

Intel® 82575 Gigabit Ethernet Controller

Design Guide V1.00

June 2007

317698-001



INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. Intel products are not intended for use in medical, life saving, life sustaining, critical control or safety systems, or in nuclear facility applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

Intel Corporation may have patents or pending patent applications, trademarks, copyrights, or other intellectual property rights that relate to the presented subject matter. The furnishing of documents and other materials and information does not provide any license, express or implied, by estoppel or otherwise, to any such patents, trademarks, copyrights, or other intellectual property rights.

IMPORTANT - PLEASE READ BEFORE INSTALLING OR USING INTEL® PRE-RELEASE PRODUCTS.

Please review the terms at http://www.intel.com/netcomms/prerelease_terms.htm carefully before using any Intel® pre-release product, including any evaluation, development or reference hardware and/or software product (collectively, "Pre-Release Product"). By using the Pre-Release Product, you indicate your acceptance of these terms, which constitute the agreement (the "Agreement") between you and Intel Corporation ("Intel"). In the event that you do not agree with any of these terms and conditions, do not use or install the Pre-Release Product and promptly return it unused to Intel.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. See http://www.intel.com/products/processor_number for details.

This document contains information on products in the design phase of development. The information here is subject to change without notice. Do not finalize a design with this information.

The 82575 Gigabit Ethernet Controller may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Hyper-Threading Technology requires a computer system with an Intel® Pentium® 4 processor supporting HT Technology and a HT Technology enabled chipset, BIOS and operating system. Performance will vary depending on the specific hardware and software you use. See http://www.intel.com/products/ht/Hyperthreading_more.htm for additional information.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an order number and are referenced in this document, or other Intel literature may be obtained by calling 1-800-548-4725 or by visiting Intel's website at <http://www.intel.com>.

Intel and Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

*Other names and brands may be claimed as the property of others.

Copyright © 2005-2007, Intel Corporation. All Rights Reserved.



Contents

1.0	Introduction	1
1.1	Scope	1
1.2	Reference Documents	2
2.0	PCI Express Port Connection to the Device	3
2.1	PCI Express Reference Clock	3
2.2	Other PCI Express Signals	3
2.3	Physical Layer Features	3
2.3.1	Link Width Configuration	3
2.3.2	Polarity Inversion	4
2.3.3	Lane Reversal	4
2.4	PCI Express Routing	5
3.0	Ethernet Component Design Guidelines	7
3.1	General Design Considerations for Ethernet Controllers	7
3.1.1	Clock Source	7
3.1.2	Magnetics for 1000 BASE-T	7
3.2	Designing with the 82575/EB/ES Gigabit Ethernet Controller	8
3.2.1	LAN Disable for 82575 Ethernet Controller Gigabit Ethernet Controller	9
3.2.2	Serial EEPROM	10
3.2.3	EEPROM Map Information	11
3.2.4	FLASH	12
3.3	SMBus and NC-SI	14
3.4	Power Supplies for the 82575 Ethernet Controller Controllers	15
3.4.1	82575 Ethernet Controller Power Sequencing	17
3.4.2	82575 Ethernet Controller Device Power Supply Filtering	19
3.4.3	82575 Ethernet Controller Controller Power Management and Wake Up	19
3.4.4	Power Management	20
3.5	82575 Ethernet Controller Device Test Capability	22
3.6	PHY Functionality	22
3.6.1	Auto Cross-over for MDI and MDI-X resolution	22
3.6.2	Smartspeed	23
3.6.3	Flow Control	23
3.6.4	Low-Power Link Up	23
3.6.5	Link Energy Detect	24
3.6.6	Polarity Correction	24
3.6.7	Auto-Negotiation differences between PHY, SerDes and SGMII	25
3.6.8	Copper PHY Link Configuration	25
3.7	Copper/Fiber Switch	26
3.8	Device Disable	27
3.8.1	BIOS handling of Device Disable	28
3.9	Software-Definable Pins (SDPs)	28
4.0	Frequency Control Device Design Considerations	30
4.1	Frequency Control Component Types	30
4.1.1	Quartz Crystal	30
4.1.2	Fixed Crystal Oscillator	30
4.1.3	Programmable Crystal Oscillators	31
4.1.4	Ceramic Resonator	31
5.0	Crystal Selection Parameters	32
5.1	Vibrational Mode	32
5.2	Nominal Frequency	32



- 5.3 Frequency Tolerance32
- 5.4 Temperature Stability and Environmental Requirements.....32
- 5.5 Calibration Mode.....33
- 5.6 Load Capacitance.....33
- 5.7 Shunt Capacitance.....34
- 5.8 Equivalent Series Resistance34
- 5.9 Drive Level34
- 5.10 Aging34
- 5.11 Reference Crystal34
 - 5.11.1 Reference Crystal Selection35
 - 5.11.2 Circuit Board.....35
 - 5.11.3 Temperature Changes35
- 6.0 Oscillator Support37**
 - 6.1 Oscillator Solution.....37
- 7.0 Ethernet Component Layout Guidelines.....39**
 - 7.1 Layout Considerations for 82575 Ethernet Controllers39
 - 7.1.1 Guidelines for Component Placement.....39
 - 7.1.2 Crystals and Oscillators42
 - 7.1.3 Board Stack Up Recommendations43
 - 7.1.4 Differential Pair Trace Routing for 10/100/1000 Designs44
 - 7.1.5 Signal Trace Geometry for 1000 BASE-T Designs45
 - 7.1.6 Trace Length and Symmetry for 1000 BASE-T Designs.....45
 - 7.1.7 Routing 1.8 V to the Magnetics Center Tap.....46
 - 7.1.8 Impedance Discontinuities46
 - 7.1.9 Reducing Circuit Inductance46
 - 7.1.10 Signal Isolation46
 - 7.1.11 Power and Ground Planes47
 - 7.1.12 Traces for Decoupling Capacitors47
 - 7.1.13 Light Emitting Diodes for Designs Based on the 82575 Controller47
 - 7.1.14 Thermal Design Considerations.....48
 - 7.2 Physical Layer Conformance Testing.....48
 - 7.2.1 Conformance Tests for 10/100/1000 Mbps Designs48
 - 7.3 Troubleshooting Common Physical Layout Issues48
- 8.0 Thermal Management50**
- 9.0 Reference Design Bill of Materials50**
- 10.0 Design and Layout Checklists50**
- 11.0 Reference Schematics50**
- 12.0 Symbol.....50**



Revision History

Date	Revision	Description
0.25	Jan 2006	Initial publication of preliminary design guide information.
0.50	July 2006	Added features listings, NC-SI, LED, strapping, pull-up/pull-down information.
0.75	March 2007	Changed classification to "Confidential"; updated crystal layout guidance; removed thermal sensor references; removed password requirements for schematic, checklist, and symbol files; updated EEPROM selection information.
1.00	June 2007	Changed classification to unclassified; removed information regarding Smart Power Down feature; changed signal name from LAN_PWR_GOOD to Internal_Power_On_Reset.



This page intentionally left blank.



1.0 Introduction

The Intel® 82575 Ethernet Controller is a single, compact component that offers two fully-integrated Gigabit Ethernet Media Access Control (MAC) and physical layer (PHY) ports. This device uses the PCI Express* (PCIe) architecture (Rev. 1.1RD). The 82575 enables two-port implementation in a relatively small area and can be used for server and workstation network designs with critical space constraints.

The 82575 provides:

- a standard IEEE 802.3 Ethernet interface for 1000BASE-T, 100BASE-TX, and 10BASE-T applications (802.3, 802.3u, and 802.3ab).
- a Serializer-Deserializer (SERDES) to support 1000Base-SX/LX (optical fiber) and Gigabit backplane applications. Information concerning SERDES can be found in the *82575 Ethernet Controller SERDES Application Note*.
- SGMII for SFP/external PHY
- management of MAC and PHY Ethernet layer functions
- management of PCI Express packet traffic across its transaction, link, and physical/logical layers.
- I/O Acceleration Technologies (I/OAT2). to accelerate the data transactions by hardware means optimizing the TCP flow and reducing the load on the CPU.

In addition, the 82575's on-board System Management Bus (SMB) ports enable network manageability implementations required by information technology personnel for remote control and alerting via the LAN. With SMB, management packets can be routed to or from a management processor. The SMB ports enable industry standards, such as Intelligent Platform Management Interface (IPMI) and Alert Standard Forum (ASF) 2.0, to be implemented using the 82575. In addition, on-chip ASF 2.0 circuitry provides alerting and remote control capabilities with standardized interfaces. The 82575 Ethernet Controller contains a dedicated microcontroller for manageability with with NC-SI and DMTF support.

The 82575 with PCIe architecture is designed for high-performance and low-host-memory access latency. The device connects directly to a system Memory Control Hub (MCH) or I/O Controller Hub (ICH) using one, two, or four PCI Express lanes.

Wide internal data paths eliminate performance bottlenecks by efficiently handling large address and data words. Combining a parallel and pipelined logic architecture optimized for ethernet and independent transmit and receive queues, the 82575 efficiently handles packets with minimum latency. The 82575 includes advanced interrupt handling features. It uses efficient ring buffer descriptor data structures, with up to 64 packet descriptors cached on chip. A large 48 KByte on-chip packet buffer maintains superior performance. In addition, using hardware acceleration, the controller offloads tasks from the host, such as TCP/UDP/IP checksum calculations and TCP segmentation.

The 82575 is packaged in 25mm x 25mm, 576-ball grid array.

1.1 Scope

This application note contains Ethernet design guidelines applicable to LOM designs based on PCI Express-supported chipsets.



1.2 Reference Documents

This application assumes that the designer is acquainted with high-speed design and board layout techniques. The following documents provide additional information:

- *82575 Ethernet Controller Product Datasheet*. Intel Corporation.
- *PCI Express Base Specification, Revision 1.1*. PCI Special Interest Group.
- *PCI Express Card Electromechanical Specification, Revision 1.0a*. PCI Special Interest Group.
- *PCI Bus Power Management Interface Specification, Revision 1.1*. PCI Special Interest Group.
- *IEEE Standard 802.3, 2002 & 2005 Edition*. Institute of Electrical and Electronics Engineers (IEEE).
- Incorporates various IEEE standards previously published separately.
- *Intel Ethernet Controllers Timing Device Selection Guide, AP-419*. Intel Corporation.
- *Intel 82575 Ethernet Controller Thermal Design Considerations*. Intel Corporation

Note: Intel documentation is subject to frequent revision. Verify with your local Intel sales office that you have the latest information before finalizing a design.



2.0 PCI Express Port Connection to the Device

PCI Express (PCIe*) is a dual simplex point-to-point serial differential low-voltage interconnect. The signaling bit rate is 2.5 Gbps per lane per direction. Each port consists of a group of transmitters and receivers located on the same chip. Each lane consists of a transmitter and a receiver pair. A link between the ports of two devices is a collection of lanes. The device supports up to four lanes on the PCIe interface.

Each signal is 8b/10b encoded with an embedded clock.

The PCI Express topology consists of a transmitter (Tx) located on one device connected through a differential pair connected to the receiver (Rx) on a second device. The controller may be located on the motherboard or on an add-in card using a connector specified by PCI Express.

The lane is AC-coupled between its corresponding transmitter and receiver. The AC-coupling capacitor is located on the board close to transmitter side. Each end of the link is terminated on the die into nominal 100 Ω differential DC impedance. Board termination is not required.

For more information on PCI Express, refer to the *PCI Express* Base Specification, Revision 1.1* and *PCI Express* Card Electromechanical Specification, Revision 1.1RD*.

For information about PCIe power management with the 82575, refer to section 3.4 in this document.

2.1 PCI Express Reference Clock

The device uses a 100 MHz differential reference clock, denoted PE_CLK_P and PE_CLK_N. This signal is typically generated on the system board and routed to the PCI Express port. For add-in cards, the clock will be furnished at the PCI Express connector.

The frequency tolerance for the PCI Express reference clock is +/- 300 ppm.

2.2 Other PCI Express Signals

The device also implements other signals required by the PCI Express specification. The Ethernet controller signals power management events to the system using the PE_WAKE# signal, which operates very similarly to the familiar PCI PME# signal. Finally, there is a PE_RST# signal which serves as the familiar reset function for the controller.

2.3 Physical Layer Features

2.3.1 Link Width Configuration

The device supports a maximum link width of x4, x2, or x1 as determined by the EEPROM Lane_Width field in PCIe init configuration.

The max link width is loaded into the Maximum Link Width field of the PCIe capability Register (LCAP[11:6]). The 82575 Ethernet Controller default is x4 link.

During link configuration, the platform and the 82575 Ethernet Controller negotiate on a common link width. The link width must be one of the supported PCIe link widths (1x, 2x, 4x), such that:

- If Maximum Link Width = x4, then the 82575 Ethernet Controller negotiates to either x4, x2 or x1



- If Maximum Link Width = x2, then the 82575 Ethernet Controller negotiates to either x2 or x1
- If Maximum Link Width = x1, then the 82575 Ethernet Controller only negotiates to x1

2.3.2 Polarity Inversion

If polarity inversion is detected the Receiver must invert the received data.

During the training sequence, the Receiver looks at Symbols 6-15 of TS1 and TS2 as the indicator of lane polarity inversion (D+ and D- are swapped). If lane polarity inversion occurs, the TS1 Symbols 6-15 received will be D21.5 as opposed to the expected D10.2. Similarly, if lane polarity inversion occurs, Symbols 6-15 of the TS2 ordered set will be D26.5 as opposed to the expected D5.2. This provides the clear indication of lane polarity inversion.

2.3.3 Lane Reversal

The following lane reversal modes are supported (see Figure below):

- Lane configuration of x4, x2, and x1
- Lane reversal in x4 and in x2
- Degraded mode (downshift) from x4 to x2 to x1 and from x2 to x1, with one restriction - if lane reversal is executed in x4, then downshift is only to x1 and not to x2.

These restrictions require that a x2 interface to the 82575 Ethernet Controller must connect to lanes 0 & 1 on the 82575 Ethernet Controller. The PCI Express Card Electromechanical specification does not allow routing a x2 link to a wider connector. Therefore, the system designer is not allowed to connect a x2 link to lanes 2 and 3 of a PCI Express connector. It is also recommended that, when using x2 mode on a network interface card, the 82575 Ethernet Controller be connected to lanes 0 & 1 of the card.

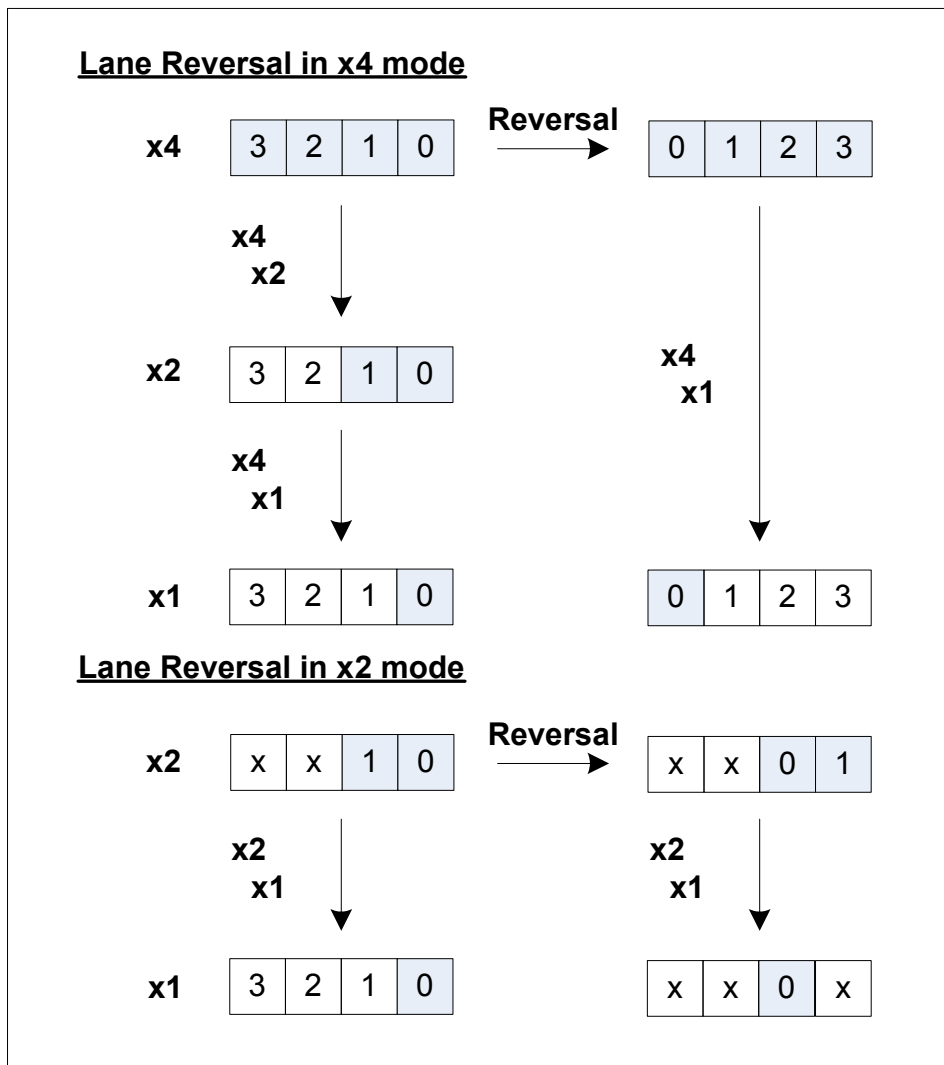


Figure 1. Lane Reversal supported modes

Configuration bits: EEPROM "Lane reversal disable" bit - disables lane reversal altogether

2.4 PCI Express Routing

For information regarding the PCIe signal routing, please refer to the Intel PCIe Design Guide. Contact your Intel representative for information.



This page left intentionally blank.



3.0 Ethernet Component Design Guidelines

These sections provide recommendations for selecting components and connecting special pins.

For 1000 BASE-T designs, the main design elements are the 82575 Gigabit Ethernet Controller, an integrated discrete or magnetics module with RJ-45 connector, an EEPROM, and a clock source.

3.1 General Design Considerations for Ethernet Controllers

Follow good engineering practices with respect to unused inputs by terminating them with pull-up or pull-down resistors, unless the datasheet, design guide or reference schematic indicates otherwise. Do not attach pull-up or pull-down resistors to any balls identified as No Connect. These devices may have special test modes that could be entered unintentionally.

3.1.1 Clock Source

All designs require a 25 MHz clock source. The 82575 Gigabit Ethernet Controller uses the 25 MHz source to generate clocks up to 125 MHz and 1.25 GHz for the PHY circuits, and 1.25 GHz for the SERDES. For optimum results with lowest cost, connect a 25 MHz parallel resonant crystal and appropriate load capacitors at the XTAL1 and XTAL2 leads. The frequency tolerance of the timing device should be 30 ppm or better. Refer to the application note, Intel Fast Ethernet Controllers Timing Device Selection Guide, AP-419, for more information on choosing crystals.

For further information regarding the clock for the 82575, see the sections about frequency control, crystals, and oscillators later in this document.

3.1.2 Magnetics for 1000 BASE-T

Magnetics for the 82575 can be either integrated or discrete.

The magnetics module has a critical effect on overall IEEE and emissions conformance. The device should meet the performance required for a design with reasonable margin to allow for manufacturing variation. Occasionally, components that meet basic specifications may cause the system to fail IEEE testing because of interactions with other components or the printed circuit board itself. Carefully qualifying new magnetics modules prevents this problem.

When using discrete magnetics it is necessary to use Bob Smith termination: Use four 75 Ω resistors for cable-side center taps and unused pins. This method terminates pair-to-pair common mode impedance of the CAT5 cable.

Use an EFT capacitor attached to the termination plane. Suggested values are 1500 pF/2KV or 1000 pF/3KV. A minimum of 50-mil spacing from capacitor to traces and components should be maintained.

3.1.2.1 Magnetics Module Qualification Steps

The steps involved in magnetics module qualification are similar to those for crystal qualification:

1. Verify that the vendor's published specifications in the component datasheet meet or exceed the required IEEE specifications.
2. Independently measure the component's electrical parameters on the test bench, checking samples from multiple lots. Check that the measured behavior is



consistent from sample to sample and that measurements meet the published specifications.

3. Perform physical layer conformance testing and EMC (FCC and EN) testing in real systems. Vary temperature and voltage while performing system level tests.

3.1.2.2 Modules for 1000 BASE-T Ethernet

Magnetics modules for 1000 BASE-T Ethernet are similar to those designed solely for 10/100 Mbps, except that there are four differential signal pairs instead of two. Use the following guidelines to verify specific electrical parameters:

1. Verify that the rated return loss is 19 dB or greater from 2 MHz through 40 MHz for 100/1000 BASE-TX.
2. Verify that the rated return loss is 12 dB or greater at 80 MHz for 100 BASE-TX (the specification requires greater than or equal to 10 dB).
3. Verify that the rated return loss is 10 dB or greater at 100 MHz for 1000 BASE-TX (the specification requires greater than or equal to 8 dB).
4. Verify that the insertion loss is less than 1.0 dB at 100 kHz through 80 MHz for 100 BASE-TX.
5. Verify that the insertion loss is less than 1.4 dB at 100 kHz through 100 MHz for 1000 BASE-T.
6. Verify at least 30 dB of crosstalk isolation between adjacent channels (through 150 MHz).
7. Verify high voltage isolation to 15000 Vrms. (Does not apply to discrete magnetics.)
8. Transmitter OCL should be greater than or equal to 350 μ H with 8 mA DC bias.

3.1.2.3 Third-Party Magnetics Manufacturers

The following magnetics modules have been used successfully in previous designs..

Manufacturer	Part Number
Pulse	H5007
Bel (discrete)	Bel 0344FLA

3.1.2.4 Layout Guidelines for Use with Integrated and Discrete Magnetics

Layout requirements are slightly different when using discrete magnetics.

These include:

- Ground cut for HV installation (not required for integrated magnetics)
- A maximum of two (2) vias
- Turns less than 45°
- Discrete terminators

3.2 Designing with the 82575/EB/ES Gigabit Ethernet Controller

This section provides design guidelines specific to the 82575/EB/ES controller.



3.2.1 LAN Disable for 82575 Ethernet Controller Gigabit Ethernet Controller

The 82575 Ethernet Controller device has three signals that can be used for disabling Ethernet functions from system BIOS. LAN0_DIS_N and LAN1_DIS_N are the separated port disable signals and DEV_OFF_N is the device disable signal. Each signal can be driven from a system output port. Choose outputs from devices that retain their values during reset. For example, ICH7 resumes GPIO outputs (GP24, 25, 27, 28) transition high during reset. It is important not to use these signals to drive LAN0_DIS_N or LAN1_DIS_N because these inputs are latched upon the rising edge of PE_RST_N or an inband reset end. The DEV_OFF_N input is completely asynchronous and does not have this restriction.

Each PHY may be disabled if its LAN function's LAN Disable input indicates that the relevant function should be disabled. Since the PHY is shared between the LAN function and manageability, it may not be desired to power down the PHY in LAN Disable. The PHY_in_LAN_Disable EEPROM bit determines whether the PHY (and MAC) are powered down when the LAN Disable pin is asserted. Default is not to power down.

A LAN port may also be disabled through EEPROM settings. If the LAN_DIS EEPROM bit is set, the PHY enters power down. Note, however, that setting the EEPROM LAN_PCI_DIS bit does not bring the PHY into power down.

Table 1. PCI/LAN Function Index

PCI Function #	LAN Function Select	Function 0	Function 1
Both LAN functions are enabled	0	LAN 0	LAN 1
LAN 0 is disabled	0	Dummy	LAN1
LAN 1 is disabled	0	LAN 0	-
LAN 0 is disabled	1	LAN 1	-
Both LAN functions are enabled	1	LAN 1	LAN 0
LAN 1 is disabled	1	Dummy	LAN 0
Both LAN functions are disabled	Don't Care	All PCI functions are disabled Whole Device is at deep PD	

Table 2. Strapping Options for LAN Disable

Symbol	Ball #	Name and function
LAN1_DIS_N	A15	This pin is a strapping option pin always active. This pin has an internal weak pull-up resistor. In case this pin is not connected or driven hi during init time, LAN 1 is enabled. In case this pin is driven low during init time, LAN 1 function is disabled. This pin is also used for testing and scan.
LAN0_DIS_N	B13	This pin is a strapping option pin always active. This pin has an internal weak pull-up resistor. In case this pin is not connected or driven hi during init time, LAN 0 is enabled. In case this pin is driven low during init time, LAN 0 is disabled. This pin is also used for testing and scan.

Table 3. Control Options for LAN Disable

Function	Default	Control options
LAN 0	1	Strapping Option + EEPROM word 20h bit 13 (full/PCI only disable in case of strap)
LAN 1	1	Strapping Option + EEPROM word 10h bit 13 (full/PCI only disable in case of strap)/ EEPROM Word 10h bit 11 (full disable) / EEPROM word 10h bit 10 (PCI only disable)

3.2.2 Serial EEPROM

The 82575 Ethernet Controller Gigabit Ethernet Controller uses an Serial Peripheral Interface (SPI)* EEPROM. Several words of the EEPROM are accessed automatically by the device after reset to provide pre-boot configuration data before it is accessed by host software. The remainder of the EEPROM space is available to software for storing the MAC address, serial numbers, and additional information. This information is available to the 82575 Ethernet Controller also and is part of the pre-boot configuration data.

The 82575 has a thermal sensor that can send alerts. Trip points for the sensor are set in the EEPROM. For information regarding the use of the sensor and the programming its function in the EEPROM, please refer to the EEPROM Programming Information and Map Application Note.

3.2.2.1 General Regions

The EEPROM is divided into four regions based on the type of access:

- Hardware accessed--this region is accessed by other hardware
- Alert (ASF) accessed--this region is accessed by alert routines
- PT accessed--this region is accessed by the pass-through routines
- Software accessed--this region is accessed by applications

3.2.2.2 EEPROM-less Operation

The 82575 can be operated without an EEPROM, however the following conditions apply:

- Non-manageability mode only



- Legacy Wake On LAN (magic packets) is not supported
- All the initializations normally loaded from the EEPROM will be loaded by the host driver.

For more information, see the *82575 Gigabit Ethernet Controller Software Developer's Manual* and the *82575 EEPROM Information Guide Application Note AP-499*.

3.2.2.3 SPI EEPROMs for 82575 Ethernet Controller Controller

SPI EEPROMs that have been found to work satisfactorily with the 82575 device are listed in [Table 4](#). SPI EEPROMs must be rated for a clock rate of at least 2 MHz.

Table 4. SPI EEPROMs for 82575 Ethernet Controller Controller

Manufacturer	Size	Manufacturer's Part Number
Catalyst	32Kb	25C32S 0113A
Catalyst	8Kb	25C08S
Catalyst	64Kb	25C64S 0139B
STM	256Kb	95256W6 K350V
STM	64Kb	95640W6
STM	32Kb	95320W6
STM	16Kb	95160W6
STM	8Kb	95080W6
Motorola	64Kb	25AA640
Motorola	32kb	25AA320
Motorola	16Kb	25AA160A

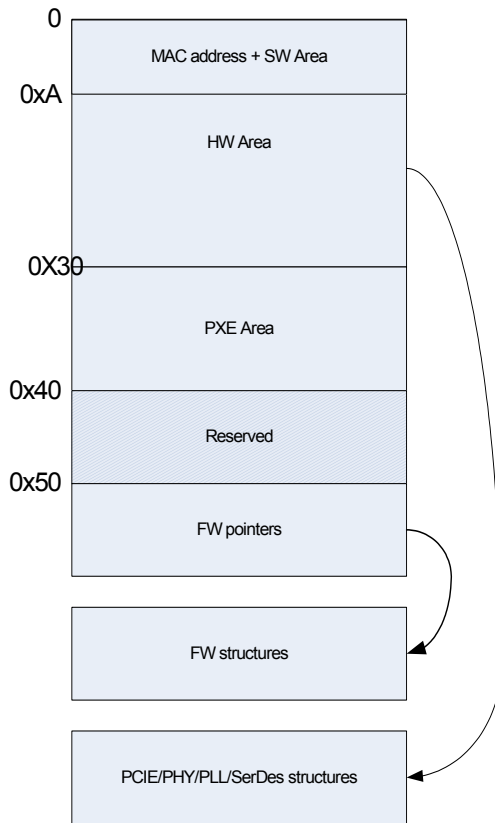
Note: Use a 128 kbit EEPROM for all applications until an appropriate size for each application is determined. Recommended manufacturer and part numbers are Atmel's AT25128N or Microchip's 25LC128.

For more information on the various management options refer to Intel's Application Note 459, *82573/82572/82571/ESB2/82575 LAN Total Cost of Ownership (TCO) System Management Bus Interface*.

3.2.3 EEPROM Map Information

The table below summarizes the EEPROM map for the 82575 Ethernet Controller Gigabit Ethernet Controller. For more about the using an EEPROM, see the 82575 Ethernet Controller EEPROM Map and Programming Information Guide, Application Note (AP-*nnn*).

Table 5. 82575 Ethernet Controller EEPROM Memory Layout



3.2.3.1 EEUPDATE

Intel has an MS-DOS* software utility called EEUPDATE, which can be used to program EEPROM images in development or production line environments. To obtain a copy of this program, contact your Intel representative.

3.2.4 FLASH

The 82575 Ethernet Controller provides two different methods for software access to the Flash.

- Using the legacy Flash transactions the Flash is read from, or written to, whenever the host CPU performs a read or a write operation to a memory location that is within the FLASH address mapping.
- Upon boot via accesses in the space indicated by the Expansion ROM Base Address Register.

All accesses to the Flash require the appropriate command sequence for the device used. Refer to the specific Flash data sheet for more details on reading from or writing to Flash. Accesses to the Flash are based on a direct decode of CPU accesses to a memory window defined in either:

1. 82575 Flash Base Address Register (PCIe Control Register at offset 14h or 18h).



2. A particular address range of the IOADDR register defined by the IO Base Address Register (PCIe Control Register at offset 18h or 20h).
3. The Expansion ROM Base Address Register (PCIe Control Register at offset 30h).

The 82575 controls accesses to the Flash when it decodes a valid access.

Note: Flash read accesses must always be assembled by the 82575 whenever the access is greater than a byte-wide access.

Note: Byte reads or writes to the Flash take on the order of 2us. The device will continue to issue retry accesses during this time.

Note: The 82575 Ethernet Controller supports only byte writes to the Flash.

Another way for SW to access the Flash is directly using the Flash's 4-wire interface through the Flash Access Register (FLA). It can use this for reads, writes, or other Flash operations (accessing the Flash status register, erase...).

To directly access the Flash, software should follow these steps:

1. Write a 1 to the Flash Request bit (FLA.FL_REQ)
2. Read the Flash Grant bit (FLA.FL_GNT) until it becomes 1. It will remain 0 as long as there are other accesses to the Flash.
3. Write or read the Flash using the direct access to the 4-wire interface as defined in the Flash Access Register (FLA). The exact protocol used depends on the Flash placed on the board and can be found in the appropriate datasheet.
4. Write a 0 to the Flash Request bit (FLA.FL_REQ).

3.2.4.1 Flash Write Control

The Flash is write controlled by the FWE bits in the EEPROM/FLASH Control and Data Register (EEC). Note that attempts to write to the Flash device when writes are disabled (FWE10) should not be attempted. Behavior after such an operation is undefined, and may result in component and/or system hangs.

After sending one byte write to the flash, the software can check if it can send the next byte to write (check if the write process in the Flash had finished) by reading the Flash Access Register. If bit (FLA.FL_BUSY) in this register is set, the current write did not finish. If bit (FLA.FL_BUSY) is clear, then the software can continue and write the next byte to the Flash.

3.2.4.2 Flash Erase Control

When software wants to erase the Flash, it should set bit FLA.FL_ER in the Flash Access Register to one (Flash erase and set bits EEC.FWE in the EEPROM/Flash Control Register to zero).

The hardware will get this command and send the erase command to the Flash. The erase process will finish by itself. Software should wait for the end of the erase process before any further access to the flash. This can be checked by using the Flash Write control mechanism described earlier.

The op-code used for erase operation is defined in the FLASHOP register.



Note: Sector erase by SW is not supported. In order to delete a sector, the serial (bit bang) interface should be used.

3.2.4.3 FLASH Device Information

While Intel does not make specific recommendations regarding FLASH devices, the following devices have been used successfully in previous designs:

Manufacturer	Device	Notes
Intel	Blanshard	Please contact your Intel representative for information
Atmel	AT25F1024	
Atmel	AT25F2048	

3.3 SMBus and NC-SI

SMB and NC-SI are optional interfaces for pass-through and/or configuration traffic between the BMC and the 82575 Ethernet Controller. Please refer to the *82575 TCO/ System Management Interface Guide* for more information.

Note: Intel recommends that the SMBus be connected to the ICH or BMC for the EEPROM recovery solution. If the connection is to a BMC, it will be able to send the EEPROM release command.

The 82575 Ethernet Controller NC-SI interface is a connection to an external BMC. It operates in one of two modes:

- NC-SI-SMB mode - In conjunction with an SMB interface, where pass-through traffic passes through NC-SI and configuration traffic passes through SMB
- NC-SI mode - As a single interface with an external BMC, where all traffic (other than header redirection) between the 82575 Ethernet Controller and the BMC flows through the interface.

The Clock-out (if enabled) is provided in all power states (unless the device is disabled).

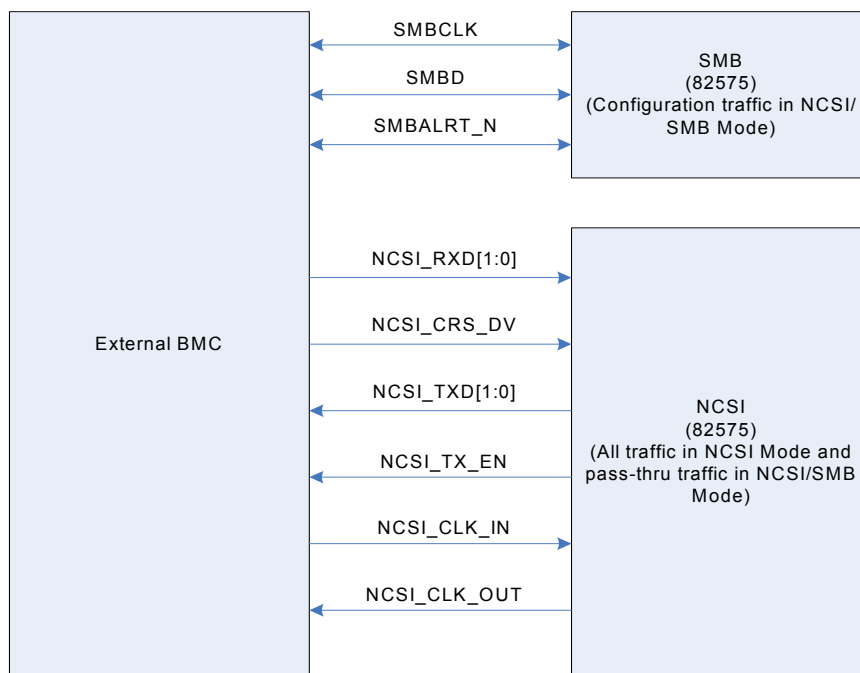


Figure 2. External BMC Connections with NC-SI and SMB

The 82575 Ethernet Controller also supports the DMTF protocol.

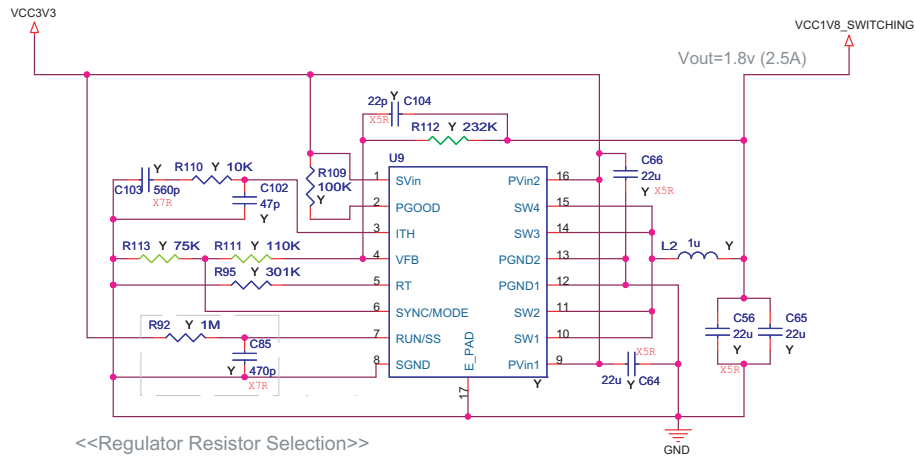
For more information about NC-SI and DMTF, see the *82575 Family System Management Application Note*.

3.4 Power Supplies for the 82575 Ethernet Controller Controllers

The 82575 Ethernet Controller Gigabit Ethernet Controllers require three power rails: 3.3 V, 1.8 V and 1.0 V. (See the 82575 Ethernet Controller Product Datasheet for power requirements.) A central power supply can provide all the required voltage sources, or the power can be derived from the 3.3 V supply and regulated locally using external regulators. If the LAN wake capability will be used, all voltages must remain present during system power down. Local regulation of the LAN voltages from system 3.3 V_{main} and 3.3 V_{aux} voltages is recommended.

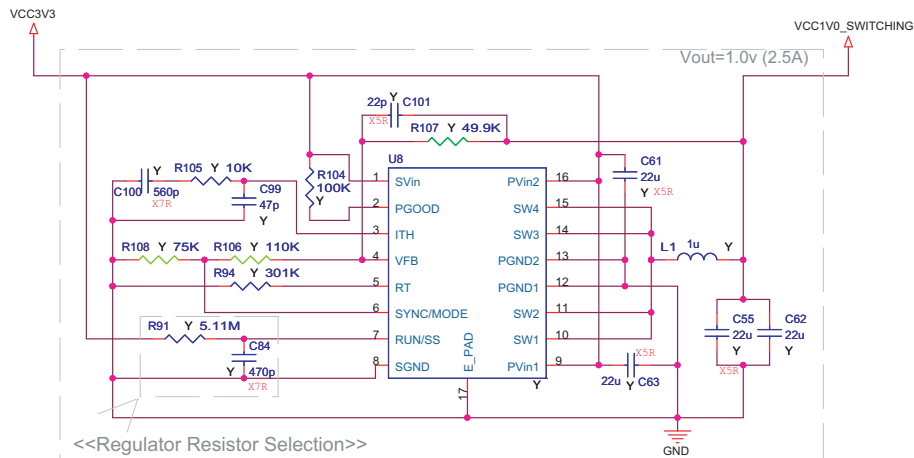
External voltage regulators need to generate the proper voltage, supply current requirements (with adequate margin), and provide the proper power sequencing.

Due to the current demand, a Switching Voltage Regulator (SVR) is highly recommended for the 1.0 V power rail. Figure 2 shows an example of a compact, low-part count, SVR that can be used for both the 1.0 V and 1.8 V power supplies.



<<Regulator Resistor Selection>>
 $V_{out} = 0.8 * (1 + R_{up} / (R_{right} + R_{left}))$
 $R_{up} = ((V_{out} / 0.8) - 1) * (R_{right} + R_{left})$
 $R_{up} = ((1.8 / 0.8) - 1) * (110k + 75k) = 231.25k = 232Kohm$
 3.3v -> 1.8v Step-Down Switching Regulator

VCC1V0_SWITCHING
 Rise after
 VCC1V8_SWITCHING



<<Regulator Resistor Selection>>
 $V_{out} = 0.8 * (1 + R_{up} / (R_{right} + R_{left}))$
 $R_{up} = ((V_{out} / 0.8) - 1) * (R_{right} + R_{left})$
 $R_{up} = ((1.0 / 0.8) - 1) * (110k + 75k) = 46.25k = 49.9Kohm$
 3.3v -> 1.0v Step-Down Switching Regulator

Figure 3. Example Switching Voltage Regulator for 1.0 V and 1.8 V



The 1.8 V rail has a lower current requirement; however, the use of a SVR is still recommended for adequate margin. Using an LVR in this application is acceptable as long as adequate margin exists in the design, and sequencing can be controlled. Figure 3 shows an example of a compact low-part-count LVR that could be used for the 1.8 V supply.




Figure 4. Example of Linear Voltage Regulator for 1.8 V power rail

3.4.1 82575 Ethernet Controller Power Sequencing

Regardless of which type of regulator used, all regulators need to adhere to the sequencing shown in Figure 5 to avoid latch-up and forward-biased internal diodes.




Figure 5. Proper power sequencing for 82575 Ethernet Controller




Figure 6. Power On Flowchart

In addition, the following limitations exist:



- 1.8 V must not exceed 3.3 V.
- 1.0 V must not exceed 3.3 V.
- 1.0 V must not exceed 1.8 V.

The power supplies are all expected to ramp during a short power-up internal (approximately 20ms or better). Do not leave the device in a prolonged state were some, but not all, voltages are applied.

3.4.1.1 Using Regulators With Enable Pins

The use of regulators with enable pins is very helpful in controlling sequencing. Connecting the enable of the 1.8 V regulator to 3.3 V will allow the 1.8 V to ramp as shown in Figure 3. Connecting the enable of the 1.0 V regulator to the 1.8 V output assures that the 1.0 V rail will ramp after the 1.8 V rail. This provides a quick solution to power sequencing. Make sure to check design parameters for inputs with this configuration.

3.4.2 82575 Ethernet Controller Device Power Supply Filtering

Provide several high-frequency bypass capacitors for each power rail (see table below), selecting values in the range of 0.01 μ F to 0.1 μ F. If possible, orient the capacitors close to the device and adjacent to power pads. Decoupling capacitors should connect to the power planes with short, thick (18 mils or more) traces and 14 mil vias. Long and thin traces are more inductive and would reduce the intended effect of decoupling capacitors.

Power Rail	4.7 μ F or 10 μ F	0.1 μ F
3.3 V	1	2
1.8 V	1	4
1.0 V	1	6

Table 6. Minimum Number of Bypass Capacitors per Power Rail.

Furnish approximately 4.7 μ F to 10 μ F of bulk capacitance for all the power rails; placement should be as close to the device power connection as possible.

3.4.3 82575 Ethernet Controller Controller Power Management and Wake Up

The 82575 Ethernet Controller Gigabit Ethernet Controller supports low power operation as defined in the PCI Bus Power Management Specification. There are two defined power states, D0 and D3. The D0 state provides full power operation and is divided into two sub-states: D0u (uninitialized) and D0a (active). The D3 state provides low power operation and is also divided into two sub-states: D3hot and D3cold.

To enter the low power state (D3), the software driver must stop data transmission and reception. Either the operating system or the driver must program the Power Management Control/Status Register (PMCSR) and the Wakeup Control Register (WUC). If wakeup is desired, the appropriate wakeup LAN address filters must also be set. The initial power management settings are specified by EEPROM bits.

When the 82575 Ethernet Controller transitions to either of the D3 low power states, the 1.0 V, 1.8 V, and 3.3 V sources must continue to be supplied to the device. Otherwise, it will not be possible to use a wakeup mechanism. The AUX_PWR signal is a

logic input to the 82575 Ethernet Controller that denotes auxiliary power is available. If AUX_PWR is asserted, the 82575 Ethernet Controller device will advertise that it supports wake up from a D3cold state.

The 82575 Ethernet Controller device supports both Advanced Power Management (APM) wakeup and Advanced Configuration and Power Interface (ACPI) wakeup. APM wakeup has also been known in the past as "Wake on LAN" and as "Magic Packet Wake-up".

Wakeup uses the PE_WAKEn signal to wake the system up. PE_WAKEn is an active low signal typically connected to a GPIO port on the chipset that goes active in response to receiving a "Magic Packet", a network wakeup packet, or link status change indication. PE_WAKEn remains asserted until PME status is cleared in the 82575 Power Management Control/Status Register.

3.4.4 Power Management

3.4.4.1 PCIe Power Management

The 82575 supports D0 and D3 power states defined in the PCI Power Management and PCI Express Specifications. D0 is divided into two sub-states: D0u (D0 Un-initialized), and D0a (D0 active). In addition, the 82575 supports a Dr state that is entered when PE_RST_N is asserted (including the D3cold state).

The diagrams below show the power states and transitions between them.

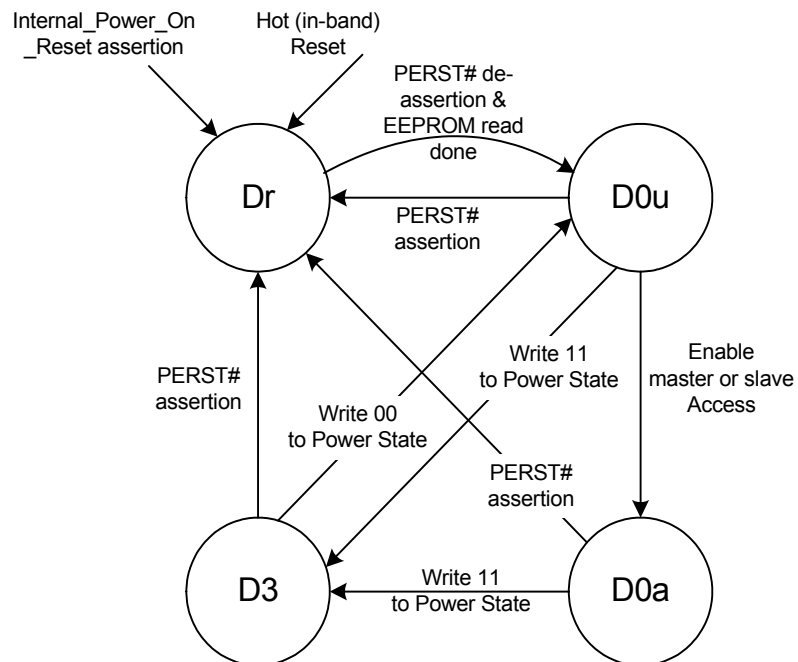


Figure 7. Power Management State Diagram

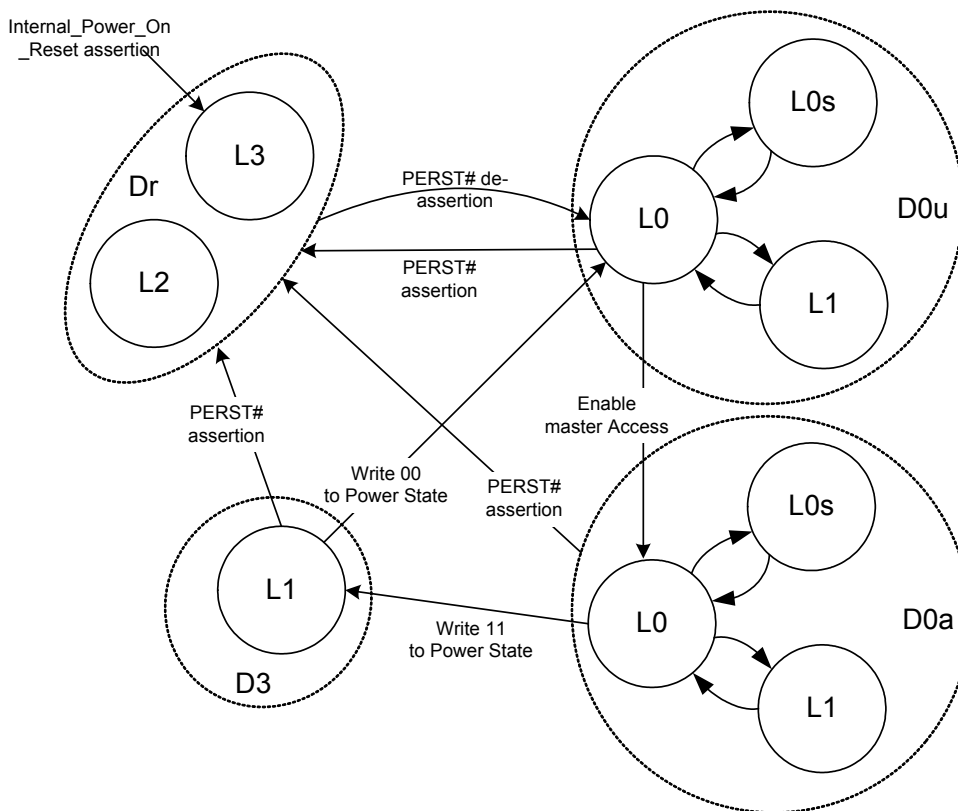


Figure 8. PCIe Power Management Flow/State Diagram

3.4.4.2 82575 Ethernet Controller Power Management

If DisableD3Cold=0, the 82575 uses the AUX_PWR indication that auxiliary power is available to the controller, and therefore advertises D3cold Wake Up support. The amount of power required for the function (which includes the entire network interface card) is advertised in the Power Management Data Register, which is loaded from the EEPROM.

If D3cold is supported, the PME_En and PME_Status bits of the Power Management Control/Status Register (PMCSR), as well as their shadow bits in the Wake Up Control Register (WUC) will be reset only by the power up reset (detection of power rising).

The only effect of setting AUX_PWR to 1 is advertising D3cold Wake Up support and changing the reset function of PME_En and PME_Status. AUX_PWR is a strapping option in the 82575.

3.5 82575 Ethernet Controller Device Test Capability

The 82575 Ethernet Controller Gigabit Ethernet Controller contains a test access port (3.3 V only) conforming to the IEEE 1149.1a-1994 (JTAG) Boundary Scan specification. To use the test access port, connect these balls to pads accessible by your test equipment.

A BSDL (Boundary Scan Definition Language) file describing the 82575 Ethernet Controller device is available for use in your test environment.

The controller also contains an XOR test tree mechanism for simple board tests. Details of XOR tree operation are available from your Intel representative.

Information about how to obtain test models is available from your Intel representative.

3.6 PHY Functionality

This section describes various functions of the PHY.

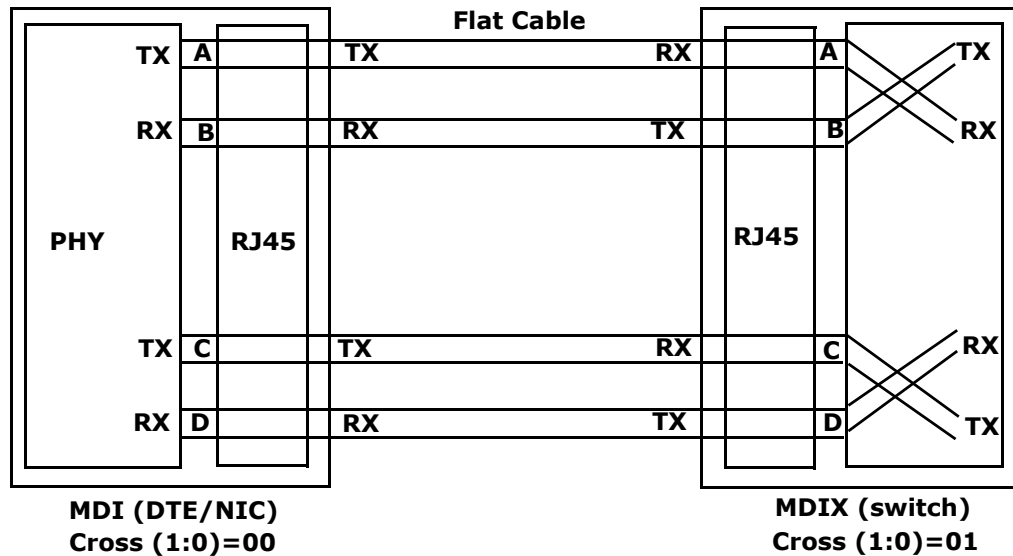
3.6.1 Auto Cross-over for MDI and MDI-X resolution

Twisted pair Ethernet PHY's must be correctly configured for MDI or MDI-X operation to interoperate. The PHY supports the automatic MDI/MDI-X configuration; manual (non-automatic) configuration is still possible by special cable, etc.

For 1000BASE-T links, pair identification is determined automatically in accordance with the standard.

For 10/100 links and during auto-negotiation, pair usage is determined by Bits 12 and 13 in the Port Control Register (PHYREG18).

In addition, the PHY has an Automatic Crossover Detection function. If Bit 18.12 = 1, the PHY automatically detects which application is being used and configures itself accordingly





3.6.2 Smartspeed

SmartSpeed is an enhancement to auto-negotiation that allows the PHY to react to network conditions that are preventing a 1000BASE-T link, such as cable problems. These problems may allow auto-negotiation to complete, but then inhibit completion of the training phase. Normally, if a 1000BASE-T link fails, the PHY returns to the auto-negotiation state with the same speed settings indefinitely.

With SmartSpeed enabled, after a configurable number (1-5, Register 27.8:6) of failed attempts, the PHY automatically downgrades the highest ability it advertises to the next lower speed: from 1000 to 100 to 10. Once a link is established, and if it is later broken, the PHY automatically upgrades the capabilities advertised to the original setting. This allows the PHY to automatically recover once the problem is corrected.

3.6.2.1 Using SmartSpeed

SmartSpeed is enabled by setting PHYREG.16.7 = 1. When SmartSpeed downgrades the PHY advertised capabilities, it sets Bit PHYREG.19.5. When link is established, its speed is indicated in PHYREG.17.15:14. SmartSpeed automatically resets the highest-level auto-negotiation abilities advertised, if the link is established and then lost for more than two seconds.

Number of failed attempts allowed is configured by Register 27.8:6.

Note: When SmartSpeed is enabled, the M/S (Master-Slave) resolution is not given seven attempts to try to resolve M/S status (see IEEE 802.3 clause 40.5.2), this is because SmartSpeed will downgrade the link after five attempts.

Note: The time to link with Smart Speed in most cases is approximately 2.5 seconds, in other cases it could take more than 2.5 seconds, depending on configuration and other factors.

3.6.3 Flow Control

Flow control allows congested nodes to pause traffic. Flow control is essentially a MAC-to-MAC function. MACs indicate their ability to implement flow control during auto-negotiation. This ability is communicated through two bits in the auto-negotiation registers (PHYREG.4.10 and PHYREG.4.11).

Prior to auto-negotiation, the MAC indicates its flow control capabilities via PHYREG.4.10 (Pause) and PHYREG.4.10 (ASM_DIR). After auto-negotiation, the link partner's flow control capabilities are indicated in PHYREG.5.10 and PHYREG.5.11.

There are two forms of flow control that can be established via auto-negotiation: symmetric and asymmetric. Symmetric flow control is for point-to-point links; asymmetric for hub-to-end-node connections. Symmetric flow control allows either node to flow-control the other. Asymmetric flow-control allows a repeater or switch to flow-control a DTE, but not vice versa.

It is the responsibility of the MAC to implement the correct function. The PHY merely allows the two MACs to communicate their abilities to each other.

3.6.4 Low-Power Link Up

Normally, PHY speed negotiation tries to establish a link at the highest possible speed. The PHY supports an additional mode of operation, where the PHY drives to establish a link at a low speed. The link-up process allows a link to come up at the lowest possible speed in cases where power is valued over performance. Different behavior is defined for the D0 state and the other non-D0 states.



The table below summarizes link speed as function of power management state, link speed control, and gigabit speed enabling:

Power Management State	Low Power Link Up (reg 25.1 & 2)	Gigabit disable bits		PHY speed negotiation
		Disable 1000 (reg 25.6)	Disable 1000 in non-D0a (reg 25.3)	
D0a	0	0	X	PHY negotiates to highest speed advertised ("normal operation")
		1		PHY negotiates to highest speed advertised ("normal operation"), excluding 1000
	1	0	X	PHY goes through Low Power Link Up (LPLU) procedure, starting with advertised values
		1		PHY goes through LPLU procedure, starting with advertised values. Does not advertise 1000
Non-D0a	0	0	0	PHY negotiates to highest speed advertised
		0	1	PHY negotiates to highest speed advertised, excluding 1000
		1	X	PHY negotiates to highest speed advertised, excluding 1000
	1	0	0	PHY goes through LPLU procedure, starting at 10
		0	1	PHY goes through LPLU procedure, starting at 10. Does not advertise 1000

3.6.5 Link Energy Detect

The PHY de-asserts the Link Energy Detect Bit (PHYREG 25.4) whenever energy is not detected on the link. This bit provides an indication of a cable becoming plugged or unplugged.

This bit is valid only if auto-negotiation is enabled.

In order to correctly deduce that there is no energy, this bit must read as zero for three consecutive reads each second.

3.6.6 Polarity Correction

The PHY automatically detects and corrects for the condition where the receive signal (MDI_PLUS_0/MDI_MINUS_0) is inverted. Reversed polarity is detected if eight inverted link pulses, or four inverted end-of-frame markers, are received consecutively. If link pulses or data are not received for 96-130 ms, the polarity state is reset to a non-inverted state.

Automatic polarity correction may be disabled by setting Bit PHYREG.27.5



3.6.7 Auto-Negotiation differences between PHY, SerDes and SGMII

SGMII protocol includes an auto-negotiation process in order to establish the MAC - PHY connection. This auto-negotiation process is not dependent on the SRDS0/1_SIG_DET signal, as this signal indicates the status of the PHY signal detection (usually used in Optical PHY).

The following shows the outcome of this auto-negotiation process:

- Link status
- Speed
- Duplex.

This information is used by the hardware to configure the MAC, when operating in SGMII mode.

For SerDes and SGMII modes, bits FD and LU of the Device Status register (STATUS), and bits in the PCS_LSTS register provide status information regarding the negotiated link.

- Auto-Negotiation may be initiated by the following:
 - LRST transition from 1 to 0
 - PCS_LCMD.AN_ENABLE transition from 0 to 1
 - Receipt of /C/ ordered set during normal operation
 - Receipt of different value of the /C/ ordered set during the negotiation process
 - Transition from loss of synchronization to synchronized state (if AN_ENABLE is set).
 - PCS_LCMD.AN_RESTART transition from 0 to 1

Resolution of the negotiated link determines device operation with respect to speed and duplex settings. These negotiated capabilities override advertised and software-controlled device configuration.

When working in SGMII mode, there is no need for setting of the PCAS_ANADV register, as the MAC advertisement word is fixed. The result of the SGMII level auto-negotiation can be read from the PCS_LPAB register.

3.6.8 Copper PHY Link Configuration

When operating with the internal PHY, link configuration is generally determined by PHY Auto-Negotiation. The driver must intervene in cases where a successful link is not negotiated or the user desires to manually configure the link.

PHY Auto-Negotiation (Speed, Duplex, Flow-Control) when using a copper PHY, the PHY performs the Auto-Negotiation function. Auto-Negotiation provides a method for two link partners to exchange information in a systematic manner in order to establish a link configuration providing the highest common level of functionality supported by both partners. Once configured, the link partners exchange configuration information to resolve link settings such as:

- Speed: 10/100/1000 Mb/s
- Duplex: Full- or Half-
- Flow Control Operation

PHY specific information required for establishing the link is also exchanged.

Note: If flow control is enabled, the settings for the desired flow control behavior must be set by software in the PHY registers and Auto-Negotiation restarted. After Auto-Negotiation



is complete, the driver must read the PHY registers to determine the resolved flow control behavior of the link and reflect these in the MAC register settings (CTRL.TFCE and CTRL.RFCE).

Note: Once PHY Auto-negotiation is complete, the PHY will assert a link indication (LINK) to the MAC. Software must have set the "Set Link Up" bit in the Device Control Register (CTRL.SLU) before the MAC recognizes the LINK indication from the PHY and can consider the link to be up.

3.7 Copper/Fiber Switch

The 82575 Ethernet Controller provides significant amount of flexibility in pairing a LAN device with a particular type of media (copper or fiber-optic) as well as the specific transceiver/interface used to communicate with the media. Each MAC, representing a distinct LAN device, can be coupled with an internal copper PHY (the default) or SERDES interface independently. The link configuration specified for each LAN device may be specified in the LINK_MODE field of the Extended Device Control Register (CTRL_EXT) and initialized from the EEPROM Initialization Control Word 3 associated with each LAN device.

In some applications, the software may need to be aware of the presence of a link on the connection not currently active. In order to supply such an indication, any of the the 82575 Ethernet Controller ports may set the AUTOSENSE_EN bit in the CONNSW register (address 0x00034) in order to enable sensing of the non-active connection activity. When in SerDes detect mode, the software should define which indication is used to detect the energy change in SerDes/SGMII mode. It can be either the external signal detect pin or the internal signal detect. This is done using the CONNSW.ENRGSRC bit.

The software can then enable the OMED interrupt in ICR in order to get an indication of any detection of energy in the non active connection.

The following procedure should be followed in order to enable the auto-sense mode:

SerDes-Detect Mode (PHY is active)

- Set CONNSW.ENRGSRC to determine the sources for the signal detect indication (1- external SIG_DET, 0- internal SerDes electrical idle). The default of this bit is set by EEPROM.
- Set CONNSW.AUTOSENSE_EN.
- When signal is detected on the SerDes link, the 82575 Ethernet Controller will set the interrupt bit OMED in ICR and, if enabled, issue an interrupt. The CONNSW.AUTOSENSE_EN will be cleared unless CONNSW.ASCLR_DIS is set. In such a case, the host driver is responsible for the clearing of the AUTOSENSE_EN bit.

PHY-Detect Mode

- Set CONNSW.AUTOSENSE_CONF = 1.
- Reset the PHY by assertion and de-assertion of CTRL.PHY_RST.
- Wait until EEMNGCTL.CFG_DONE is set.
- Enter the PHY to Link-Disconnect mode by setting why-reg25.5 via MDIC register.
- Set CONNSW.AUTOSENSE_EN = 1 and clear CONNSW.AUTOSENSE_CONF.
- When signal is detected on the PHY link, the hardware will set the interrupt bit OMED in ICR and, if enabled, issue an interrupt.



- The 82575 will put the PHY in power down unless CONNSW.ASCLR_DIS is set. In such a case the host driver is responsible for the clearing of the AUTOSENSE_EN bit

According to the result of the interrupt, the software can then decide to switch to the other core.

The following procedures need to be followed to actually switch between the two modes:

Internal PHY-to-SerDes Transition

- Disable Receiver by clearing RCTL.RXEN
- Disable Transmitter by clearing TCTL.EN
- Verify the device has stopped processing outstanding cycles and is idle.
- Modify LINK mode to SER/DES or SGMII by setting CTRL_EXT.LINK_MODE to 10b or 11b respectively.
- Enable/Disable flow control values within the MAC.
- Set up Tx and Rx queues and enable Tx and Rx processes.

SerDes-to-Internal PHY Transition

- Disable Receiver by clearing RCTL.RXEN
- Disable Transmitter by clearing TCTL.EN
- Verify the 82575 has stopped processing outstanding cycles and is idle.
- Modify LINK mode to PHY mode by setting CTRL_EXT.LINK_MODE to 00b.
- Set Link Up indication by setting CTRL.SLU
- Reset the PHY by setting CTRL.PHY_RST, waiting 10 ms and clearing CTRL.PHY_RST.
- Set up PHY with desired auto-negotiation parameters
- Set up Tx and Rx queues and enable Tx and Rx processes.

The device's link mode is controlled by the Extended Device Control register --

CTRL_EXT (0x00018) bits 23:22. The default value for the LINK_MODE setting is directly mapped from the EEPROM's initialization Control Word 3 (bits 1:0). Software can modify the LINK_MODE indication by writing the corresponding value into this register.

Note: Before dynamically cycling a mode, ensure via the software device driver that the current mode of operation is not in the process of transmitting or receiving data. This is achieved by disabling the transmitter and receiver, waiting until the device is in an idle state, and then beginning the process for changing the link mode.

Note: The mode switch in this method, is only valid until the next hardware reset of the chip. After hardware reset the link mode is restored to the default set by the EEPROM. To get a permanent change of the link mode, the default in the EEPROM should be changed.

3.8 Device Disable

For a LOM design, it may be desirable for the system to provide BIOS-setup capability for selectively enabling or disabling LOM devices. This may allow the end-user more control over system resource-management, avoid conflicts with add-in NIC solutions, etc. the 82575 Ethernet Controller provides support for selectively enabling or disabling it.



Note that if the device is configured to provide a 50MHz NC-SI clock (via the NC-SI Output Clock EEPROM bit), then the NC-SI clock must be provided in Device Disable mode as well the device should not be disabled.

Device Disable is initiated by asserting the asynchronous DEV_OFF_N pin. The DEV_OFF_N pin has an internal pull-up resistor, so that it can be left not connected to enable device operation.

The EEPROM's "Device Disable Power Down En" bit enables device disable mode (hardware default is that the mode is disabled).

While in device disable mode, the PCI Express link is in L3 state. The PHY is in power down mode. Output buffers are tri-stated.

Assertion or deassertion of PCI Express PE_RST_N does not have any effect while the device is in device disable mode (that is, the device stays in the respective mode as long as DEV_OFF_N is asserted). However, the device may momentarily exit the device disable mode from the time PCI Express PE_RST_N is de-asserted again and until the EEPROM is read.

During power-up, the DEV_OFF_N pin is ignored until the EEPROM is read. From that point, the device may enter Device Disable if DEV_OFF_N is asserted.

Note: The DEV_OFF_N pin should maintain its state during system reset and system sleep states. It should also insure the proper default value on system power-up. For example, one could use a GPIO pin that defaults to '1' (enable) and is on system suspend power (i.e., it maintains state in S0-S5 ACPI states).

3.8.1 BIOS handling of Device Disable

Assume that in the following power up sequence the DEV_OFF_N signal is driven high (or it is already disabled)

1. The PCIe is established following the GIO_PWR_GOOD
2. BIOS recognizes that the whole Device should be disabled
3. The BIOS drive the DEV_OFF_N signal to the low level.
4. As a result, the device samples the DEV_OFF_N signals and enters either the device disable mode.
5. The BIOS could put the Link in the Electrical IDLE state (at the other end of the PCI Express link) by clearing the LINK Disable bit in the Link Control Register.
6. BIOS may start with the Device enumeration procedure (the whole Device functions are invisible)
7. Proceed with Nominal operation
8. Re-enable could be done by driving hi the DEV_OFF_N signal, followed later by bus enumeration.

3.9 Software-Definable Pins (SDPs)

The 82575 has four software-defined pins (SDP) per port that can be used for miscellaneous hardware or software-controllable purposes. These pins and their function are bound to a specific LAN device (eight SDPs may not be associated with a single LAN device, for example). These pins can each be individually configured to act as either input or output pins. The default direction of each of the four pins is configurable via EEPROM, as well as the default value of any pins configured as outputs.



Note: To avoid signal contention, all four pins are set as input pins until after EEPROM configuration has been loaded.

In addition to all four pins being individually configurable as inputs or outputs, they may be configured for use as general-purpose interrupt (GPI) inputs. To act as GPI pins, the desired pins must be configured as inputs. A separate GPI interrupt-detection enable is then used to enable rising-edge detection of the input pin (rising-edge detection occurs by comparing values sampled at 62.5 MHz, as opposed to an edge-detection circuit). When detected, a corresponding GPI interrupt is indicated in the Interrupt Cause register.

The use, direction, and values of SDPs are controlled and accessed using fields in the Device Control Register (CTRL) and Extended Device Control Register (CTRL_EXT).



4.0 Frequency Control Device Design Considerations

This section provides information regarding frequency control devices, including crystals and oscillators, for use with all Intel Ethernet controllers. Several suitable frequency control devices are available; none of which present any unusual challenges in selection. The concepts documented herein are applicable to other data communication circuits, including Platform LAN Connect devices (PHYs).

The Intel Ethernet controllers contain amplifiers, which when used with the specific external components, form the basis for feedback oscillators. These oscillator circuits, which are both economical and reliable, are described in more detail in “[Crystal Selection Parameters](#)”.

The Intel Ethernet controllers also have bus clock input functionality, however a discussion of this feature is beyond the scope of this document, and will not be addressed.

The chosen frequency control device vendor should be consulted early in the design cycle. Crystal and oscillator manufacturers familiar with networking equipment clock requirements may provide assistance in selecting an optimum, low-cost solution.

4.1 Frequency Control Component Types

Several types of third-party frequency reference components are currently marketed. A discussion of each follows, listed in preferred order.

4.1.1 Quartz Crystal

Quartz crystals are generally considered to be the mainstay of frequency control components due to their low cost and ease of implementation. They are available from numerous vendors in many package types and with various specification options.

4.1.2 Fixed Crystal Oscillator

A packaged fixed crystal oscillator comprises an inverter, a quartz crystal, and passive components conveniently packaged together. The device renders a strong, consistent square wave output. Oscillators used with microprocessors are supplied in many configurations and tolerances.

Crystal oscillators should be restricted to use in special situations, such as shared clocking among devices or multiple controllers. As clock routing can be difficult to accomplish, it is preferable to provide a separate crystal for each device.

For Intel Ethernet controllers, it is acceptable to overdrive the internal inverter by connecting a 25 MHz external oscillator to the XTAL1 lead, leaving the XTAL2 lead unconnected. The oscillator should be specified to drive CMOS logic levels, and the clock trace to the device should be as short as possible. Device specifications typically call for a 40% (minimum) to 60% (maximum) duty cycle and a ± 50 ppm frequency tolerance.

Note: Please contact your Intel Customer Representative to obtain the most current device documentation prior to implementing this solution.



4.1.3 Programmable Crystal Oscillators

A programmable oscillator can be configured to operate at many frequencies. The device contains a crystal frequency reference and a phase lock loop (PLL) clock generator. The frequency multipliers and divisors are controlled by programmable fuses.

A programmable oscillator's accuracy depends heavily on the Ethernet device's differential transmit lines. The Physical Layer (PHY) uses the clock input from the device to drive a differential Manchester (for 10 Mbps operation), an MLT-3 (for 100 Mbps operation) or a PAM-5 (for 1000 Mbps operation) encoded analog signal across the twisted pair cable. These signals are referred to as self-clocking, which means the clock must be recovered at the receiving link partner. Clock recovery is performed with another PLL that locks onto the signal at the other end.

PLLs are prone to exhibit frequency jitter. The transmitted signal can also have considerable jitter even with the programmable oscillator working within its specified frequency tolerance. PLLs must be designed carefully to lock onto signals over a reasonable frequency range. If the transmitted signal has high jitter and the receiver's PLL loses its lock, then bit errors or link loss can occur.

PHY devices are deployed for many different communication applications. Some PHYs contain PLLs with marginal lock range and cannot tolerate the jitter inherent in data transmission clocked with a programmable oscillator. The American National Standards Institute (ANSI) X3.263-1995 standard test method for transmit jitter is not stringent enough to predict PLL-to-PLL lock failures, therefore, the use of programmable oscillators is generally not recommended.

4.1.4 Ceramic Resonator

Similar to a quartz crystal, a ceramic resonator is a piezoelectric device. A ceramic resonator typically carries a frequency tolerance of $\pm 0.5\%$, – inadequate for use with Intel Ethernet controllers, and therefore, should not be utilized.



5.0 Crystal Selection Parameters

All crystals used with Intel Ethernet controllers are described as “AT-cut,” which refers to the angle at which the unit is sliced with respect to the long axis of the quartz stone. Table 4 lists crystals which have been used successfully in other designs (however, no particular product is recommended):

Table 7. Crystal Manufacturers and Part Numbers

Manufacturer	Part No.
RALTRON	AS-25.000-20-F-SMD-T
CITIZEN AMERICA CORP	HCM4925.000MBBKTR
NDK AMERICA INC	41CD25.0S11005020
TXC CORPORATION - USA	6C25000131

For information about crystal selection parameters, see the chart in the *82575 Ethernet Controller Advance Information Datasheet*. Contact your Intel representative to obtain this document.

5.1 Vibrational Mode

Crystals in the above-referenced frequency range are available in both fundamental and third overtone. Unless there is a special need for third overtone, use fundamental mode crystals.

At any given operating frequency, third overtone crystals are thicker and more rugged than fundamental mode crystals. Third overtone crystals are more suitable for use in military or harsh industrial environments. Third overtone crystals require a trap circuit (extra capacitor and inductor) in the load circuitry to suppress fundamental mode oscillation as the circuit powers up. Selecting values for these components is beyond the scope of this document.

5.2 Nominal Frequency

Intel Ethernet controllers use a crystal frequency of 25.000 MHz. The 25 MHz input is used to generate a 125 MHz transmit clock for 100BASE-TX and 1000BASE-TX operation – 10 MHz and 20 MHz transmit clocks, for 10BASE-T operation.

5.3 Frequency Tolerance

The frequency tolerance for an Ethernet Platform LAN Connect is dictated by the IEEE 802.3 specification as ± 50 parts per million (ppm). This measurement is referenced to a standard temperature of 25° C. Intel recommends a frequency tolerance of ± 30 ppm.

5.4 Temperature Stability and Environmental Requirements

Temperature stability is a standard measure of how the oscillation frequency varies over the full operational temperature range (and beyond). Several optional temperature ranges are currently available, including -40° C to +85° C for industrial environments. Some vendors separate operating temperatures from temperature stability. Manufacturers may also list temperature stability as 50 ppm in their data sheets.

Note: Crystals also carry other specifications for storage temperature, shock resistance, and reflow solder conditions. Crystal vendors should be consulted early in the design cycle to discuss the application and its environmental requirements.

5.5 Calibration Mode

The terms “series-resonant” and “parallel-resonant” are often used to describe crystal oscillator circuits. Specifying parallel mode is critical to determining how the crystal frequency is calibrated at the factory.

A crystal specified and tested as series resonant oscillates without problem in a parallel-resonant circuit, but the frequency is higher than nominal by several hundred parts per million. The purpose of adding load capacitors to a crystal oscillator circuit is to establish resonance at a frequency higher than the crystal’s inherent series resonant frequency.

Figure 3 illustrates a simplified schematic of the internal oscillator circuit. Pin X1 and X2 refers to XTAL1 and XTAL2 in the Ethernet device, respectively. The crystal and the capacitors form a feedback element for the internal inverting amplifier. This combination is called parallel-resonant, because it has positive reactance at the selected frequency. In other words, the crystal behaves like an inductor in a parallel LC circuit. Oscillators with piezoelectric feedback elements are also known as “Pierce” oscillators.




Figure 9. Internal Oscillator Circuit

5.6 Load Capacitance

The formula for crystal load capacitance is as follows:

$$C_L = \frac{(C1 \cdot C2)}{(C1 + C2)} + C_{\text{stray}}$$

where $C1 = C2 = 27 \text{ pF}$

and C_{stray} = allowance for additional capacitance in pads, traces and the chip carrier within the Ethernet device package



An allowance of 3 pF to 7 pF accounts for lumped stray capacitance. The calculated load capacitance is 16 pF with an estimated stray capacitance of about 5 pF.

Individual stray capacitance components can be estimated and added. For example, surface mount pads for the load capacitors add approximately 2.5 pF in parallel to each capacitor. This technique is especially useful if Y1, C1 and C2 must be placed farther than approximately one-half (0.5) inch from the device. It is worth noting that thin circuit boards generally have higher stray capacitance than thick circuit boards. Consult the PCIe Design Guide for more information.

The oscillator frequency should be measured with a precision frequency counter where possible. The load specification or values of C1 and C2 should be fine tuned for the design. As the actual capacitance load increases, the oscillator frequency decreases.

Note: C1 and C2 may vary by as much as 5% (approximately 1 pF) from their nominal values.

5.7 Shunt Capacitance

The shunt capacitance parameter is relatively unimportant compared to load capacitance. Shunt capacitance represents the effect of the crystal's mechanical holder and contacts. The shunt capacitance should equal a maximum of 7 pF.

5.8 Equivalent Series Resistance

Equivalent Series Resistance (ESR) is the real component of the crystal's impedance at the calibration frequency, which the inverting amplifier's loop gain must overcome. ESR varies inversely with frequency for a given crystal family. The lower the ESR, the faster the crystal starts up. Use crystals with an ESR value of 50 Ω or better.

5.9 Drive Level

Drive level refers to power dissipation in use. The allowable drive level for a Surface Mounted Technology (SMT) crystal is less than its through-hole counterpart, because surface mount crystals are typically made from narrow, rectangular AT strips, rather than circular AT quartz blanks.

Some crystal data sheets list crystals with a maximum drive level of 1 mW. However, Intel Ethernet controllers drive crystals to a level less than the suggested 0.5 mW value. This parameter does not have much value for on-chip oscillator use.

5.10 Aging

Aging is a permanent change in frequency (and resistance) occurring over time. This parameter is most important in its first year because new crystals age faster than old crystals. Use crystals with a maximum of ± 5 ppm per year aging.

5.11 Reference Crystal

The normal tolerances of the discrete crystal components can contribute to small frequency offsets with respect to the target center frequency. To minimize the risk of tolerance-caused frequency offsets causing a small percentage of production line units to be outside of the acceptable frequency range, it is important to account for those shifts while empirically determining the proper values for the discrete loading capacitors, C1 and C2.



Even with a perfect support circuit, most crystals will oscillate slightly higher or slightly lower than the exact center of the target frequency. Therefore, frequency measurements (which determine the correct value for C1 and C2) should be performed with an ideal reference crystal. When the capacitive load is exactly equal to the crystal's load rating, an ideal reference crystal will be perfectly centered at the desired target frequency.

5.11.1 Reference Crystal Selection

There are several methods available for choosing the appropriate reference crystal:

- If a Saunders and Associates (S&A) crystal network analyzer is available, then discrete crystal components can be tested until one is found with zero or nearly zero ppm deviation (with the appropriate capacitive load). A crystal with zero or near zero ppm deviation will be a good reference crystal to use in subsequent frequency tests to determine the best values for C1 and C2.
- If a crystal analyzer is not available, then the selection of a reference crystal can be done by measuring a statistically valid sample population of crystals, which has units from multiple lots and approved vendors. The crystal, which has an oscillation frequency closest to the center of the distribution, should be the reference crystal used during testing to determine the best values for C1 and C2.
- It may also be possible to ask the approved crystal vendors or manufacturers to provide a reference crystal with zero or nearly zero deviation from the specified frequency when it has the specified CLoad capacitance.

When choosing a crystal, customers must keep in mind that to comply with IEEE specifications for 10/100 and 10/100/1000Base-T Ethernet LAN, the transmitter reference frequency must be precise within ± 50 ppm. Intel® recommends customers to use a transmitter reference frequency that is accurate to within ± 30 ppm to account for variations in crystal accuracy due to crystal manufacturing tolerance..

5.11.2 Circuit Board

Since the dielectric layers of the circuit board are allowed some reasonable variation in thickness, the stray capacitance from the printed board (to the crystal circuit) will also vary. If the thickness tolerance for the outer layers of dielectric are controlled within ± 17 percent of nominal, then the circuit board should not cause more than ± 2 pF variation to the stray capacitance at the crystal. When tuning crystal frequency, it is recommended that at least three circuit boards are tested for frequency. These boards should be from different production lots of bare circuit boards.

Alternatively, a larger sample population of circuit boards can be used. A larger population will increase the probability of obtaining the full range of possible variations in dielectric thickness and the full range of variation in stray capacitance.

Next, the exact same crystal and discrete load capacitors (C1 and C2) must be soldered onto each board, and the LAN reference frequency should be measured on each circuit board.

The circuit board, which has a LAN reference frequency closest to the center of the frequency distribution, should be used while performing the frequency measurements to select the appropriate value for C1 and C2.

5.11.3 Temperature Changes

Temperature changes can cause the crystal frequency to shift. Therefore, frequency measurements should be done in the final system chassis across the system's rated operating temperature range.



This page is intentionally left blank.



6.0 Oscillator Support

The 82575 clock input circuit is optimized for use with an external crystal. However, an oscillator can also be used in place of the crystal with the proper design considerations:

- The clock oscillator has an internal voltage regulator of 1.2 V to isolate it from the external noise of other circuits to minimize jitter. If an external clock is used, this imposes a maximum input clock amplitude of 1.2 V.
- The input capacitance introduced by the 82575 (approximately 20 pF) is greater than the capacitance specified by a typical oscillator (approximately 15 pF).
- The input clock jitter from the oscillator can impact the 82575 clock and its performance.

Table 8. 82575 Clock Oscillator Specifications

Symbol	Parameter	Specifications			Units
		Min	Typical	Max	
f ₀	Frequency (@25°C)	-	25	-	MHz
V _{p-p}	External Oscillator Supply Swing	3.0	3.3	3.6	V
V _{scp-p}	XTAL1 Swing	1.1	1.2	1.3	V
Df/f ₀	Frequency Tolerance (@ -20 to +70 °C)	-	±50		ppm
T _{opr}	Operating Temperature	-	-20 to +70		°C
Δf/f ₀	Aging	-	±5		ppm/year
C _{coupling}	Coupling Capacitor	8	10	12	pF

Note: The power consumption of additional circuitry equals about 1.5 mW.

The following table lists oscillators that have been used successfully in past designs (please note that no particular product is recommended):

Table 9. Oscillator Manufacturers and Part Numbers

Manufacturer	Part No.
RALTRON	CO4305-25.000-TR
CITIZEN AMERICA CORP	CSX750FBB25.000MTR

6.1 Oscillator Solution

This solution involves capacitor C1, which forms a capacitor divider with C_{stray} of about 20 pF. This attenuates the input clock amplitude and adjusts the clock oscillator load capacitance.

$$V_{in} = VDD * (C1 / (C1 + C_{stray}))$$

$$V_{in} = 3.3 * (C1 / (C1 + C_{stray}))$$

This enables load clock oscillators of 15 pF to be used. If the value of C_{stray} is unknown, C1 should be adjusted by tuning the input clock amplitude to approximately 1 V_{ptp}. If C_{stray} equals 20 pF, then C1 is 10 pF ±10%.

A low capacitance, high impedance probe ($C < 1 \text{ pF}$, $R > 500 \text{ K}\Omega$) should be used for testing. Probing the parameters can affect the measurement of the clock amplitude and cause errors in the adjustment. A test should also be done after the probe has been removed for circuit operation.

If jitter performance is poor, a lower jitter clock oscillator can be implemented.

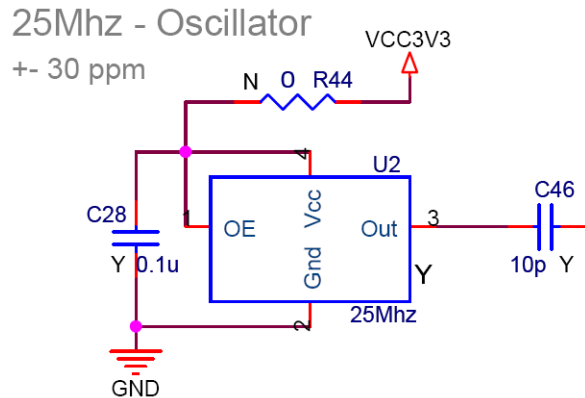


Figure 10. Reference Oscillator Circuit

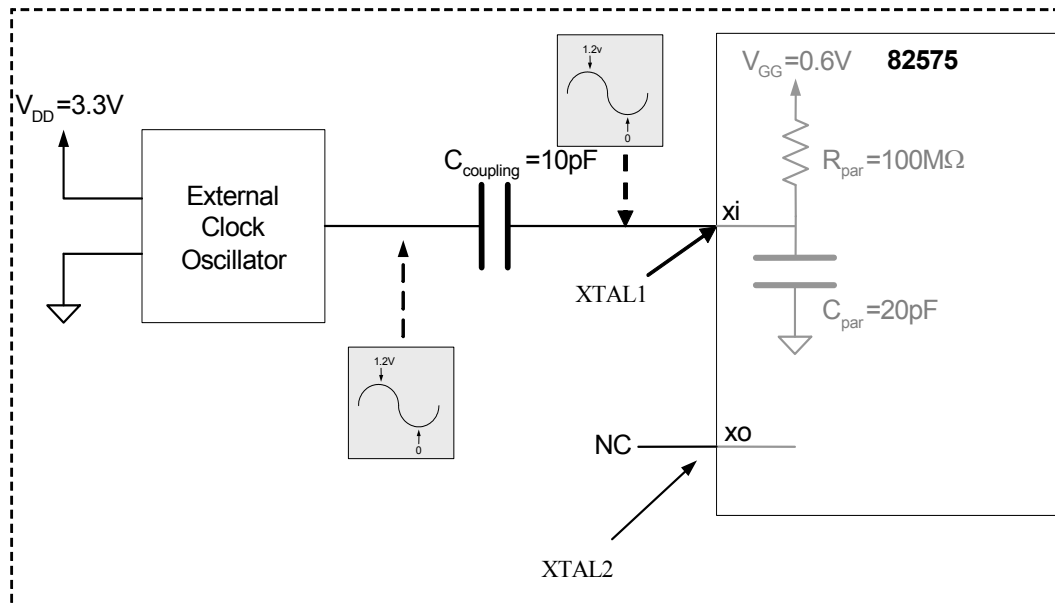


Figure 11. External Clock Oscillator Connectivity



7.0 Ethernet Component Layout Guidelines

These sections provide recommendations for performing printed circuit board layouts. Good layout practices are essential to meet IEEE PHY conformance specifications and EMI regulatory requirements.

7.1 Layout Considerations for 82575 Ethernet Controllers

Critical signal traces should be kept as short as possible to decrease the likelihood of being affected by high frequency noise from other signals, including noise carried on power and ground planes. Keeping the traces as short as possible can also reduce capacitive loading.

Since the transmission line medium extends onto the printed circuit board, special attention must be paid to layout and routing of the differential signal pairs.

Designing for 1000 BASE-T Gigabit operation is very similar to designing for 10 and 100 Mbps. For the 82575 Gigabit Ethernet controller, system level tests should be performed at all three speeds.


7.1.1 Guidelines for Component Placement

Component placement can affect signal quality, emissions, and component operating temperature. This section provides guidelines for component placement.

Careful component placement can:

- Decrease potential problems directly related to electromagnetic interference (EMI), which could cause failure to meet applicable government test specifications.
- Simplify the task of routing traces. To some extent, component orientation will affect the complexity of trace routing. The overall objective is to minimize turns and crossovers between traces.

Minimizing the amount of space needed for the Ethernet LAN interface is important because other interfaces will compete for physical space on a motherboard near the connector. The Ethernet LAN circuits need to be as close as possible to the connector.



NOTE: This figure represents a 10/100 diagram. Use the same design considerations for the two differential pairs not shown for gigabit implementations.

Figure 12. General Placement Distances for 1000 BASE-T Designs

Figure 5 shows some basic placement distance guidelines. The figure shows two differential pairs, but can be generalized for a Gigabit system with four analog pairs. The ideal placement for the Ethernet silicon would be approximately one inch behind the magnetics module.

While it is generally a good idea to minimize lengths and distances, this figure also illustrates the need to keep the LAN silicon away from the edge of the board and the magnetics module for best EMI performance.

The following figures illustrate a reference layout for discrete and integrated magnetics.




Figure 13. Layout for Integrated Magnetics

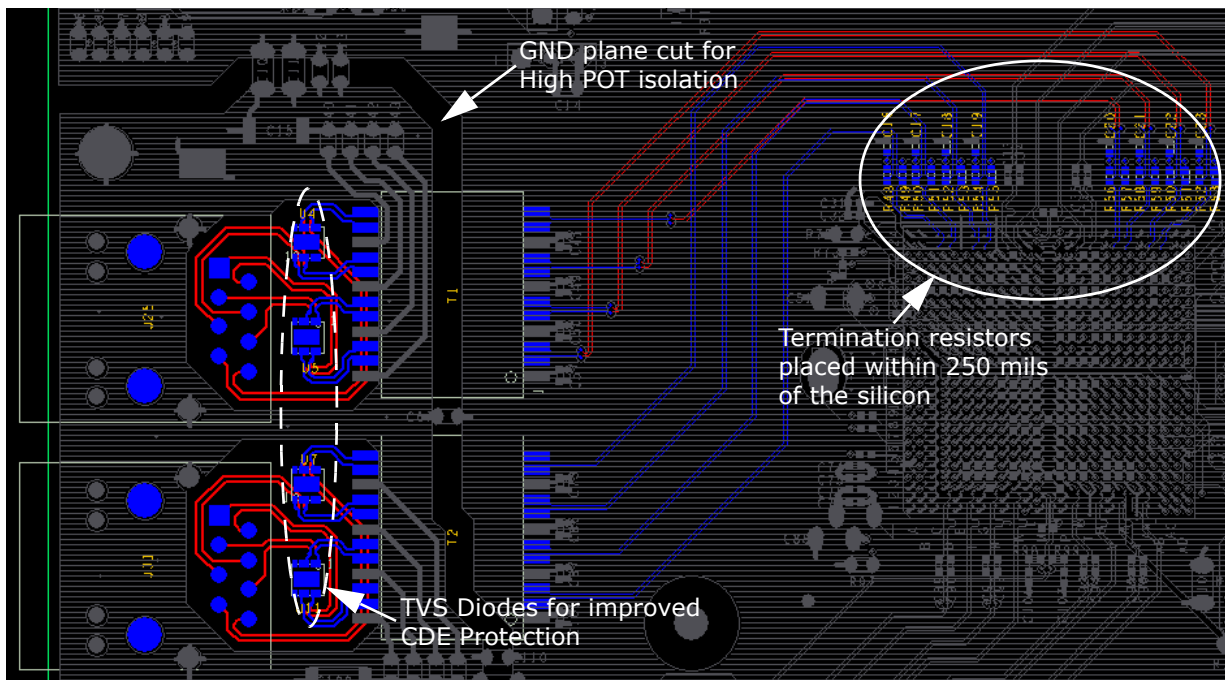


Figure 14. Layout for Discrete Magnetics



7.1.2 Crystals and Oscillators

Clock sources should not be placed near I/O ports or board edges. Radiation from these devices may be coupled into the I/O ports and radiate beyond the system chassis. Crystals should also be kept away from the Ethernet magnetics module to prevent interference.

7.1.2.1 Crystal layout considerations

Note: Failure to follow these guidelines could result in the 25 MHz clock failing to start.

When designing the layout for the crystal circuit, the following rules must be used:

- Place load capacitors as close as possible (within design-for-manufacturability rules) to the crystal solder pads. They should be no more than 90 mils away from crystal pads.
- The two load capacitors, crystal component, the Ethernet controller device, and the crystal circuit traces must all be located on the same side of the circuit board (maximum of one via-to-ground load capacitor on each Xtal trace).
- Use 27 pF (5% tolerance) 0402 load capacitors.
- Place load capacitor solder pad directly in line with circuit trace (see Figure 15, point A).
- Place a 30-ohm (5% tolerance) 0402 series resistor on Xtal2 (see Figure 15, point C). The location of the resistor along the Xtal2 trace is flexible, as long as it is between the load capacitor and the controller.
- Use 50-ohm impedance single-ended microstrip traces for the crystal circuit.
- Route traces so that electro-magnetic fields from Xtal2 do not couple onto Xtal1. No differential traces.
- Route Xtal1 and Xtal2 traces to nearest inside corners of crystal pad (see Figure 15, point B).
- Ensure that the traces from Xtal1 and Xtal2 are symmetrically routed and that their lengths are matched.
- The total trace length of Xtal1 or Xtal2 should be less than 750 mils.

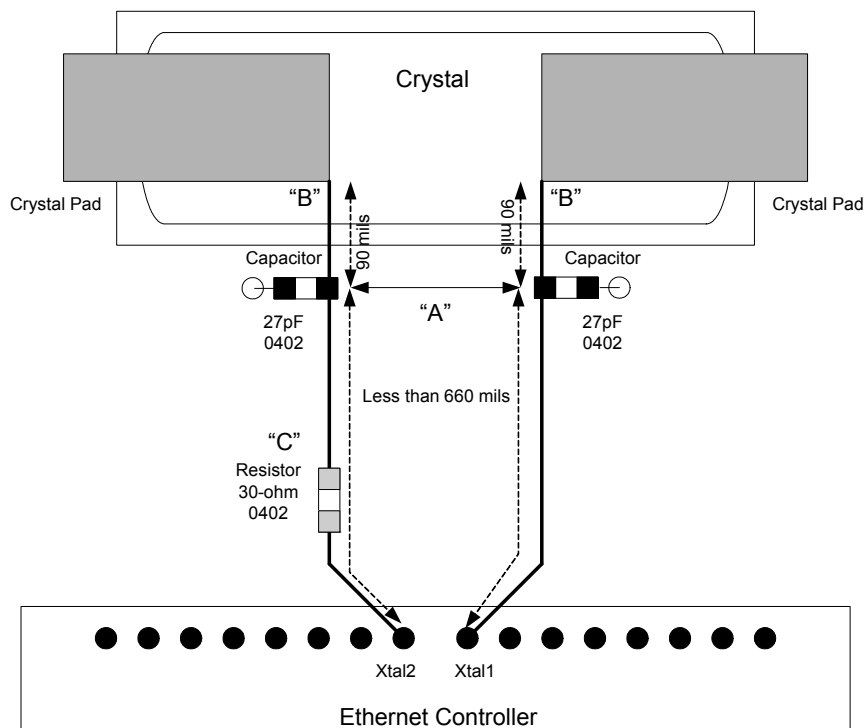


Figure 15. Recommended Crystal Placement and Layout

7.1.3 Board Stack Up Recommendations

Printed circuit boards for these designs typically have six, eight, or more layers. Although, the 82575 does not dictate the stackup, here is an example of a typical six-layer board stackup:

- Layer 1 is a signal layer. It can contain the differential analog pairs from the Ethernet device to the magnetics module, or to an optical transceiver.
- Layer 2 is a signal ground layer. Chassis ground may also be fabricated in Layer 2 under the connector side of the magnetics module.
- Layer 3 is used for power planes.
- Layer 4 is a signal layer.
- Layer 5 is an additional ground layer.
- Layer 6 is a signal layer. For 1000 BASE-T (copper) Gigabit designs, it is common to route two of the differential pairs (per port) on this layer.

This board stack up configuration can be adjusted to conform to your company's design rules

7.1.4 Differential Pair Trace Routing for 10/100/1000 Designs

Trace routing considerations are important to minimize the effects of crosstalk and propagation delays on sections of the board where high-speed signals exist. Signal traces should be kept as short as possible to decrease interference from other signals, including those propagated through power and ground planes. Observe the following suggestions to help optimize board performance:

- Maintain constant symmetry and spacing between the traces within a differential pair.
- Minimize the difference in signal trace lengths of a differential pair.
- Keep the total length of each differential pair under 4 inches. Although possible, designs with differential traces longer than 5 inches are much more likely to have degraded receive BER (Bit Error Rate) performance, IEEE PHY conformance failures, and/or excessive EMI (Electromagnetic Interference) radiation.
- Do not route a pair of differential traces closer than 100 mils to another differential pair.
- Do not route any other signal traces parallel to the differential traces, and closer than 100 mils to the differential traces (300 mils is recommended).
- Keep maximum separation within differential pairs to 7 mils.
- For high-speed signals, the number of corners and vias should be kept to a minimum. If a 90° bend is required, it is recommended to use two 45° bends instead. Refer to [Figure 16](#).

Note:

In manufacturing, vias are required for testing and troubleshooting purposes. The via size should be a 17-mil (± 2 mils for manufacturing variance) finished hole size (FHS).

- Traces should be routed away from board edges by a distance greater than the trace height above the reference plane. This allows the field around the trace to couple more easily to the ground plane rather than to adjacent wires or boards.
- Do not route traces and vias under crystals or oscillators. This will prevent coupling to or from the clock. And as a general rule, place traces from clocks and drives at a minimum distance from apertures by a distance that is greater than the largest aperture dimension




Figure 16. Trace Routing

- The reference plane for the differential pairs should be continuous and low impedance. It is recommended that the reference plane be either ground or 1.8 V (the voltage used by the PHY). This provides an adequate return path for and high frequency noise currents.
- Do not route differential pairs over splits in the associated reference plane as it may cause discontinuity in impedances.



7.1.4.1 Signal Termination and Coupling

The four differential pairs of each port are terminated with 49.9 Ω (1% tolerance) resistors, placed near the 82575 controller. One resistor connects to the MDI+ signal trace and another resistor connects to the MDI- signal trace. The opposite ends of the resistors connect together and to ground through a single 0.1 μ F capacitor. The capacitor should be placed as close as possible to the 49.9 ohm resistors, using a wide trace. Stubs created by the 49.9 Ω (1% tolerance) termination resistors should be kept at a minimum.

Do not vary the suggested component values. Be sure to lay out symmetrical pads and traces for these components such that the length and symmetry of the differential pairs are not disturbed.

7.1.5 Signal Trace Geometry for 1000 BASE-T Designs

The key factors in controlling trace EMI radiation are the trace length and the ratio of trace-width to trace-height above the reference plane. To minimize trace inductance, high-speed signals and signal layers that are close to a reference or power plane should be as short and wide as practical. Ideally, this trace width to height above the ground plane ratio is between 1:1 and 3:1. To maintain trace impedance, the width of the trace should be modified when changing from one board layer to another if the two layers are not equidistant from the neighboring planes.

Each pair of signal should have a differential impedance of 100 Ω . +/- 15%. If a particular tool cannot design differential traces, it is permissible to specify 55-65 Ω single-ended traces as long as the spacing between the two traces is minimized. As an example, consider a differential trace pair on Layer 1 that is 8 mils (0.2 mm) wide and 2 mils (0.05 mm) thick, with a spacing of 8 mils (0.2 mm). If the fiberglass layer is 8 mils (0.2mm) thick with a dielectric constant, E_R , of 4.7, the calculated single-ended impedance would be approximately 61 Ω and the calculated differential impedance would be approximately 100 Ω .

When performing a board layout, do not allow the CAD tool auto-router to route the differential pairs without intervention. In most cases, the differential pairs will have to be routed manually.

Note: Measuring trace impedance for layout designs targeting 100 Ω often results in lower actual impedance. Designers should verify actual trace impedance and adjust the layout accordingly. If the actual impedance is consistently low, a target of 105 – 110 Ω should compensate for second order effects.

It is necessary to compensate for trace-to-trace edge coupling, which can lower the differential impedance by up to 10 Ω , when the traces within a pair are closer than 30 mils (edge to edge).

7.1.6 Trace Length and Symmetry for 1000 BASE-T Designs

As indicated earlier, the overall length of differential pairs should be less than four inches measured from the Ethernet device to the magnetics.

The differential traces (within each pair) should be equal in total length to within 50 mils (1.25mm) and as symmetrical as possible. Asymmetrical and unequal length traces in the differential pairs contribute to common mode noise. If a choice has to be made between matching lengths and fixing symmetry, more emphasis should be placed on fixing symmetry. Common mode noise can degrade the receive circuit's performance and contribute to radiated emissions.



7.1.6.1 Signal Detect

Each port of the 82575 controller has a Signal Detect pin for connection to optical transceivers. For designs without optical transceivers, these signals can be left unconnected because they have internal pull-up resistors. Signal Detect is not a high-speed signal and does not require special layout.

7.1.7 Routing 1.8 V to the Magnetics Center Tap

The central-tap 1.8 V should be delivered as a solid supply plane (1.8 V) directly to the magnetic module or, if this is not possible, by a short and thick trace (lower than 0.2ohm DC resistance). The decoupling capacitors for the central tap pins should be placed as close as possible to the magnetic component. This improves both EMI and IEEE compliance.

7.1.8 Impedance Discontinuities

Impedance discontinuities cause unwanted signal reflections. Minimize vias (signal through holes) and other transmission line irregularities. If vias must be used, a reasonable budget is two per differential trace. Unused pads and stub traces should also be avoided.

7.1.9 Reducing Circuit Inductance

Traces should be routed over a continuous reference plane with no interruptions. If there are vacant areas on a reference or power plane, the signal conductors should not cross the vacant area. This causes impedance mismatches and associated radiated noise levels. Noisy logic grounds should be separated from analog signal grounds to reduce coupling. Noisy logic grounds can sometimes affect sensitive DC subsystems such as analog to digital conversion, operational amplifiers, etc. All ground vias should be connected to every ground plane; and similarly, every power via, to all power planes at equal potential. This helps reduce circuit inductance. Another recommendation is to physically locate grounds to minimize the loop area between a signal path and its return path. Rise and fall times should be as slow as possible. Because signals with fast rise and fall times contain many high frequency harmonics, which can radiate significantly. The most sensitive signal returns closest to the chassis ground should be connected together. This will result in a smaller loop area and reduce the likelihood of crosstalk. The effect of different configurations on the amount of crosstalk can be studied using electronics modeling software.

7.1.10 Signal Isolation

To maintain best signal integrity, keep digital signals far away from the analog traces. A good rule of thumb is no digital signal should be within 300 mils (7.5mm) of the differential pairs. If digital signals on other board layers cannot be separated by a ground plane, they should be routed perpendicular to the differential pairs. If there is another LAN controller on the board, take care to keep the differential pairs from that circuit away.

Some rules to follow for signal isolation:

- Separate and group signals by function on separate layers if possible. If possible, maintain a gap of 100 mils between all differential pairs (Ethernet) and other nets, but group associated differential pairs together. Note: Over the length of the trace run, each differential pair should be at least 0.3 inches away from any parallel signal traces.
- Physically group together all components associated with one clock trace to reduce trace length and radiation.



- Isolate I/O signals from high-speed signals to minimize crosstalk, which can increase EMI emission and susceptibility to EMI from other signals.
- Avoid routing high-speed LAN traces near other high-frequency signals associated with a video controller, cache controller, processor, or other similar devices.

7.1.11 Power and Ground Planes

Good grounding requires minimizing inductance levels in the interconnections and keeping ground returns short, signal loop areas small, and power inputs bypassed to signal return, will significantly reduce EMI radiation.

The following guidelines help reduce circuit inductance in both backplanes and motherboards:

- Route traces over a continuous plane with no interruptions. Do not route over a split power or ground plane. If there are vacant areas on a ground or power plane, avoid routing signals over the vacant area. This will increase inductance and EMI radiation levels.
- Separate noisy digital grounds from analog grounds to reduce coupling. Noisy digital grounds may affect sensitive DC subsystems.
- All ground vias should be connected to every ground plane; and every power via should be connected to all power planes at equal potential. This helps reduce circuit inductance.
- Physically locate grounds between a signal path and its return. This will minimize the loop area.
- Avoid fast rise/fall times as much as possible. Signals with fast rise and fall times contain many high frequency harmonics, which can radiate EMI.
- The ground plane beneath a magnetics module should be split. The RJ45 connector side of the transformer module should have chassis ground beneath it.

7.1.12 Traces for Decoupling Capacitors

Traces between decoupling and I/O filter capacitors should be as short and wide as practical. Long and thin traces are more inductive and would reduce the intended effect of decoupling capacitors. Also for similar reasons, traces to I/O signals and signal terminations should be as short as possible. Vias to the decoupling capacitors should be sufficiently large in diameter to decrease series inductance.

7.1.13 Light Emitting Diodes for Designs Based on the 82575 Controller

The 82575 controller provides four programmable high-current push-pull (active high) outputs per port to directly drive LEDs for link activity and speed indication. Each LAN device provides an independent set of LED outputs; these pins and their function are bound to a specific LAN device. Each of the four LED outputs can be individually configured to select the particular event, state, or activity, which will be indicated on that output. In addition, each LED can be individually configured for output polarity, as well as for blinking versus non-blinking (steady-state) indication.

Since the LEDs are likely to be integral to a magnetics module, take care to route the LED traces away from potential sources of EMI noise. In some cases, it may be desirable to attach filter capacitors.

The LED ports are fully programmable through the EEPROM interface.



7.1.14 Thermal Design Considerations

The 82575 Gigabit Ethernet Controller contains a thermal sensor that is accessible through the SMBus. Trip points can be set in the EEPROM for the device.

IcePak* and FlowTherm* models are available for the 82575 Ethernet Controller; contact your Intel representative for information.

Refer to the application note: *Intel® 82575 Ethernet Controller Thermal Design Considerations* for more information.

7.2 Physical Layer Conformance Testing

Physical layer conformance testing (also known as IEEE testing) is a fundamental capability for all companies with Ethernet LAN products. PHY testing is the final determination that a layout has been performed successfully. If your company does not have the resources and equipment to perform these tests, consider contracting the tests to an outside facility.

7.2.1 Conformance Tests for 10/100/1000 Mbps Designs

Crucial tests are as follows, listed in priority order:

- Bit Error Rate (BER). Good indicator of real world network performance. Perform bit error rate testing with long and short cables and many link partners. The test limit is 10^{-11} errors.
- Output Amplitude, Rise and Fall Time (10/100Mbps), Symmetry and Droop (1000Mbps). For the 82575 controller, use the appropriate PHY test waveform.
- Return Loss. Indicator of proper impedance matching, measured through the RJ-45 connector back toward the magnetics module.
- Jitter Test (10/100Mbps) or Unfiltered Jitter Test (1000Mbps). Indicator of clock recovery ability (master and slave for Gigabit controller).

7.3 Troubleshooting Common Physical Layout Issues

The following is a list of common physical layer design and layout mistakes in LAN On Motherboard Designs.

1. Lack of symmetry between the two traces within a differential pair. Asymmetry can create common-mode noise and distort the waveforms. For each component and/or via that one trace encounters, the other trace should encounter the same component or a via at the same distance from the Ethernet silicon.
2. Unequal length of the two traces within a differential pair. Inequalities create common-mode noise and will distort the transmit or receive waveforms.
3. Excessive distance between the Ethernet silicon and the magnetics. Long traces on FR4 fiberglass epoxy substrate will attenuate the analog signals. In addition, any impedance mismatch in the traces will be aggravated if they are longer than the four inch guideline.
4. Routing any other trace parallel to and close to one of the differential traces. Crosstalk getting onto the receive channel will cause degraded long cable BER. Crosstalk getting onto the transmit channel can cause excessive EMI emissions and can cause poor transmit BER on long cables. At a minimum, other signals should be kept 0.3 inches from the differential traces.
5. Routing one pair of differential traces too close to another pair of differential traces. After exiting the Ethernet silicon, the trace pairs should be kept 0.3 inches or more away from the other trace pairs. The only possible exceptions are in the vicinities



where the traces enter or exit the magnetics, the RJ-45 connector, and the Ethernet silicon.

6. Use of a low-quality magnetics module.
7. Re-use of an out-of-date physical layer schematic in a Ethernet silicon design. The terminations and decoupling can be different from one PHY to another.
8. Incorrect differential trace impedances. It is important to have $\sim 100 \Omega$ impedance between the two traces within a differential pair. This becomes even more important as the differential traces become longer. To calculate differential impedance, many impedance calculators only multiply the single-ended impedance by two. This does not take into account edge-to-edge capacitive coupling between the two traces. When the two traces within a differential pair are kept close to each other, the edge coupling can lower the effective differential impedance by 5Ω to 20Ω . Short traces will have fewer problems if the differential impedance is slightly off target.



8.0 Thermal Management

Please see the 82575 Thermal Application Note, available on the Intel Developer site.

9.0 Reference Design Bill of Materials

The bill of materials for Intel's reference designs is available on the Intel Developer site.

10.0 Design and Layout Checklists

Design and Layout checklists are available on the Intel Developer site; please contact your Intel representative to obtain these documents.

11.0 Reference Schematics

The reference schematics describing typical designs for the 82575 Ethernet Controller are available on the Intel Developer site.

The schematics are updated as needed when changes are made in the device, so please check to ensure you have the latest version for use with your design.

12.0 Symbol

The symbol for the 82575 Ethernet Controller is available on the Intel Developer site.

The symbol is updated as needed, so please check to ensure you have the latest version for use with your design.

§ §