Protocol Region Target Device Number '0' (Internal registers) '101' (BUFFER_REQUEST message) Source CPU Device Number (bits [12:8] in 0x140030) '0'
Data 0	
[31] [30] [29:28] [27:25] [24:22] [21] [20:10] [9:0]	Sniffer (0-Sniffer type message) Unknown (0-Unknown message) '0' Reserved Target Port Number Multicast/Unicast (0-Unicast) Byte Count Reserved

11.3.5 START_OF_PACKET Message (GalNet to CPU)

A message from a target GalNet device to the CPU which contains the target Buffer address.

Table 15: START_OF_PACKET Message (GalNet to CPU)

PCI Bits	Description
Address	
[31:8]	START_OF_PACKET Base Address (0x140038)
[7:3]	Offset Counter (32 total entries)
[2:0]	'000'
Data 0	
[31]	Sniffer (0 - Sniffer type message)
[30]	'0'
[29:22]	Target Port Number (1bit per each port; bit 22 - Port 0, bit 23 - Port 1 etc.)
[21]	Multicast/Unicast (0-Unicast)
[20:10]	Byte Count
[9:0]	Target Buffer Address (divided by 0x600)
Data 1	
[31:0]	'0'

The START_OF_PACKET message must be transferred in a burst of two words over the PCI.

11.3.6 START_OF_PACKET Message (CPU to GalNet)

A message from the target CPU to the source GalNet device which contains the Empty Buffer address. This message will only be sent when the GalNet devices needs to forward a unicast packet to the CPU via Intervention mode.

Table 16: START_OF_PACKET Message (CPU to GalNet)

PCI Bits	Description
Address	
[31:27]	GalNet Protocol Region
[26:22]	Source Device Number
[21]	'0' (Internal registers)
[20:18]	'110' (START_OF_PACKET message)
[17:0]	'0'
Data 0	
[31]	Sniffer (0 - Sniffer type message)
[30]	'0'
[29:22]	Reserved
[21]	Multicast/Unicast (0-Unicast)
[20:10]	Byte Count
[9:0]	Target Buffer Address (divided by 0x600)
Data 1	
[31:18]	'0'
[17:15]	Source Port Number
[14:5]	Source Buffer Address (divided by 0x600)
[4:0]	Target CPU Device Number (bits [12:8] in 0x140030)

The START_OF_PACKET message must be transferred in a burst of two words over the PCI.

11.3.7 PACKET_TRANSFER Message (GalNet to CPU 16 Block Buffer)

A burst of up to 8 32-bit words from the source GalNet device to the target CPU buffer which contains the packet. This message is used when the GalNet device forwards one of the following packets directly to the CPU 16 block buffer:

- Unicast packets destined for the CPU (Device Number in the address table is equal to the CPU number)
- Multicast packets
- Unknown packets (if set by bit 6 in the Global Control Register, 0x140028)
- BPDU messages
- Sniffer packets when the CPU is the target sniffer
- EASE packets

The data is entered starting at the 8th word (33rd byte) of the next free block. Words 1 to 7 are left empty for user purposes. The address of this packet transfer is based upon the CPU Buffer Base Address.

Table 17: PACKET_TRANSFER Message (GalNet to CPU)

PCI Bits	Description
Address	
[31:15]	CPU Buffer Base Address (0x140034)
[14:11]	Offset Counter (16 total blocks)
[10:2]	Address within the buffer
[1:0]	'0'
Data 0	
[31:0]	Data 0
----	----
----	----
Data 7	
[31:0]	Data 7

The PACKET_TRANSFER message can be transferred using any burst size over the PCI. The GalNet device uses a burst of 8 words.

11.3.8 PACKET_TRANSFER Message (GalNet to CPU in Unicast Intervention Mode)

A burst of up to 8 32-bit words from the source GalNet device to the target CPU buffer which contains the unicast packet whose address is tagged for intervention mode. The first data word of the packet will be written to second word of the buffer. The first word is left empty for the END_OF_PACKET message.

Table 18: PACKET_TRANSFER Message (GalNet to CPU)

PCI Bits	Description
Address	
[31:27]	GalNet Protocol Region
[26:22]	CPU Device Number (bits [12:8] in 0x140030)
[21]	'1' (DRAM)
[20:0]	DRAM location
Data 0	
[31:0]	Data 0
----	----
----	----

PCI Bits	Description
Address	
Data 7	
[31:0]	Data 7

The PACKET_TRANSFER message can be transferred using any burst size over the PCI. The GalNet device uses a burst of 8 words.

11.3.9 PACKET_TRANSFER Message (CPU to GalNet)

A burst of up to 8 32-bit words from the source CPU Buffer to the target GalNet device which contains the packet.

Table 19: PACKET_TRANSFER Message (CPU to GalNet)

PCI Bits	Description
Address	
[31:27]	GalNet Protocol Region
[26:22]	Target Device Number
[21]	'1' (DRAM)
[20:0]	DRAM location
Data 0	
[31:0]	Data 0
----	----
----	----
Data 7	
[31:0]	Data 7

The PACKET_TRANSFER message can be transferred using any burst size over the PCI. It is recommended to use a burst of 8 words.

11.3.10 END_OF_PACKET Message (GalNet to CPU 16 Block Buffer)

A message from the source GalNet device to the CPU which indicates the end of the packet. This message is used when the GalNet device forwards one of the following packets directly to the CPU 16 block buffer:

- Unicast packets destined for the CPU (Device Number in the address table is equal to the CPU number)
- Multicast packets
- Unknown packets (if set by bit 6 in the Global Control Register, 0x140028)
- BPDU messages
- Sniffer packets when the CPU is the target sniffer
- EASE packets

The END_OF_PACKET message is written to the first word of the buffer.

Table 20: END_OF_PACKET Message (GalNet to CPU 16 Block Buffer)

PCI Bits	Description
Address	
[31:15]	CPU Buffer Base Address (0x140038)
[14:11]	Offset Counter (16 total blocks)
[10:0]	'0' (Always to the First word)
Data 0	
[31]	Sniffer Packet
[30:24]	Reserved
[23:16]	EASE sample (port0, bit 16; port7, bit 23)
[15]	EASE sample is an original packet to CPU
[14:12]	Source Channel Number
[11:1]	Byte Count
[0]	Valid Bit

11.3.11 END_OF_PACKET Message (GalNet to CPU in Unicast Intervention Mode)

A message from the source GalNet device to the CPU which indicates the end of a unicast packet transfer to the CPU in intervention mode. The END_OF_PACKET message is written to the first word of the buffer.

Table 21: END_OF_PACKET Message (GalNet to CPU in Unicast Intervention Mode)

PCI Bits	Description
Address	
[31:27]	GalNet Protocol Region
[26:22]	CPU Device Number (bits [12:8] in 0x140030)
[21]	'1' (DRAM)
[20:0]	'0' (Always to the First word)
Data 0	
[31]	Sniffer Packet
[30:24]	Reserved
[23:16]	EASE sample (port0, bit 16; port7, bit 23)
[15]	EASE sample is an original packet to CPU
[14:12]	Source Channel Number
[11:1]	Byte Count
[0]	Valid Bit

11.3.12 END_OF_PACKET Message (CPU to GalNet)

A message from the source CPU to the target GalNet device which indicates the end of the packet.

Table 22: END_OF_PACKET Message (CPU to GalNet)

PCI Bits	Description
Address	
[31:27]	GalNet Protocol Region
[26:22]	Target Device Number
[21]	'0' (Internal registers)
[20:18]	'111' (END_OF_PACKET message)
[17:0]	'0'
Data 0	
[31]	CRC (0 - Do not append, 1 - Append)
[30]	'0'
[29:22]	Target Port Number (1bit per each port; bit 22 - Port 0, bit 23 - Port 1 etc.)
[21]	Multicast/Unicast
[20:10]	Byte Count
[9:0]	Target Buffer Address (divided by 0x600)

12. Fast PCI Bus Operation

The GT-48004A can act as either a master initiating a PCI bus operation, or as a target responding to a PCI bus operation. The GT-48004A is enabled to both master the PCI bus and respond to memory accesses after RESET.

The GT-48004A has two sets of internal registers that are accessible through the PCI interface for each FEU. Internal control registers for the GT-48004A are memory mapped and accessible through standard PCI read/write operations. PCI control registers are accessible through PCI configuration cycles.

12.1 Separate Logical PCI Interfaces for Each FEU

Each FEU within a GT-48004A has a separate logical PCI interface. This means that there are separate Req/Gnt* pairs and a separate IdSel for each FEU. This also means that there are two separate Configuration Headers, Register banks, etc. From both the software and hardware standpoint, it is as if the GT-48004A has two physical PCI devices in a single package.

12.2 Interfacing Management Processors to Fast PCI

The Fast PCI bus required by the GT-48004A runs up to 66MHz. To connect a management processor, you must have a PCI interface chip capable of running at this speed. Galileo currently offers two system controller products which incorporate Fast PCI interfaces as well as a memory controller (and all other core logic):

- The GT-64120 interfaces all 64-bit bus MIPS processors (R4700 through R5000) to the GT-48004A. The GT-64120 includes two 66MHz PCI interfaces as well as a 75MHz CPU interface and SDRAM controller.
- The GT-64111 interfaces 32-bit bus MIPS processors (R4640, RM5230 and Vr4300) to the GT-48004A. The GT-64111 includes a single 66MHz PCI interfaces as well as a 66MHz CPU interface and EDO DRAM controller. The GT-64111 is the more appropriate choice for cost sensitive applications.

It is possible to run the GT-48004A at 33MHz on the PCI bus, however, doing so will not realize any performance benefit over the GT-48002A.

12.3 PCI Configuration Header Registers

The GT-48004A's PCI configuration registers are located in the standard PCI configuration header locations. Access to these registers is through PCI configuration reads/writes with the specific IdSel0/1 signal asserted. PCI configuration registers control PCI functions and return PCI information (device/vendor ID, etc.) Note that there are two separate configuration headers in the GT-48004A; one for FEU0 and the other for FEU1.

12.4 Accessing DRAM and Internal Registers through the PCI Interface

All GT-48004A internal control registers and the local DRAM are mapped into PCI memory space. Each FEU in the GT-48004A looks at PCI address bits 31:22 to determine if it is the target for the current cycle. This results in a decode region of 4Mbytes per FEU in a GT-48004A. Bit 21 of the address is used by the FEU to subdecode on-chip between accesses to DRAM and accesses to the internal registers. Bits 20:0 are used to select a specific register or memory location. Register addresses given in this document refer to the value of bits 20:0 in the PCI address. The GT-48004A expects data in little-endian format, as is required by the PCI specification.

The GT-48004A acts as a "medium speed" decode device, returning DevSel* in 2 clocks.

12.5 Fast PCI Bandwidth/Performance Issues

The fast PCI bus has an ideal maximum bandwidth of just under 264Mbytes/sec at 66MHz (about 2 gigabits per second.) In systems with many GalNet devices the bandwidth available on the PCI bus can become a performance limitation.

There are several factors that reduce the maximum achievable bus bandwidth:

- Aggressiveness of the PCI arbiter. For example, the use of hidden arbitration can save cycles between adjacent accesses and improve overall bandwidth. Simpler arbiters will degrade bandwidth as cycles that could be used for data transfer are used for arbitration.
- The GalNet protocol does involve some overhead when compared to raw data transfer.

While performance in any given system is impossible to estimate due to design differences, Galileo Technology has

determined the theoretical maximum PCI bus loading for each 100Mbps port (see Table 23.) The bandwidth numbers shown in this table take into account overhead for the GalNet protocol, as well as overhead for the PCI bus itself (arbitration, clocks/transfer, etc.)

Table 23: PCI Bandwidth Estimates

Packet Length	Maximum PCI Bandwidth Required per FULL DUPLEX 100Mbps Port	Maximum Number of 100Mbps Ports per PCI Bus with 0% Packet Loss at Worst Case Load ¹
64	TBD	TBD
128	TBD	TBD
256	TBD	TBD
512	TBD	TBD
1024	TBD	TBD

1. Maximum number of ports to guarantee 0% packet loss at worst case loading conditions.

Please note that the DRAM bandwidth limitations seen in some GT-48002A systems are lessened (or eliminated) in the GT-48004A due to the double speed memory interface.

12.6 Plug-and-Play Considerations In PCI Systems

The GT-48004A is not suited to Plug and Play configurations where the GT-48004A sits directly on the x86 PCI bus. Customers wanting to build PC based switches are urged to use either a GT-64120 (aka GT-146H) or Intel i960RP device as a “front end”

12.7 PCI Bus in Stand-Alone Systems

Stand-alone applications still require use of the PCI bus for forwarding packets between FEUs. The PCI bus pins must be connected as shown in Table 24 to insure proper operation. All pull-up and pull-down resistors should have a value of 4.7KΩ. Be sure to have your own reset and PCI clock on-board.

Table 24: Pin Strapping Requirements for PCI Signals in 4-Port Applications

Pin Name	Strapping
DevSel*	Pull up
Stop*	Pull up
Par*	Pull up
PErr*	Pull up
Frame*	Pull up
IRdy*	Pull up
TRdy*	Pull up
Gnt0/1*	Connect to PAL arbiter. NOTE: One of the two Gnt* signals MUST be asserted at all times. This ensures that one of the two FEUs is “parked” on the PCI bus, per the PCI bus spec.
IdSel0/1	Pull down
SErr*	Float
Req0/1*	Connect to PAL arbiter
Int*	Float
AD[31:0]	Float
CBE[3:0]	Float

12.8 PCI Bus Arbiter

All GT-48004A systems will require a PCI bus arbiter. The arbiter examines the bus request signals from each FEU and determines which FEU or external device is granted the bus. Galileo provides reference designs for PCI bus arbiters (implemented in inexpensive PALs) on our website.

13. Fast Ethernet Interfaces

Each FEU in the GT-48004A interfaces directly to two MII (Media Independent Interface) ports (4 ports total) which are compliant with the IEEE standard (please see 802.3u Fast Ethernet standard for detailed interface information and timing parameters). Each MII port has the following characteristics:

- Capable of supporting both 10 Mbps and 100 Mbps data rates in half or full duplex modes
- Data and delimiters are synchronous to clock references
- Provides independent 4-bit wide transmit and receive paths
- Uses TTL signal levels
- Provides a simple management interface (common to all ports)
- Capable of driving a limited length of shielded cable

The GT-48004A incorporates all the required digital circuitry to interface to 100BaseTX, 100BaseT4, and 100BaseFX. Two Fast Ethernet ports are integrated in the GT-48004A and only a small amount of external logic is needed to implement the standard physical interfaces.

13.1 10/100 MII Compatible Interface

The GT-48004A MAC allows it to be connected to a 10Mbps or 100Mbps network. The GT-48004A interfaces to an IEEE 802.3 10/100 Mbps MII compatible PHY device. The data path consists of a separate nibble-wide stream for both transmit and receive activities. The GT-48004A can switch automatically between 10 or 100 Mbps operation depending on the speed of the network. Data transfers are clocked by the 25 MHz transmit and receive clocks in 100 Mbps operation, or by 2.5 MHz transmit and receive clocks in 10 Mbps operation. The clock inputs are driven by the PHY, which controls the clock rate based on auto-negotiation.

13.2 Media Access Control (MAC)

The GT-48004A MAC performs all of the functions of the 802.3 protocol such as frame formatting, frame stripping, collision handling, deferral to link traffic, etc. The GT-48004A ensures the any outgoing packet complies with the 802.3 specification in terms of preamble structure. The GT-48004A transmits 56 preamble bits before Start of Frame Delimiter (SFD). The GT-48004A operates in half-duplex or full-duplex modes. In half-duplex mode, the GT-48004A checks that there is no competitor for the network bus before transmission. In addition to listening for a clear line before transmitting, the GT-48004A handles collisions in a pre-determined way. If two nodes attempt to transmit at the same time, the signals collide and the data on the line is garbled. The GT-48004A listens while it is transmitting, and it can detect a collision. If a collision is detected, the GT-48004A transmits a 'JAM' pattern and then delays its re-transmission for a random time period determined by the backoff algorithm. In full-duplex mode, the GT-48004A transmits unconditionally.

13.3 Auto-negotiation

13.3.1 Disabled

When EnAutoNeg* is HIGH, auto-negotiation is disabled both ports and each port can be selected to be in half- or full-duplex mode independently. Following RESET the port duplex mode is set by the state sampled on DAddr[6] for Port 0 and DAddr[7] for Port 1 (for each FEU). This value can be overridden in each port's Port Control register. The speed that each port operates in (10Mbps or 100Mbps) is determined by the frequency of TxClk[x] and RxClk[x] generated by the PHY. When the port is operating at 10Mbps, the PHY generates a 2.5MHz clock for both TxClk and RxClk. When the port is operating at 100Mbps, the PHY generates a 25MHz clock for both TxClk and RxClk.

13.3.2 Enabled

When EnAutoNeg* is LOW, auto-negotiation is enabled for both ports and the GT-48004A decodes the duplex mode for each port from the values of the Auto-Negotiation Advertisement register and the Auto-Negotiation Link Partner Ability registers at the end of the Auto-Negotiation process. (Note: The auto-negotiation feature on the 48002A is used ONLY to tell the 48002A-P-2 the duplex that each port is operating in. The SPEED (10/100) of each port is determined ONLY by RxClk and TxClk.) Once the duplex mode is resolved, the GT-48004A updates the port control registers with that duplex mode. The GT-48004A will continuously perform the following operations for each port (PHY addresses 1

and 2 alternately), implemented as READ commands issued via the MDC/MDIO interface:

1. Read the PHY Auto-Negotiation Complete status. As long as PHY bit 1.5 (Register 1, bit 5) is '0' switch to Half-Duplex mode and continue to read PHY register bit 1.5. Continue to step 2 when PHY bit 1.5 is '1', signaling Auto-negotiation is complete.

Steps 2 through 6 are performed once for every transition of PHY bit 1.5 from '0' to '1'. Once PHY bit 1.5 remains '1' and PHY registers 4 and 5 have already been read, the GT-48004A will continue to read PHY register 1, and monitor PHY bit 1.5. Steps 2 to 6 are performed once, if after Rst* de-assertion, the PHY bit 1.5 is read as '1', in order to update the GT-48004A duplex mode.

NOTE: PHY bit 1.2 (Link Status) is read and latched during this same register read operation, regardless of the Auto-Negotiation status.

2. Read the Auto-Negotiation Advertisement register, PHY Register 4. Continue to step 3.
3. Read the Auto-Negotiation Link Partner Ability register, PHY Register 5. Continue to step 4.
4. Resolve the highest common ability of the two link partners in the following manner (according to the 802.3u Priority Resolution clause 28B.3):

```
if (bit 4.8 AND bit 5.8) == '1' then ability is 100BASE-TX Full Duplex
else if (bit 4.9 AND bit 5.9) == '1' then ability is 100BASE-T4 Half Duplex
else if (bit 4.7 AND bit 5.7) == '1' then ability is 100BASE-TX Half Duplex
else if (bit 4.6 AND bit 5.6) == '1' then ability is 10BASE-T Full Duplex
else ability is 10BASE-T Half Duplex;
```

Continue to step 5.

5. Resolve the duplex mode of the two link partners in the following manner:

```
if ( (ability == "100BASE-TX Full Duplex") or (ability == "10BASE-T Full Duplex") ) then
duplex mode = FULL DUPLEX
else duplex mode = HALF DUPLEX;
```

NOTE: the value of the duplex mode indication should change only after reading both PHY registers 4 and 5. Continue to step 6.

6. Update the Port Control Register by writing the correct duplex mode bit. Continue with step 1.

13.3.3 Auto-negotiation Control Per Port

Since EnAutoNeg* pin enables or disables auto-negotiation for both ports, applications which require one port to auto-negotiate and one port to be forced into a certain mode require an additional PLD to interface between the GT-48004A's MDIO and MDC pin and the PHY.

While the GT-48004A waits for auto-negotiation to complete on a port, it forces that particular port to half-duplex mode. If this port is not auto-negotiation capable (but supports full duplex mode), or the port needs to be forced to full-duplex while the other is left to auto-negotiate, the PLD will monitor the MDIO line and make some modifications to the data read by GT-48004A from the PHYs. Therefore, instead of defaulting to half-duplex, the modified data read back by the GT-48004A will place the port in full-duplex. In essence, the PLD mimics the completion of the auto-negotiation cycle with a device that is not capable to auto-negotiate. For detailed information about implementing auto-negotiation on a per port basis, including the required PLD equations, please download the application note located on Galileo's website (<http://www.galileoT.com>).

13.3.4 Auto-Negotiation, Software Detection

Auto-negotiation is enabled or disabled via a hardware pin. The status of auto-negotiation (enabled or disabled) can be indirectly evaluated via software. When auto-negotiation is enabled, the CPU does not have write control to the duplex mode bit in the port control register of both ports. The CPU only has read access of these bits. Therefore, if a CPU can write modify these bits, and then read these bits to detect a change, it can conclude that auto-negotiation is DISABLED. Otherwise, if it cannot change the value of these bits, it can conclude that auto-negotiation is ENABLED. To avoid detecting changes in this bit during an auto-negotiation cycle, this write/modify/read should be done several times.

13.4 Backoff Algorithm Options

The GT-48004A implements the truncated exponential backoff algorithm defined by the 802.3 standard. Aggressiveness of the backoff algorithm used by all of the ports is controlled by the Limit4 pin. Limit4 controls the number of consecutive packet collisions that will occur before the consecutive collision counter is reset. When Limit4 is LOW, the GT-48004A resets the collision counter after 16 consecutive retransmit trials, restarts the backoff algorithm, and continues to try and retransmit the frame. A packet which is endlessly colliding on re-transmits will continue to be re-transmitted forever, only changing backoff intervals. The GT-48004A supports port partitioning on consecutive collisions, a mode which must be enabled by the CPU. The retransmission is done from the data already stored in the DRAM. In the case of a successful transmission, the GT-48004A is ready to transmit any other frames queued in its transmit FIFO within the minimum IPG of the link.

When Limit4 is HIGH, the GT-48004A will reset its collision counter and restarts the backoff algorithm after 4 consecutive transmit trials. This results in the GT-48004A being more aggressive in acquiring the media following a collision. This will result in better overall switch throughput (less packet loss) in standardized tests. Limit4 can be toggled during switch operation.

13.5 Data Blinder

The data blinder field (DataBlind in the Serial Parameters register) sets the period of time during which the port does not look at the wire to decide to transmit (inhibit time.) The default value is 32 bit times.

13.6 Inter-Packet Gap (IPG)

IPG is the idle time between any two successive packets from the same port. The default (from the standard) is 9.6uS for 10Mbps Ethernet and 960nsec for 100-Mbps Fast Ethernet. Note that the IPG can be made smaller or larger than the Ethernet standards by programming. Making the IPG smaller can improve test scores at the cost of Ethernet compatibility (a trick used by many vendors during head-to-head magazine tests.) We do not recommend this mode of operation, however, as it violates IEEE standards.

IPG is programmable in the Serial Parameters register.

13.7 10/100 Mbps MII Transmission

When the GT-48004A has a frame ready for transmission, it samples the link activity. If the RxDV signal is inactive (no activity on the link), and the Inter-packet gap (IPG) counter has expired, frame transmission begins. The data is transmitted via pins TxD[3:0] of the transmitting port, clocked on the rising edge of TxClk. The signal TxEn is asserted at this same time. In the case of collision, the PHY asserts the CoL signal on the GT-48004A which will then stop transmitting the frame and append a jam sequence onto the link. After the end of a collided transmission, the GT-48004A will back off and attempt to retransmit once the backoff counter expires.

A waveform of the signals which are synchronous to TxClk (TxD0[3:0], TxD1[3:0], TxEn[1:0]) is shown in Figure 10. The actual delay times of the GT-48004A are tighter than the IEEE 802.3u standard, clause 22.3.1, as shown in Table 25.

Figure 10: MII Transmit Signal Timing




Table 25: MII Signal Timings Synchronous to TxClk

IEEE 802.3u Spec.						
Name	Parameter	MIN	MAX	MIN	MAX	Units
vd	Valid Delay after Rising TxClk	2	14	0	25	ns

13.8 10/100 Mbps MII Reception

Frame reception starts with the assertion of RxDV (while the GT-48004A is not transmitting) by the PHY. Once RxDV is asserted, the GT-48004A will begin sampling incoming data on pins RxDV[3:0] on the rising edge of RxClk. Reception ends when the RxDV is deasserted by the PHY. The last nibble sampled by the GT-48004A is the nibble present on RxD[3:0] on the last RxClk rising edge in which RxDV is still asserted. During reception, the RxDV is asserted. If, while RxDV is asserted, the GT-48004A detects the assertion of RxEr, it will designate this packet as corrupted. While no reception is taking place, RxDV should remain deasserted.

A waveform of the signals which are synchronous to RxClk (RxD0[3:0], RxD1[3:0], RxDV[1:0], RxEr[1:0]) is shown in Figure 11. The setup and hold times of the GT-48004A are tighter than the IEEE 802.3u standard, clause 22.3.2, as shown in Table 26.

Figure 11: MII Receive Signal Timing




Table 26: MII Signal Timings Synchronous to RxClk

IEEE 802.3u Spec.						
Name	Parameter	MIN	MAX	MIN	MAX	Units
ts	Setup Time to Rising RxClk	6		10		ns
th	Hold Time after Rising RxClk	1		10		ns

13.9 10/100 Mbps Full-Duplex Operation

When operating in Full-duplex mode the GT-48004A can transmit and receive frames simultaneously. In full-duplex mode, the CrS signal is associated with received frames only and has no effect on transmitted frames. The Col signal is ignored by the GT-48004A while in Full-duplex mode. Transmission starts when TxEn goes active. Transmission starts regardless of the state of RxDV. Reception starts when the RxDV signal is asserted indicating traffic on the receive port of the PHY.

13.10 Illegal Frames

The GT-48004A will discard all illegal frames and increment the appropriate error MIB counters. Examples include: runts (less than 64 bytes), oversize (greater than 1518 or 1522 bytes), and bad FCS (bad CRC.)

13.11 Partition Mode

A port enters Partition Mode when more than 61 consecutive collisions are seen on the port (whether the port is running at 10 or 100 Mbps). In Partition Mode the port continues to transmit but it will not receive. The PaEn bit in the corresponding Port Control register is set when a port is partitioned. A port is returned to normal operation mode when a good packet is seen on the wire.

13.11.1 Enabling Partition Mode

Partition is enabled (for all ports) by setting the enable bit in the GT-48004A Control Register. The default value is Partition disabled for all ports. You must have a CPU in the system to enable partition mode, there is no pin strapping option.

13.11.2 Entering Partition State

When Partition is enabled, a port will enter Partition state when either of the following two situations occur:

- The port detects a collision on every one of 61 consecutive retransmit attempts of the same packet.
- The port detects a single collision which occurs for more than 512 bit times

While in Partition state:

- If the interrupt is not masked, the GT-48004A will issue an interrupt to the CPU upon entering Partition state, and will set the partition bit of that port in the Interrupt Cause register.
- The port will continue to transmit its pending packet, regardless of the collision detection, and will not follow the usual Backoff Algorithm. Additional packets pending for transmission, will be transmitted, while ignoring the internal collision indication. This frees the port's transmit buffers which would otherwise be filled up at the expense of the other ports' buffers. The assumption is that Partition is a system failure situation (bad connector/cable/station), thus dropping packets is a small price to pay vs. the cost of halting the switch due to full buffers.
- The Partition Indication is available via the LED interface (both the status led - blinking twice, and a dedicated led - on constantly).

13.11.3 Exiting from Partition State

The port will exit from Partition state, following the end of a successful packet transmission or packet reception. A successful packet transmission will be declared, if no collisions were detected during packet transmit or receive. If the interrupt is not masked, the GT-48004A will issue an interrupt to the CPU upon exiting from Partition state, and will clear the partition bit of that port in the Interrupt Cause register.

13.12 Back-pressure

Back-pressure is not supported by the GT-48004A. Back-pressure can greatly deteriorate real network throughput in favor of better standardized test scores. Back-pressure works is by "jamming" an entire network segment when the switch cannot accept new transmissions. This blocks *all* traffic, including traffic between nodes on the same network segment (like blocking local phone calls when the long distance circuits are busy.) Simply dropping packets is a better solution for overall network throughput, although this is not reflected in some simplistic standardized tests. (Dropped

packets are detected by higher-layer protocols and regenerated.)

13.13 VLAN Tagging Support

The GT-48004A has the ability to receive and transmit Ethernet frames up to 1522 bytes in length, thereby accommodating the standard 802.1Q VLAN tagging bytes (four extra bytes.) This feature is enabled/disabled by bit 9 in the Port Control Register (0x040200 - 0x040204). The default is to have this feature disabled (i.e. bytes longer than 1522 are discarded as over-size frames.)

14. MII Management Interface (SMI)

The GT-48004A MAC contains an MII Management Interface (SMI) for an MII compliant PHY devices. This allows control and status parameters to be passed between the GT-48004A and the PHY (parameters specified by the CPU) by one serial pin (MDIO) and a clocking pin (MDC), reducing the number of control pins required for PHY mode control. Typically, the GT-48004A will continuously query the PHY devices for their link status, without CPU intervention. The predefined PHY addresses for the link query are 1 and 2 (out of possible 32 addresses). This protocol complies with the National DP83840 PHY device as well as other available PHYs.

A CPU connected to the GT-48004A has access to all of the PHY addresses/registers, by writing and reading to/from a dedicated set of GT-48004A SMI control registers. The SMI allows the CPU to have direct control over an MII-compatible PHY device via the GT-48004A SMI control register. This allows the driver software to place the PHY in specific modes such as Full Duplex, Loopback, Power Down, 10/100 speed selection as well as control of the PHY device's Auto-Negotiation function, if it exists. The CPU writes commands to the GT-48004A SMI register and the GT-48004A reads or writes control/status parameters to the PHY device via a serial, bi-directional data pin called MDIO. These serial data transfers are clocked by the GT-48004A MDC clock output.

14.1 SMI Cycles

The SMI protocol consists of a bit stream that is driven or sampled by the GT-48004A on each rising edge of the MDC clock. The bit stream format of the SMI frame is described in Table 27.

Table 27: SMI Bit Stream Format

	PRE	ST	OP	PhyAd	RegAd	TA	Data	IDLE
READ	1...1	01	10	AAAAA	RRRRR	Z0	D.D(16)	Z
WRITE	1...1	01	01	AAAAA	RRRRR	10	D.D(16)	Z

- PRE (Preamble). At the beginning of each transaction, the GT-48004A sends a sequence of 32 contiguous logic one bits on MDIO with 32 corresponding cycles on MDC to provide the PHY with a pattern that it can use to establish synchronization.
- ST (Start of Frame). A Start of Frame pattern of 01.
- OP (Operation Code). 10 - Read; 01 - Write
- PhyAd (PHY Address). A 5 bit address of the PHY device (32 possible addresses). The first PHY address bit transmitted by the GT-48004A is the MSB of the address.
- RegAd (Register Address). A 5 bit address of the PHY register (32 possible registers in each PHY). The first register address bit transmitted by the GT-48004A is the MSB of the address. The GT-48004A always queries the PHY device for status of the link by reading register 1, bit 2.
- TA (Turn Around). The turnaround time is a 2 bit time spacing between the Register Address field and the Data field of the SMI frame to avoid contention during a read transaction. During a Read transaction the PHY should not drive MDIO in the first bit time and drive '0' in the second bit time. During a write transaction, the GT-48004A drives a '10' pattern to fill the TA time.
- Data (Data). The data field is 16 bits long. The PHY drives the data field during Read transactions. The GT-48004A drives the data field during write transactions. The first data bit transmitted and received shall be bit 15 of the PHY register being addressed.
- IDLE (Idle). The IDLE condition on MDIO is a high impedance state. The MDIO driver is disabled and the PHY should pull-up the MDIO line to a logic one.

14.1.1 SMI Timing Requirements

Figure 12 shows a waveform of the MDC line which is driven by the GT-48004A. Table 28 shows typical MDC timings.

Figure 12: GT-48004A MDC Waveform




Table 28: Typical MDC Timings

Name	Parameter	TYPICAL	Units
t1	MDC High Time	480	ns
t2	MDC Low Time	480	ns
t3 ¹	MDC Period	960	ns

1. MDC is generated internally by dividing the Clk clock input by 32. Clk clock frequency of 33MHz is assumed.

A waveform of the MDIO line when it is driven by the GT-48004A is shown in Figure 13. The actual delay times of the GT-48004A are tighter than the IEEE 802.3u standard, clause 22.3.4, as shown in Table 29.

Figure 13: GT-48004A MDIO Output Delay




Table 29: MDIO Signal Timings (GT-48004A Driving MDIO)

Name	Parameter	IEEE 802.3u Spec.				Units
		MIN	MAX	MIN	MAX	
t4	Rising MDC to Valid MDIO	10	50	0	300	ns

A waveform of the MDIO line when it is driven by the PHY is shown in Figure 13. The setup and hold times requirements of the GT-48004A are the same as the IEEE 802.3u standard, clause 22.3.4, as shown in Table 30.

Figure 14: GT-48004A MDIO Setup and Hold Time




Table 30: MDIO Signal Timings (PHY Driving MDIO)

Name	Parameter	MIN	MAX	IEEE 802.3u Spec.		Units
				MIN	MAX	
t5	Setup Time to Rising MDC	10		10		ns
t6	Hold Time after Rising MDC	10		10		ns

14.2 Link Detection and Link Detection Bypass (ForceLinkPass*)

Typically, the GT-48004A will continuously query the PHY devices for its link status, without CPU intervention. The pre-defined PHY addresses for the link query are 1 and 2 (out of possible 32 addresses). The GT-48004A will alternately read register 1 from PHY1 and PHY2 and update the internal link bits according to the value of bit 2 of register 1. In the case of “link is down” (i.e. bit 2 is ‘0’), that port will enter link test fail state. In this state, all of the port’s logic is reset. The port will exit from link test fail state only when the “link is up” i.e. bit 2 of register 1 is read from the port’s PHY as ‘1’.

The GT-48004A offers the option to disable the link detection mechanism by forcing the link state of both ports to the link test pass state. This is done with the ForceLinkPass* pin, which is sample at RESET. When ForceLinkPass* is LOW on RESET, the link status of both ports remains in the “link is up” state regardless of the PHY’s link bit value. When ForceLinkPass* is HIGH on RESET, the link status of the ports is read through the SMI from the PHY devices (register 1, bit 2).

15. Network Management Support

The GT-48004A supports the following management features:

- HP-EASE packet sampling technology
- Repeater MIB and PCI counters
- Port monitoring (sniffer) mode

This section describes the MIB counter, connectivity matrix and port mirroring functions. The HP-EASE functions are described in the following section.

15.1 Repeater MIB and PCI Counters

The GT-48004A incorporates a full set of Repeater MIB counters for each Ethernet port, as well as counters for activity on the PCI interface. Counters are accessed by the management CPU through the PCI interface.

The Repeater MIB counters include the following:

- Bytes Received
- Bytes Sent
- Frames Received
- Frames Sent
- Total Bytes Received (Good and Bad)
- Total Frames Received (Good and Bad)
- Multicast Frames Received
- Broadcast Frames Received
- CRC + Alignment Error
- Oversize Frames
- Fragments
- Jabber Frames
- Collision
- Late Collision
- Frames with length of 64 Bytes
- Frames with length of between 65-127 Bytes
- Frames with length of between 128-255 Bytes
- Frames with length of between 256-511 Bytes
- Frames with length of between 512-1023 Bytes
- Frames with length of between 1024-1522 Bytes
- MAC Receive Error (received packets with RxEr asserted)
- Dropped Frames

The global PCI counters are:

- PCI Frames Received
- PCI Frames Sent

Please see the register description section below for more information on the repeater MIB registers.

15.2 Monitoring (Sniffer) Mode

The CPU can program the GT-48004A to work in Monitoring mode for one of the two Fast Ethernet ports. In Monitoring mode, the GT-48004A sends all receive (including local traffic) and transmit packets to the CPU or to a port in one of the GT-48004A devices which was assigned to be the Sniffer. The packets that are forwarded to the Sniffer are not necessarily in a linear time order.

Monitoring Mode is enabled by setting bit 2 in the Port Control register. The target sniffer is written into the CPU and Sniffer Numbers register. Only one port of a single GT-48004A device can work in monitoring mode at a time.

15.3 Spanning Tree Support

The GT-48004A provides the hardware assistance for Bridge Spanning Tree Algorithm implementation. The Spanning

Tree algorithm itself is performed by a management CPU.

The GT-48004A includes a SpanEn bit in the Global Control register and additional SpanEn bits in each of the 8 Port Control registers. Table 31 summarizes the hardware assistance for the Spanning Tree algorithm.

Table 31: Spanning Tree Enable Bit Definition

SpanEn (Global)	SpanEn (Port)	Logic State	Remarks
0	x	Port Enable	No Spanning Tree. Treat BPDUs as regular Multicast.
1	1	Blocking, Listening, Learning	Transfer BPDUs to CPU. All receive/transmit packets are rejected, except BPDU messages from the CPU. Address learning disabled.
1	0	Forward	Transfer BPDU to the CPU. Accept all packets. Address learning enabled.

Note 1: The GT-48004A does not learn MAC addresses during the Spanning Tree ‘Learning’ stage (it is ‘learning’ the bridge topology while in this mode.) The GT-48004A only learns MAC addresses in the Forward mode.

Note 2: The CPU can send BPDU messages to a port of the GT-48004A which is disabled. The mechanism to send BPDUs from the CPU to a locked port is to send a BUFFER_REQUEST message like the format shown in Section 11.3.4, but with the LSB bits of the address as 0x58, instead of 0x0. This BUFFER_REQUEST message will cause the GT-48001A to allocate a buffer regardless of the state of that port.

15.4 Broadcast Storm Filtering

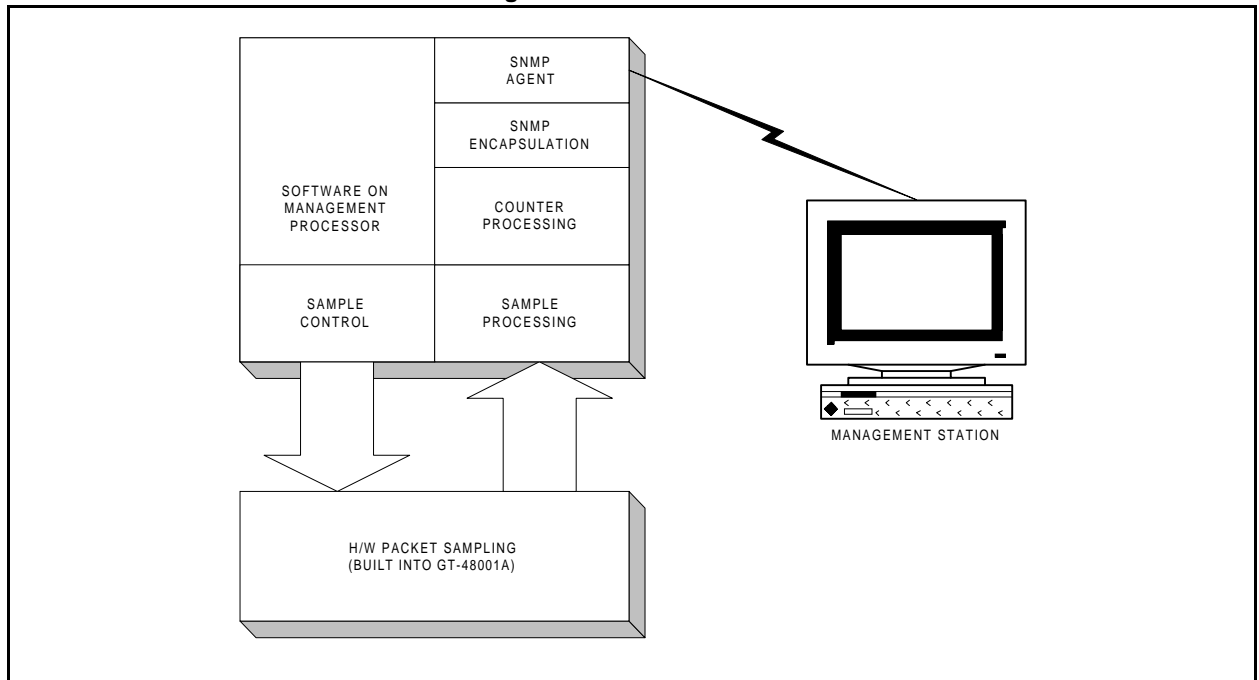
Excessive broadcast packets (broadcast “storms”) can be filtered in a managed switch by setting the DisBroad bit in the port control register. Broadcast packets can be re-enabled once the loops causing the broadcast storm are eliminated via the spanning tree algorithm.

16. HP-EASE Packet Sampling Technology

Hewlett-Packard's Embedded Advanced Sampling Environment (HP-EASE) is supported directly by custom hardware in the GT-48004A device. HP-EASE provides many of the management functions that RMON provides, but at a greatly reduced implementation cost. With the GT-48004A device, the switch OEM has the freedom to implement RMON, sampled RMON, EASE, or any of the above.

The HP-EASE functionality is best understood using the familiar stack model, as shown in Figure 15. An HP-EASE agent running on the management CPU, combines samples of packets passing through the switch, with snapshots of the MIB counters and forwards these to a management console through SNMP "trap" messages. This approach is different from RMON, as RMON keeps all data within the switch until polled by the management console. The GT-48004A provides the hardware packet sampling portion of the HP-EASE protocol. This capability can be used to implement a "true" HP-EASE compliant product, or as a method to implement sampled management such as statistical RMON.

Figure 15: HP-EASE Stack



A system implementing HP-EASE provides network management functions similar to RMON. However, a network device (or CPU) implementing HP-EASE is not required to have the extensive CPU or memory resources needed to fully implement all groups of RMON. EASE significantly reduces the requirement for these resources by statistically sampling data on network segments while RMON requires the processing of every network event. The philosophy of EASE is to off-load the intelligence and processing power required for network monitoring to the network management station rather than the network device.

16.1 HP EASE Technology Overview

A system with HP EASE provides network management functions similar to RMON. However, a network device (or CPU) implementing HP EASE is not required to have the extensive CPU or memory resources needed to fully implement all groups of RMON. EASE significantly reduces the requirement for these resources by statistically sampling data on network segments. RMON requires the processing of every network event. The philosophy of EASE is to off-load the intelligence and processing power required for network monitoring to the network management station rather than the network device.

HP EASE sampling requires a counter for each network segment. This counter indicates the number of packets to be skipped before a sample is taken. When the counter reaches zero, the next packet on the network segment is captured by the network device. Software then truncates the sampled packet, to some small fixed length, and appends a snapshot of specific MIB counters for that segment. The counter snapshot does not have to be taken simultaneously with the sample. Software may introduce a delay of some milliseconds after the packet is sampled by hardware, however minimizing this delay makes EASE more accurate. The newly created datagram is sent off to the network management station as an SNMP trap. The network management station records the sample and counters in a database, and uses the information to obtain traffic load estimates, top talker matrices, high-level protocol flows, and other useful sets of information. After the sample has been taken, the CPU loads the count-down counter with the next skip count to capture the next sampled packet. The skip count is a random value loaded by software.

EASE software in the network device must keep track of the last receive error sources and the associated error conditions. The network device keeps track of errors associated with received packets and informs the CPU of the source address (SA) of these error packets.

16.2 EASE Functionality on the GT-48004A

Support for EASE sampling is directly integrated in the GT-48004A chip, but requires the presence of a CPU in order to function, for enabling the EASE support as well as the sample packet processing. Each GT-48004A chip supports two network segments (one per each port) as well as a PCI system bus. Sampling will occur only on the network segments, and sampled packets will be sent to the CPU via the PCI system bus. Sampling is not performed on the PCI bus. It may, however, be performed on packets received from the PCI bus, but only as a function of the counters for the destination ports (i.e. packets entering the GT-48004A via the PCI bus and being transmitted through one or more ports). There is no counter for the PCI interface itself. Only good packets of valid length are sampled. All other packets are not sampled and do not affect the skip count. All counters and registers implemented in the GT-48004A chip in order to support EASE functionality, may be accessed by the CPU from the PCI system bus.

16.3 Ease_Register

A register is defined for each external port supported on the GT-48004A device. This register is used by the CPU to load the internal count down counter, described above, with a random skip count. The count-down counter is 20 bits in length and is used to actually determine when a sample is to be made. The GT-48004A implements a shadow register for each of the Ease_Registers. The shadow register address is the same address as the Ease_Register address. After a value has been written to the Ease_Register it is transferred to an internal 1 word deep FIFO (the shadow register) or directly to the actual count-down counter if that counter is currently idle and empty. If the value can not be transferred to the count-down counter, the value will be held in the Ease_Register shadow register until space becomes available (i.e. a sample has been taken). If the Ease_Register shadow register was written and the CPU does attempt to write a new value, the new value will silently replace the existing value. If the Ease_Register is empty at the time a new value needs to be loaded into the internal counter or the shadow register, the GT-48004A will simply wait, indefinitely, for the CPU to write a value into EASE_register. In this situation, EASE is effectively disabled on that port.

EASE (8 Registers), Offset: 0x040228 - 0x04022c

Bits	Field Name	Function	Initial Value
19:0	Ease_Register	value loaded to the internal count-down counter of port 0	0x0
31:20	-	Reserved.	-

16.4 EASE Interrupts

A status bit indicating the full/empty status of the Ease_Register for each external port supported on the GT-48004A, is maintained as part of the Interrupt Cause register. When a value is moved from the Ease_Register into an internal counter or shadow register, a bit is reset in the Interrupt Cause register indicating that the Ease_Register is now empty.

Setting this bit should also generate a processor interrupt. The Interrupt Cause register may be read to determine the state of the Ease_Registers, and may be written to clear the interrupt condition described above. It is possible for the CPU to mask the interrupt condition as well as clear the interrupt condition. The GT-48004A implements a mask bit in the Interrupt Mask register for each EASE status bit in the Interrupt Cause register. Masking and clearing the interrupts are executed in a way that is consistent with the other interrupts supported by the GT-48004A.

Interrupt Cause Register, Offset: 0x044

Bits	Field Name	Function	Initial Value
19	EaseReg0Empty	Ease_Register of port 0 is empty	0x0
20	EaseReg1Empty	Ease_Register of port 1 is empty	0x0
27	ErrorSASent	Error_Source message sent to CPU	0x0
other	various interrupt cause bits	left unchanged. See register section for details	

Interrupt Mask Register, Offset: 0x048

Bits	Field Name	Function	Initial Value
27:19	MaskBits	Mask the CPU interrupt line for the appropriate bits in the Interrupt Cause register.	0x0
other	various interrupt mask bits	left unchanged. See register section for details	0x0

16.5 Sampled Packet Indication

Sampled packets are copied into the CPU's receive buffers using the same mechanism as normal receive packets. The only difference, from the CPU's point of view, is that the GT-48004A will put an indication in the first word of the receive buffer which identifies the packet as a sample. The sample indication bits specify which ports on the particular GT-48004A the sample is associated with. It is possible for a single sample to be associated with more than one port at a time. For example, a broadcast packet flooded to all ports may be sampled on several ports if each of their skip counters had previously been decremented to zero.

Each GT-48004A device operates independently, so it is possible for the CPU to receive the same sample from different GT-48004A devices. For example, a broadcast packet flooded to all ports in the system may be sampled by several GT-48004As at the same time. Each sample will result in a separate copy of the packet being sent to the CPU. It is also possible to sample a packet which would normally be received by the CPU. In this case, only a single copy of the packet can be sent to the CPU. The CPU should be responsible for determining if a sampled packet should also be accepted as a normal receive packet. In the case where a normally received packet is also a sample from multiple GT-48004A devices (e.g. a broadcast packet), the GT-48004A must provide an indication which allows the CPU to avoid processing duplicate packets. This indication is provided by the GT-48004A which actually received the packet from an external link.

An additional bit in the packet header indicates that the sample packet was originally received from an external link to the CPU as opposed to the PCI system bus. Other GT-48004A devices which sampled the flooded packet only because it was received from the PCI interface and is being transmitted on a port whose internal counter was decremented to zero will not have this indication. These samples are "pure" samples and the CPU will know that it should not process the packet as a normally received packet.

The first word in each 2K block holding the packet to CPU contains the following bits:

- Sniffer (bit 31) (active HIGH)
- n/a (bits [30:18])
- EASE sample for port 1 (bit 17) (active HIGH)
- EASE sample for port 0 (bit 16) (active HIGH)
- EASE sample is an original packet to CPU (bit 15) (active HIGH)
- Source Channel number (bit 12, 1 - port1; 0 - port0;)
- Byte Count (bits [11:1], bit 11 is MSB)
- Valid bit (bit 0, active HIGH)

These parameters are written at the end of the packet transfer.

16.6 Error Source Indications

EASE software in the network device must keep track of the last receive error sources and the associated error conditions. The GT-48004A informs the CPU of error source conditions by writing the Error_Source message to a new Error_Source buffer area in CPU memory. Operations in the Error_Source buffer area are similar to those in the NEW_ADDRESS, START_OF_PACKET and Intervention buffer areas. There is an Error_Source Base Address Register in the GT-48004A in which the CPU writes a pointer to the Error_Source buffer area. The Error_Source buffer area is able to hold 32 entries. Two types of errors are defined for this procedure: FCS error and frames too long. When the GT-48004A receives a packet with any of the above conditions, it will generate and write an Error_Source message to the CPU's buffer area. The Error_Source message will contain the 48-bit source address of the error packet, the source port number and an indication of the error type. The CPU may poll the Error_Source buffer area for new messages. However, the GT-48004A includes a separate bit in the Interrupt Cause register which indicates that the GT-48004A has written an Error_Source message into the CPU's memory. An appropriate mask bit is defined in the Interrupt Mask register.

CPU Error Source Base Address, Offset: 0x140050

Bits	Field Name	Function	Initial Value
31:8	ErrorSourceBaseAdd	Contains a pointer to the CPU 'Error_Source' area. The area includes 32 entries (2 32-bit words each) for the GT-48004A's 'Error_Source' messages.	0x0
7:0	-	Reserved. Must be 0x0 when written.	-

'Error_Source': The data written by the GT-48004A device to the 'Error_Source' messages buffer area that contains information about an Error Source Address. The data format is as follows:

PCI Bits	Description
Address	
[31:8] [7:3] [2:0]	ErrorSourceBaseAdd offset pointer to entry '000'
Data 0	
[31:3] [2] [1] [0]	MAC address [19:47] 1 - n/a 1 - FCS Error 1 - Over Count Error
Data 1	
[31:25] [24] [23:19] [18:0]	reserved 0 - port 0; 1 - port 1; Device# (bit 23 is MSB) MAC address [0:18]

16.7 Enabling/Disabling EASE Functionality

An explicit HP EASE enable/disable bit is provided in the Global Control register for the GT-48004A device. When HP EASE is disabled using this bit, no EASE samples nor Error Source messages are sent to the CPU. HP EASE packet sampling can be disabled on a port anytime the internal counter can not be reloaded with a new skip count because the CPU has not provided any new values via the Ease_Register. Interrupt conditions generated by an empty Ease_Register can be masked by appropriate bits in the Ease_Full_Mask and/or Interrupt Cause registers.

16.8 Interaction With Other GT-48004A Features

Some GT-48004A features are incompatible with HP EASE, and will require that HP EASE be disabled. In most cases, it can be the responsibility of software to assure that HP EASE is disabled when used with other GT-48004A features.

- Broadcast Intervention Mode: HP EASE is independent of broadcast intervention mode.
- Unicast Intervention Mode: HP EASE should be disabled when using this mode.
- Sniffer Mode: HP EASE should be disabled on the GT-48004A device which one of its ports has been configured to work in monitoring mode (e.g. that port's Rx and Tx traffic are sent to the Sniffer).
- RMON Station-to-Station Matrix: HP EASE is independent of the Station-to-Station Matrix feature.
- Spanning Tree Support: HP EASE may only be enabled on GT-48004A which has its ports configured in the forwarding state. If the global Spanning Tree enable bit is set and the port is blocked, listening or learning, HP EASE should not be enabled. If the global Spanning Tree enable bit is clear or the port is in the forwarding state, HP EASE can be enabled. It can be the responsibility of software to assure that EASE is enabled correctly when used with Spanning Tree features.
- Address Learning: HP EASE in no way impacts the learning process of the GT-48004A device.
- LED Serial Interface: EASE packet sample indications are accessible via the serial LED interface mainly for debug purpose. LedData bit # 125: EASE sample indication for port 0. LedData bit # 126: EASE sample indication for port 1. (LedData bit#1 is the bit on which LedStb is HIGH). The LED circuit implements a "monostable" stretching function to enable viewing these dynamic signals.

17. DRAM Interface and Usage

The GT-48004A includes direct support two separate 1-2Mbyte arrays of 35nS EDO DRAM. The performance of 35nS EDO satisfies the required bandwidth for data transfer, address recognition and Tx descriptor fetch/update. The DRAM interface is entirely glueless. All accesses are performed as 32-bits. The DRAM interface is designed for 35ns EDO DRAMs and all timings are guaranteed to work with these devices. Refresh is performed automatically by the GT-48004A.

NOTE: 35ns EDO DRAMs are available from several vendors including Silicon Magic (www.simagic.com) and Mosel-Vitec (www.mosel-vitec.com).

Each FEU within the GT-48004A requires its own separate DRAM array. This is more efficient than a single 64-bit wide array as it allows address lookups and packet transfers to occur independently.

Each FEU requires about 300Kbytes of the DRAM for the address table and other private data structures. The remainder is used for packet buffers. Following power-up or system RESET, the GT-48004A device creates the MAC Address Table in DRAM, and initializes all locations in the table to indicate that invalid entries exist in all locations.

Galileo recommends using DRAM with 256K x 16 configuration. When using this configuration, 2 DRAM chips are required for 1 MByte, and 4 DRAM chips are required for 2 MBytes. If 1 MByte is selected, RAS0* should be connected to 2 DRAM chips while RAS1* should be left unconnected.

If 2 MBytes is selected, RAS0* will control the first 1MB bank, while RAS1* will activate the second 1MB bank. DData[31:0], DAddr[8:0], CAS*, and WE* should be connected to both banks.

Using 1 or 2 MBytes of DRAM is entirely up to the architect. 2MBytes increases the size of the Rx Buffer space. This performance advantage must be weighed against the cost of additional memory.

18. LED Support

The serial LED interface is similar to the 3-pin LED interface of the GT-48001A device which requires a PAL to interpret the LED bit stream. Galileo also provides reference designs and example PAL equations in the LED interface application note available on our website.

NOTE: The GT-48004A does not implement the parallel LED interface found on the GT-48002A.

18.1 LED Indications Interface Description

Table 32 shows the data accessible on the LED Indications Serial Interface for each of the GT-48004A ports.

Table 32: LED Signals Available

Data Description	Symbolic Signal Name	Type
Primary Port Status LED	primary_port_status	n/a
Transmit data in progress	transmit	dynamic
Receive data in progress	receive	dynamic
Collision active	collision	dynamic
Forwarding of unknown packets enabled	unknown_enable	static
The port is configured as Sniffer	port_is_sniffer	static
Full/Half duplex	full_duplex	static
Receive Buffer Full	rx_buffer_full	dynamic

18.2 Detailed LED Signal Description

18.2.1 Primary Port Status LED

The Primary Port Status LED indicates the port status in two operation modes, selectable via the LEDMode input.

18.2.1.1 Primary Port Status LED in Mode 0: (LEDMode input is LOW)

In this mode, the Port Status LED provides the following information:

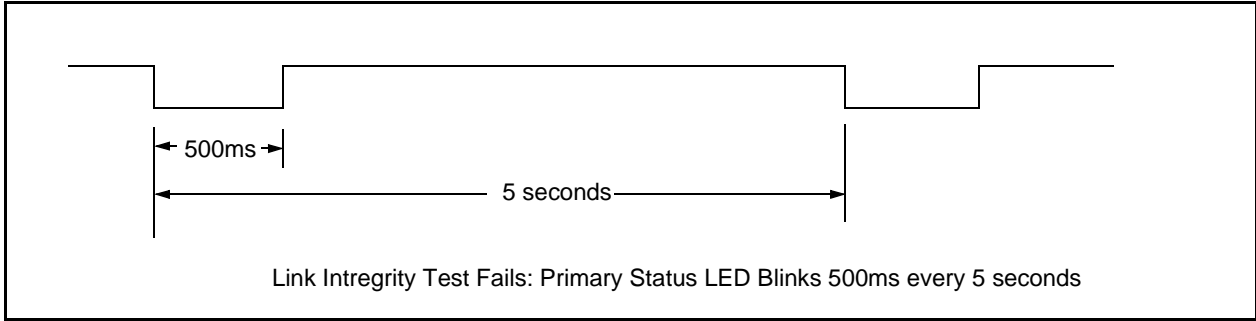
```

if Port is disabled
    Port Status LED is OFF;
else if Link Integrity test failed
    Port Status LED blinks once;
else if Partition State detected
    Port Status LED blinks twice;
else Everything is OK (Port Status LED is ON)

```

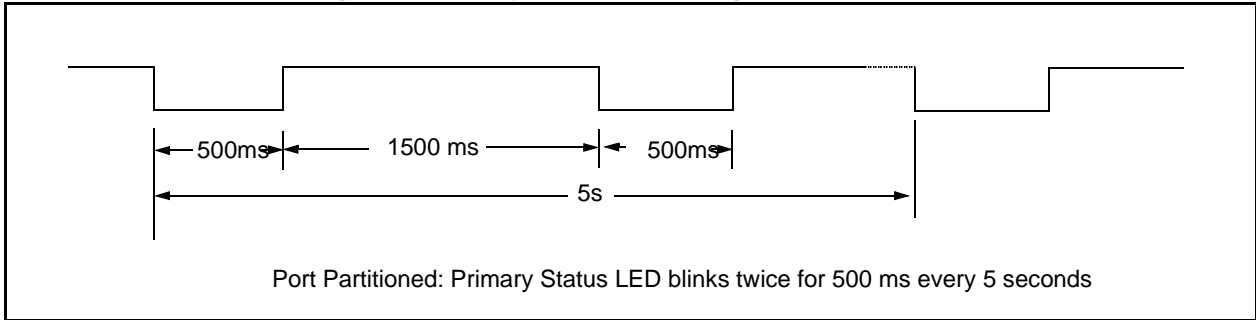
18.2.1.2 Status LED blink timing (Mode 0)

Link Integrity test failed, Status LED blinks once. Primary status bit is active for 500 ms every 5s.



Partition, Status LED blinks twice. Primary status bit is activated twice every 5seconds for 500 ms each time, with a period of 1500 ms between two consecutive activations.

Figure 16: Primary Status LED Timing (Port Partitioned)



18.2.1.3 Primary Port Status LED (Mode 1): (LEDMode input is HIGH)

In this mode, the Port Status LED initially displays the port link status information (indication type a) and then switches to reflect the port traffic (Transmit or Receive) activity (indication type b). The switching between the two types of indications is as follows:

Indication type a:

The Port Status LED indicates the port link status for a period of about 3 to 3.5 seconds (active - link test passes, inactive - link test fails) following any of the events below:

1. Rst* deassertion
2. Transition between link test fail to link test pass

Following this time period, the Port Status LED will switch to indication type b. In the case that the link test fails, the Port Status LED will remain inactive and will not switch to indication type b.

Indication type b:

The Port Status LED indicates the port traffic (Transmit or Receive) activity which is a logical OR of the TxEn active and RxDV active dynamic signals. The “monostable” function is applied to this indication type so the LED can be viewed for a period of about 62 ms for LEDMode 0, or for about 7.5 ms for LEDMode 1, per each traffic activity. The Port Status LED will switch to indication type a on the following cases:

1. Rst* assertion

2. Transition from link test pass to link test fail

18.2.2 Transmit data in progress

This signal indicates the port is transmitting data.

18.2.3 Receive data in progress

This signal indicates port receive activity.

18.2.4 Collision active

This signal indicates the port collision event detected by the port.

18.2.5 Full/Half duplex

This signal indicates the port duplex: active - full duplex, inactive - half duplex.

18.2.6 Receive Buffer Full

In order for this LED to be active, Rx Buffer Threshold must be enabled (see Table 4). This signal indicates the port receive buffer status: active - the buffer exceeds it's programmed threshold, inactive otherwise.

18.2.7 Forwarding of unknown packets enabled

This signal indicates the port mode of forwarding unknown packets: active - forwarding unknown packets enabled, inactive otherwise.

18.2.8 The port is configured as Sniffer

This signal indicates if the port mode is configured as a sniffer target port: active - port is a target sniffer, inactive otherwise.

18.2.9 Link Fail State

This signal indicates the port link status: active - link is down, inactive - link is up.

18.2.10 Partition State

This signal indicates the port partition status: active - port entered partition state, inactive otherwise.

18.2.11 Secondary Port Status LED

Indicates the secondary status mode as per the inverted value of LEDMode input.

18.2.12 Pure Port Status LED

This signal will be inactive for any of the following events:

- Port is disabled
- Link Integrity Test Failed
- Partition State Detected

Otherwise, this signal is active.

18.3 LED Signals Timing Type

18.3.1 Static LED Signals

These signals are stable for relatively long time periods. The LED indication directly reflects their current value. The static signals are:

- Port Status (LEDMode 0)
- Forwarding of Unknown packets enabled
- Port Configured as Sniffer
- Full/Half Duplex

18.3.2 Dynamic Internal Signals:

These signals are typically active for short time periods. In order to be visible through the LED Indication Interfaces, the GT-48004A includes a “monostable” function per each of these dynamic signals so they can be viewed on the LED indication output for a period of about 62 ms in LEDMode 0 and 7.5 ms in LEDMode 1. The dynamic signals are:

- Port Status (LEDMode 1)
- Transmit data in progress (TxEn)
- Receive data in progress (RxDV)
- Collision active (Col)
- Receive Buffer Full

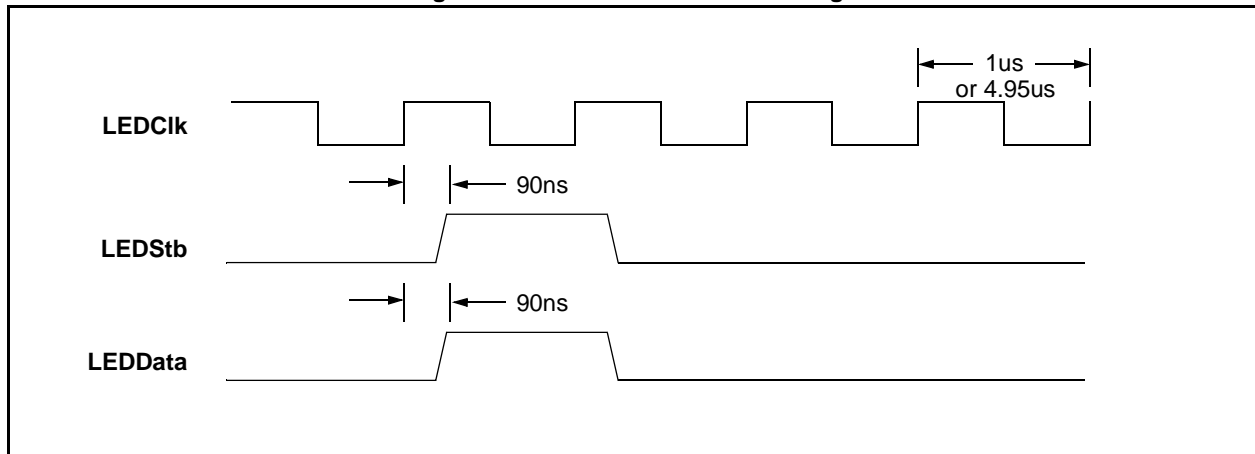
18.4 Serial LED Interface Description

The LED serial interface consists of three outputs:

- **LEDClk**: LEDClk is the primary timebase of the LED Indications Interface. It is a 50% duty cycle free running clock at a fixed frequency selectable via the LEDMode input. If LEDMode is LOW, LEDClk will be 1 MHz. If LEDMode is HIGH, LEDClk will be 202 KHz. LEDClk is HIGH-Z when Rst* is asserted.
- **LEDStb**: LEDStb (active HIGH) indicates the beginning of the data frame. LEDStb is activated for a duration of one LEDClk cycle once every 128 LEDClk cycles, starting from Rst* deactivation. This signal marks the beginning of the 128 bit long LED data frame. LEDStb transitions occur 90 ns after LEDClk rising edge.
- **LEDData**: The internal signals are multiplexed on the LEDData output for every data frame. LEDStb activation signals the presence of data bit #1 (out of 128 bits) on the LEDData output. LEDData transitions occur 90 ns after LEDClk rising edge. All internal signals accessible via LEDData are active HIGH internally and are inverted on the LEDData output (i.e. when an internal signal is active, the data bit on the LEDData output will be LOW). For example: If port 0 transmits data, the internal_event_transmit[0] signal is active HIGH and the corresponding bit 9 in the LEDData serial stream is LOW.

The timings for the LED serial interface are shown in Figure 17.

Figure 17: Serial LED Interface Timings



18.4.1 Table of Internal Activities/Status Driven via the Serial LED Interface

The following table defines a bit by bit description of the internal signals driven through the LED Indications Serial Interface. The bit number refers to the activation of LEDStb. LEDStb is active for bit# 1. RESERVED bits contents is not defined (i.e. can be either HIGH or LOW).

Table 33: LED Signals

Bit Number	Signal	Bit Number	Signal
1	primary_port_status[0]	22	port_is_sniffer[1]
2	primary_port_status[1]	23	full_duplex[1]
3-8	PRIVATE	24-72	-reserved-
9	transmit[0]	73	link_test_fail[0]
10	receive[0]	74	link_test_fail[1]
11	collision[0]	75-80	-reserved-
12	rx_buffer_full[0]	81	partition[0]
13	unknown_enable[0]	82	partition[1]
14	port_is_sniffer[0]	83-88	-reserved-
15	full_duplex[0]	89	secondary_port_status[0]
16	-reserved-	90	secondary_port_status[1]
17	transmit[1]	91-96	-reserved-
18	receive[1]	97	pure_port_status[0]
19	collision[1]	98	pure_port_status[1]
20	rx_buffer_full[1]	99-128	-reserved-
21	unknown_enable[1]		

19. Interrupts

The GT-48004A signals interrupts to a management CPU via the PCI INTA# pin. Interrupts are maskable through the Interrupt Mask register and the interrupt source is determined through the Interrupt Cause register. The Interrupt Mask register defaults to masking all interrupts. A '0' in the appropriate bit means that particular interrupt will be masked. A '1' in the appropriate bit means that particular interrupt will not be masked. The default is that all interrupts are masked.

Interrupts are cleared by writing '0' to the corresponding bit in the Interrupt Cause register. Writing '1' to a bit in the Cause register has no effect.

NOTE: There is an Interrupt Cause and Mask register in each FEU. The interrupts requested by each FEU are ORed before being brought out to the package pin.

20. RESET Configuration

The GT-48004A uses several pins as configuration inputs to set certain parameters following a RESET. The definition of the configuration pins changes immediately after RESET to their usual function.

20.1 Configuration Pins

Configuration pins must be pulled up or down externally at RESET to select the desired operational parameter. The recommended value of the pull-up/down resistors is 4.7K ohms. Table 34 shows the configuration pins for the GT-48004A.

Table 34: RESET Pin Strapping Options

Pin	Configuration Function
A/BDAddr[4:0]	Device Number: Each FEU must have a different device number.
A/BDAddr[5]	DRAM Size for FEU
0-	2Mbyte
1-	1Mbyte
A/BDAddr[6]	Half/Full Duplex Mode for Port 0/2
0-	Half Duplex
1-	Full Duplex
A/BDAddr[7]	Half/Full Duplex Mode for Port 1/3
0-	Half Duplex
1-	Full Duplex
A/BDAddr[8]	DRAM Type for each FEU
0-	Reserved
1-	EDO
LEDMode*	LED Mode
0-	LEDMode 0
1-	LEDMode 1
ForceLinkPass*	Force Link Pass
0-	Read Link Status from the PHY via SMI
1-	Force Link Status to "link is up"

20.2 Configuration Input Timings

The configuration inputs have two timing requirements:

- setup/hold time to clock (as any synchronous input)
- setup of at least 10 clock cycles before RESET de-assertion (rising edge).

You can guarantee these parameters by using resistors to strap the configuration pins and delaying RESET de-asser-

tion until least 10 clock cycles after the clock is stable.

21. Switch Expansion

The GT-48004A is designed for switches of up to 16 Fast Ethernet ports. In such a configuration, there are 4 GT-48004A devices and a single GT-64120 device on the local Fast PCI bus. At 50MHz (66MHz target) board design will require careful attention to trace lengths, loading effects, etc. Please refer to Galileo's 16-port GT-48004A reference design for tips. Expanding beyond 16 ports may be possible through 66MHz PCI-to-PCI bridges, although such a design will be challenging from a board design standpoint.

For customers not concerned with ultimate performance, it is possible to expand beyond 16 ports by tying the MII interfaces of 2 segments together and building a "bridge in a box". This is, of course, a blocking architecture but is adequate for lower performance switches, or for those building systems to compete with 10/100 autonegotiating repeaters (segmented repeaters.)

22. Development Tools

Galileo Technology and a number of third party vendors offer development tools for the GT-48004A. Vendors wishing to join Galileo's third party developers program should contact Galileo's Ethernet marketing group (ethernet@galileoT.com.)

22.1 Evaluation Platforms/Reference Designs

Galileo will provide an evaluation board/reference design for the GT-48004A. Details will be included in a future revision of this document.

22.2 Verilog Models

Galileo provides Verilog models for all components, including the GT-48004A. Please contact your local sales representative for pricing information.

22.3 Complimentary Products

Galileo manufactures a line of CPU core logic chips for MIPs processors that incorporate PCI interfaces. The GT-64120 provides all core logic functions for 64-bit bus MIPs processors: RM5260, R4600, R4650, R4700 and R5000 as well as a two 32-bit Fast PCI interfaces.

23. Register Tables

The GT-48004A incorporates the required PCI Configuration registers, Command registers and various counters for management purposes. The GT-48004A can work in stand-alone mode, in which there is no requirement for CPU intervention (a system with no CPU) in which case the default values of the control registers are used.

The actual address of an internal register is the sum of the GT-48004A base address and the particular register's offset. The CPU defines the base address by writing to the Internal Registers Base Address register in the PCI Configuration area.

Note that the internal register map of each FEU is separate and the base address is controlled by the device number assigned to each FEU.

23.1 Register Map

Description	Offset
Internal Control Registers	
Global Control	0x140028
Port Control 0	0x040200
Port Control 1	0x040204
Status	0x14002c
CPU and Sniffer Numbers	0x140030
Interrupt Cause	0x044
Interrupt Mask	0x048
Serial Parameters	0x040220
Receive Buffers Threshold	0x040224
CPU Buffer Base Address	0x140034
CPU Start Of Packet Base Address	0x140038
CPU New Address Base Address	0x14003c
CPU Intervention Base Address	0x140048
Device Table	0x40
Timeout Counter	0x04c
Port 0 Counter Block	0x040000 - 0x040054
Port 1 Counter Block	0x040058 - 0x0400ac
PCI Global Counters	0x140040 - 0x140044
SMI Register	0x14004c
PCI Configuration	
Device and Vendor ID	0x000
Status and Command	0x004
Class Code and Revision ID	0x008
Header Type, Latency Timer, Cache Line	0x00c
Internal/DRAM Base Address Register	0x010
Interrupt Pin and Line	0x03c

23.2 Internal Control Registers

Global Control, Offset: 0x140028

Bits	Field Name	Function	Initial Value
0	DisLearnPro	Disable Learning Process. 0 - Learning process is enabled 1 - Learning process is disabled. The GT-48004A will not learn any new addresses from the Ethernet ports	0x0
1	RMONEEn	RMON Enable. 0 - RMON disabled 1 - The GT-48004A enters the RMON mode (Station-to-Station connectivity matrix) and asserts the ChipSel* pin when it reads the packet's Byte Count and Source and Destination Addresses from the DRAM	0x0
2	DevTabMod	Device Table Mode. 0 - The GT-48004A updates the appropriate bits in the table upon master abort 1 - The CPU updates the Device Table	0x0
3	-	Reserved.	0x1
4	DRAMArbPri	DRAM Arbiter Priority. This bit indicates the DRAM arbitration scheme of the four GT-48004A internal units as follows: 0 - 1) Frame Control unit, 2) Switching Core unit, 3) PCI and InterPCI Control units in round robin scheme. 1 - 1) Frame Control unit, 2) Switching Core unit, 3) InterPCI Control unit, 4) PCI Control unit	0x0
5	DescArbPri	Descriptor Control Arbiter Priority. This bit indicates the descriptor control arbitration scheme as follows: 0 - Round robin between the PCI side and the 2 Ethernet ports. The 2 ports have equal priority. 1 - PCI has higher priority than the 2 ports. The priority is: PCI, Port 0; PCI, Port 1, PCI, Port 0; PCI, Port 1...	0x1
6	ForwUnk	Forward Unknown Packets. It defines whether the GT-48004A will forward Unknown packets to the CPU or not. 0 - Do not forward 1 - Forward	0x0
7	ForwNewAdd	Forward New Address. It defines whether the GT-48004A will forward new address messages to the CPU or not. 0 - Do not forward 1 - Forward	0x0
8	RecEn	Recovery Enable. It defines whether the recovery process is enabled or not. 0 - Disabled 1 - Enabled	0x1

Bits	Field Name	Function	Initial Value
9	SnifTyp	Sniffer Type. This bit indicates the Sniffer type. The Sniffer can be a CPU or a dedicated port in one of the GT-48004A devices which was assigned to be the Sniffer. 0 - CPU type 1 - GT-48004A type	0x1
10	CPUEn	CPU Enable. This bit indicates that there is a CPU in the system. 0 - CPU does not exist 1 - CPU exists	0x0
11	RMONTToPCI	RMON to PCI Enable. Meaningful only when RMONEEn bit is set. 0 - The GT-48004A reads the DA/SA of the packet that is forwarded only to the local ports. 1 - The GT-48004A reads the DA/SA of the packet that is forwarded to the local ports and the PCI.	0x0
19:12	-	Reserved.	0x0
20	BufThrEn	Buffer Threshold Enable. 0 - There is no limitation on the buffers' allocation (other than physical memory size.) 1 - The buffers allocated to the ports and the PCI are limited to the number which is written in the Rx Buffers Threshold register. This bit is meaningful only when DisBufThr* pin is disabled.	0x1
21	-	Reserved.	0x0
22	ForwMulti	Forward Multicast. 0 - The GT-48004A forwards Multicast packets to all the ports and to the CPU. 1 - Multicast packets forwarded only to the CPU.	0x0
23	ParEn	Partition Enable. When more than 61 collisions occur while transmitting, the GT-48004A enters the Partition mode. It waits for the first good packet from the wire, and then goes back to Normal mode. Under Partition mode it continues transmitting, but not receiving. 0 - Normal mode 1 - Partition mode	0x0
24	SpanEn	Spanning Tree Enable. 0 - The BPDU (Bridge Protocol Data Unit) packets are treated as Multicast packets, and therefore are forwarded to all ports. 1 - The GT-48004A forwards BPDU packets only to the CPU.	0x0
25	EnEASE	EASE sampling enable/disable. 0 - EASE sampling disabled. 1 - EASE sampling enabled.	0x0

Bits	Field Name	Function	Initial Value
26	MIBCtrMode	0 - MIB counters reflect forwarded packets only. 1 - MIB counters reflect local and forwarded packets. NOTE: A local packet is a packet which is destined to a station on the same port and is not switched.	0x0
27	EnableDevice	Used to enable or disable the GT-48004A. 0 - The GT-48004A is disabled. 1 - The GT-48004A is enabled/disabled by the EnDev* pin.	0x0
28	RstQueues	Reset Queues. Used to reset the transmit and receive queues via software. 0 - The Tx and Rx queues of the GT-48004A are reset according to RstQueue*. 1 - Reset the Tx and Rx queues immediately.	0x0
29	CRCGenEn	Used to enable or disable the CRC generation during transmit of CPU generated packets to the GT-48004A. 0 - CRC generation is disabled 1 - CRC generation is enabled.	0x0
30	MIBClrMode	MIB Counter Clear-on-Read mode 0 - MIB counters will be cleared after they have been read. 1 - MIB counters will not cleared after they have been read.	0x0
31	Reserved	This bit is reserved and must be written to 0.	0x0

Port Control (2 Registers), Offset: 0x040200 - 0x040204

Bits	Field Name	Function	Initial Value
0	PortEn	Port Enable. 0 - Port is disabled 1 - Port is enabled	0x1
1	FullDx	Half/Full Duplex. 0 - Port works in half-duplex mode 1 - Port works in full-duplex mode	DAddr[6] for port 0 or DAddr[7] for port 1 (at RESET)
2	MonMode	Monitoring Mode. 0 - Port works in normal mode 1 - Port works in monitoring mode; all Rx and Tx packets are sent to the Sniffer	0x0
5:3	Reserved	Reserved.	0x1
6	FilBroad	Filter Broadcast. 0 - Broadcast packets are forwarded to all ports. 1 - The GT-48004A discards Broadcast packets.	0x0
7	ForwUnk	Forward Unknown. 0 - Unknown packets are forwarded. 1 - The GT-48004A does not forward Unknown packets to this port.	0x0
8	SpanEn	Spanning Tree Enable. Meaningful only when SpanEn bit in the Global Control register is set. 0 - All packets are accepted. 1 - The GT-48004A discards all incoming/outgoing packets except for BPDU packets.	0x0
9	VTagEn	Virtual LAN Tagging Packet Extension Enable 0 - Accept packets up to 1518 bytes in length 1 - Accept packets up to 1522 bytes in length	0x0

Note: Modifying bits [9:1] of any Port Control Register must be preceded by disabling that port via bit[0], the EnDev* input or bit[27] of the Global Control Register (EnableDevice).

Status, Offset: 0x14002c (Read-Only)

Bits	Field Name	Function	Initial Value
4:0	DevNum	Device Number. Indicates the GT-48004A number chosen by the designer.	DAddr[4:0] at RESET
5	DRAMSize	DRAM Size. Indicates the DRAM size. 0- 2Mbyte 1- 1Mbyte	DAddr[5] at RESET
6	Port0Par	Port0 Partition. This bit indicates the port 0 Partition status. 0 - No Partition (Normal mode) 1 - Partition	0x0
7	Port0LTF	Port0 Link Test Fail. This bit indicates the port 0 Link Test status. 0 - Link Test Pass 1 - Link Test Fail	0x0
8	Port1Par	Port1 Partition. This bit indicates the port 1 Partition status. 0 - No Partition (Normal mode) 1 - Partition	0x0
9	Port1LTF	Port1 Link Test Fail. This bit indicates the port 1 Link Test status. 0 - Link Test Pass 1 - Link Test Fail	0x0

CPU and Sniffer Numbers, Offset: 0x140030

Bits	Field Name	Function	Initial Value
4:0	SnifDevNum	Sniffer Device Number. These bits specify the Sniffer Device Number chosen by the designer. The same value should be programmed to all GT-48004A/GT-48001A devices sharing the same PCI bus.	0x0 (bit 4 is MSB)
7:5	SnifPortNum	Sniffer Port Number. These bits specify the Sniffer Port Number chosen by the designer. The same value should be programmed to all GT-48004A/GT-48001A devices sharing the same PCI bus.	0x0 (bit 7 is MSB)
12:8	CPUDevNum	CPU Device Number. These bits specify the Device Number of the CPU as chosen by the designer. This number must not be the same as any other GalNet device in the system. The same value should be programmed to all GT-48004A/GT-48001A devices sharing the same PCI bus. The corresponding bit for the CPUDevNum value should be set in the Device table (offset 0x40) for all GT-48004A/GT-48001A devices sharing the same PCI bus.	0x0 (bit 12 is MSB)

Interrupt Cause, Offset: 0x044

Bits	Field Name	Function	Initial Value
0	IntSumm	Interrupt Summary. Logical 'OR' of all the interrupt bits.	0x0
1	NewAdd	New Address. This bit is set by the GT-48004A when a new address is received.	0x0
2	TxEnd	Tx End of Packet. This bit is set by the GT-48004A upon transferring a END_OF_PACKET message to the CPU main memory.	0x0
3	RxStart	Rx Start of Packet. This bit is set by the GT-48004A upon sending a START_OF_PACKET message to the CPU.	0x0
4	AddrRecF	Address Recognition Failed. This bit is set by the GT-48004A when the address recognition cycle fails (due to a large number of MAC addresses).	0x0
5	FlushTxQ	Flush Tx Queue. This bit is set by the GT-48004A when one of the Tx queues is flushed due to the Watchdog Timer.	0x0
6	MastRdPar	Master Read Parity. This bit is set by the GT-48004A upon master read parity error on the PCI.	0x0
7	MastWrPar	Master Write Parity. This bit is set by the GT-48004A upon master write error on the PCI.	0x0
8	AddPar	Address Parity. This bit is set by the GT-48004A upon address parity error on the PCI.	0x0
9	MastAbort	Master Abort. the This bit is set by the GT-48004A upon master abort on PCI.	0x0
10	TarAbort	Target Abort. This bit is set by the GT-48004A upon target abort on the PCI.	0x0
11	LinkChange	Link State Change. This bit is set by the GT-48004A upon a change in the link state (down->up, up->down) for any pin.	0x0
12	Part	Partition. This bit is set by the GT-48004A upon entering Partition state in one of the ports.	0x0
13	BufWrap	Buffer Wrap-Around. This bit is set by the GT-48004A upon transferring 16 packets to the CPU main memory.	0x0

Bits	Field Name	Function	Initial Value
14	Interv	Intervention. This bit is set by the GT-48004A upon transferring a BUFFER_REQUEST message in Intervention mode.	0x0
15	IntervWrap	Intervention Wrap-Around. This bit is set by the GT-48004A upon transferring 256 BUFFER_REQUEST messages in Intervention mode.	0x0

Interrupt Mask, Offset: 0x048

Bits	Field Name	Function	Initial Value
15:0	MaskBits	Mask to the CPU interrupt line for the appropriate bits in the Interrupt Cause register. The default is to mask all interrupts. 0 - Mask Interrupt 1 - Do not mask Interrupt	0x0000

Serial Parameters, Offset: 0x040220

Bits	Field Name	Function	Initial Value
11:0	-	Reserved.	0x823
16:12	IPGData	Inter-Packet Gap (IPG): The IPG varies between 12 bit-times (0x3) and 124 bit-times (0x1F). The step is 40ns@100mbs or 400 ns@10mbs (4 bit-times). The default value is 18 hex, or 0.96us@100mbs (9.6uS@10mbs). Value should be written in hexadecimal format. Note: these bits may be changed only when PortEn bits are set to 0 in all Port Control Registers (Port is disabled).	0x18 = 24 decimal
21:17	DataBlind	Data Blinder: The number of nibbles from the beginning of the IFG, in which the GT-48004A will restart the IFG counter when detecting a carrier activity. Following this value, the GT-48004A will enter the Data Blinder zone and will not reset the IFG counter to ensure fair access to the medium. Value should be written in hexadecimal format. The default is 10 hex (64 bit times - 2/3 of the default IFG). The step is 40ns (4 bit-times). Valid range is 3 to 1F hex nibbles. Note: these bits may be changed only when PortEn bits are set to 0 in all Port Control Registers (Port is disabled).	0x10 = 16 decimal
25:22	Reserved	Reserved	0xB

Rx Buffers Threshold, Offset: 0x040224

Bits	Field Name	Function	Initial Value
3:0	TxWatTim	Tx Watchdog Timer. For 100 Mbps operation, the default value of the timer is 63msec and the range is between 10.5mSec and 168msec, in 10.5 ms steps. For 10 Mbps operation, the default value of the timer is 630msec and the range is between 105ms to 1680ms, in 105 ms steps. Valid values: 1 to F hex. Value should be written in hexadecimal format.	0x6 (63 ms @ 100Mbps; 630 ms @ 10 Mbps)
8:4	-	Reserved.	0x04
12:9	RxBufThr	Rx Buffer Threshold. These bits determine the threshold of the receive buffers. Meaningful only when the BufThrEn bit in the Global Control register is set. The Rx Buffer Threshold equals to: (RxBufThr x 20 + 20), the step is 20, and the value varies between 20 (0x0) and 320 (0xf). The default is 80 (0x3) for 1MB DRAM and 140 (0x6) for 2MB DRAM. Value should be written in hexadecimal format.	0x3 (80) @ 1M DRAM 0x6 (140) @ 2M DRAM
15:13	-	Reserved. During writes to this register, bits 15:13 must be written to their default values.	0x1 @ 1M DRAM 0x3 @ 2M DRAM
18:16	PCIBufThr	PCI Buffer Threshold. These bits determine the threshold of the PCI receive buffers. The PCI Buffer Threshold equals to: (PCIBufThr x 50 +100). step is 50, and the value varies between 100 (0x0) and 320 (0xf). The default is 150 (0x1) for 1MB DRAM and 200 (0x2) for 2MB DRAM. Value should be written in hexadecimal format.	0x1 (150) @ 1M DRAM 0x2 (200) @ 2M DRAM

CPU Buffer Base Address, Offset: 0x140034

Bits	Field Name	Function	Initial Value
31:15	BaseAdd	Contains a pointer to the CPU buffer area	0x0
14:0	Reserve	Reserved. Must be 0x0 when written.	-

Note: When reading the CPU Buffer Base Address Register (0x140034), the address will be read out 1 bit shifted left.

CPU START_OF_PACKET Base Address, Offset: 0x140038

Bits	Field Name	Function	Initial Value
31:8	StBaseAdd	Contains a pointer to the CPU START_OF_PACKET area. The area includes 32 entries (2 32-bit words each) for the GT-48004A's START_OF_PACKET messages.	-
7:0	Reserve	Reserved. Must be 0x0 when written.	-

CPU NEW_ADDRESS Base Address, Offset: 0x14003c

Bits	Field Name	Function	Initial Value
31:8	NABaseAdd	Contains a pointer to the CPU NEW_ADDRESS area. The area includes 32 entries (2 32-bit words each) for the GT-48004A's NEW_ADDRESS messages.	0x0
7:0	Reserve	Reserved. Must be 0x0 when written.	-

CPU Intervention Base Address, Offset: 0x140048

Bits	Field Name	Function	Initial Value
31:11	IntBaseAdd	Contains a pointer to the CPU BUFFER_REQUEST area for unicast packets that are marked for intervention. The area includes 256 entries (2 32-bit words each) for the GT-48004A's BUFFER_REQUEST messages.	0x0
10:0	Reserve	Reserved. Must be 0x0 when written.	-

Device Table, Offset: 0x40

Bits	Field Name	Function	Initial Value
31:0	DevTab	GT-48004A Device Table. Each bit represents a GT-48004A device in the system.	0xffffffff

Time Out Counter, Offset: 0x4c

Bits	Field Name	Function	Initial Value
7:0	TimeOut0	Specifies in PCI clock units the number of clocks the GT-48004A holds the PCI bus before the generation of Retry termination. Used for the first data transfer.	0x0f (16 clocks)
15:8	TimeOut1	Specifies in PCI clock units the number of clocks the GT-48004A holds the PCI bus before the generation of Retry termination. Used for data transfers following the first data.	0x07 (8 clocks)
23:16	RetryCounter	Specifies the number of retries the GT-48004A attempts to do in the PCI before aborting the transaction. Value of 0x0 disable this counter (unlimited retries).	0xFF (256 times)

23.3 Port MIB Counters (2 Blocks), Offset: 0x040000 - 0x0400ac

The CPU must read all of the MIB counters during initialization in order to reset the counters to '0'. All counters are 32-bits. The CPU must access the counters using single datum transactions (burst reads/writes are not allowed.) MIB counters can be cleared by reading, or left intact after a read based on the MIB Clear Mode (bit 30 in the Global Command Register.) During initialization, the CPU must read all of the MIB counters in order to reset the counters to '0'. The counters will only be reset to '0' if MIBClrMode (bit 30 of the Global Control Register) is set to '0' (default). If MIBClrMode bit is '1', reading the MIB counters will have no effect.

Table 35 lists definitions for terms used in the counter descriptions.

Table 35: Definitions Used in Counter Descriptions

Term	Definition
Packet Data Section	All data bytes in the packet following the SFD until the end of the packet
Packet Data Length	The number of data bytes in the packet data section
Data Octet	A single byte from the packet data section
Nibble	4 bits (half byte) of a data octet
Misaligned Packet	A packet with an odd number of nibbles
Received Good Packet	A received packet which is not rejected and enters the switching core to be transmitted later
Transmitted Good Packet	Any transmitted packet from the GT-48004A
Collision Event	A collision has been detected until than 576 bit times into the transmitted packet after TxEn.
Late Collision Event	A collision has been detected later than 576 bit times into the transmitted packet after TxEn.
Rx Error Event	The input RX_ERR has been asserted
Dropped Packet	A received packet which is ignored due to lack of available receive buffers (port is in buffer_full state)
Local Packet	A received packet whose destination address is mapped to the receiving port
Rejected Packet	A received packet which is not forwarded due to error such as bad CRC, Rx Error Event, Invalid size (too short or too long).
MIBCtrMode	Bit 26, MIBCtrMode, of the Global Control Register (offset 0x140028)

Term	Definition
VTagEn	Bit 9, VTagEn, of the Port Control Register (offset 0x040200-0x040204)
MAXFRAMESIZE	1518 for VTagEn = 0 (default) or 1522 for VTagEn = 1

Table 36: Port MIB Counters

Address for Port 0	Counter Name	Function	Initial Value
0x040000	Bytes Received	MIBCtrMode = 0: This counter is incremented once for every data octet of good packets (unicast + multicast + broadcast) received. MIBCtrMode = 1: This counter is incremented once for every data octet of good packets (unicast + multicast + broadcast packets) and for LOCAL and DROPPED packets.	-
0x040004	Bytes Sent	This counter is incremented once for every data octet of a transmitted good packet.	-
0x040008	Frames Received	MIBCtrMode = 0: This counter is incremented once for every good packet (unicast + multicast + broadcast) received. MIBCtrMode = 1: This counter is incremented once for every good packet (unicast + multicast + broadcast packets) and for LOCAL and DROPPED packets received.	-
0x04000c	Frames Sent	This counter is incremented once for every transmitted good packet.	-
0x040010	Total Bytes Received	This counter is incremented once for every data octet of all received packets. This includes data octets of rejected and local packets which are NOT forwarded to the switching core for transmission. This counter should reflect all the data octets received on the line. NOTE: A nibble is NOT counted as a whole byte.	-
0x040014	Total Frames Received	This counter is incremented once for every received packets. This includes rejected and local packets which are NOT forwarded to the switching core for transmission. This counter should reflect all packets received on the line.	-
0x040018	Broadcast Frames Received	MIBCtrMode = 0: This counter is incremented once for every good broadcast packet received. MIBCtrMode = 1: This counter is incremented once for every good broadcast packet received and for LOCAL and DROPPED broadcast packets.	-
0x04001c	Multicast Frames Received	MIBCtrMode = 0: This counter is incremented once for every good multicast packet received. MIBCtrMode = 1: This counter is incremented once for every good multicast packet received and for LOCAL and DROPPED broadcast packets. This counter does not include Broadcast packets.	-

Table 36: Port MIB Counters

Address for Port 0	Counter Name	Function	Initial Value
0x040020	CRC Error	This counter is incremented once for every received packet which meets all the following conditions (i.e. logical AND of the following conditions): 1. Packet data length is between 64 and MAXFRAMESIZE bytes inclusive (i.e. valid packet data length per IEEE std). 2. Packet has invalid CRC (non-A also counted packets with odd number of nibbles). 3. Collision Event has not been detected. 4. Late Collision Event has not been detected. 5. Rx Error Event has not been detected.	-
0x040024	Oversize Frames	This counter is incremented once for every received packet which meets all the following conditions (i.e. logical AND of the following conditions): 1. Packet data length is greater than and MAXFRAMESIZE. 2. Packet has valid CRC. 3. Rx Error Event has not been detected.	-
0x040028	Fragments	This counter is incremented once for every received packet which meets all the following conditions (i.e. logical AND of the following conditions): 1. Packet data length is less than 64 bytes -OR- packet without SFD and is less than 64 bytes in length. 2. Collision Event has not been detected. 3. Late Collision Event has not been detected. 4. Rx Error Event has not been detected.	-
0x04002c	Jabber	This counter is incremented once for every received packet which meets all the following conditions (i.e. logical AND of the following conditions): 1. Packet data length is greater than MAXFRAMESIZE. 2. Packet has invalid CRC. 3. Rx Error Event has not been detected.	-
0x040030	Collision	This counter is incremented once for every received packet which meets both of the following conditions (i.e. logical AND of the following conditions): 1. Collision Event has been detected. 2. Rx Error Event has not been detected.	-
0x040034	Late Collision	This counter is incremented once for every received packet which meets both of the following conditions (i.e. logical AND of the following conditions): 1. Late Collision Event has been detected. 2. Rx Error Event has not been detected.	-
0x040038	Frames 64 Bytes	This counter is incremented once for every received and transmitted packet with size of 64 bytes. This counter includes rejected received and transmitted packets.	-
0x04003c	Frames 65-127 Bytes	This counter is incremented once for every received and transmitted packet with size of 65 to 127 bytes. This counter includes rejected received and transmitted packets.	-

Table 36: Port MIB Counters

Address for Port 0	Counter Name	Function	Initial Value
0x040040	Frames 128-255 Bytes	This counter is incremented once for every received and transmitted packet with lsize of 128 to 255 bytes. This counter includes rejected received and transmitted packets.	-
0x040044	Frames 256-511 Bytes	This counter is incremented once for every received and transmitted packet with lsize of 256-511 bytes. This counter includes rejected received and transmitted packets.	-
0x040048	Frames 512-1023 Bytes	This counter is incremented once for every received and transmitted packet with lsize of 512-1023 bytes. This counter includes rejected received and transmitted packets	-
0x04004c	Frames 1024-1522 Bytes	This counter is incremented once for every received and transmitted packet with lsize of 1024 to MAXFRAMESIZE bytes. This counter includes rejected received and transmitted packets.	-
0x040050	MAC Rx Error	This counter is incremented once for every received and transmitted packet where the Rx Error Event has been detected. This counters that will not be incremented when this counter increments include: CRC Error, Oversize Frames, Fragments, Jabber, Collision and Late Collision.	-
0x040054	Dropped Frames	This counter is incremented once for every received dropped packet.	-

23.4 PCI Global Counters, Offset: 0x140040 - 0x140044

The CPU must read all counters during initialization in order to reset the counters to '0'. All counters are 32-bits.

Table 37: PCI MIB Counters

Address	Counter Name	Function	Initial Value
0x140040	PCIFraRec	Good Frames Received from the PCI	-
0x140044	PCIFraSent	Good Frames Sent to the PCI	-

23.5 SMI Register, Offset: 0x14004c

Bits	Field Name	Function	Initial Value
15:0	Data	<p>For SMI READ operation: Two PCI transactions are required: (1) PCI write to the SMI register with OpCode = 1, PhyAd, RegAd with the Data being any value. (2) PCI read from the SMI register. When reading back the SMI register, the Data is the addressed Phy register contents if the ReadValid bit (#27) is 1. The Data remains undefined as long as ReadValid is 0. No PCI write to the SMI register should take place until the ReadValid is returned with the value of '1'.</p> <p>For SMI WRITE operation: One PCI transaction is required: PCI write to the SMI register with OpCode = 0, PhyAd, RegAd with the Data to be written to the addressed Phy register. No PCI write to the SMI register should take place until a minimum of 320 MDC clock cycles have passed.</p>	N/A
20:16	PhyAd	PHY device address	0x0
25:21	RegAd	PHY device register address	0x0
26	OpCode	0 - Write 1 - Read	0x1
27	ReadValid	1 - Indicates that the Read operation has been completed for the addressed RegAd register, and the data is valid on the Data field. Once set, this bit is cleared following the PCI read operation.	0x0
31:28	N/A	This bits should be driven 0x0 during any write to the SMI register.	0x0

23.6 PCI Configuration Registers

The GT- 48001 contains the required PCI configuration registers. These registers are accessed from the PCI.

Device and Vendor ID, PCI Offset: 0x000 (Read Only)

Bits	Field Name	Function	Initial Value
15:0	VenId	Vendor ID. Provides the manufacturer of the PCI device. For the GT-48004A this is Galileo's PCI vendor ID (0x11ab.)	0x11ab
31:16	DevId	Provides the unique GT- 48002A device ID number	0x4802

Status and Command, PCI Offset: 0x004

Bits	Field Name	Function	Initial Value
1	MemEn	Memory Enable. Controls the GT-48004A's response to memory accesses. 0 - Disable, the GT-48004A will ignore all memory accesses on the PCI bus. 1 - Enable, the GT-48004A will respond to memory accesses on the PCI bus.	0x1
2	MasEn	Master Enable. Controls the GT-48004A's ability to act as a master on the PCI bus. 0 - Disable 1 - Enable	0x1
4	MemWrInv	Memory Write and Invalidate. Controls the GT-48004A's ability to generate Memory Write and Invalidate commands on the PCI bus. 0 - Disable 1 - Enable	0x0
5	Reserved	Reserved. Read only.	0x0
6	ParEn	Parity Enable. Controls the GT-48004A's ability to respond to parity errors on the PCI. 0 - Disable 1 - Enable	0x0
8	SysErrEn	System Error Enable. Controls the GT-48004A's ability to assert the SErr* pin. 0 - Disable 1 - Enable	0x0
22:9	Reserve	Reserved. Read only.	0x0
23	TarFastBB	Target Fast Back-to-Back. This indicates that the GT-48004A is capable of accepting Fast Back-to-Back transactions on the PCI bus. Read-only bit.	0x1
24	DataParDet	Data Parity Detected. This bit is set by the GT-48004A when it detects a Data Parity Error during a master operation. This bit is cleared by writing '1' to it.	0x0
26:25	DevSelTim	Device Select Timing. These bits indicate the GT-48004A's DevSel* timing. The GT-48004A's DevSel* timing is always set at medium (01), as defined in the PCI specification. Read only.	0x1
28	TarAbort	Target Abort. This bit is set upon Target Abort. This bit is cleared by writing '1' to it.	0x0

Bits	Field Name	Function	Initial Value
29	MastAbort	Master Abort. This pin is set upon Master Abort. This bit is cleared by writing '1' to it.	0x0
30	SysError	System Error. This pin is set upon System Error. This bit is cleared by writing '1' to it.	0x0
31	DetParErr	Detected Parity Error. This pin is set upon detection of Parity Error (in both master and slave operations). This bit is cleared by writing '1' to it.	0x0

Class Code and Revision ID, Offset: 0x008 (Read-Only)

Bits	Field Name	Function	Initial Value
7:0	RevID	Revision ID. Indicates the GT-48004A revision number.	0x10
23:16	SubClass	SubClass. Indicates the GT-48004A Subclass (0x0 - Ethernet).	0x0
31:24	BaseClass	Base Class. Indicates the GT-48004A Base Class (0x2 - Network Device).	0x2

Header Type, Latency Timer, Cache Line, Offset: 0x00C

Bits	Field Name	Function	Initial Value
7:0	CacheLine	Cache Line. Specifies the GT-48004A's cache line size (size=8). Read only.	0x7
15:8	LatTimer	Latency Timer. Specifies in units of PCI bus clocks the value of the latency timer of the GT-48004A. Default is 256 cycles (0xff).	0xff
23:16	HeadType	Header Type. Specifies the layout of bytes 10 hex through 3f hex.	0x0

For more information on these fields, please refer to the PCI specification.

Internal Register/DRAM Base Address Register, Offset: 0x010

Bits	Field Name	Function	Initial Value
21:0	-	These bits are cleared.	0x0
21	DramAccess	0 - Access internal registers 1 - Access DRAM array	0x0
26:22	DevNum	Device Number. These bits specify the GalNet Device Number.	DAddr[4:0] at RESET
31:27	BaseAdd	Galnet Protocol Region (GPR) Base Address. These bits set the base address for the GalNet Protocol Region. All GalNet devices in a system must have the same GPR base address (see text.)	0x1

Note: This register has special behavior in a plug and play environment (see Section 12.6 on page 49). The device number can be changed in this register after RESET (see Section 8.3 on page 30).

Interrupt Pin and Line, Offset: 0x03c

Bits	Field Name	Function	Initial Value
7:0	IntLine	Interrupt Line. Provides interrupt line routing information.	0x0
15:8	IntPin	Interrupt Pin. Indicates which interrupt pin the GT-48004A uses. The GT-48004A uses INTA in the PCI slot.	0x1

24. GT-48004A PINOUT TABLE, 329 pin BGA (sorted by ball number)

TABLE 38.

Ball #	Signal Name	Ball #	Signal Name	Ball #	Signal Name
A01	NC	B16	Vref	D08	GND
A02	NC	B17	AD[18]	D09	NC
A03	NC	B18	AD[21]	D10	VCC
A04	NC	B19	CBE[3]*	D11	NC
A05	NC	B20	AD[26]	D12	GND
A06	Req0*	B21	AD[30]	D13	NC
A07	AD[0]	B22	Req1*	D14	VCC
A08	AD[3]	B23	Int1*	D15	Rst*
A09	AD[6]	C01	ADData[29]	D16	GND
A10	AD[8]	C02	ADData[28]	D17	AD[19]
A11	AD[11]	C03	NC	D18	VCC
A12	AD[15]	C04	NC	D19	AD[28]
A13	PErr*	C05	NC	D20	IdSel1
A14	Trdy*	C06	Gnt0*	D21	BDData[28]
A15	Frame*	C07	AD[2]	D22	BDData[26]
A16	AD[16]	C08	AD[5]	D23	BDData[27]
A17	AD[20]	C09	CBE[0]*	E01	ADData[23]
A18	AD[22]	C10	AD[10]	E02	ADData[22]
A19	AD[24]	C11	AD[12]	E03	ADData[24]
A20	AD[27]	C12	CBE[1]*	E04	NC
A21	AD[31]	C13	Par	E20	NC
A22	Gnt1*	C14	Stop*	E21	BDData[25]
A23	NC	C15	Irdy*	E22	BDData[23]
B01	ADData[31]	C16	CBE[2]*	E23	BDData[24]
B02	ADData[30]	C17	AD[17]	F01	ADData[20]
B03	NC	C18	AD[23]	F02	ADData[19]
B04	NC	C19	AD[25]	F03	ADData[21]
B05	Int0*	C20	AD[29]	F04	GND
B06	IdSel0	C21	BDData[31]	F20	GND
B07	AD[1]	C22	BDData[29]	F21	BDData[22]
B08	AD[4]	C23	BDData[30]	F22	BDData[20]
B09	AD[7]	D01	ADData[26]	F23	BDData[21]
B10	AD[9]	D02	ADData[25]	G01	ADData[18]
B11	AD[13]	D03	ADData[27]	G02	ADData[17]
B12	AD[14]	D04	NC	G03	ADData[16]
B13	SErr*	D05	NC	G04	NC
B14	DevSel*	D06	VCC	G20	NC
B15	NC	D07	NC	G21	BDData[17]

TABLE 38.

Ball #	Signal Name	Ball #	Signal Name	Ball #	Signal Name
G22	BDData[18]	L14	GND	P14	GND
G23	BDData[19]	L20	NC	P20	GND
H01	ADDData[15]	L21	BDData[6]	P21	BWE*
H02	ADDData[14]	L22	BDData[5]	P22	BRAS[1]*
H03	ADDData[13]	L23	BDData[7]	P23	BRAS[0]*
H04	VCC	M01	ADDData[1]	R01	ADAddr[4]
H20	VCC	M02	ADDData[0]	R02	ADAddr[5]
H21	BDData[14]	M03	ADDData[2]	R03	ADAddr[6]
H22	BDData[15]	M04	VCC	R04	NC
H23	BDData[16]	M10	GND	R20	BDAddr[6]
J01	ADDData[12]	M11	GND	R21	BDAddr[8]
J02	ADDData[11]	M12	GND	R22	BDAddr[7]
J03	ADDData[10]	M13	GND	R23	BDAddr[5]
J04	NC	M14	GND	T01	ADAddr[1]
J20	NC	M20	VCC	T02	ADAddr[2]
J21	BDData[11]	M21	BDData[2]	T03	ADAddr[3]
J22	BDData[12]	M22	BDData[4]	T04	VCC
J23	BDData[13]	M23	BDData[3]	T20	VCC
K01	ADDData[9]	N01	ARAS[1]*	T21	BDAddr[4]
K02	ADDData[8]	N02	ACAS*	T22	BDAddr[3]
K03	ADDData[7]	N03	AWE*	T23	BDAddr[2]
K04	GND	N04	NC	U01	MDIO0
K10	GND	N10	GND	U02	NC
K11	GND	N11	GND	U03	ADAddr[0]
K12	GND	N12	GND	U04	NC
K13	GND	N13	GND	U20	NC
K14	GND	N14	GND	U21	BDAddr[1]
K20	GND	N20	BCAS*	U22	BDAddr[0]
K21	BDData[8]	N21	PClk	U23	MDC
K22	BDData[9]	N22	BDData[1]	V01	TxCIk[0]
K23	BDData[10]	N23	BDData[0]	V02	TxEn[0]
L01	ADDData[5]	P01	ADAddr[7]	V03	TxD0[3]
L02	ADDData[4]	P02	ADAddr[8]	V04	GND
L03	ADDData[3]	P03	ARAS[0]*	V20	GND
L04	ADDData[6]	P04	GND	V21	TxCIk[2]
L10	GND	P10	GND	V22	MDIO1
L11	GND	P11	GND	V23	TxEn[2]
L12	GND	P12	GND	W01	TxD0[1]
L13	GND	P13	GND	W02	TxD0[2]

TABLE 38.

Ball #	Signal Name	Ball #	Signal Name	Ball #	Signal Name
W03	TxD0[0]	AA05	RxD1[2]	AB15	LEDData1
W04	NC	AA06	RxEr[1]	AB16	LEDRxBufFull*[3]
W20	NC	AA07	LEDRxBufFull*[1]	AB17	CrS[3]
W21	TxD2[1]	AA08	LedClk0	AB18	RxEr[3]
W22	TxD2[3]	AA09	EnDev*	AB19	RxD3[2]
W23	TxD2[2]	AA10	Test0 ¹	AB20	TxD3[1]
Y01	RxD0[2]	AA11	Limit4	AB21	TxCk[3]
Y02	RxD0[3]	AA12	DisWD*	AB22	RxEr[2]
Y03	RxD0[1]	AA13	JTDI	AB23	RxCk[2]
Y04	NC	AA14	JTRST*	AC01	CrS[0]
Y05	TxD1[0]	AA15	LEDStb1	AC02	Col[0]
Y06	VCC	AA16	BForceLinkPass*	AC03	TxD1[3]
Y07	Col[1]	AA17	Col[3]	AC04	TxCk[1]
Y08	GND	AA18	RxD3[1]	AC05	RxD1[1]
Y09	LEDStb0	AA19	TxD3[0]	AC06	RxCk[1]
Y10	VCC	AA20	TxD3[3]	AC07	CrS[1]
Y11	VLAN	AA21	Col[2]	AC08	AForceLinkPass*
Y12	GND	AA22	RxD2[1]	AC09	LEDData0
Y13	JTMS	AA23	RxD2[0]	AC10	RstQueue*
Y14	VCC	AB01	RxDV[0]	AC11	LEDMode
Y15	LedClk1	AB02	TxEr[1]	AC12	SkipInIt*
Y16	GND	AB03	TxD1[2]	AC13	JTCLK
Y17	RxDV[3]	AB04	RxD1[3]	AC14	EnAutoNeg1*
Y18	VCC	AB05	RxD1[0]	AC15	LEDRxBufFull*[2]
Y19	NC	AB06	RxDV[1]	AC16	NC
Y20	NC	AB07	NC	AC17	RxCk[3]
Y21	RxD2[2]	AB08	LEDRxBufFull*[0]	AC18	RxD3[0]
Y22	TxD2[0]	AB09	EnAutoNeg0*	AC19	RxD3[3]
Y23	RxD2[3]	AB10	Priority	AC20	TxD3[2]
AA01	RxEr[0]	AB11	DisBufThr*	AC21	TxEr[3]
AA02	RxD0[0]	AB12	Scan*	AC22	CrS[2]
AA03	RxCk[0]	AB13	Tristate*	AC23	RxDV[2]
AA04	TxD1[1]	AB14	JTDO		

1. This pin should be tied to GND through a 4.7k ohm resistor.

25. DC Characteristics - *PRELIMINARY/SUBJECT TO CHANGE*

25.1 Absolute Maximum Ratings

Symbol	Parameter	Min.	Max.	Unit
Vdd	Supply Voltage	-0.3	4	V
Vi	Input Voltage	-0.3	Vdd+0.3	V
Vo	Output Voltage	-0.3	Vdd+0.3	V
Io	Output Current		24	mA
Iik	Input Protect Diode Current		+/-20	mA
Iok	Output Protect Diode Current		+/-20	mA
Tc	Operating Case Temperature	0	70	C
Tstg	Storage Temperature	-40	125	C
ESD			2000	V

25.2 Recommended Operating Conditions

Symbol	Parameter	Min.	Typ.	Max.	Unit
Vdd	Supply Voltage	3.15	3.3	3.45	V
Vi	Input Voltage	0		Vdd	V
Vo	Output Voltage	0		Vdd	V
Tc	Case Operating Temperature	0		70	C
Cin	Input Capacitance		7.2		pF
Cout	Output Capacitance		7.2		pF

25.3 DC Electrical Characteristics Over Operating Range

(Tc=0-70 C; Vdd=+3.3V, +/-5%)

Symbol	Parameter	Test Condition	Min.	Typ.	Max.	Unit
Vih	Input HIGH level	Guaranteed Logic HIGH level	2.0		Vdd + 0.5V	V
Vil	Input LOW level	Guaranteed Logic LOW level	-0.5		0.8	V
Voh	Output HIGH Voltage	IoH = 2 mA IoH = 4 mA IoH = 8 mA IoH = 12 mA IoH = 16 mA IoH = 24 mA	2.4			V
Vol	Output LOW Voltage	IoL = 2 mA IoL = 4 mA IoL = 8 mA IoL = 12 mA IoL = 16 mA IoL = 24 mA			0.4	V

Symbol	Parameter	Test Condition	Min.	Typ.	Max.	Unit
lih	Input HIGH Current				+ -1	uA
lil	Input LOW Current				+ -1	uA
lozh	High Impedance Output Current				+ -1	uA
lozl	High Impedance Output Current				+ -1	uA
Vh	Input Hysteresis					
Icc	Operating Current	Vdd = 3.45V f=66MHz		TBD		mA

NOTE: Pullup/Pulldown resistors are 45KOhm minimum, 65KOhm typical, 80KOhm maximum.

25.4 Thermal Data

Table 39 shows the package thermal data for the GT-48004A. Galileo recommends the use of heatsink for most systems, especially those with little or no airflow. Please check with Galileo if you are in doubt as to thermal considerations for your system.

Table 39: 329 BGA Thermal Data

Parameter	Definition	Value
θ_{ja}	Thermal resistance: junction to ambient, 0 ft/s airflow	22 C/W
θ_{jc}	Thermal resistance: junction to case, 0 ft/s airflow	9.6 C/W
T _j	Operating Junction temperature	100C
T _j	Maximum Junction temperature ¹	125C

1. Operating the device for extended periods of time at the maximum junction temperature may cause permanent damage to this device, resulting in functional or reliability type failures.

26. AC Timing - PRELIMINARY/SUBJECT TO CHANGE

(Tc= 0-70°C; VDD= +3.3V, +/- 5%)

Symbol	Signals	Description	Min	Max	Unit
	Clk	System Clock	66		MHz
	Clk	Rise/Fall Time (PCI Specification Rev. 2.1)	1	4	V/ns
t4	AD[31:0], CBE[3:0]*			4	ns
t4	Rst*, Frame*, IRdy*, TRdy*, DevSel*, Stop*, PErr*, Par, Int*, Gnt*, IdSel, Req*, SErr*			5	ns
t5	AD[31:0], CBE[3:0]*		1		ns
t5	Rst*, Frame*, IRdy*, TRdy*, DevSel*, Stop*, PErr*, Par, Int*, Gnt*, IdSel, Req*, SErr*		1		ns
t3	DAddr[8:0], DData[31:0], CAS*, RAS*, WE*	Delay from Clock Rising or Falling Edge	2	12	ns
t4	RstQueue*, EnDev*	Setup	10		ns
t5	RstQueue*, EnDev*	Hold	1		ns
t6	DData[31:0]	Float Delay	2	12	ns
t7	DData[31:0]	Drive Delay	2	12	ns
t8	DData[31:0]	Setup	2		ns
t9	DData[31:0]	Hold	3		ns
t4	RxD[3:0], RxDv, RxEr, Col, Crs, MDio	Setup	10		ns
t5	RxD[3:0], RxDv, RxEr	Hold	1		ns
t3	Txd[3:0], TxEn, Led-Clk, LedStb, LedData, MDC, MDio	Delay from Clock Rising or Falling Edge	2	10	ns

Notes:

1. All Delays, Setup, and Hold times are referred to Clk rising edge, unless stated otherwise.
2. All outputs are specified for 50pF load.
3. "All Inputs" and "All Outputs" also refer to I/O signal behavior.

Figure 18: Output Delay from Clock Rising Edge




Figure 19: Input Setup and Hold to Clock rising Edge




Figure 20: Output Delay from Clock rising edge




Figure 21: Output Float and Drive Delay


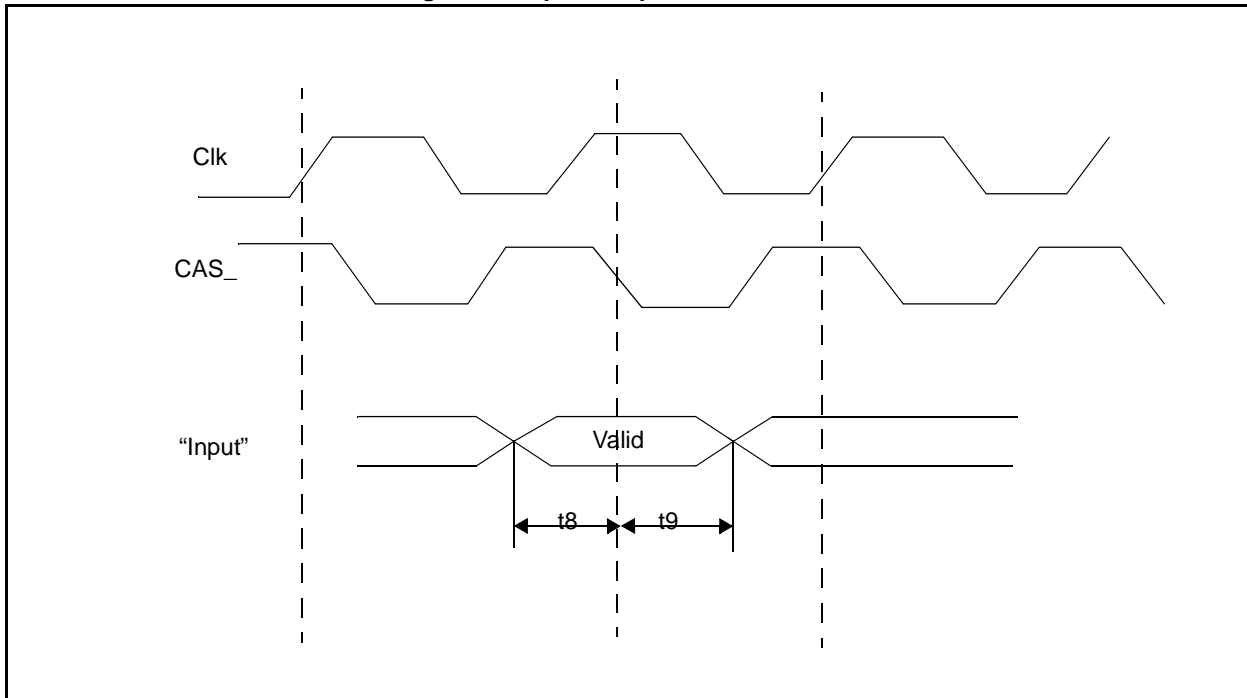


Figure 22: Input Setup and Hold to CAS_



27. Functional Waveforms

Figure 23: EDO DRAM Read

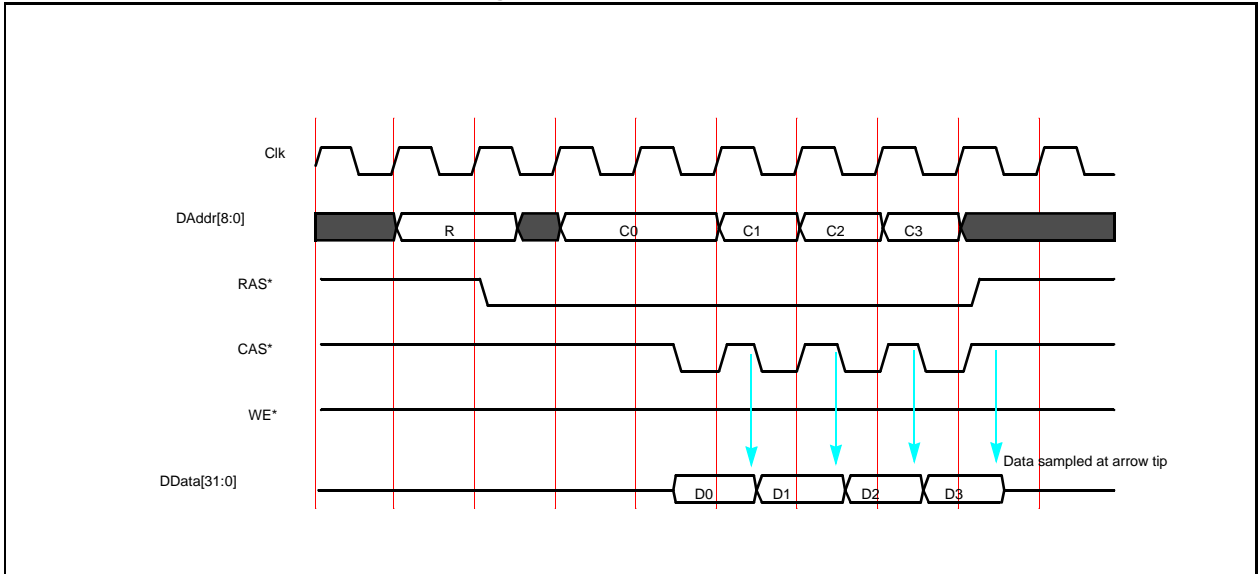
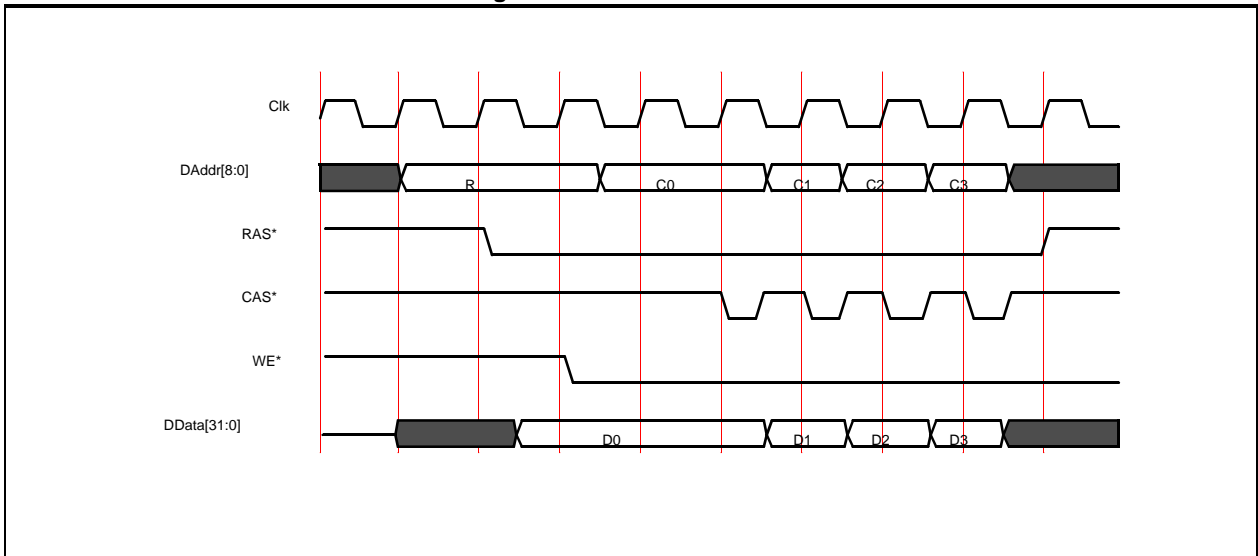


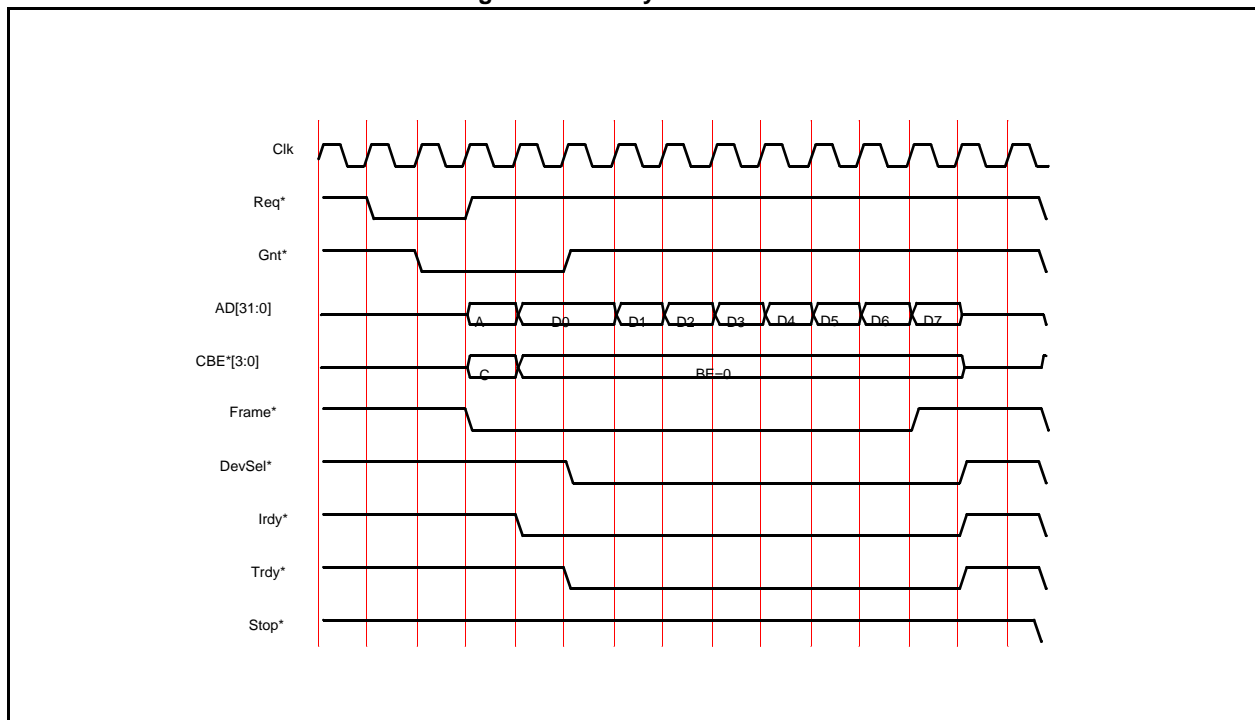
Figure 24: EDO DRAM Write




27.1 PCI Read/Write Cycle

The GT-48004A uses standard PCI read and write cycles. The longest burst that is executed by at GT-48004A device is eight words.


Figure 25: PCI Cycle Waveform




28. Packaging



BOTTOM VIEW
329 SOLDER BALLS




TOP VIEW



4	2.33±0.21	0.56±0.06
2	2.33 ^{+0.19} / _{-0.21}	0.56 ^{+0.04} / _{-0.06}
NO. LAYERS	DIM "A"	DIM "B"

PBGA THICKNESS SCHEDULE B



SIDE VIEW

Table 40: Document History

Document Type	Rev. Number	Date	Comments
PRODUCT PREVIEW	0.1	7/13/97	First rev. Derived from GT-48002A Rev 1.2 in progress data sheet. [MK]
PRODUCT PREVIEW	0.2	8/19/97	Fixed pinout (deleted multiple FEU pins when not existant.) Fixed typos. Put in correct BGA size.
PRODUCT PREVIEW	0.3	2/10/98	Added VLAN and priority support, final pinout and package mechanicals. Very limited release.
PRELIMINARY	1.0	2/12/98	Final cleanup. First non-NDA release.