Intel® Ethernet Controller I225/I226

Design Considerations and Guidelines

Client Connectivity Division

September 2022

Revision 2.6.2

Revision History

Revision	Date	Comments
1.0	July, 2019	Initial Release.
1.1	August, 2019	Fix typo. The Vct Capacitors for the discrete magnetics designs should be 68 pF (it was listed 0.68 uF).
1.2	October, 2019	Remove MDI Lane Swap. Update Vct cap value. Suggest using 0.1uF (instead of 68pF).
2.0	October, 2020	Add MDI Lane Swap (it is now supported). Remove Magnetics Common-Mode Chokes locations requirement. Update Maximum MDI Trace Length table. Add Leakage Information Update Design Notes
2.1	November, 2020	Add Voltage Shifter Information
2.5	September 2021	Added I226 Added System Thermal Solution
2.6	February 2022	Add more info for the SMBUS Section
2.6.1	April 2022	Chapter 1.6: update statement of JRST_N signals
2.6.2	September 2022	Update table 1-3 MDI Pin Assignment for MDI LAN SWAP

1.0 Design Considerations and Guidelines

This section provides guidelines for connecting interfaces, using special pins, selecting components, and layout guidance. Unused interfaces should be terminated with pull-up or pull-down resistors. Pin #8 is reserved and should pull-down with 0 Ω resistor. Please check the I225/I226 collateral package or ask your Intel representative if there is question about schematic or layout designs.

1.1 Connecting the PCIe Interface

The I225/I226 connects to the host system using a PCIe interface. The lane must be AC-coupled between its corresponding transmitter and receiver, with the AC-coupling capacitor located close to the transmitter side (within 1 inch). Each end of the link is terminated on the die into nominal 100 differential DC impedance. Board termination is not required.

The I225/I226 supports the PCIe Gen 2, v3.1 interface. Refer to PCI Express* Card Electromechanical Specification, rev2.0 and the PCI Express* Gen2 Specification, ver3.1.

1.1.1 Link Width Configuration

The I225/I226 is a MAC/PHY PCIe device that supports link width x1. During link configuration, the platform and the I225/I226 negotiates on a common link width. The link width result must be x1.

1.1.2 Polarity Inversion and Lane Reversal

To ease routing, designers have the flexibility to use the lane reversal modes supported by the I225/I226. During the link training sequence, the polarity inversion of differential pairs is detected and automatically corrected.

1.1.3 PCIe Reference Clock

The I225/I226 requires a 100 MHz differential reference clock called PE_CLK_p and PE_CLK_n. This signal is typically generated from a Peripheral Controller Hub (PCH) and provided to the I225/I226. For Network Interface Cards (NICs), the clock is furnished at the PCIe connector and the signal must be routed to the PCIe connector.

Note: The frequency tolerance for the PCIe reference clock is ±300 ppm.

1.1.4 PCIe Bias Resistor

For proper biasing of the PCIe analog interface, a 200 Ω 1% pull-down resistor must be connected between the RBIAS (pin #20) pin and ground in order to calibrate the PCIe analog modules. To avoid noise coupled onto this reference signal, place the bias resistor close to the I225/I226 and keep traces as short as possible. Do not change this resistor value.

1.1.5 Miscellaneous PCIe Signals

PE_WAKE_N: The I225/I226 signals the power management events to the system by pulling low the PE_WAKE_N signal. Note that somewhere in the system, this signal must be pulled high to the auxiliary 3.3 V supply rail.

PE_RST#: The PE_RST# signal, which serves as the familiar reset function for the I225/I226, needs to be connected to the host system's corresponding signal. On Motherboard, it should be Platform Reset.

1.1.6 PCIe Layout Recommendations

For information regarding the PCIe signal routing, refer to the Intel PCIe Design Guide.

1.2 Connecting The 2.5G BASE-T MDI Interfaces

The MDI line interface on the I225/I226 is capable of driving up to +100 meters of CAT-5E or CAT-6 unshielded twisted pair (100 Ω differential impedance) on 10 Mb/s, 100 Mb/s, 1 Gb/s, or 2.5 Gb/s. It is designed to drive this via a quad, 50 Ω center tapped 1:1 transformer connected to an RJ-45 PCB-mount jack. Solutions that combine the transformer and RJ-45 jack into a single device are also supported.

The MDI line interface on the I225/I226 supports automatic A/B and C/D pair swaps and inversions (called the MDI-X feature). It also supports provisioned ABCD to DCBA pair reversal (called the MDI Lane Swap feature) for ease of routing via a hardware strap. Please see the MDI Lane Swap Configuration Feature section for details.

1.2.1 PHY Calibration Resistor

For proper biasing of the MDI analog interface, a 22 KΩ 1% pull-down resistor must be connected between the RCAL pin (pin#42) and ground in order to calibrate the I225/I226 MDI analog front-end. To avoid noise coupled onto this reference signal, place the calibration resistor close to the I225/I226 and keep traces as short as possible. Do not change this resistor value.

1.2.2 MDI Circuit Guidelines

The MDI discrete design and integrated magnetic components were chosen for inclusion in the reference design and Bill of Material (BOM). Refer to Section 1.9 for more details. These components are capable of delivering the performance required for this demanding application.

For systems with enhanced ESD performance requirements, the additional enhanced ESD parts might be needed. Consult with the ESD part vendors for their latest technology and techniques. Also consult with your Intel representative for the schematic and layout implementation.

1.2.3 Magnetics Component

The magnetics module has a critical effect on overall IEEE and emissions conformance. The I225/I226 should meet the performance required for a design with reasonable margin to allow for manufacturing variation. Carefully qualifying new magnetics modules prevents problems that might arise because of interactions with other components or the Printed Circuit Board (PCB) itself. The magnetics specified should comply with the specifications listed in the Intel® 2.5G BASE-T Magnetic Specification Electrical/ Mechanical Requirements for 2.5G BASE-T Magnetic Components.

The steps involved in magnetics module qualification are:

1. Verify that the vendor's published specifications in the component datasheet meet or exceed the required IEEE 802.3an specifications and the internal specifications listed in the Intel® 2.5-G BASE-T Magnetic Specification Electrical/Mechanical Requirements for 2.5G BASE-T Magnetic Components.
2. Independently measure the component's electrical parameters on a 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.

Magnetics modules for 2.5G BASE-T Ethernet as used by the I225/I226 are similar to those designed for 1G BASE-T, except that the electrical requirements for the board layout and magnetics are more stringent.

1.2.4 MDI Layout Guidance

The magnetics can be implemented either as a discrete module or an integrated solution combining the magnetics and the RJ-45 jack.

For the discrete magnetics, the I225/I226 designs consist of an RJ-45 jack, Bob Smith termination and discrete magnetics all separated parts (see Figure 1-2).

For the integrated magnetics, the I225/I226 designs is simple. The RJ-45 connector has Bob Smith termination and discrete magnetics integrated inside (see Figure 1-6).

Minimizing the amount of space needed for the PHY is important because other interfaces compete for physical space on the LAN on Motherboard (LOM) near the connector. The PHY circuits need to be as close as possible to the connector.

Figure 1-2 illustrates some basic placement distance guidelines. It shows four differential pairs and the layout can be generalized for a 2.5 GbE system with four analog pairs. The ideal placement for the I225/I226 is approximately two inches behind the magnetics module for both the discrete and integrated solutions.

1.2.4.1 Discrete Magnetics

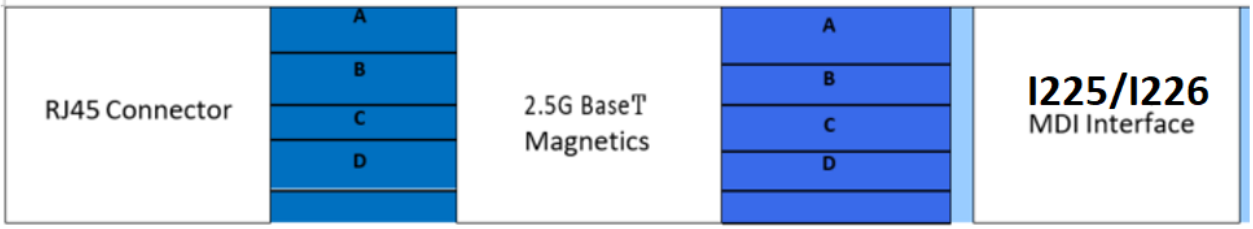

Figure 1-2 Discrete Magnetics

The I225/I226 impedance of the Bob Smith termination is 75 Ω Channel A, B, C and D, whereas the PHY input impedance is 100 Ω. The I225/I226, referred as a LAN silicon in Figure 1-3 and Figure 1-5, must be at least 1.5 inches from the I/ O back panel. To help reduce EMI interference, the following recommendations should be followed:

- Minimize the length of the MDI interface.
- Place the MDI traces no closer than 0.5 inch (1.3 cm) from the board edge.

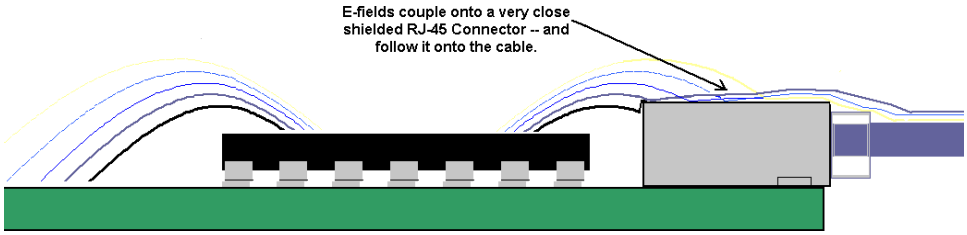

Figure 1-3 Effect of Device Placed Less Than One Inch from the RJ-45 Connector

- The MDI implementation with the Discrete Magnetics Connector is shown in Figure 1-4. The optional stuffing options (bridge capacitors) can be tuned for improving the EMI performance. Some Designs cam successfully pass EMI tests when using one single Ground Plane (CHASSIS GND and SYSTEM GND are connected together, no stuffing option).
- The TERM-PLANE ground (isolating island) is still required.

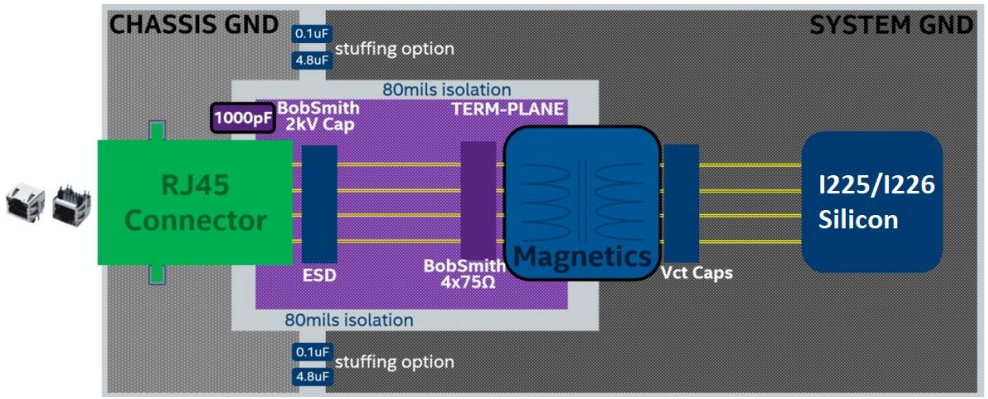

Figure 1-4 Discrete Magnetics Connector Layout Topology

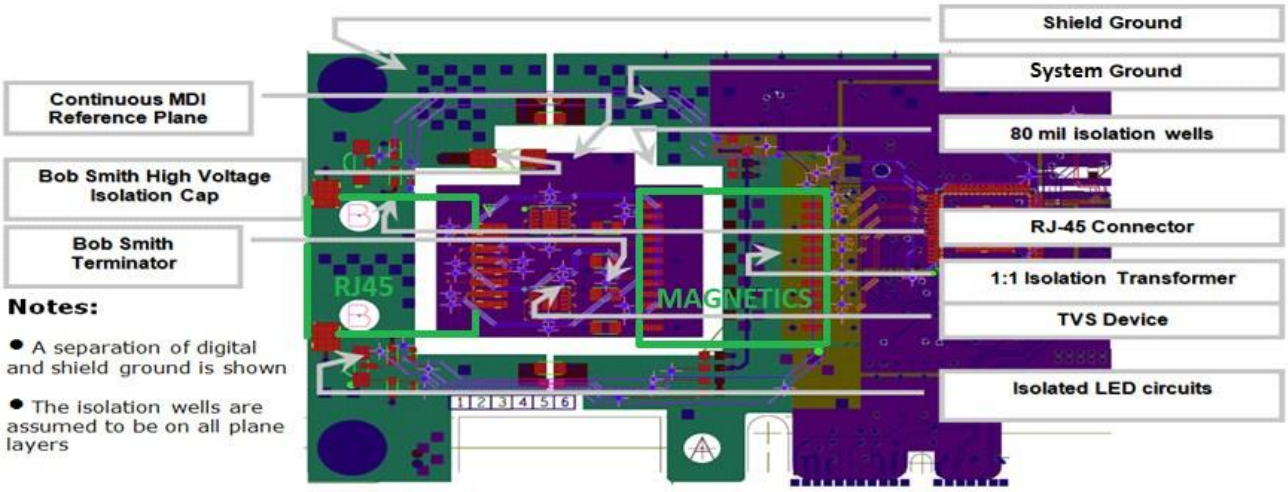

Figure 1-5 Example of a Discrete Magnetics Layout

Figure 1-5 explanation:

- The ground is split underneath the magnetics, with the MDI Reference Plane GND (the purple island) being present under the front-end (Figure 1-5), and the circuit system ground under the I225/I226 side of the magnetics. The MDI traces on the line side are referenced to the MDI reference plane. Splitting the ground planes beneath the transformer minimizes noise coupling between the primary and secondary sides of the transformer and between the adjacent coils in the magnetics. This arrangement also improves the common mode choke functionality of magneticsmodule.
- Between the RJ45 connector and the magnetics device, the MDI Reference Plane GND (the purple island) can be voided. When using this voiding technique, this MDI segment (cable side of the magnetics) will have no reference plane. Please make sure using the trace width/separation to target 100-ohm differential impedance. The MDI segment on the silicon side of the magnetics still references to the system ground.
- The Bob Smith termination (4x 75 Ω) to be implemented from the line side transformer center taps to the reference plane.
- Bob Smith termination and the shield of the RJ-45 jack are both connected to chassis ground, which are the two large circular metallization in the MDI (see Figure 1-6).
- The resistors in the Bob Smith termination are required to be 0805's to handle the cable discharge voltages.
- Similarly, the magnetics should be placed as close as possible to the RJ-45 jack.
- In Figure 1-5, the Hi-POT clearance for cable discharge is shown and designed with ≥80 mils of clearance from the ground. Therefore, the MDI reference plan has 80 mils clearance away from any other signals (including shield ground and system ground). Note that the breakdown voltage in FR-4 is lowest in the x-y axes planar to the PCB and highest in the z-axis between layers.
- The MDI reference plane is coupling to the shield ground via the Bob Smith high voltage isolation capacitor. This helps ESD performance.
- The shield ground can be coupled to the system ground via 2-4 small value capacitors. Spread and distribute these capacitors out. The capacitor value can be tuned for improving the EMI performance.

1.2.4.2 Integrated Magnetics

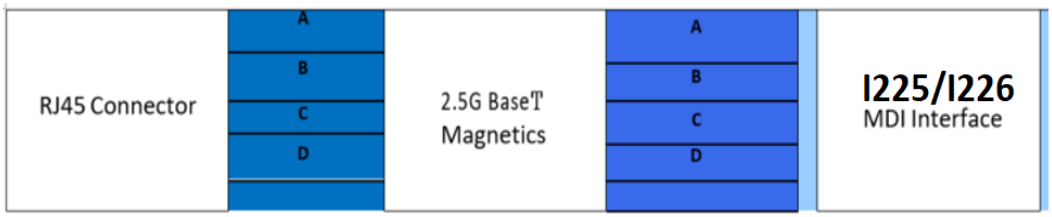

Figure 1-6 Integrated Magnetics

At high level, the guidance on MDI trace implementation between the I225/I226 and the integrated module remains the same as with discrete magnetics, including length requirements and target impedance.

However, the RJ45 connector has the magnetics and the Bob Smith termination network integrated inside. Therefore, there is no need to take care of the high-voltage isolation routing. Do not split the connector shield ground with the system ground. Instead, connect the RJ45 shield ground directly to the system ground.

The integrated magnetics RJ45**+USB ports**, do not slit the Ground plane.

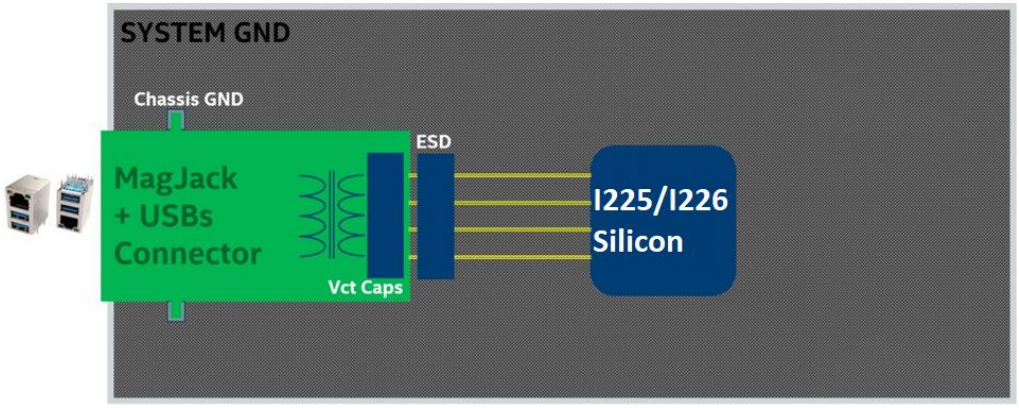

Figure 1-7 Integrated Magnetics/USB Connector Layout Topology

The MagJack/ICM integrated magnetics connector can have the slit ground implementation for the EMI improvement. The stuffing cap values can be adjusted to gain 1-2dBs.

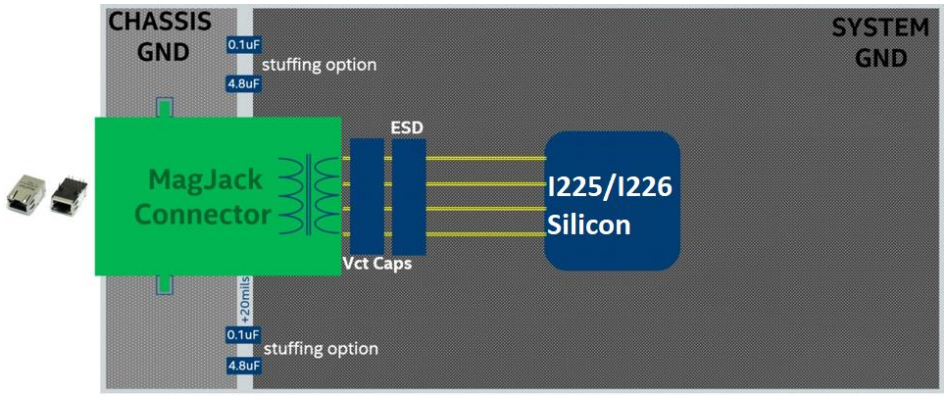

Figure 1-8 Integrated Magnetics Connector Layout Topology (with the ground slit for the EMI improvement -optional)

1.2.4.3 MDI Differential Pair Trace Routing

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.

1.2.4.4 Signal Trace Geometry

One of 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, the trace-width to trace-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 signals should target a differential impedance of 100 Ω ±15%.

A set of trace length calculation tools are made available from Intel to aid with MDI topology design. Contact your Intel representative for tool availability.

When designing a board layout, the automatic router feature of the CAD tool must not route the differential pairs without intervention. In most cases, the differential pairs require manual routing.

Note: Measuring trace impedance for layout designs targeting 100 Ω often results in lower actual impedance due to over-etching. Designers should verify actual trace impedance and adjust the layout accordingly. If the actual impedance is consistently low, a target of 105 Ω to 110 Ω should compensate for over-etching.

1.2.4.4.1 Matching Traces within a Pair (P and N)

P and N for each MDI pair should be matched to within 5 mils on the PCB to prevent common-to-differential and differential-to-common conversion due to the length mismatch.

If in-pair length matching is not possible using bends or small loops, serpentine routing (zig-zag of a shorter trace) is acceptable if only one to three meanders are routed within 200 mils of the source of the skew or the end(s) of any otherwise unmatched lengths of a differential trace segment. For example, near device pins and/or at or near the connector pins and possibly at differential signal vias. Refer to the Intel® Ethernet Controller I225/I226 Checklists for more details.

Table 1-1 MDI Routing Summary

Parameter	Main Route Guidelines	Breakout Guidelines[1]	Notes
Signal group	MDI_P[3:0] MDI_N[3:0]		
Microstrip*/Stripline* uncoupled single-ended impedance specification	50 Ω ±10%		
Microstrip/Stripline uncoupled differential impedance specification	100 Ω ±15%		2,3
Microstrip nominal trace width	Design dependent	Design dependent	
Microstrip nominal trace space	Design dependent	Design dependent	3
Microstrip/Stripline trace length	< 4 inches (recommended)		
Microstrip/Stripline pair-to-pair space (edge-to-edge)	≥ 7 times the dielectric thickness		
Microstrip/Stripline bus-to-bus spacing	≥ 7 times the dielectric thickness		
Matching traces within a pair (P and N)	< 5 mils		
Keep pair-to-pair length differences	< 2 inches		

1. Pair-to-pair spacing ≥ 7 times the dielectric thickness for a maximum distance of 500 mils from the pin. The phase tolerance between MDI_P and MDI_N is <5mils.
2. Board designers should ideally target 100 Ω ±10%. If it's not feasible (due to board stackup) it is recommended that board designers use a 95 Ω ±10% target differential impedance for MDI with the expectation that the center of the impedance is always targeted at 95 Ω. The ±10% tolerance is provided to allow for board manufacturing process variations and not lower target impedances. The minimum value of impedance cannot be lower than 90 Ω.
3. Simulation shows 80 differential trace impedances degrade MDI return loss measurements by approximately 1 dB from that of 90 Ω

Table 1-2 Maximum Trace Lengths Based on Trace Geometry and Board Stackup

Dielectric Thickness (mils)	Dielectric Constant (DK) @1MHz	Width/Space/Width (mils)	Pair-to-Pair Space (mils)	Nominal Impedance (Ω)	Impedance Tolerance (+/-%)	Maximum Trace Length (inches)[1]
2.7	4.05	4 / 10 / 4	19	95[2]	17[2]	3.5
2.7	4.05	4 / 10 / 4	19	95[2]	15[2]	4
2.7	4.05	4 / 10 / 4	19	95	10	5
3.3	4.1	4.2 / 9 / 4.2	23	100[2]	17[2]	4

1. Longer MDI trace lengths can be achievable but that might make it more difficult to achieve IEEE conformance. Simulations have shown deviations are possible if traces are kept short. Longer traces are possible; use cost considerations and stack-up tolerance for differential pairs to determine length requirements.
2. Deviations from 100 Ω nominal and/or tolerances greater than 15% decrease the maximum length for IEEE conformance.

Note: Use the MDI differential trace calculator to determine the maximum MDI trace length for your trace geometry and board stackup. Contact your Intel representative for access.

The following factors can limit the maximum MDI differential trace lengths for IEEE conformance:

- Dielectric thickness
- Dielectric constant
- Nominal differential trace impedance
- Trace impedance tolerance
- Copper trace losses
- Additional devices, such as switches, in the MDI path might impact IEEE conformance. Board geometry should also be factored in when setting trace length.

1.2.5 PHY MDI Lane Swap Configuration

In general, the RJ45 connector on the system PCB can have two different structures with the tap up or the tap down as shown in Figure below. The difference between tap-up and tap-down is a swap of positioning ABCD and DCBA orders. The I225/I226 pin#18 strapping allows system designers to perform the MDI lane swap configuration. As a result, a PCB layout can avoid the crossing routing problem.

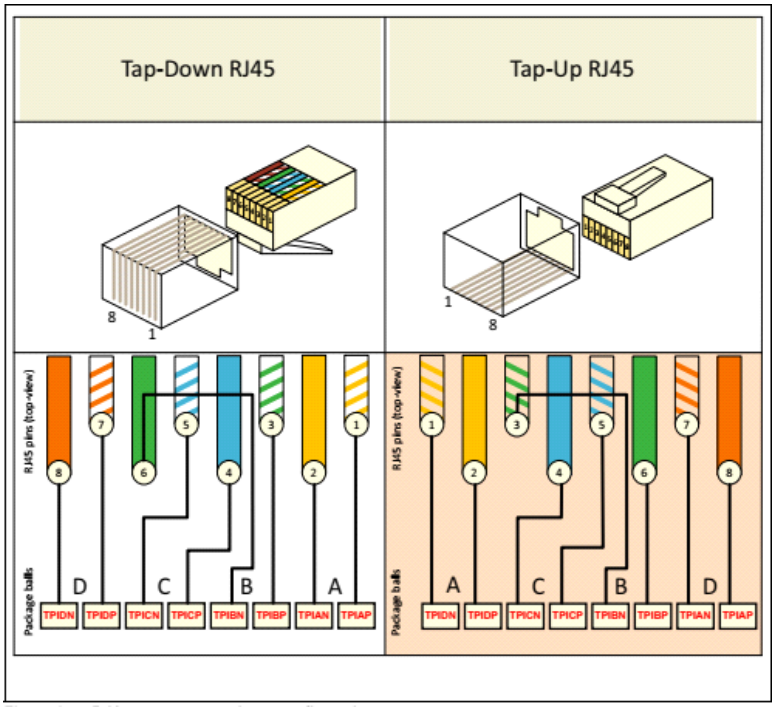

Figure 1-9 Tap-Down RJ45 vs Tap-Up RJ45

MDI Crossing Issue: When selecting RJ45 connectors, the RJ45 MDI lanes order might not line up with the Silicon MDI lanes order. This causes the MDI crossing routing problem in layout.

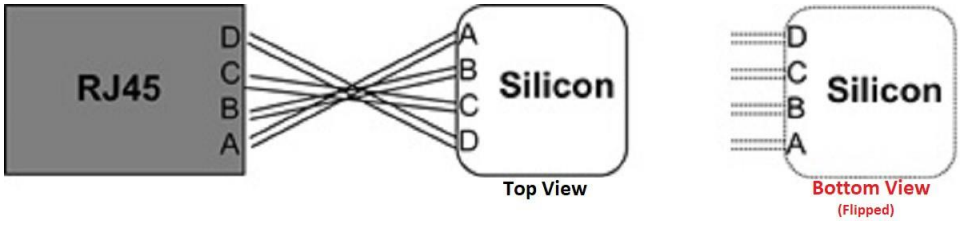

Figure 1-10 Crossing Routing Issue Between RJ45 and Silicon

Silicon: Pin#18 allows the MDI lane swap configuration. This is a hardware strap.

Default Configuration (Unstrap): Pin#18 has its internal pull-up resistor. The pin can leave unconnected.

MDI Swap Configuration (Strapped): Pin 18 needs a 10 K ohms pull-down. This hardware strap activates the MDI lane swap feature. ABCD order are reversed.

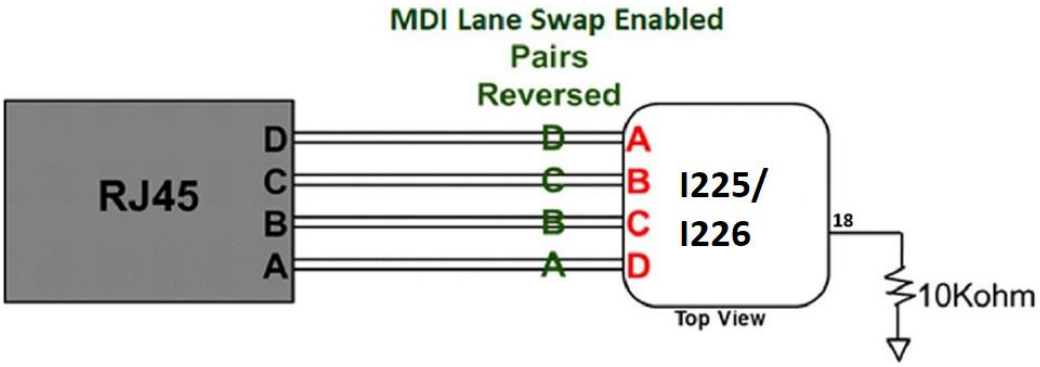

Figure 1-11 MDI Lane Swap Enabled by Strapping

Table 1-3 MDI Pin Assignment

Pin #	Pin Name	Pin Type	MDI_LANE_ SWAP = 1 Functionality	MDI_LANE_ SWAP = 0 Functionality	Function
45	MDI_A_P	AI / AO	MDI_A_P	MDI_D_P	Media Dependent Interface. Connect four MDI pairs directly to the magnetics/RJ45 component. No external terminations are needed.
46	MDI_A_N	AI / AO	MDI_A_N	MDI_D_N	
48	MDI_B_P	AI / AO	MDI_B_P	MDI_C_P	
49	MDI_B_N	AI / AO	MDI_B_N	MDI_C_N	
51	MDI_C_P	AI / AO	MDI_C_P	MDI_B_P	
52	MDI_C_N	AI / AO	MDI_C_N	MDI_B_N	
54	MDI_D_P	AI / AO	MDI_D_P	MDI_A_P	
55	MDI_D_N	AI / AO	MDI_D_N	MDI_A_N	
42	PHY_CAL	Calibration for all GPHY Ethernet ports.			

1.2.6 Center Tap Connection Via Capacitors to Ground

The I225/I226 has a voltage-mode driver. Therefore, it is required that a center tap be decoupled to ground via capacitors.

When using an integrated magnetic, a 0.1 µF capacitor should be connected from each center tap pin to ground. Most of integrated magnetic connector will have 4x 0.1 µF capacitor inside. Therefore, the I225/I226 does not need any external capacitor. Please add 1x 0.1 µF capacitor (empty) as a place holder.

When using a discrete magnetic, add a 0.1 µF capacitor to each center tap pin of the magnetics. The I225/I226 should have a total of 4x 0.1 µF capacitors at the magnetics center tap. Tie all 4x0.1uF together might improve EMI performance. _Please see I225/I226 Reference Schematic for details.

1.3 Connecting the SMBus Interface

The I225/I226 SMB_DATA and SMB_CLK signals must be connected to the PCH's SMLINK0 for supporting vPro. These I225/I226 SMBUS pins require pull-up resistors.

For PCH's SMLINK0 power rail is 3.3V: Use only one PU = 499 Ω between PCH <-> I225/I226.
For PCH's SMLINK0 power rail is 1.8V: Use the Voltage-Level Shifter. Both sides of this Voltage Shifter should have one PU. These PU values are validated and defined in Intel RVP Schematics.

The SMBUS PU value for the Votlage-Level Shifter should be within 1K-2K Ω, depend on the total bus capacitances and PCH's internal setting. Please use the value setting on Intel RVP Platform Schematics.

If the interface is not used, add 10Kohm pull-up resistors.

Refer to the I225/I226 Reference Schematic and I225/I226 Design Checklist for more details.

1.4 Connecting the Flash Interface

1.4.1 Connecting the Flash

Intel recommends having a Flash socket for test & validation purpose. However, the socket is not required for production board. The following picture shows a typical SPI connection for the I225/I226. Please see the Datasheet and I225/I226 Reference Schematic for details of the strapping options.

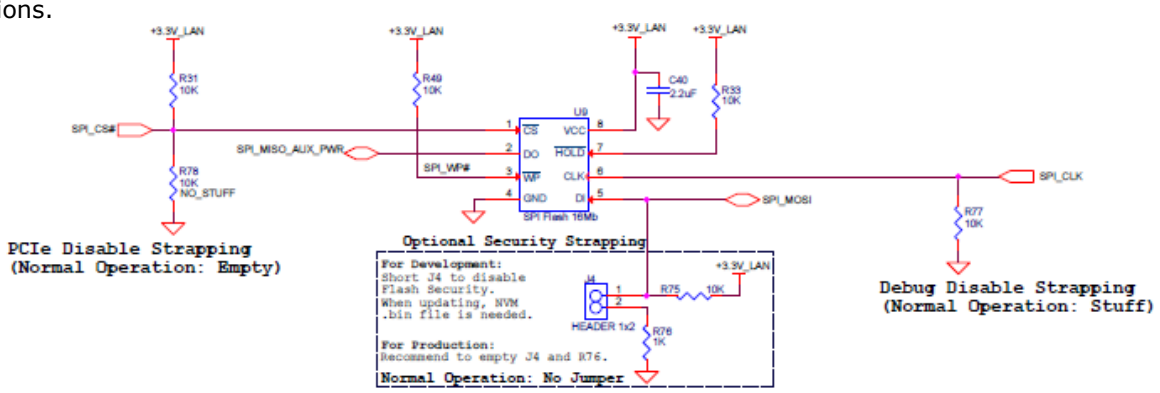

Figure 1-12 Typical SPI Flash Connection

The I225/I226 provides support for a SPI Flash device that is made accessible to the system through the following:

- Flash Base Address register. The PCIe Control register at an offset that will be identified in a later release of this document.
- An address range of the IOADDR register, defined by the I/O Base Address register (PCIe Control register at an offset that will be identified in a later release of this document.
- Expansion ROM Base Address register. The PCIe Control register is at an offset that will be identified in a later release of this document.

1.4.2 Supported Flash Devices

The I225/I226 requires an external SPI Flash. Several words of the Flash are accessed automatically by the I225/I226 after reset to provide pre-boot configuration data before it is accessed by host software. The remainder of the Flash space is available to software for storing the MAC address, serial numbers, and additional information.

See Datasheet for more details about Flash.

1.5 Connecting the Light Emitting Diodes (LEDs)

The I225/I226 provides three programmable outputs to directly drive the LEDs for link activity and speed indications. Each of the three LED outputs can be individually configured to select the particular event, state, or activity, which is 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.

The LED ports are fully programmable through the Flash interface (LEDCTL register). In addition, the hardware-default configuration for all LED outputs can be specified via a Flash field, thus supporting LED displays configurable to a particular OEM preference.

System designers need to provide separate current limiting resistors for each LED connected.

Since the LEDs are likely to be placed close to the board edge and to external interconnect, take care to route the LED traces away from potential sources of EMI noise. In some cases, it might be desirable to attach filter capacitors.

1.5.1 LED Configuration #1

This configuration has been used widely in Sever/Workstation/Desktop/AIC/etc.

Expectation:

Speed LED is a dual-color. Link-Activity LED is a single-color.

Highest speed LED is Green. Second highest speed LED is Yellow. All other-speed LEDs are OFF.

Table 1-4 LED Configuration #1 Expected Behavior

Speed LED		
10 Mbps	OFF	
100 Mbps	OFF	
1000 Mbps	Yellow	←2nd Highest Speed
2500 Mbps	Green	←Highest Speed
Link-Activity LED		
Link-up	Green	
Tx/Rx Activity	Blinking Green	

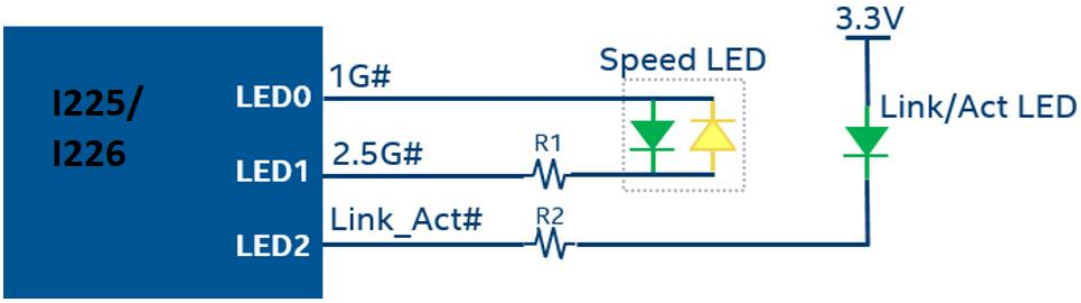

Figure 1-13 Typical Schematic Setup for LED Configuration #1

Design Notes:

- All LED pins are active-low signals.
- R1/R2 are current-limit resistors.
- Good practice for server. Just by looking at the back of server racks, IF all LEDS are green, users know all ports are running at the highest speeds.

Table 1-5 I225/I226 NVM Setting for Configuration #1: Offset 0x1C = 0x0508 Offset 0x1F = 0x8427

NVM word offset 0x1C = 0x0508 (for LED1 Setting)															
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	0	0	0	0	1	0	1	0	0	0	0	1	0	0	0
Reserved					TXAMP			Blink1	Invert1	Reserved		LED1=0x8 (for 2.5G#)			

NVM word offset 0x1F = 0x8427 (for LED0 and LED2 Setting)															
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	0	1	0	0	0	0	1	0	0	1	1	1
Blink2	Invert2	Reserved		LED2=0x4 (for Link_Act#)				Blink0	Invert0	Global Blink	Reserved	LED0=0x7 (for 1G#)			

Design Notes:

Invert = 0 (LED is an active low signal)

Blink2 = 1 (LED2 is solid = link-up and blinking = Tx/Rx)

Global Blink = 1 (fast blinking mode for all LEDs)

1.5.2 LED Configuration #2

This configuration is for the I225/I226 design, which needs individual color for each Ethernet speed.

Expectation:

The speed LED can display three different colors. The Link-Activity LED is a single-color.

Each speed LED has a different color. This has an advantage for distinct displays of all speeds at assigned colors. The RJ45 cost expected to be higher.

Table 1-6 LED Configuration #2 Expected Behavior

Speed LED		
10 Mbps	OFF	
100 Mbps	Green	← A
1000 Mbps	Yellow	← Mixed A+B color, both LEDs turned on at a same time.
2500 Mbps	Red	← B
Link-Activity LED		
Link-up	Green	
Tx/Rx Activity	Blinking Green	

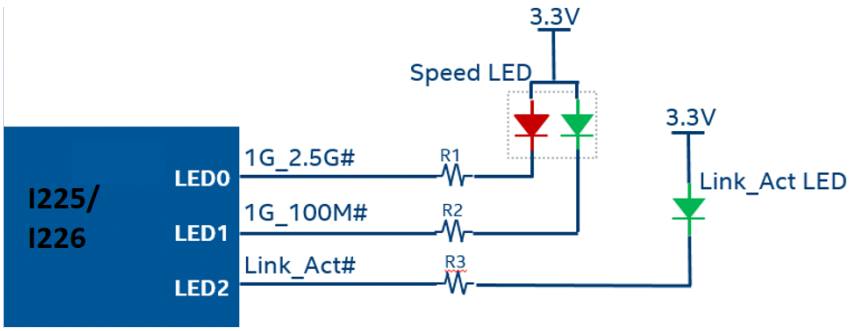

Figure 1-14 Typical Schematic Setup for LED Configuration #2

Design Notes:

- All LED pins are active-low signals.
- R1/R2/R3 are current-limit resistors.
- Each speed can have its own color by this 3-pin LED. Traditionally, 1 GbE is yellow and 100 Mb/s is green. This setup creates a new red indication for 2.5 GbE.

Table 1-7 Mixed-color Examples for the LED Configuration #2

	Example#	Color A	Color B	Mixed A+B
Recommended →	1	Red	Green	= Yellow
	2	Red	Yellow	= Orange
	3	Orange	Green	= Yellow

Table 1-8 I225/I226 NVM Setting for Configuration #2: Offset 0x1C = 0x050D Offset 0x1F = 0x842E

NVM word offset 0x1C = 0x050D (for LED1 Setting)															
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	0	0	0	0	1	0	1	0	0	0	0	1	1	0	1
Reserved					TXAMP			Blink1	Invert1	Reserved		LED1=0xD (for 1G_100M#)			

NVM word offset 0x1F = 0x842E (for LED0 and LED2 Setting)															
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	0	1	0	0	0	0	1	0	1	1	1	0
Blink2	Invert2	Reserved		LED2=0x4 (for Link_Act#)				Blink0	Invert0	Global Blink	Reserved	LED0=0xE (for 1G_2.5G#)			

Design Notes:

- Invert = 0 (LED is an active low signal)
- Blink2 = 1 (LED2 is solid = Link-up and blinking = Tx/Rx)
- Global Blink = 1 (fast blinking mode for all LEDs)

1.5 Connecting LAN_Disable_N Signal

The I225/I226 has the LAN_DISABLE_N pin (pin #2) that can be used for disabling Ethernet functions (disable the silicon). In a typical normal operation design, pull this LAN_DISABLE_N pin up with a 10 K Ω resistor to LAN 3.3 Vdc. The silicon will be always enabled and ready whenever the power is available.

For supporting RTD3 and Ultra Low Power state, LAN_DISABLE_N must be connected to PCH's GPIO. When driving the LAN_DISABLE_N signal low, the I225/I226 will enter a ULP state if enabled in the NVM.
While in the ULP state, the Host can be setup and de-assert the LAN_DISABLE_N signal to exit.
See I225/I226 Datasheet for detail.

1.6 Connecting the JTAG Port

The I225/I226 contains a test access port (3.3 V only) conforming to the IEEE 1149.1-2001 Edition (JTAG) specification. To use the test access port, connect these balls to pads accessible by specific test equipment.

For normal operation, pull-up JTCK, JTDO, JTMS and JTDI signals with 10 K Ω resistors to 3.3 Vdc. JRST_N signals can be left open or pull-low, which hold the JTAG interface in reset (test disabled).

1.7 Crystal Design Considerations

All designs require an external clock. The only option for this clock source is a 25 MHz external crystal. The I225/I226 uses the crystal to generate clocks for the high-speed interfaces.

The chosen crystal vendor should be consulted early in the design cycle. Crystal manufacturers familiar with networking equipment clock requirements might provide assistance in selecting an optimum, low-cost solution.

The following parameters are required when selecting the 25 MHz external crystal component for I225/I226.

External Crystal Specifications:

Parameter Name	Symbol	Recommended Value	Max/Min Range	Conditions
Frequency	f_o	25 [MHz]		@25 [°C]
Vibration Mode		Fundamental		
Frequency Tolerance @25 °C	Df/f_o @25°C	±30 [ppm]		@25 [°C]
Temperature Tolerance	Df/f_o	±30 [ppm]		0 to +70 [°C]
Series Resistance (ESR)	R_s		60 [Ω] max	@25 [MHz]
Crystal Load Capacitance	C_{load}	18 [pF]		
Shunt Capacitance	C_o		5 [pF] max	
Drive Level	D_L		100 [µW] max	
Aging	Df/f_o	±5 ppm per year	±5 ppm per year max	
Calibration Mode		Parallel		
Insulation Resistance			500 [MΩ] min	@ 100 Vdc

*Note: The selected Crystal Component must meet or exceed the specified drive Level (D_L). The crystal Component datasheet should show 100uW or more.

1.7.1 Quartz Crystal

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

1.7.2 Vibrational Mode

Crystals are available in both fundamental and third overtone. Unless there is a special need for third overtone, use fundamental mode crystals. At any 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.

1.7.3 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.

1.7.4 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). The Commercial operating temperature range is 0 °C to +70 °C. The Industrial operating temperature range requires -40 °C to +85 °C. Some vendors separate operating temperatures from temperature stability. Manufacturers can also list temperature stability as 50 ppm in their Datasheet.

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.

1.7.5 Calibration Mode

The terms series-resonant and parallel-resonant are 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.

1.7.6 Reference Crystal Circuit

Figure 1-15 shows a typical schematic of the I225/I226 oscillating circuit. Pin X1 and X2 refers to XTAL1 and XTAL2 in the I225/I226, 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.

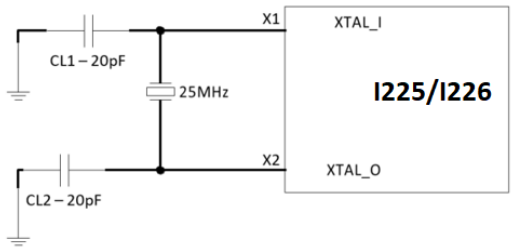

Figure 1-15 Crystal Circuit

1.7.7 Crystal Load Capacitance

The formula for crystal load capacitance is as follows:

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

where C1 = C2 = 20 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 10 pF accounts for lumped stray capacitance.

The calculated crystal load capacitance is 18 pF with C1 = C2 = 20pF and an estimated stray capacitance of about 8 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. Thin circuit boards generally have higher stray capacitance than thick circuit boards.

The oscillating 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 oscillating frequency decreases.

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

1.7.8 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 5 pF.

1.7.9 Equivalent Series Resistance (ESR)

ESR is the actual 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 60 Ω or better.

1.7.10 Driver 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.

I225/I226 Ethernet controllers can drive crystals up to 100 mW value. Therefore, the crystal datasheets should support at least of 100 mW. Otherwise, the silicon will over drive the crystal.

1.7.11 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.

1.7.12 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 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 is perfectly centered at the desired target frequency.

1.7.13 Reference Crystal Selection

There are several methods available for choosing the appropriate reference crystal. These are listed below.

- If a Saunders and Associates (S&A) crystal network analyzer is available, 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 is 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, 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 might 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 C_{Load} capacitance.

When choosing a crystal, keep in mind that to comply with IEEE specifications for 100/ 1000/ 2.5G BASE-T Ethernet LAN, the transmitter reference frequency must be precise within ±50 ppm. Intel recommends using a transmitter reference frequency that is accurate to within ±10 ppm to account for variations in crystal accuracy due to crystal manufacturing tolerance.

1.7.14 Circuit Board

Since dielectric layers of the circuit board are allowed some reasonable variation in thickness, stray capacitance from the printed board (to the crystal circuit) also vary. If the thickness tolerance for the outer layers of dielectric are controlled within ±17 percent of nominal, 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 increases 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.

1.7.15 Temperature Changes

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

1.8 Leakage Protection Circuit

Normally, board designers use the LAN_DISABLE_n pin to disable the I225/I226. However, there are some
cases where board designers prefer to use a FET to cut off the +3.3V_LAN power supply.
When the I225/I226 is completely powered down, there might be leakage on these five pins:

- Pin 3 — PE_CLKREQ_N
- Pin 4 — ULP_WAKE_N
- Pin 10 — SMB_DATA
- Pin 11 — SMB_CLK
- Pin 12 — PE_WAKE_N

To avoid leakage, a basic logic MOSFET can be setup to prevent the leakage path. Refer to the I225/I226 Design Checklist for more detail.

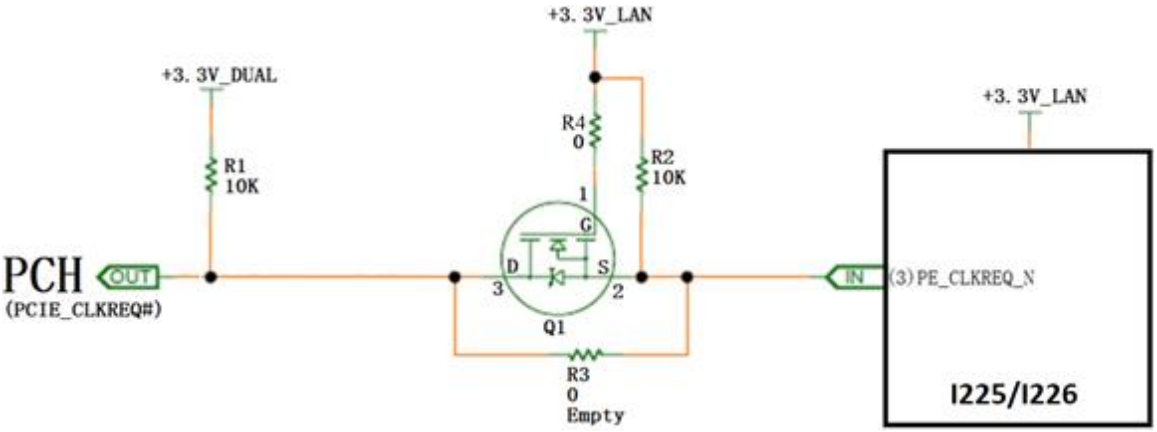

Figure 1-16: Leakage Protection Circuit

1.9 Voltage Shifter Circuit

Intel PCH I/O might require the new lower 1.8V level. However, I225/I226 I/O can only support 3.3V voltage level. We will need the voltage level shifter to connect these I/O pins together.

There are 2 different voltage level shifter solutions:

1. **Using the basic logic MOSFET device**: This circuit is very similar to the leakage current protection (as shown on Figure 1-16. Leakage Protection Circuit)
2. **Using the Voltage Shifter Bus isolation device**.

See Intel RVP Platform Schematic or I225/I226 Reference Schematic for details.

1.10 System Thermal Solution

In the I225/I226 Datasheet, we have the power consumption table. System Designers should measure all speeds/conditions based on the product requirements. Analyze the collected data to define the system thermal solution for I225/I226 silicon.

By specification, the Operating Tjunction should never access the Tjunction Maximum listed in the datasheet. The most common way is to use the Surface Thermocouples device to collect the silicon Top Case Temperature, and then calculate to project the junction temperatures. As a good rule of thumb, the measured Case Temp should not go over 100°F in all operating conditions for the production quality. If the design accessed this condition, designers should consider improving the system thermal solution, by adding heatsink and/or fan.

Please consult with thermal engineers for further suggestions.

1.11 PCB Guidelines

This section describes the general PCB design guidance targeted as supplementary information in addition to the specific requirements and recommendations for individual interfaces. For items not directly covered in the specific interface sections, these guidance recommendations apply to the design.

1.11.1 Via Usage

Use vias to optimize signal integrity. Figure 1-17 shows correct via usage. Figure 1-18 shows the type of topology that should be avoided.

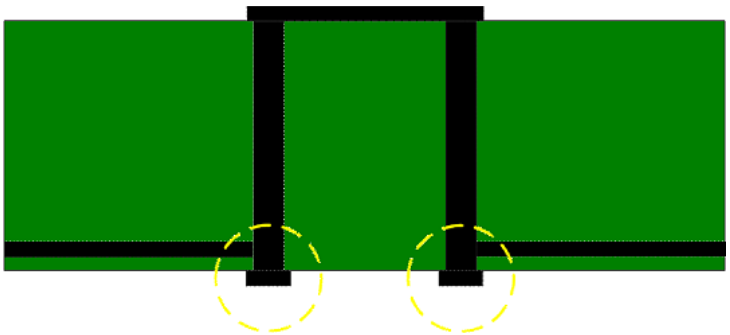

Figure 1-17 Correct Via Usage

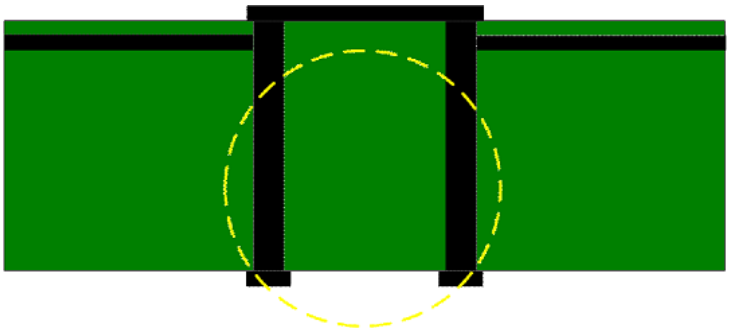

Figure 1-18 Incorrect Via Usage

Any via stubs on the MDI differential signal traces must be less than 35 mils in length. Keeping MDI signal via stubs less than or equal to 20 mils is preferable.

Place ground vias adjacent to signal vias used for the MDI interface. DO NOT embed vias between the high-speed signals, but place them adjacent to the signal vias. This helps to create a better ground path for the return current of the AC signals, which also helps address impedance mismatches and EMC performance.

It is recommended that, in the breakout region between the via and the capacitor pad, target a Z0 for the via to capacitor trace equal to 50 Ω. This minimizes impedance imbalance.

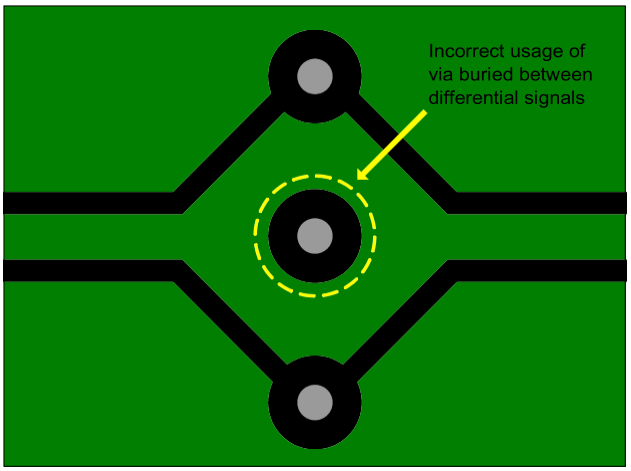

Figure 1-19 Incorrect Via Usage for Differential Pair

1.11.2 Reference Planes

Do not cross plane splits with the MDI high-speed differential signals. This causes impedance mismatches and negatively affects the return current paths for the board design and layout. See Figure 1-20.

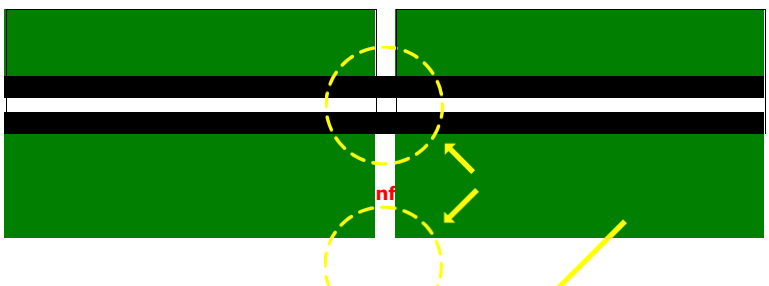

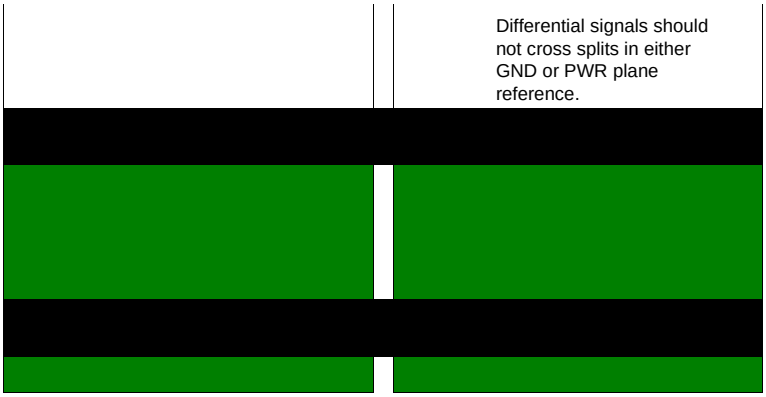

Figure 1-20 Improper Differential Signal Routing - Plane Split

Traces should not cross power or ground plane splits if at all possible. Traces should stay seven times the dielectric height away from plane splits or voids. If traces must cross splits, capacitive coupling should be added to stitch the two planes together to provide a better AC return path for the high-speed signals. To be effective, the capacitors should have low ESR and low equivalent series inductance.

Note: Even with plane split stitching capacitors, crossing plane splits is extremely high risk for 2.5G BASE-T designs

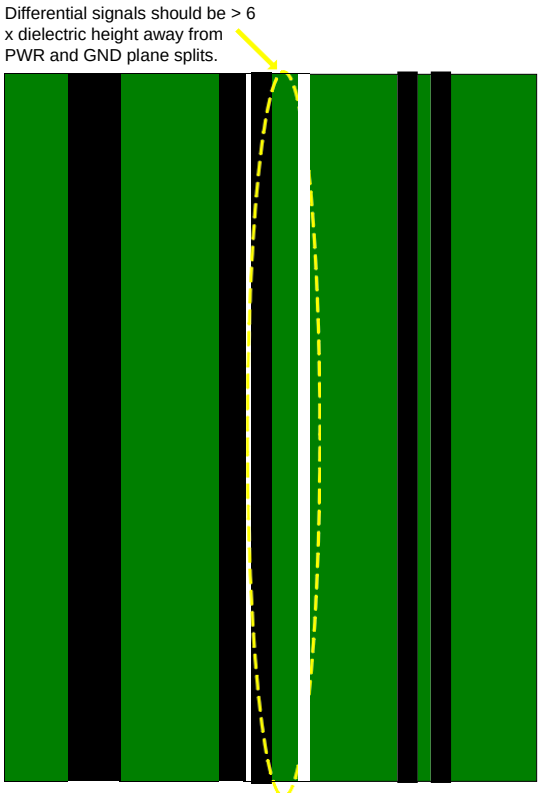

Figure 1-21 Differential Signal Routing - Plane Split and Void Proximity

It is recommended that the MDI signals stay at least seven times the dielectric height away from any power or ground plane split. This improves impedance balance and return current paths.

If a high-speed signal needs to reference a power plane, ensure that the height of the secondary (power) reference plane is at least 3 x the height of the primary (ground) reference plane.

1.11.3 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. Routing over a void in the reference plane causes impedance mismatches and usually increases radiated noise levels. Noisy logic grounds should NOT be located near or under high-speed signals or near sensitive analog pin regions of the LAN silicon. If a noisy ground area must be near these sensitive signals or IC pins, ensure sufficient decoupling and bulk capacitance in these areas. Noisy logic and switching power supply 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 should be equally potential power planes. 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 while still meeting the relevant electrical requirements 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 results in a smaller loop area and reduces the likelihood of crosstalk. The effect of different configurations on the amount of crosstalk can be studied using electronics modeling and simulation software.

1.11.4 Signal Isolation

To maintain the best signal integrity, keep digital signals far away from the analog traces. A good rule to follow is no digital signal should be within 7x to 10x dielectric height of the differential pairs. If digital signals on other board layers cannot be separated by a ground plane, they should be routed at a right angle (90 degrees) to the differential signal traces. If there is another Ethernet controller on the board, take care to keep the differential pairs away from that circuit. The same thing applies to switching regulator traces.

Rules to follow for signal isolation:

- Separate and group signals by function on separate board layers if possible. Maintain a separation that is at least seven times the thinnest adjacent dielectric height between all differential pairs (Ethernet) and other nets, but group associated differential pairs together.
- Over the length of the trace run, each differential pair should be at least seven times the thinnest adjacent dielectric height away from any parallel signal traces.
- Physically group together all components associated with one clock trace to reduce trace length and radiation.
- Isolate other I/O signals from high-speed signals to minimize crosstalk. Crosstalk can increase radiated EMI and can also increase 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.

1.11.5 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 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.

1.11.6 Power and Ground Planes

Good grounding requires minimizing inductance levels in the interconnections and keeping ground returns short, signal loop areas small, and locating decoupling capacitors at or near power inputs to bypass to the signal return. This significantly reduces EMI radiation.

These general following guidelines help reducing circuit inductance:

- 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. Routing signals over power or ground voids increases inductance and increases radiated EMI levels.
- Use distance and/or extra decoupling capacitors to separate noisy digital grounds from analog grounds to reduce coupling. Noisy digital grounds can 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 minimizes 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.
- Do not route high-speed signals near switching regulator circuits.
- It's acceptable to put ground fill or thieving on the trace layers, but preferably not closer than 50 mils to the differential traces and the connector pins.
- If differential traces must be routed on another layer, the signal vias should carry the signal to the opposite side of the PCB (to be near the top of the PCB), AND if the high-speed signals are being routed between two connectors on the same board, before the signal traces reach the second connector, they must return to the original signal layer (before reaching the connector pin). This strategy keeps via stubs short without requiring back drilling.
- Each time differential traces make a layer transition (pass through a pair of signal vias), there must be at least one ground via located near each signal via. Two ground vias near each signal via is better. See Figure 1-22 and Figure 1-23.

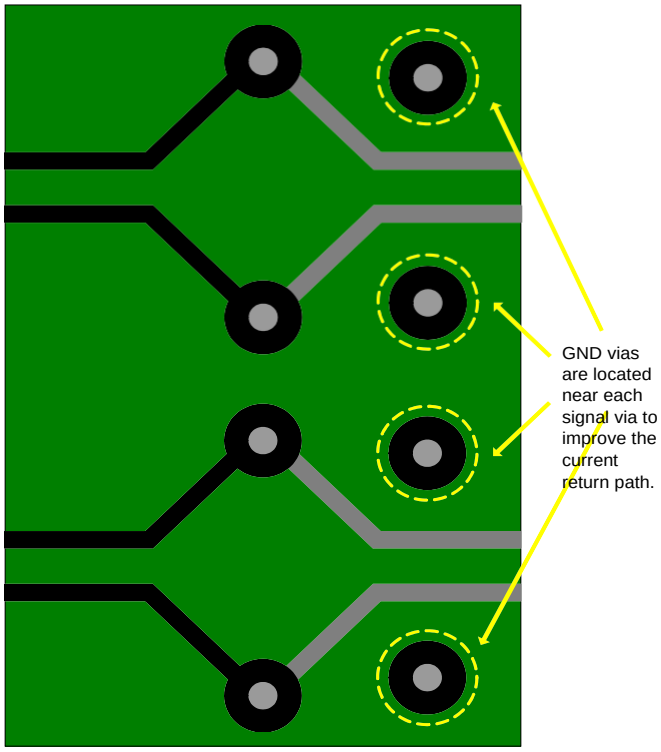

Figure 1-22 Return Path Vias for Differential Signals - Acceptable Example

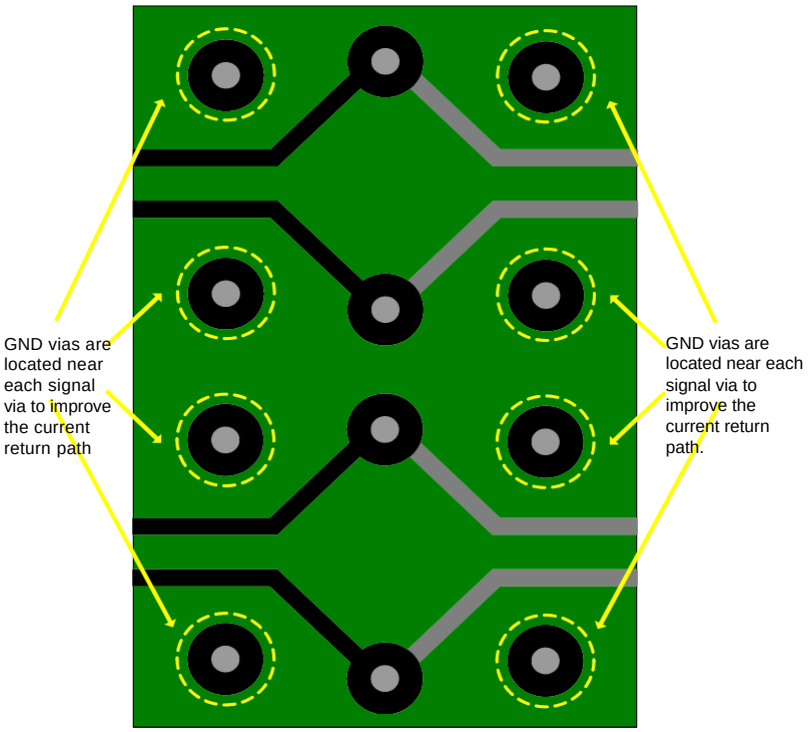

Figure 1-23 Return Path Vias for Differential Signals - Optimal Example

If the PCB fabrication process permits it, it's best to remove signal via pads on unconnected metal layers. See Figure 1-24 and Figure 1-25.

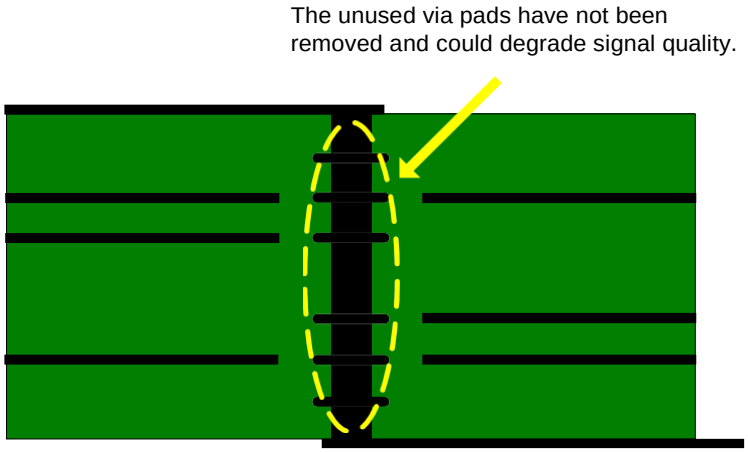

Figure 1-24 Signal Vias: Improper Padstack Example

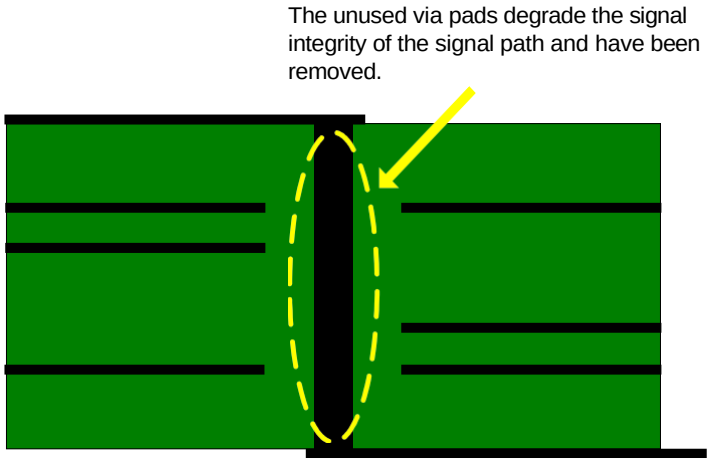

Figure 1-25 Signal Vias: Optimal Padstack

On metal layers where signal vias need to have via pads, it is desirable to reduce capacitance between the signal vias and ground-plane layers. The anti-pad diameters should be up to 20 mils larger than the via pad diameters. See Figure 1-26. Clearance between the pad and the surrounding metal should be ≥ 10 mils.

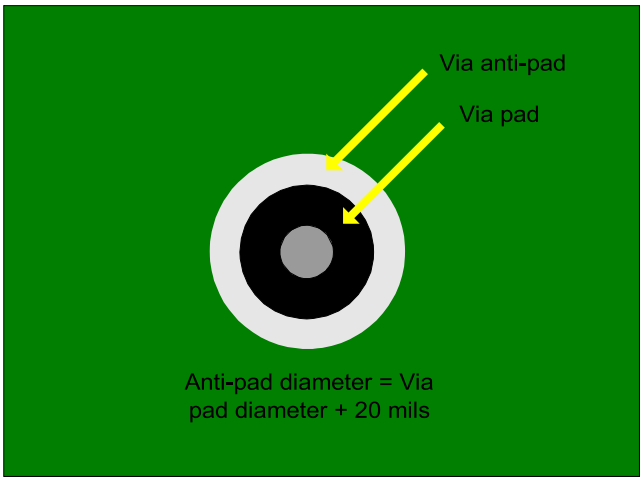

Figure 1-26 Anti-Pad Geometry

Each time differential signal vias pass through a plane layer, within each differential pair, the anti-pads should overlap. See Figure 1-27, Figure 1-28 and Figure 1-29.

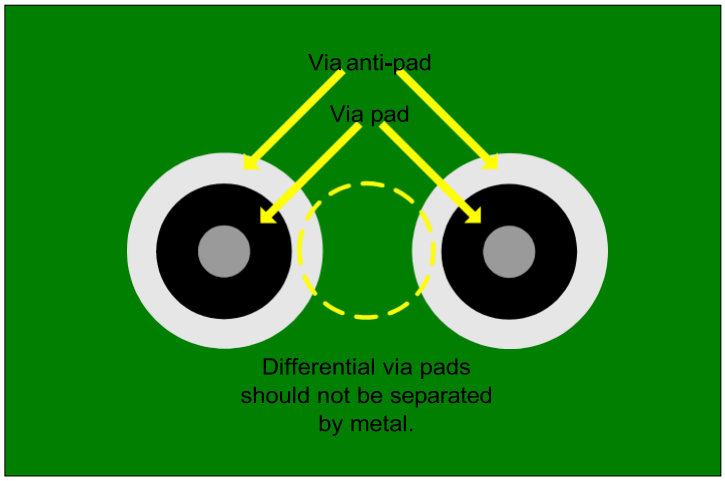

Figure 1-27 Differential Signal Vias: Improper Anti-pad Geometry

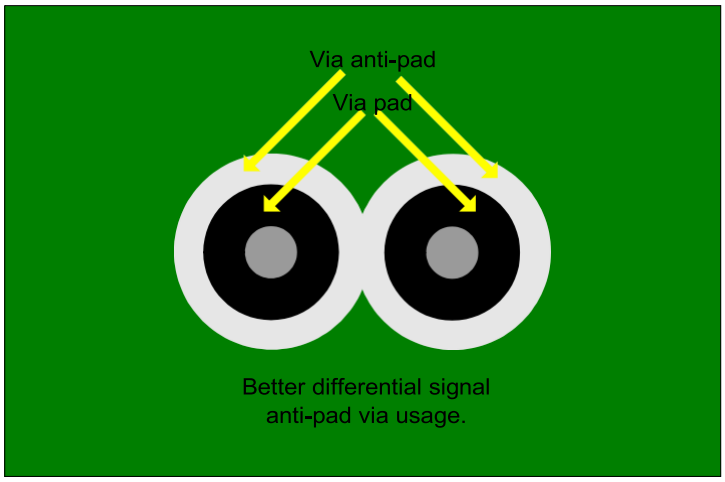

Figure 1-28 Differential Signal Vias: Acceptable Example

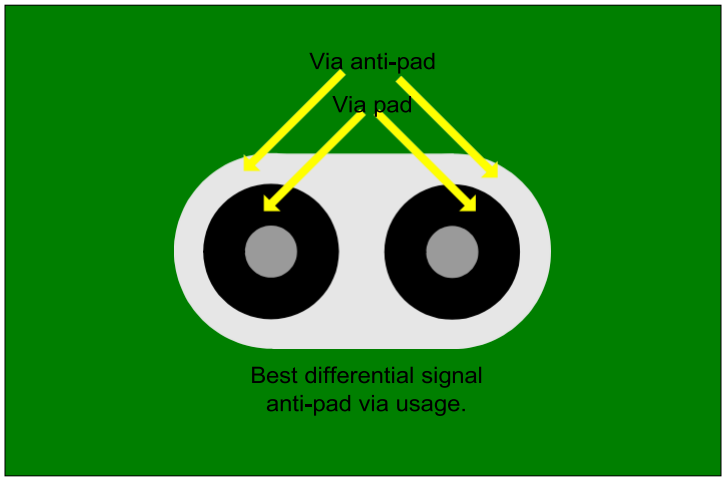

Figure 1-29 Differential Signal Vias: Optimal Example

1.11.7 Differential High Speed Layout Guideline

The following shows the Good vs Bad Practices for typical differential high-speed signals. This guideline should be applied when routing the 4x MDI pairs and PCIe TX/RX pairs of I225/I226.

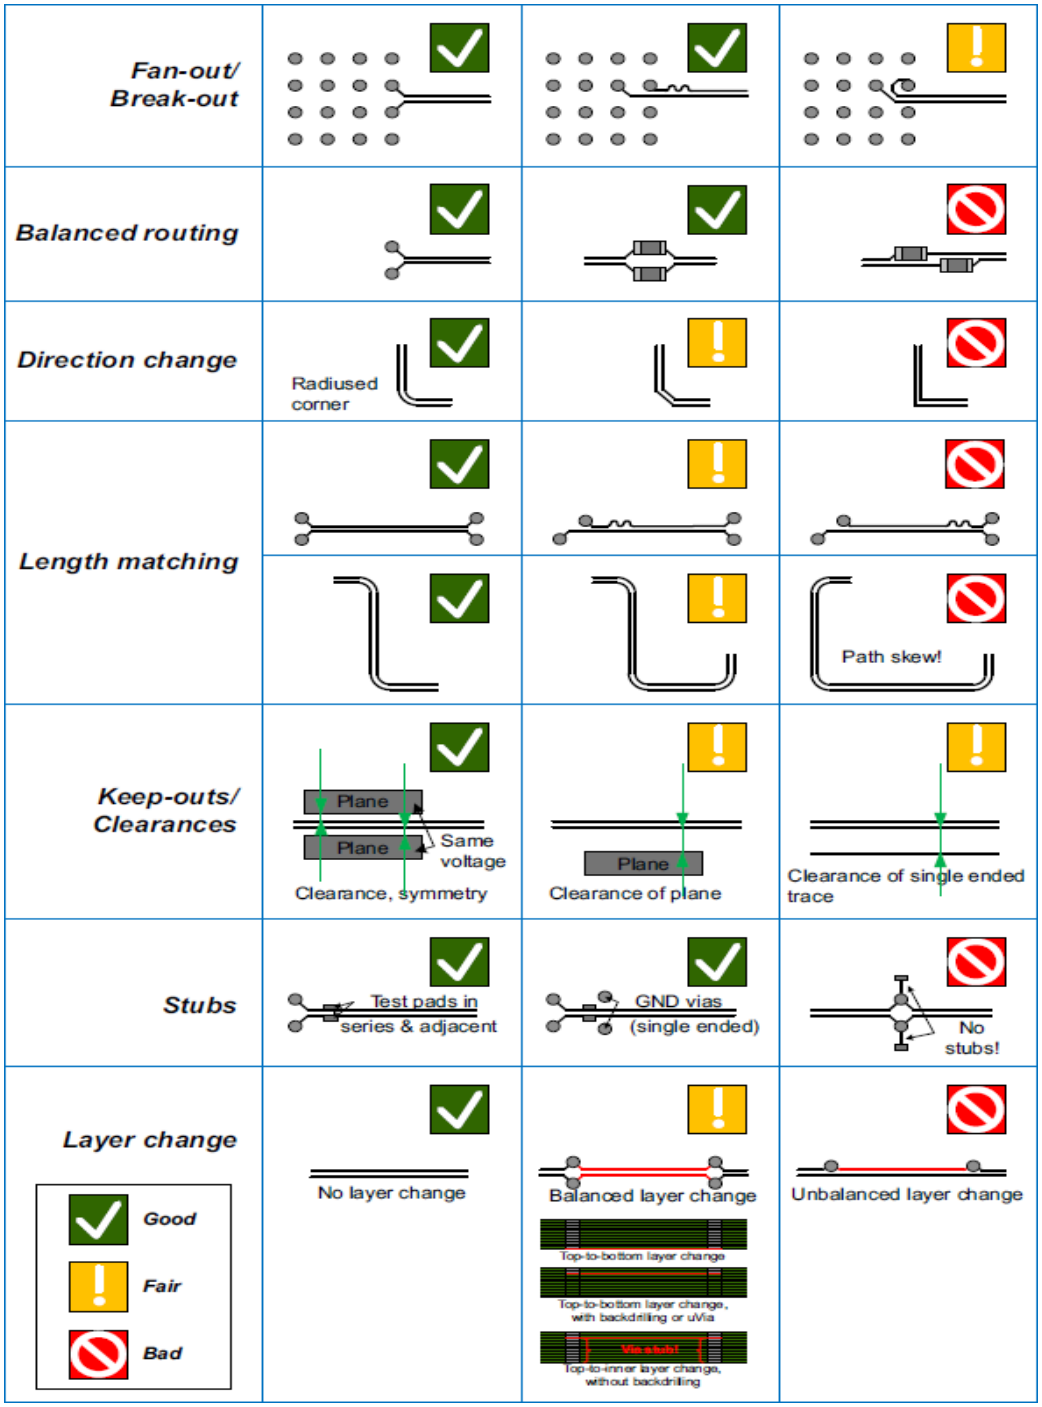

1.12 Component Selection (Bill Of Material)

Listed parts have been used successfully in the past designs. Note that no particular part is required to use with the I225/I226 silicon. It is recommended to start with the known quality parts, which Intel knows that they work. At a high-level, any part within the I225/I226 specifications should work with the silicon.

Any selected parts must go through the EV validations for production design quality.

Note: Please see the recommended parts list in the I225/I226 Datasheet.

LEGAL