Reno Reference Design

For FM2224 and FM4224 Family

Preliminary Specification

July 29, 2007 (Revision 1.0)

Order Number: RENO-HS-1.0

Trademark and Copyright Notice

Contact and Support Information:

Phone: 818-871-8100
Fax: 818-871-8101
support@fulcrummicro.com
www.fulcrummicro.com

Table of Contents

Table of Tables

Table of Figures

Revision History

Revision	Date	Content
1.0	7/29/2007	• First version

1. Introduction

1.1 Scope

FocalPoint is the family name for Fulcrum's 10G Ethernet switch chip products, which includes the FM2000 L2 switch chip platform and the FM4000 multi-layer switch chip platform. This specification details the FocalPoint Evaluation Platform for those products.

1.2 Definitions

Name	Definition
FocalPoint	The family name for Fulcrum's 10G Ethernet switch chip products, which includes the Tahoe L2 switch chip platform and the Bali multi-layer switch chip platform.
Tahoe	The original member of the FocalPoint family, Tahoe is a layer-2 10G Ethernet switch chip platform which forms the basis for many FocalPoint L2 switch product variants (including the FM2224, FM2212, FM2208, FM2112, FM2104, and FM2103).
Bali	The newest member of the FocalPoint family, Bali is an enhanced multi-layer 10G Ethernet switch chip platform which forms the basis for new FocalPoint switch product variants, all of which are pin-compatible with their original Tahoe counterparts.
Packet or Frame	Packet: On a typical computer network, data is transmitted in the form of structured and modest-sized packets. Instead of transmitting arbitrary-length strings of data, structured packets allow error checking and other relevant processing to occur on smaller easier-to-retransmit data. Packetized data also helps to alleviate traffic jams on the network when multiple nodes are contending for a shared network resource. Because packets are typically smaller than the complete data stream, techniques such as time-division multiplexing (THM) can be used to share and interleave traffic, making it appear that multiple nodes are using the network resource at the same time. Frame: While a packet is a small block of data, a Frame is the definition of how packets of data are defined and transported on a specific network. When sending data over a network, both sides of the connection must agree on a common frame format (e.g., when a frame starts, when a frame ends, padding, etc.) Combining terms, an Ethernet packet is sent onto an Ethernet interface using an Ethernet frame format. This document uses both terms interchangably; for now they shall be interpreted as synonymous. In the future, this document will be edited to standardize on the term packet.
Byte	8 bits
Half-word	16 bits
Word	32 bits
Double Word	64 bits

Name	Definition
Bit numbering	Bit 0 is the least-significant bit throughout this specification (even if Ethernet standards and specifications suggest otherwise) unless specified otherwise.
SIGNAL# or SIGNAL_N	Denote signals that are active low. Both conventions are used and are synonyms.
BUS[0:N]	Denote a bus where the least significant bit is N and the most significant bit is 0.
BUS[N:0]	Denote a bus where the least significant bit is 0 and the most significant bit is N.

2. Overview

2.1 Chassis Description

The FocalPoint Evaluation Platform chassis is designed to fit into a 1U track of a standard 19" rack (1.75" high and 17" wide). This platform has the following characteristics:

- FM2224 switch
- Freescale MPC8541E control processor
- DDR-SDRAM socket
    - Unit is shipped with 512MB of DDR
- CompactFlash socket
    - Unit is shipped with 1GB of CcompactFLASH
- 8MB of BOOT ROM (NOR FLASH)
- 1GB of NAND FLASH
- One AC power inlet with integrated fuses and filter
- One console interface for management
- One Ethernet 10/100 BaseT interface for management
    - The console interface and the Ethernet interface are mounted together into one double RJ-45 assembly.
- 12 x CX4 (with power over CX4) on main board
- Optional interface modules:
    - 12 x CX4 (with power over CX4)
    - 12 x XFP
    - 12 x SFP+

Figure 1 shows the chassis front view; Figure 2 shows the rear views. Ventilation is from front to back through ventilation holes in the front and in the back using blowers. All connectors are on the front.

Figure 1: Front View

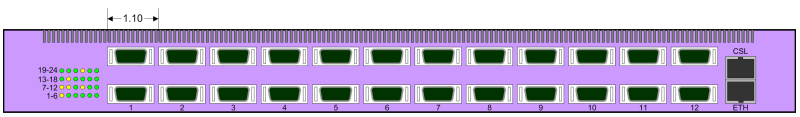

Figure 2: Back View

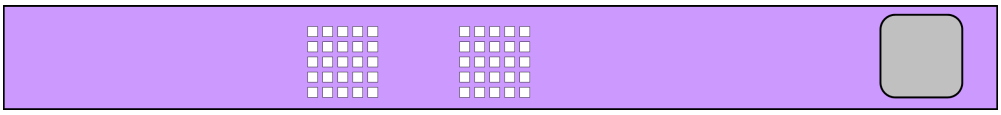

Rack mount ears can be mounted on the rear of the chassis or the front the chassis, allowing reverse installation of the chassis, with connectors in the front.

Figure 3 shows the top view of the chassis. Power is coming to the chassis through one AC 90VAC-230VAC inlet to a dedicated AC/DC power supply which produces a +12VDC connected to the motherboard.

Figure 3: Chassis Top View

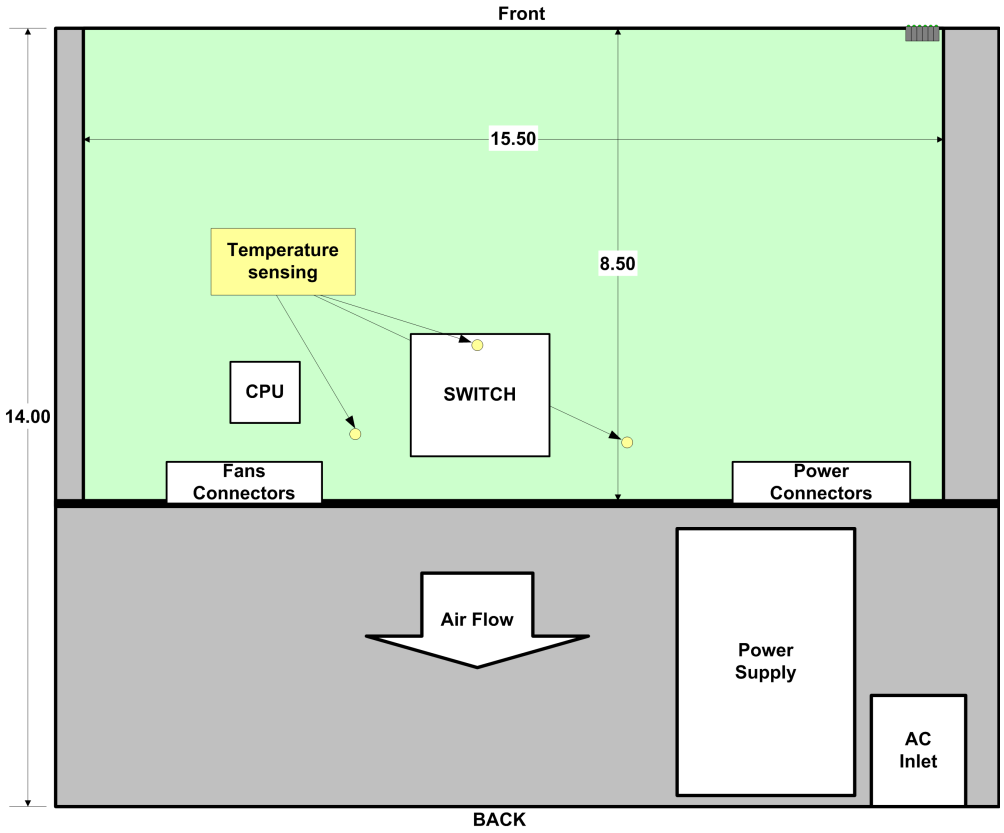

The Figure 4 shows the mechanical specification of the motherboard.

Figure 4: Motherboard Mechanical Specification

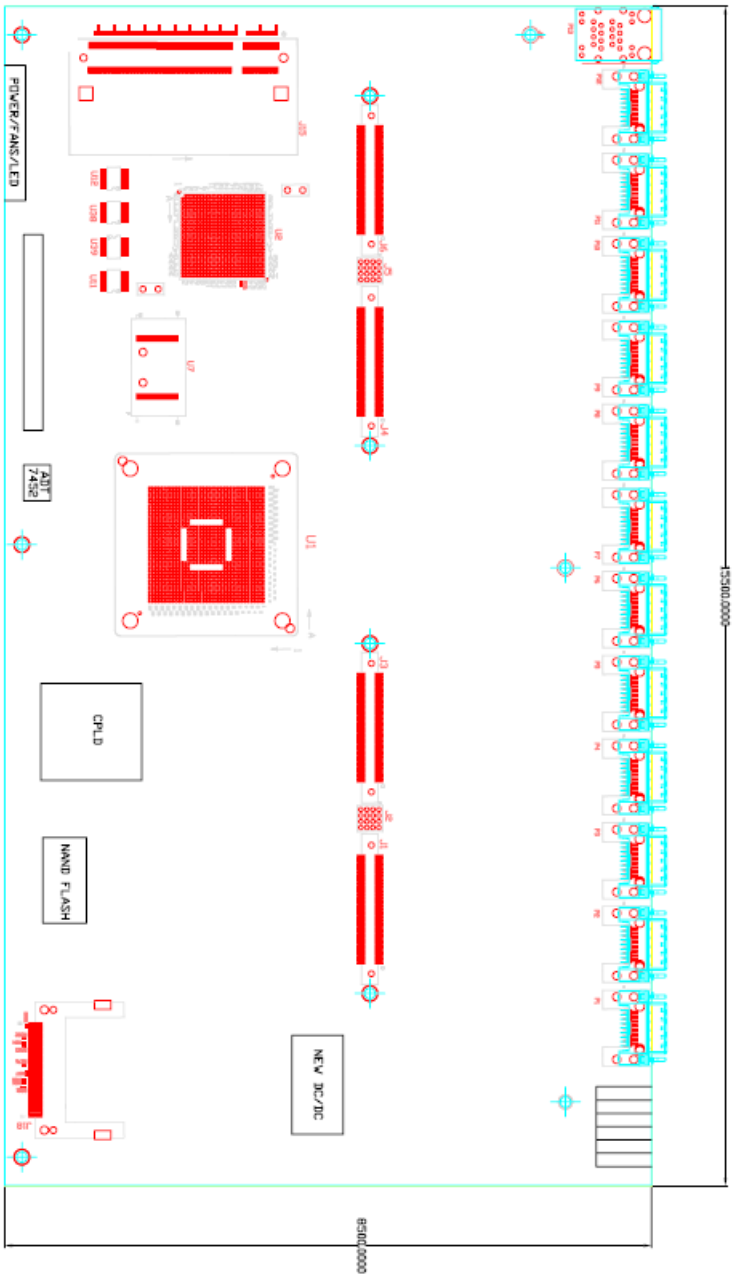

2.2 Motherboard Block Diagram

The motherboard is based on the Fulcrum Microsystems FM2224 Ethernet switch and the Freescale MPC8451E microprocessor.

The processor is connected directly to the FM2224, a DDR-SDRAM module, BOOT ROM, compactFlash, the Ethernet management port, the console port. The PCI bus is not used. The processor is also supporting

two I2C interfaces, the first one is used for management of the board while the second one used for management of the interface modules. The BOOT ROM is used to store the boot code while the compactFlash and NAND Flash are used to store the operating systems and the application software.

Figure 5 illustrates the motherboard block diagram of the design.

Figure 5: Motherboard Block Diagram

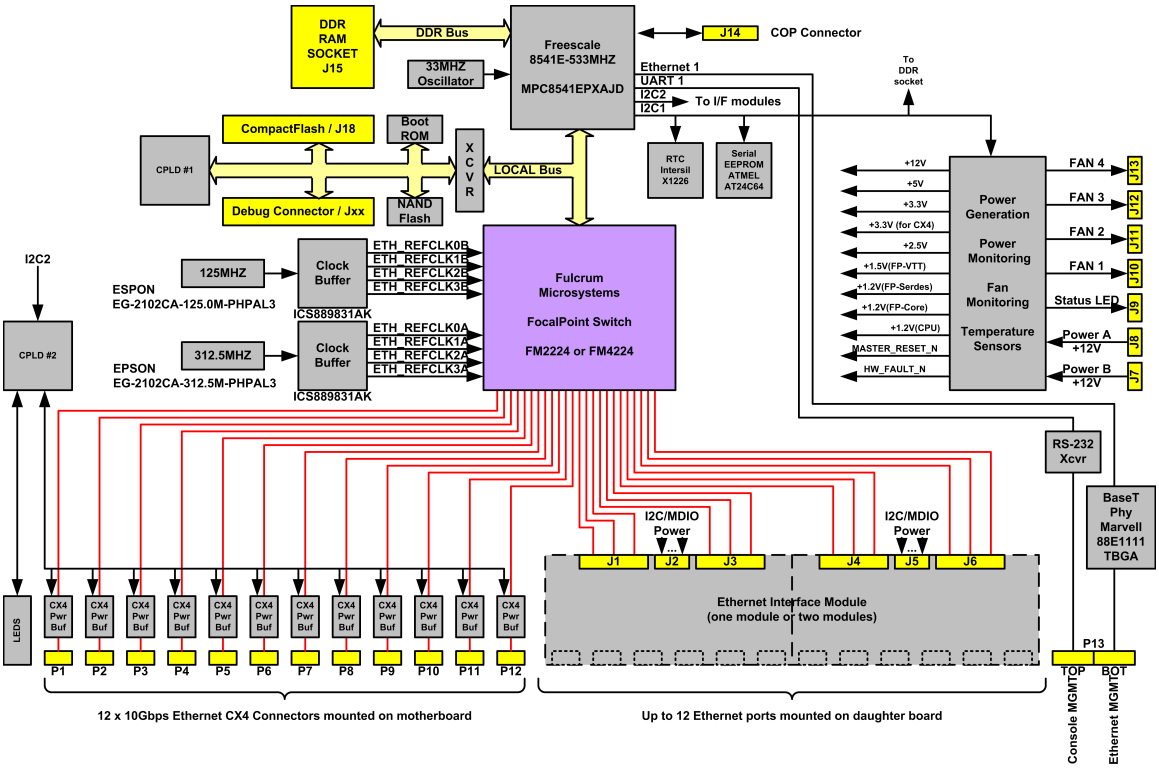

The interface modules are connected to the motherboard for high-speed differential pairs board-to-board connector. Each interface module is controlled by an I2C interface and an MDIO interface.

2.3 Connectors Pinout

2.3.1 CX4 Pinout

The CX4 pin out is shown on the next table. It is compliant to IEEE 10GBase-CX4 with the addition of power over CX4.

Table 1: CX4 Connector Pinout

Pin	Name	Direction
S1	RX Lane 0 +	Input

Pin	Name	Direction
S2	RX Lane 0 -	Input
S3	RX Lane 1 +	Input
S4	RX Lane 1 -	Input
S5	RX Lane 2 +	Input
S6	RX Lane 2 -	Input
S7	RX Lane 3 +	Input
S8	RX Lane 3 -	Input
S9	TX Lane 3 -	Input
S10	TX Lane 3 +	Input
S11	TX Lane 2 -	Input
S12	TX Lane 2 +	Input
S13	TX Lane 1 -	Input
S14	TX Lane 1 +	Input
S15	TX Lane 0 -	Input
S16	TX Lane 0 +	Input
G1	Ground	Ground
G2	TXDIS_N	Output
G3	Ground	Ground
G4	Ground	Ground
G5	Ground	Ground
G6	FAULT_N	Input
G7	DETECT	Input

Pin	Name	Direction
G8	+3.3V	Power
G9	Ground	Ground
G10	Ground	Ground
G11	Ground	Ground
G12	Ground	Ground
G13	Ground	Ground
G14	Ground	Ground
G15	Ground	Ground

2.3.2 Dual RJ45 Pinout

2.3.2.1 RS232 Console Pinout

The RS-232 on RJ45 is compliant to EIA/TIA 561 and is shown in the next table. The unit is shipped with an RJ-45 to DB9 female cross-over adapter requiring a simple straight male-female cable for connection to a standard personal computer.

Table 2: RS-232 on RJ45 Pinout

PIN	Signal	Direction	Function
1	DSR	Input	Data Set Ready
2	DCD	Input	Data Carrier Detect
3	DTR	Output	Data Terminal Ready
4	SGND	Ground	Ground
5	RD	Input	Receive Data
6	TH	Output	Transmit Data
7	CTS	Input	Clear to Send
8	RTS	Output	Request to Send

2.3.2.2 Ethernet Pinout

The Ethernet on RJ45 is compliant to IEEE 802.3 specification for 10/100/1000BASET. The interface is auto-crossover, auto-speed and auto-polarity. The pinout is shown on the next table.

Table 3: Ethernet Pinout

PIN	Signal	Direction	Function
1	TX_D1+	Auto-sense	Transmit Data
2	TX_D1-	Auto-sense	Transmit Data
3	RX_D2+	Auto-sense	Receive Data
4	BI_D3+	Auto-sense	Used in 1000BASET only
5	BI_D3-	Auto-sense	Used in 1000BASET only
6	RX_D2-	Auto-sense	Receive Data
7	BI_D4+	Auto-sense	Used in 1000BASET only
8	BI_D4-	Auto-sense	Used in 1000BASET only

2.3.3 Interface Module Pinout

Six connectors are used to interconnect to interface module; 4 are high-speed differential connectors (Fujitsu FCS-268F024-G0/D on the main board mating with Fujitsu FCS-268M024-G/3D on the interface module) and are used for the external 10G interface and 2 are 4x4 headers and are used for management of the module (Samtec SQW-104-01-F-Q on the main board mating to Samtec TW-04-03-L-Q-300-090 on the interface module).

Each high-speed differential connector carries 3 10GE ports. The pinout is shown on the next table with the pin numbering actually used on the schematic along with the pin number used by Futjisu.

Table 4: High-Speed Interface Module Connector Pinout

Pin	Fujitsu Pin #	Signal	Signal	Signal	Signal
1-4	G1,G2,S1,S2	GND	GND	RX1A-	RX1A+
5-8	G3,G4,S3,S4	GND	GND	TX1A-	TX1A+
9-12	…	GND	GND	RX1B-	RX1B+
13-16	…	GND	GND	TX1B-	TX1B+

Pin	Fujitsu Pin #	Signal	Signal	Signal	Signal
17-20	…	GND	GND	RX1C-	RX1C+
21-24	…	GND	GND	TX1C-	TX1C+
25-28	…	GND	GND	RX1D-	RX1D+
29-32	…	GND	GND	TX1D-	TX1D+
33-36	…	GND	GND	RX2A-	RX2A+
37-40	…	GND	GND	TX2A-	TX2A+
41-44	…	GND	GND	RX2B-	RX2B+
45-48	…	GND	GND	TX2B-	TX2B+
49-52	…	GND	GND	RX2C-	RX2C+
53-56	…	GND	GND	TX2C-	TX2C+
57-60	…	GND	GND	RX2D-	RX2D+
61-64	…	GND	GND	TX2D-	TX2D+
65-68	…	GND	GND	RX3A-	RX3A+
69-72	…	GND	GND	TX3A-	TX3A+
73-76	…	GND	GND	RX3B-	RX3B+
77-80	…	GND	GND	TX3B-	TX3B+
81-84	…	GND	GND	RX3C-	RX3C+
85-88	…	GND	GND	TX3C-	TX3C+
89-92	…	GND	GND	RX3D-	RX3D+
93-96	G47,G48,S47,S48	GND	GND	TX3D-	TX3D+
97-98	G49,G50	GND	GND	GND	GND

The two interface module management connectors pin out are shown in the next table. The +3.3V is provided for convenience but the power should normally be taken from the +12V inputs and converted to lower voltage

on the interface module as desired. The IO[0:3] will be sampled while the interface module is in reset to detect the interface type (both RESET_N and IO[8] are pulled down during board ID).

Table 5: J2 Interface Module Connector

PIN	SIGNAL	PIN	SIGNAL	PIN	SIGNAL	PIN	SIGNAL
A1	+12V	B1	+12V	C1	+12V	D1	+3.3V
A2	GND	B2	GND	C2	GND	D2	RSVD
A3	IO0	B3	IO1	C3	IO2	D3	IO3
A4	RSVD	B4	RSVD	C4	IO8	D4	INT#

Table 6: J5 Interface Module Connector

PIN	SIGNAL	PIN	SIGNAL	PIN	SIGNAL	PIN	SIGNAL
A1	+12V	B1	+12V	C1	+12V	D1	+3.3V
A2	GND	B2	GND	C2	GND	D2	Pull-up
A3	IO5	B3	IO6	C3	MDC	D3	IMDIO
A4	SDA	B4	SCL	C4	RESET#	D4	INT#

If a full size module is inserted, then the J2 is not used except to carry extra power if needed.

2.3.4 Power Input Pinout

There are two redundant power input connectors. Each connector has 14 pins and the pin out is listed on the next table.

Table 7: Power Connector Pinout

Pin	Usage	Pin	Usage
1	GND	2	+12V
3	GND	4	+12V
5	GND	6	+12V
7	GND	8	+12V

Pin	Usage	Pin	Usage
9	GND	10	+12V
11	GND	12	+12VSense
13	NC	14	OFF

2.3.5 Fan Connector Pinout

The fans are powered through a +12V source and support also a tachometer output which is connected to the CPU. The tachometer output is a simple momentary connection to ground every time the fan spin one turn. The PWM is controlled through an Analog ADT7462 system controller, the FAN is driven when the signal is asserted to 0.

Table 8: Fan Connector Pinout

Pin	Usage
1	PWM
2	TACHO
3	+12V

2.3.6 COP Port Pinout

See Figure 7.

2.3.7 DDR SODIMM Socket Pinout

The 200-pin DDR-SODIMM socket pinout is standard and conforms to the JEDEC standard.

2.3.8 Compact Flash Pinout

The CompactFlash socket pinout is standard and conform to the Compact Flash Standard Association.

2.4 Power Distribution

The power distribution is shown in Figure 6. The two +12V/14A inputs are combined through ORing diodes (15V VRRM, 15A, low drop out Shottkey) to produce one single +12V output. The combined +12V is then distributed to the interface modules, to the fans and to the various DC/DC converters which produces the local voltages required for the difference components on the board.

Figure 6: Local Power Generation and Distribution

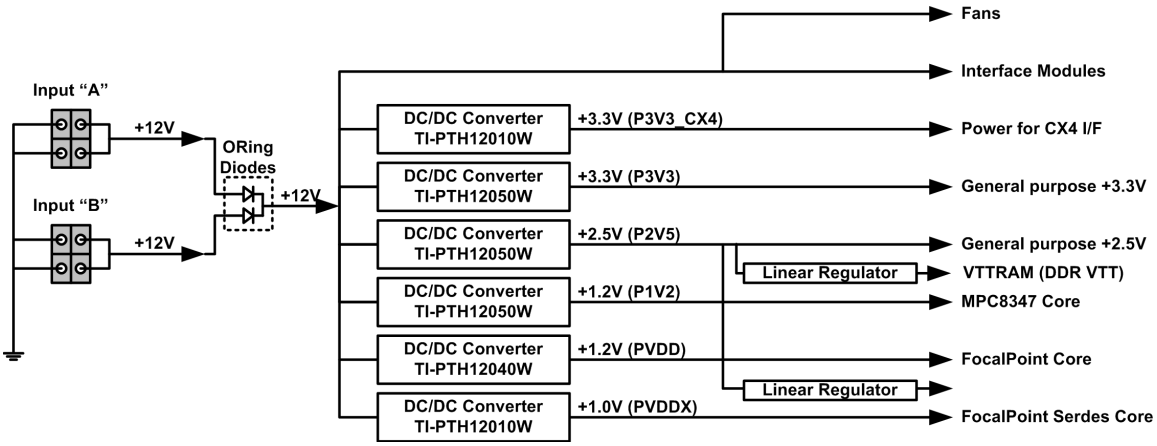

The 160W power budget is distributed along the following elements:

- 48W to the FocalPoint
    - The FocalPoint requires 1.2V for both the SERDES interface and the switch core. There are driven by two distinct DC/DC to ensure latitude in setting the power for those two critical modules.
- 36W to the Ethernet interface modules (3W per port)
- 10W for CPU subsystem
- 36W for power over CX4 (3W per port)
- 85% DC/DC efficiency
- 10W for FANs

The board contains two power monitoring circuit. The first one is a Linear Technologies LTC2900 and is used to monitor the power output for the CPU subsystem power distribution and is connected to the reset button. The reset produces a clean MASTER_RESET# signal which is de-asserted when all power is present and the reset button has been released. The three +1.2V sources are ORed together via normal diodes causing a reset if any one of these three sources becomes inoperative. The second monitoring circuit (Analog Device ADT7462) is used to control fans, measure temperature and monitor other power sources. The next figure shows the monitoring circuit.

Figure 7: Monitoring Power and Fan Controller

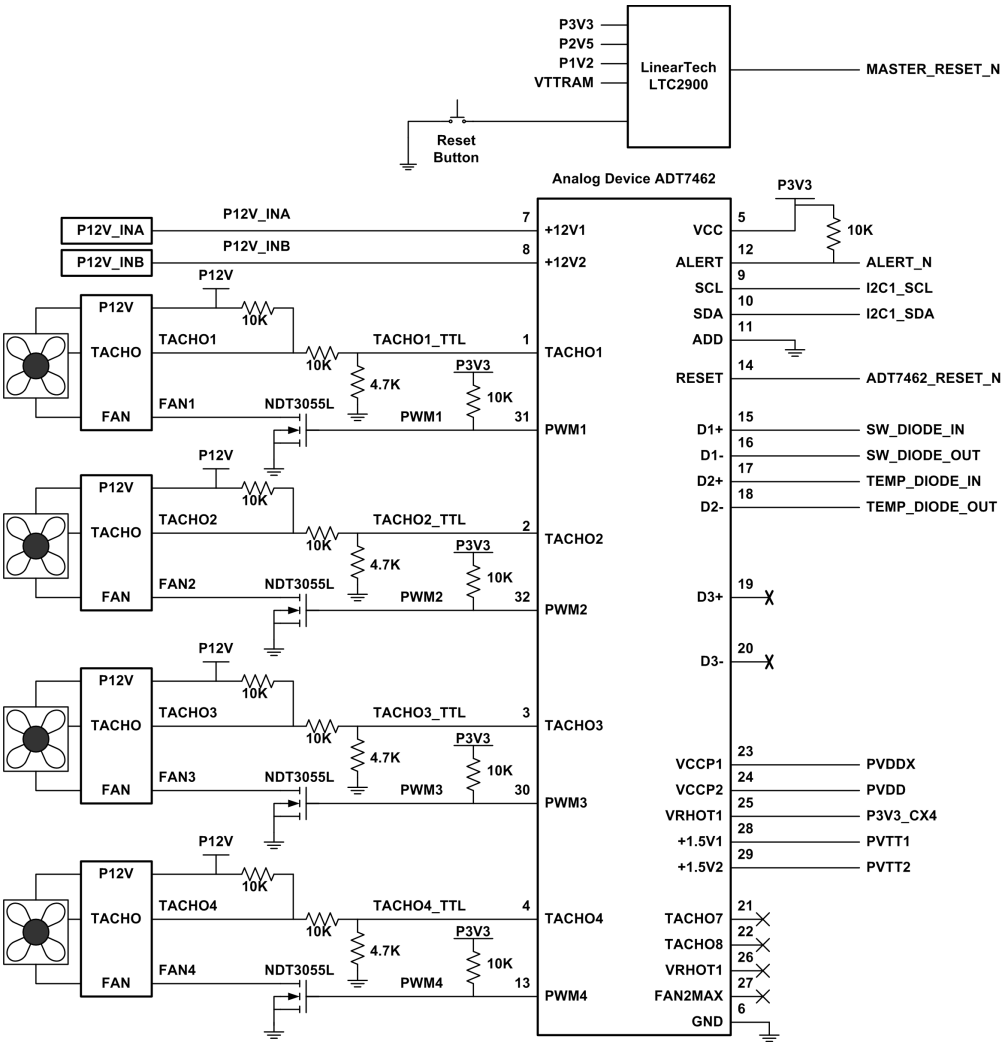

2.5 Unit Controller (MPC8541E)

The unit controller is a 533-MHz Freescale MPC8541E processor. This processor supports the following devices:

- Dual UART (only one used in this design)
- Dual Ethernet (only one used in this design)
- PCI bus
- Local bus
- DDR SDRAM bus

- Timers
- General purpose I/O
- Two I2C

Each of these elements is detailed in the next sections.

2.6 Reset and COP Port

The MASTER_RESET# from the power block is connected to the PORESET# of the MPC8541E and to the reset pins of the BOOT FLASH and CPLDs ensuring that the FLASH are ready for fetch once the master reset is released. While in reset, the MPC8541E will place all I/O pins in tri-state. Some of those I/O pins are used as GPIO pins by the software and will have external pull-up/pull-down pins to ensure a known state at power-up.

When the master reset is de-asserted, the MPC8541E detects if HRESET# is asserted externally via the COP port and, if not, detects its boot mode by reading the pins RESET_CONFIG[2:0] which are pull-down to 0 forcing the MPC8541E to retrieve its configuration from the boot ROM.

2.7 Local Bus

The local bus interconnects the MPC8541E to the FocalPoint, the BOOT ROMs, the compactFlash, the NAND flash, the CPLD and the bus expansion. The MPC8541E is the bus master of the local bus. The local bus runs at either 33MHZ/66MHZ. The local bus driven by the MPC8541 is multiplexed and must thus be de-multiplexed before being used. This is achieved using simple 3.3V latches connected on ADS. The data portion of the bus is buffered through a pair of 16-bit transceivers to reduce the load and allowing higher speed to be performed between the processor and the switch.

Note that the bit numbering has been reversed as compared to the PowerPC bit numbering to match the more conventional bit numbering of the other devices.

The following signals are part of the local bus:

- LBI_ADDR[24:2]: bussed, address
- LBI_DATA[31:0]: bussed, data, note that the CPLD is only using 8 bits
- LBI_R/W#: bussed, read, write
- LBI_CS[3:0]: point to point, chip select, 0=bootFlash, 1=compactFlash, 2=FM2224, 3=cpld
- LBI_ALE: Latch enable
- LBI_AS#: bussed, address strobe
- LBI_RDY#: bussed, data ready
- LBI_INT[1:0]: point-to-point, interrupt, 0=FM2224, 1=CPLD

Figure 8: Connection from MPC8541E to Local Bus

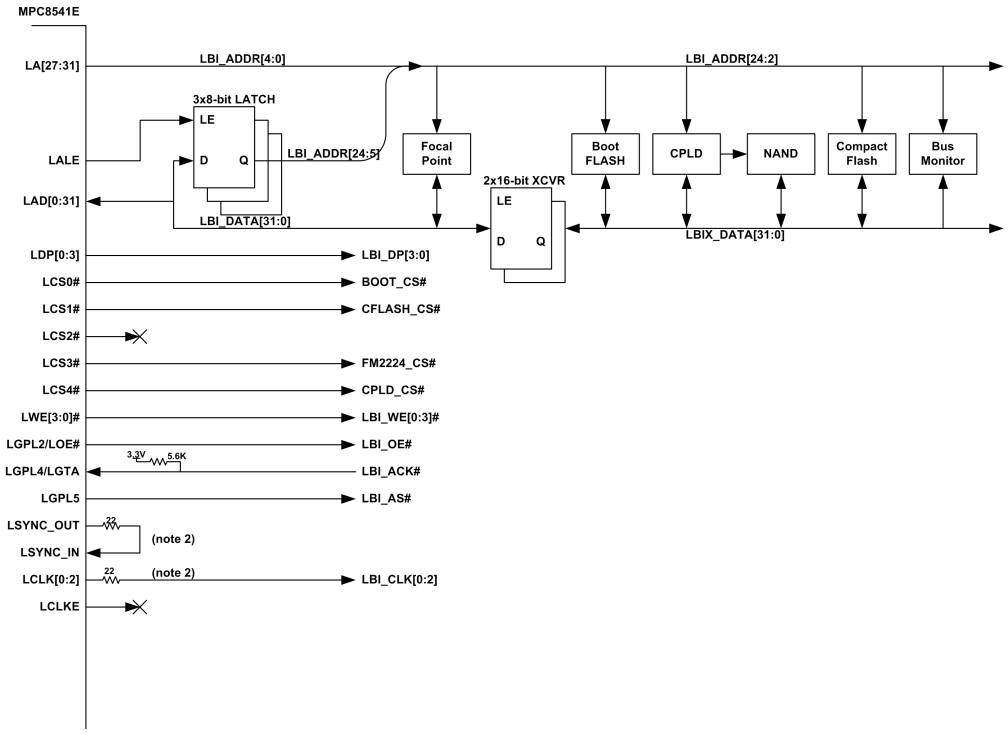

2.8 DDR SDRAM

The MPC8541E supports a PC2100 DDR SDRAM interface with or without ECC. It is connected to a 64/72-bit SO-DIMM socket (Tyco #1612507-1) and could support up to 1GB of PC2100 DDR SDRAM. The connection is shown in the next figure and follows the routing guidelines defined in [FREESCALE2].

Figure 9: DDR SDRAM Interface

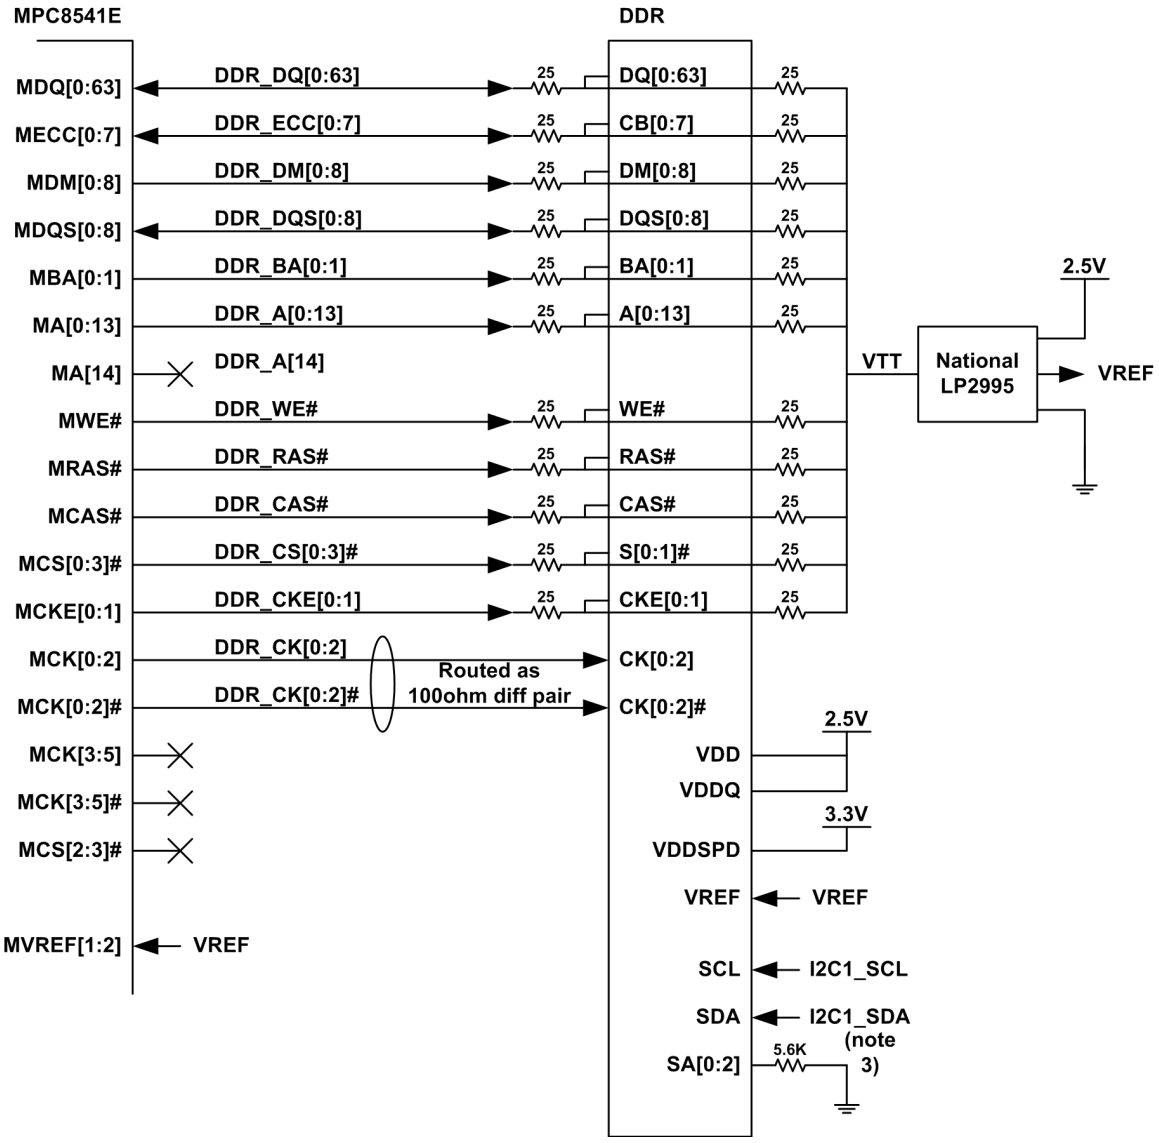

2.9 I2C Bus

The MPC8541E supports two I2C buses; I2C1 and I2C2. The first bus, I2C1, is connected on the devices listed in the next table. The second bus, I2C2, is connected to the Ethernet Interface module and the CPLD that controls the LEDs.

Table 9: Main Board I2C Devices on I2C Bus #1

Address	Function
27h	Connected to a TIPCA9555 for board identification. The lower 8 bits of the GPIO contains the board ID while the upper 8 bits contains the assembly/version ID.
2Ch, 2Dh	Connected to two Analog Device AD5173BRM50 software programmable 50K ohm resistors. Each device have two programmable resistors. They are connected as follow: • 2Ch resistor 1: controls FM2224 VDD core voltage • 2Ch resistor 2: controls FM2224 VDDX serdes core voltage • 2Dh resistor 1: controls FM2224 VTT1 (external ports 1 through 6) • 2Dh resistor 2: controls FM2224 VTT2 (external ports 7 through 12) The actual voltage programmed will be equal to: • VDD = 8000/(R+11456) + 0.8V • VDDX = 8000/(R+11440) + 0.8V • VTT = 1.21V + 0.731 * R / (2430+R) Where R is: • R = 50000 * D / 256 + 100; And D is the actual 8-bit unsigned value programmed into the device (0-255).
30h & 50h	Connected to SO-DIMM for DRAM identification. The content of the I2C is described in the JEDEC SO-DIMM specification.
51h	Connected to an ATMEL AT24C64 8Kx8 EEPROM. This device is so store bootstrap environment variables such as name and location of the file to boot, Ethernet MAC addresses, operating voltages, etc… The EEPROM contains environment definitions in an ASCII format, each definition is terminated by a NUL character (0x00) and the last line is terminated is a NULL.
57h & 6Fh	Connected to Intersil X1226 real time clock / NVRAM. The address 57h gives access to memory while address 6Fh is for real time clock. The memory is 256 bytes long. The RTC is protected by 1F super-capacitor and will provide backup for few days when power is not available.
58h	Connected to Analog Device ADT7462 system monitoring circuit.

Table 10: Main Board I2C Devices on I2C Bus #2

Address	Function
02h	Connected to the CPLD that controls equalizers and LEDs. The CPLD supports the following registers. The register address is 8-bit. Register 0x01 through 0x0C: • Control/status ports 1 through 12. Data is 8 bits. • Write: o Bit 0-1: LED color for each port - 00: OFF - 01: Green (will blink on data RX or TX) - 10: Yellow (solid) - 11: Alternating yellow/green • Read: o Bit 0-1: Color programmed per port. o Bit 2: indicate presence of an active cable (for ports 1 through 12 only) o Bit 3: indicate a fault report from the active cable (for ports 1 through 12 only) Register 0x11 through 0x1C: • Control/status ports 13 through 24. Data is 8 bits. • Write: o Bit 0-1: LED color for each port - 00: OFF - 01: Green (will blink on data RX or TX) - 10: Yellow (solid) - 11: Alternating yellow/green • Read: o Bit 0-1: Color programmed per port. Register 0x21 through 0x2C: • Configure QLM4302 equalizers ports 1 through 12. Data is 24 bits. Most significant bit sent first.

Address	Function
	• Write: o Bit 23-19: Lane 3 o Bit 18-14: Lane 2 o Bis 13-9: Lane 1 o Bits 8-4: Lane 0 o Bits 3-0: Must be 0xF • Read o The equalizers can only be written. Register 0xFF: • Global control. Data is 8 bits. • Write o Bit 0: Controls if the QLM4302 are software controlled (1) or not (0). o Bit 1: Enables Power over CX4 (1) or not (0). It set to 1, then the power will be applied if and only if an active cable is connected. If set to 0, then the power is never applied regardless of the type of cable connected. o Bit 7:2: Not used. • Read: o Returns the value written.
All others	Reserved for interface module.

2.10 Flash Memory

The board contains 3 types of flash memory:

- 8MBytes NOR FLASH used for boot code
  - Accessible through processor CS0# (CS_BOOT_N).
- 1GB NAND FLASH used to store application software
  - Accessible through processor CS3# through CPU CPLD
- CompactFLASH interface for extra storage.

o Accessible through processor CS1# (CS_FLASH_N)

2.11 UARTS & Console Port

The MPC8541E supports two UARTs. The first one is connected to the console port through an RS-232 transceiver while the second one is not used.

2.11.1 Ethernet Management Port

The MPC8541E supports two Ethernet ports. The first one is routed to the Ethernet management port through a Marvell 88E1111 Ethernet BASE transceiver configured to operate in GMII mode. The second port is not used.

Figure 10: MPC8541E to BASET PHY

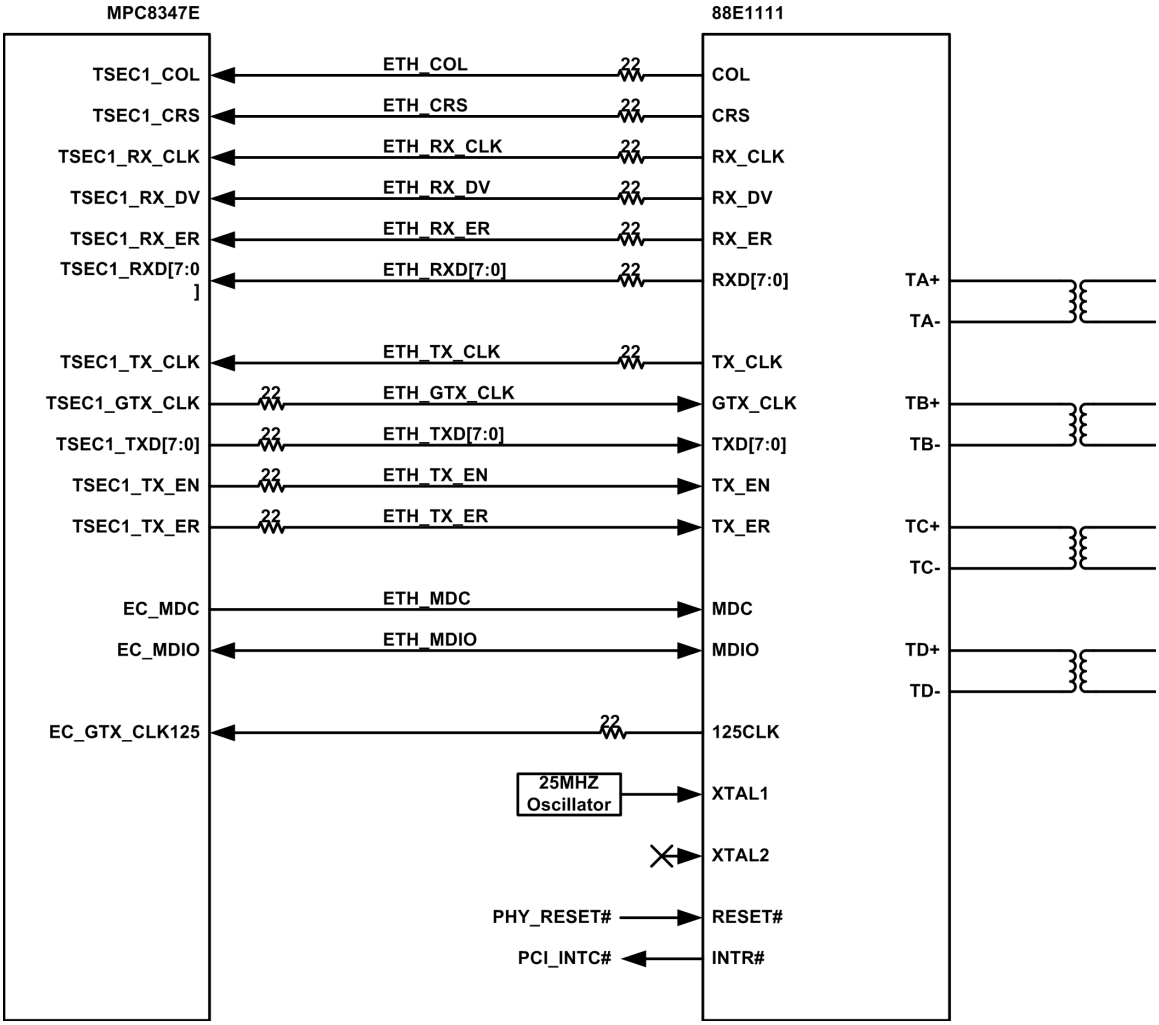

Refer to MARVELL[2] for the reference design on the 88E1111 and refer to MARVELL[3] for the bill of material for this reference design.

2.12 Ethernet Switch

The FocalPoint is the main component of the motherboard. It offers 24 x Ethernet interfaces, a bus interface for management, reference clocks, LED outputs and an EEPROM controller. The LED outputs are routed to a CPLD which may be used to override any status projected by the FocalPoint switch.

The Ethernet ports are routed to the 12 CX4 connectors on the motherboard and to the 2 Ethernet interface modules. The ports are routed as shown in Figure 14. The recommended routing is designed to reduce the number of layers required to route the interfaces by avoiding ports to cross over each other. Each Ethernet interface consists of 8 differentials pairs, 4 pairs in the transmit direction and 4 pairs in the receive direction. The figure shows also the VTTs power disposition.

Figure 11: FM2224 Ethernet Port Routing

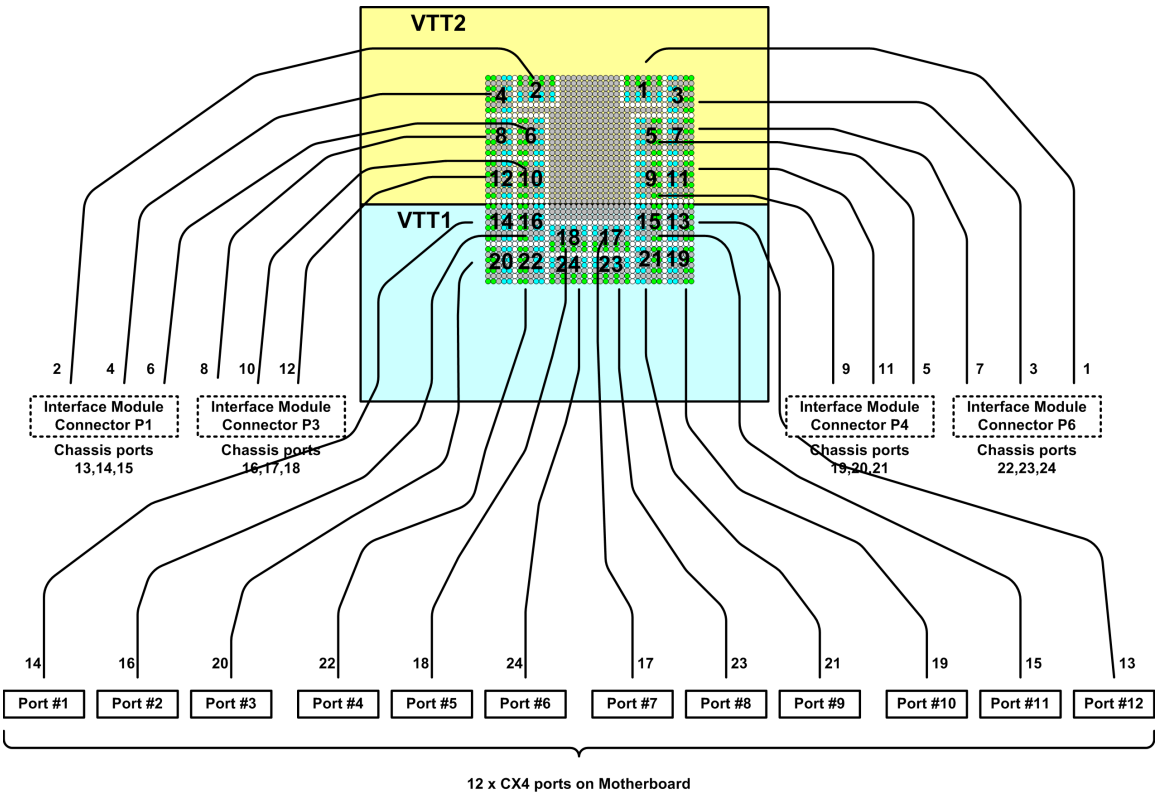

The reference clocks for the Ethernet ports are supplied by the clock distribution module as described in section 5.13. The FM2224 EEPROM controller is connected to the SPI interface of the MPC8541E and is not used by the software.

2.13 Clock Generation and Distribution Circuit

The following clocks are generated on the board:

- 32.768KHz: RTC reference clock
- 33.33MHZ: CPU reference clock
- 33.33MHZ: switch frame handler reference clock
- 312.5MHZ: switch port reference clock "A"/P>
- 125MHZ: switch port reference clock "B"

Note that the FM2224 support two independent reference clocks for each group of 6 ports for a total of eight reference clocks. Each port could be configured to select between two clocks. In this reference design, a 312.5MHZ reference clock is supplied to each group making it possible to run any port at the standard 10G interface rate and a second reference clock of 125MHZ is supplied to each group for support of the 10/100/1G BASET Interface Module.

Table 11: Reference Clocks

Reference	Clock Frequency	Usage	FM2224 Ports	External Ports
REFCLK1A	312.5MHZ	2.5/10Gbps	1,3,5,7,9,11	10,11,12,22,23,24
REFCLK1B	125.0MHZ	10/100/1Gbps	1,3,5,7,9,11	10,11,12,22,23,24
REFCLK2A	312.5MHZ	2.5/10Gbps	2,4,6,8,10,12	1,2,3,13,14,15
REFCLK2B	125.0MHZ	10/100/1Gbps	2,4,6,8,10,12	1,2,3,13,14,15
REFCLK3A	312.5MHZ	2.5/10Gbps	13,15,17,19,21,23	7,8,9,19,20,21
REFCLK3B	125.0MHZ	10/100/1Gbps	13,15,17,19,21,23	7,8,9,19,20,21
REFCLK4A	312.5MHZ	2.5/10Gbps	14,16,18,20,22,24	4,5,6,16,17,18
REFCLK4B	125.0MHZ	10/100/1Gbps	14,16,18,20,22,24	4,5,6,16,17,18

2.14

2.15 CX4 Interface

Each CX4 interface supports power over CX4 for support of active cables and support a programmable equalizer/amplifier to recover signals from very long CX4 cables. The power over CX4 is asserted by the CPLD if the power over CX4 is enabled and the presence of an active cable is detected. The circuit is shown in the next figure:

Figure 12: CX4 Interface

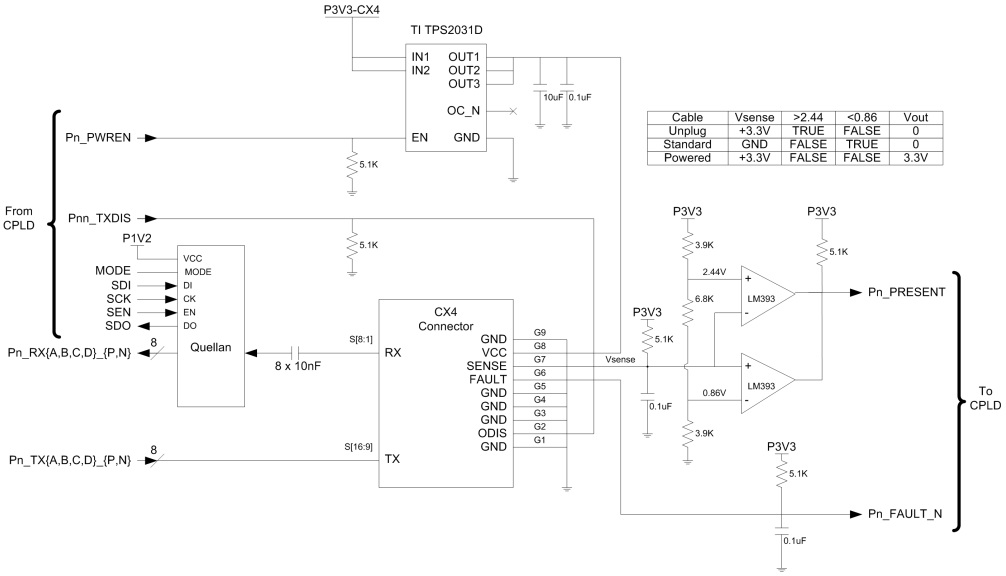

2.16 Intersil Real Time Clock

The Intersil X1226 real time clock is connected to the MPC8541E through the I2C interface. It is backed up by a 1F super pack to preserve the date and time for up to 1 week in absence of power.

2.17 Atmel Serial EEPROM

An ATMEL 24C64 is connected to the MPC8541E through the I2C interface. The device is used to store boards parameters for the boot loader.

2.18 JTAG Chains

There are 3 JTAG chains:

- Chain 1: CPU only
    - o This chain is through J14 and include the CPU only.
    - o Its main usage is to connect a COP debugger for software development.
- Chain 2: CPLDs only
    - o This chain is accessible through J9 and include the two CPLDs (CPLD_CPU and CPLD_INTRF).
    - o Its main usage is to connect a CPLD programmer.
- Chain 3: all components

- This chain is accessible through J17 and requires the JTAG external controller to drive the pin 15 (JTAG_USER0) to logic 1 to complete the chain. The elements in the chain are, in order: CPU => CPLD_CPU => CPLD_INTRF => FM2224 => Marvel 88E1111.

- Its main usage is for board testing.

The hardware allows both J14 and J9 to operate independently of each other provided J17 is not installed. J17, when used, does require both J14 and J9 to be left unconnected. The overall diagram is shown on the next figure.

Figure 13: JTAG Logic

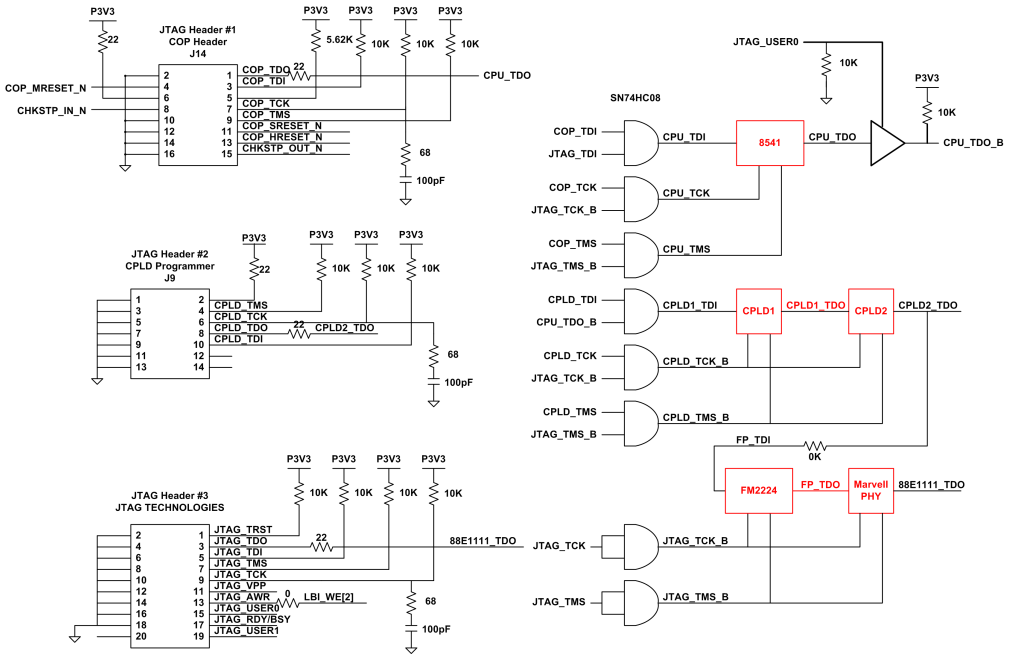

3. Boot Firmware

The boot firmware is based on U-Boot (http://sourceforge.net/projects/u-boot), a flexible boot loader available under the GNU General Public License, and has been adapted to this platform. The source is available from Fulcrum.

3.1 Memory Map

The boot firmware is stored in the BOOT flash at location 0xFF700000 and it initializes the system with the memory map shown in the next table.

Table 12: Memory Map

Address	Function
0x00000000 – 0x7FFFFFFF	DDR Memory. The unit is shipped with a 2-bank 512MB memory module mapped as follow: • 0x00000000 – 0x0FFFFFFF: Bank 0 • 0x10000000 – 0x1FFFFFFF: Bank 1
0x80000000 – 0x9FFFFFFF	PCI 1 : to PMC module
0xA0000000 – 0xBFFFFFFF	PCI 2 : not connected
0xE0000000 – 0xE00FFFFF	CCSR, Internal Memory Mapped Registers
0xE2000000 – 0xE2FFFFFF	PCI 1 : I/O (16MB)
0xE3000000 – 0xE3FFFFFF	PCI 2 : I/O (16MB)
0xE8000000 – 0xE801FFFF	L2 Cache used as SRAM. This address space is used momentary during boot initialization while the DDR controller is initialized.

0xF0000000 – 0xF001FFFF	CompactFlash. The CompactFlash is configured in tru-IDE mode and uses two chip selects out of the 8541 (CS1 and CS2). The memory map for the different registers of the compactFlash is shown below.
	Reading:
	`0xF0000006(CS1): Data Register (16 bits)`
	`0xF0000007(CS1): Error Register (8 bits)`
	`0xF000000B(CS1): Sector Count (8 bits)`
	`0xF000000F(CS1): Sector No (8 bits)`
	`0xF0000013(CS1): Cylinder Low (8 bits)`
	`0xF0000017(CS1): Cylinder High (8 bits)`
	`0xF000001B(CS1): Select Card/Head (8 bits)`

Address	Function
	0xF000001F(CS1): Status (8 bits) 0xF001001B(CS2): Alt Status (8 bits) Writing: 0xF0000006(CS1): Data Register (16 bits) 0xF0000007(CS1): Features (8 bits) 0xF000000B(CS1): Sector Count (8 bits) 0xF000000F(CS1): Sector No (8 bits) 0xF0000013(CS1): Cylinder Low (8 bits) 0xF0000017(CS1): Cylinder High (8 bits) 0xF000001B(CS1): Select Card/Head (8 bits) 0xF000001F(CS1): Command (8 bits) 0xF001001B(CS2): Device Control (8 bits)
0xF0020000 – 0xF002FFFF CPLD	Access to CPU CPLD Address map: 0xF0020000: Read or write into NAND (8 bits) 0xF0020004: Controls NAND pins: bit 0: Select (1 = selected, 0=not selected) bit 1: Address (A = bit) bit 2: Command (C = bit) bit 3: PRL (PRL = bit) bit 4: Write Protect (1 = protected, 0 = unprotected)
0xF1000000 – 0xF3FFFFFF FM2224	Reserved for FM2224
0xFF700000 – 0xFF7FFFFF CCSR.	Internal Memory Mapped Registers. This area is the default area for the internal registers after reset and gets relocated to 0xE0000000 after initial stages are completed.
0xFF800000 – 0xFFFFFFFF	Boot FLASH (8MB). The boot flash usage is used as follow: 0xFF800000 - 0xFFEFFFFF: Free (7MB) 0xFFF00000 - 0xFFFBFFFF: Boot code (768KB) Sections: .text, .data, .bss 0xFFFC0000 - 0xFFFDFFFF: Environment (128KB) This section is used only on rev A and rev B of hardware, the rev C stores the environment variable in a I2C serial EEPROM. 0xFFFFF000 - 0xFFFFFFFB: Boot page (4KB-4)

Address	Function
	Sections: .boot 0xFFFFFFFC - 0xFFFFFFFF: Start vector (4 bytes) Sections: .reset

3.2 Boot Firmware Console Access

The boot firmware starts to execute at power-up and will display information on the console port on its progress through the board initialization. The firmware will also start the application software automatically. The console is set at factory to operate at 115,200bps (8-bits no parity, no flow control).

The boot firmware autoboot process can be interrupted to allow the user access to firmware environment variables and run some basic diagnostics. By example, the default environment variables are set to boot the application from the flash system but it is possible to change this environment to boot from the network.

```
************************************************************************
***************** FULCRUM MICROSYSTEMS FM2224EP ***************************
************************* UBOOT VERSION 1.3 *******************************
************************************************************************

I2C: ready
DRAM: Initializing
DDR: 512 MB
Press any key for RAM test, otherwise, just wait 2s!
Key pressed, aborted
FLASH: 8 MB
NAND: 1024MB
L2 cache 256KB: enabled
env_relocate[214] offset = 0x20000000
env_relocate[233] malloced ENV at 1febe008
env_relocate[277] end of function
In: serial
Out: serial
Err: serial
Net: TSEC0: PHY is Marvell 88E1111S (1410cc2)
TSEC0
IDE: Bus 0: OK
Device 0: Model: LEXAR ATA FLASH Firm: V1.00 Ser#: 11154282039199099093
Type: Removable Hard Disk
Capacity: 489.2 MB = 0.4 GB (1001952 x 512)
Setting VDD: mv=1250 r=6327 d=32 set to 1.25
Setting VDDX: mv=1250 r=6337 d=32 set to 1.25
Setting VTT1: mv=1500 dev=2d addr=0 r2=963 rpot=1594 d=8 set to 1.5
Setting VTT2: mv=1500 dev=2d addr=80 r2=963 rpot=1594 d=8 set to 1.5
Setting VTT3: mv=1500 dev=2e addr=0 r2=963 rpot=1594 d=8 set to 1.5
Setting VTT4: mv=1500 dev=2e addr=80 r2=963 rpot=1594 d=8 set to 1.5
```

```
Hit any key to stop autoboot: 5
=>
```

There are two pauses in the startup process.

- The first pause is right after the DDR memory is initialized and allows to execute a RAM test if desired. The RAM test is activated by entering any key and can be aborted at any time by entering another key. The RAM test is skipped if no characters have been received within 2 seconds.
- The second pause is after the firmware has completed initialization of the board. This timer can be aborted by entering any character and will allow the user to start an interactive mode. If no characters have been received, then the firmware will automatically execute the "BOOT" command which is normally set to start the application software contained in the compactFlash.

3.3 Motherboard Layer Stacking and Routing

The main board layer stacking is shown in the next figure. Three layers are used for routing the high-speed CX4 differential pairs away from the switch; two layers are used for receiving signals and one for transmitting signals. The usage of two layers for receiving signals ensure large trace separation between adjacent ports and minimize cap-coupling between two ports; one port could have a very low level signal received from a long cable while the adjacent port could have a strong signal received from a short cable. The main board uses only one layer for the transmission as the signal strength is about the same on all ports and there is little signal corruption caused by cap-coupling.

The FocalPoint Design Guide Lines provide more details on this subject.

Table 13: Main Board Layer Stackup

Layer	Name	Usage
1 (top)	TOP	100R differential & 50 R single ended
2	GND	
3	IA	100R differential / 50R Single Ended
4	GND	
5	PWR1	1.5V/2.5V/12V and SWITCH-VTT
6	GND	
7	PWR2	1.2V and SWITCH-VDD
8	IB	50R Single Ended
9	IC	50R Single Ended

10	PWR4	3.3V and SWITCH-VDDX
11	GND	
12	ID	100R differential / 50R Single Ended
13	GND	
14	IE	100R differential / 50R Single Ended
15	GND	
16 (bottom)	BOTTOM	100R differential / 50R Single Ended

The XAUI routing shall be done using the following guidelines:

- Trace width = 5mils
- Differential trace spacing = 5mils
- Spacing to next differential pairs >= 15 mils
- Use only one layer for routing one pair
- Variation in length within one pair <= 50 mils
- Variation in length from one pair to another pair on the same port <= 300 mils
- Dielectric thickness = 10 mils

4. Interface Modules

The motherboard supports 6 sockets to connect to two single size interface modules one double size interface module. Each single interface module needs 3 connectors: two Fujitsu Micro-GIGA supporting each 3 XAUI interfaces and one 4x4 2mm header for power and control. The next figure shows the signals on each connector.

The software detects the presence of a board by reading the I/Os pin while the board is in reset. These I/Os signals have weak 5.6K pull-up on the motherboard guaranteeing a known state (11b) if the board is not present. If the I/Os are used, then the interface module shall have a tri-state buffer as shown above to drive a known board identification while the board is in reset, the other outputs from the interface module shall be in tri-state during that time. If the I/Os are not used internally on the interface module, then the interface module could connect some of these signals to ground to provide a board identification, a buffer is not required in this case.

The interface modules ID are shown in the next table:

Table 14: Interface Modules IDs

IO[3]	IO[2]	IO[1]	IO[0]	Module
0	0	1	0	CX4 Module
0	0	1	1	XFP Module
0	1	0	0	SFP Module

4.1 CX4 Interface Module

The CX4 interface module is a replication of the CX4 circuit offered on the main board. The block diagram is shown on the next figure.

Figure 14: CX4 Interface Module

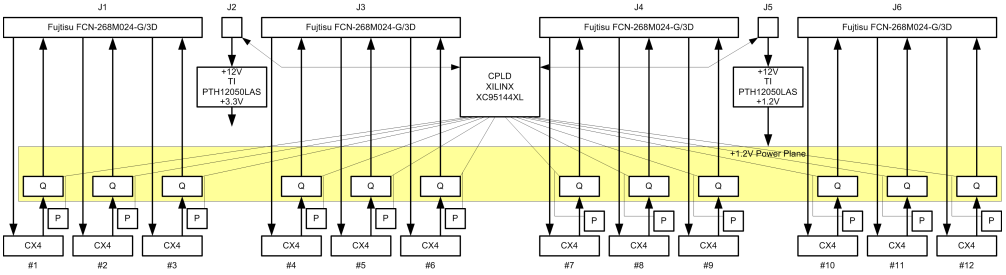

The interface module supports 12 ports, all supporting a QLM4302 equalizer and supporting power over CX4. The module carries two DC/DC modules; one for powering active cables and one for local power for the QLM4302. A CPLD is used to control the board and is accessible from the CPU through an I2C interface.

Table 15: Main Board I2C Devices on I2C Bus #2

Address	Function
02h	Reserved by main board.
04h	Connected to the CPLD that controls equalizers. The CPLD supports the following registers. The register address is 8-bit. Register 0x21 through 0x2C: • Configure QLM4302 equalizers ports 1 through 12. Data is 24 bits. Most significant bit sent first. • Write: o Bit 23-19: Lane 3 o Bit 18-14: Lane 2 o Bis 13-9: Lane 1 o Bits 8-4: Lane 0 o Bits 3-0: Must be 0xF • Read o The equalizers can only be written. Register 0xFF: • Global control. Data is 8 bits. • Write o Bit 0: Controls if the QLM4302 are software controlled (1) or not (0). o Bit 1: Enables Power over CX4 (1) or not (0). It set to 1, then the power will be applied if and only if an active cable is connected. If set to 0, then the power is never applied regardless of the type of cable connected. o Bit 7:2: Not used. • Read: o Returns the value written.
All others	Unused

4.2 XFP Interface Module

The XFP interface module supports twelve interface ports; each port consist of a XFP cage and an Aeluros AEL1004 PHY to convert the single lane 10.3125GE interface into four 3.125GE lanes. The block diagram is shown on the next figure.

Figure 15: XFP Module Block Diagram

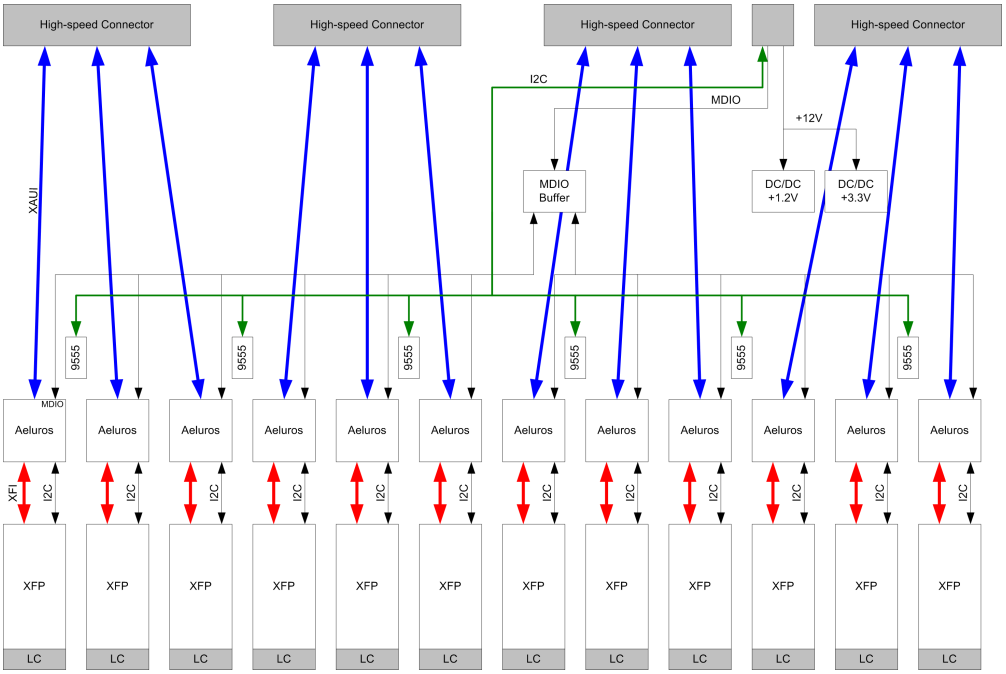

The Aeluros PHYs are controlled through an MDIO interface. The PHY address is equal to the port number (1 through 12). The MDIO registers are defined in the Aeluros data sheet.

The PHY and XFP have also extra pins (input and outputs) that are driven into an I2C GPIO device (TI PCA9555). There are six PCA9555, one for each two ports. The addresses are as follow:

- Address 0x20: ports 1 and 2
- Address 0x21: ports 3 and 4
- Address 0x22: ports 5 and 6
- Address 0x23: ports 7 and 8
- Address 0x24: ports 9 and 10
- Address 0x25: ports 11 and 12

The mapping of those signals is shown in the next table:

Table 16: XFP Module General Purpose I/Os.

PCA9555 Pins for Odd Ports (1,3,5,7,9,11)	PCA95555 Pins for Even Ports (1,3,5,7,9,11)	Function
P00	P17	CDR_LOL: • Input, positive logic • Reports lost of lock on clock recovery.
P01	P16	CMU_LOL: • Input, positive logic • Reports lost of lock on clock recovery.
P02	P15	LOS: • Input, positive logic • Reports lost of signal
P03	P14	RESET_N: • Output, negative logic. • Asserts RESET on PHY
P04	P13	TX_ON: • Output, positive logic • Enables transmitter
P05	P12	XFP_INT_N: • Input, negative logic. • Reports interrupt from XFP
P06	P11	XFP_PRESENT_N:

		• Input, negative logic. • Reports presence of XFP transceiver
P07	P10	XFP_DESEL: • Output, positive ligic • Deassert access to XFP via I2C

4.3 SFP+ Interface Module

The SFP interface module supports twelve interface ports; each port consist of a SFP cage and an AMCC QT2035S PHY to convert the single lane 10.3125GE interface into four 3.125GE lanes. The block diagram is shown on the next figure.

Figure 16: SFP+ Module Bloc Diagram

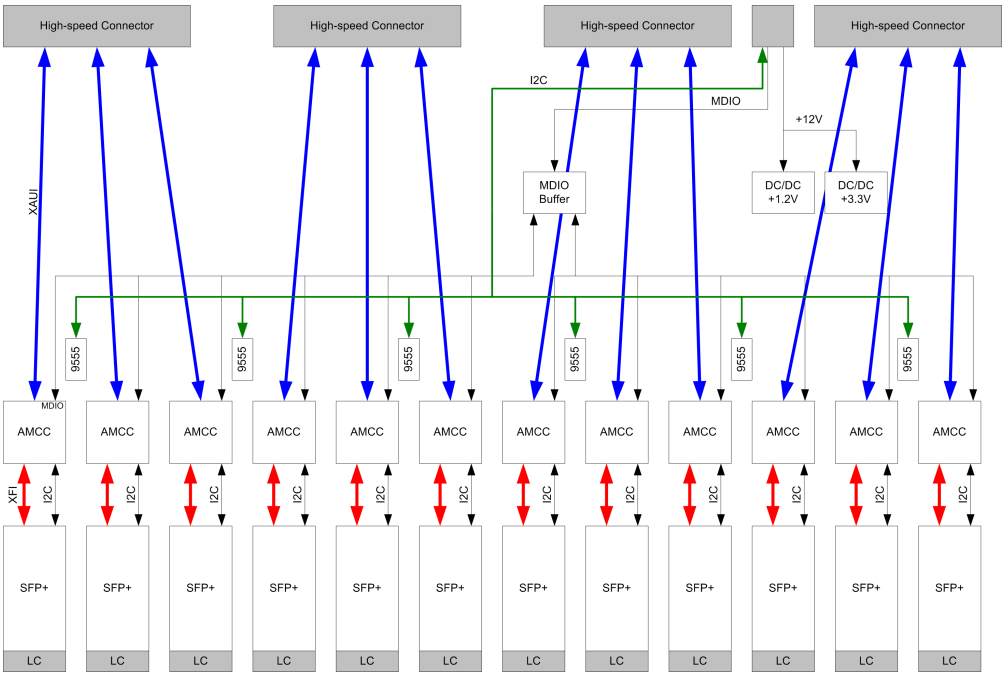

The AMCC PHYs are controlled through an MDIO interface. The MDIO interface of each PHY is shown in the following table:

Port	1	2	3	4	5	6	7	8	9	10	11	12
MDIO	8	0	12	2	10	6	14	1	9	5	13	3

The PHY and XFP have also extra pins (input and outputs) that are driven into an I2C GPIO device (TI PCA9555). There are six PCA9555, one for each two ports. The addresses are as follow:

- Address 0x20: ports 1 and 2
- Address 0x24: ports 3 and 4
- Address 0x22: ports 5 and 6
- Address 0x26: ports 7 and 8
- Address 0x21: ports 9 and 10
- Address 0x25: ports 11 and 12

The mapping of those signals is shown in the next table:

Table 17: SFP+ Module General Purpose I/Os.

PCX9555 Pins for Odd Ports (1,3,5,7,9,11)	PCA9555 Pins for Even Ports (1,3,5,7,9,11)	Function
P00	P10	Unused
P01	P11	Unused
P02	P12	Unused
P03	P13	Write Protect. • Output, positive logic. • If set to 1, then the I2C serial EEPROM storing the PHY code will not be writable.
P04	P14	LOS • Output, positive logic • Lost of signal
P05	P15	LASI • Input, positive logic. • LASI

P06	P16	RESET_N • Output, negative logic. • If set to 0, then the PHY is placed in reset. If set to 1, then the PHY runs.
P07	P17	PRESENT_N • Input, negative ligoc • Reports if the transceiver is present (0) or not (1).