COHERENT

Application Note AN-2036

Frequently Asked Questions Regarding Finisar's 1000BASE-T SFPs

Finisar's 1000BASE-T SFP transceivers are based on the SFP Multi Source Agreement (MSA). They are compatible with Gigabit Ethernet and 1000BASE-T standards as specified in IEEE Std. 802.3:2005. They can also function as 10BASE-T and 100BASE-T transceivers, but only after programming over a 2-wire bi-directional serial interface.

This Application Note covers the most commonly asked questions about Finisar's 1000BASE-T SFPs.

1. **What versions are there of the Finisar 1000BASE-T SFP and how do they differ?**
   Finisar currently offers 6 versions of 1000BASE-T SFP as listed in the following table:

Part Number	RoHS	Case Temp	PHY Rev	RxLOS	Auto-Negotation Supported?
FCMJ-8520-3	5/6	0C–85C	B0	Yes	No
FCMJ-8521-3	5/6	0C–85C	B0	No	Yes
FCLF-8520-3	6/6	0C–85C	B0	Yes	No
FCLF-8521-3	6/6	0C–85C	B0	No	Yes
FCLF8520P2BTL	6/6	-40C–85C	B2	Yes	No
FCLF8521P2BTL	6/6	-40C–85C	B2	No	Yes

   The most critical distinction is between the "8520" series products (i.e., the FCMJ-8520-3, the FCLF-8520-3, and the FCLF8520P2BTL), and the "8521" series products (i.e., the FCMJ-8521-3, the FCLF-8521-3, and the FCLF8521P2BTL). This distinction is discussed in more detail in FAQ Question 11. For topics that apply equally well to either 8520 or 8521 products, we'll refer to "1000BASE-T SFP" products or occasionally to "Cu SFP" products.

   All 6 part number variations use a Marvell 88E1111 "PHY", as discussed in FAQ Question 5. The oldest "first generation" designs use the Rev. B0 of this PHY chip; the latest "second generation" uses Rev. B2. Distinctions in performance between first and second generation products are listed in FAQ Question 24.

2. **Are there special electrical supply considerations when using the Finisar 1000BASE-T SFP?**
   Yes. The Finisar 1000BASE-T SFP consumes a maximum of 1.20W under worst-case conditions (typical is ~ 1.05 W). The maximum current draw is 375 mA, compared to 300 mA for a standard SFP. The electrical supply of the host should be examined to ensure that it can handle these increased demands. The 1000BASE-T SFP is generally more sensitive to supply voltage (3.13 to 3.47V) than most optical transceivers (e.g., 3.0 to 3.6V typical).

3. **Will the extra length of the Finisar 1000BASE-T SFP cause mechanical interference problems in my system?**
   While there is a chance of interference problems, the probability is small. If you compare the length of a Finisar 1000BASE-T SFP to a fiber SFP, with their respective cables inserted, there is little difference in length. This is due to the strain relief features on LC connectors that protrude much farther than the insulation on a Cat 5 cable. An optical and 1000BASE-T SFP with cables inserted is shown in Picture 1 below.

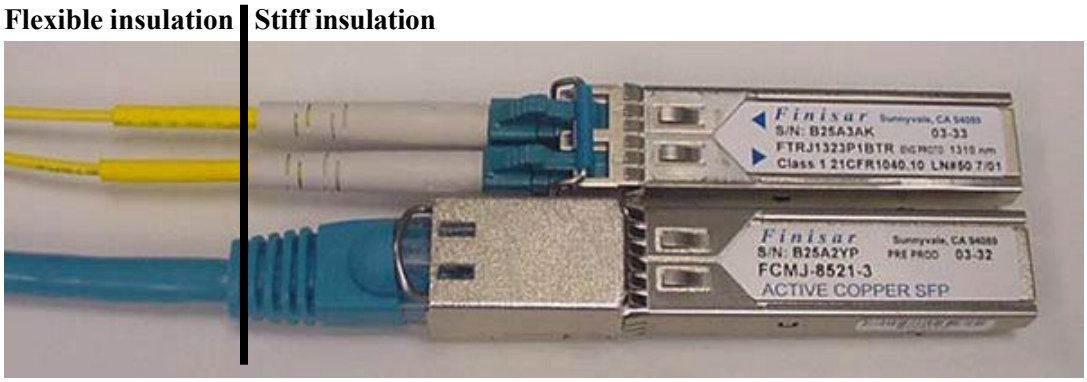

Picture 1: 1000BASE-T and optical SFP transceivers with cables inserted.

4. **What cable type is recommended for use with the Finisar 1000BASE-T SFP?**
   The Finisar 1000BASE-T SFP was designed to operate using standard Cat 5 cable that has been configured per IEEE802.3:2005. You can safely use Cat 5e or Cat 6 cables, as these are improved versions of Cat 5. For short cable runs, with intricate routing, stranded cable is recommended, as it is more flexible.

5. **What is the PHY and how can it be accessed?**
   The Finisar 1000BASE-T SFP uses the Marvell 88E1111 Physical Layer IC (PHY) to convert between the serial interface and 1000BASE-T interfaces. The first generation designs use Rev. B0 of this PHY; the second generation uses Rev. B2. This chip has a number of useful features available on internal registers that can be accessed via the 2-wire bi-directional serial interface at address 0xAC. The contents at each register address are 2 bytes wide; details for accessing the register can be found at http://www.marvell.com .

6. **What is SGMII mode?**
   SGMII stands for "Serial Gigabit Media Independent Interface." SGMII is a mode of communication between the MAC and PHY to allow for 10/100/1000BASE-T operation. In 100BASE-TX mode, the MAC still transmits to the PHY at 1.25 Gb/sec, but each byte is repeated 10 times. The PHY then converts this repeated data to 100BASE-TX format. The process is the same in 10BASE-T mode but each byte is repeated 100 times. SGMII is a Cisco Systems specification, document EN-46158, and is available for download at no cost over the internet: ftp://ftp-eng.cisco.com/smii/sgmii.pdf. At the time this FAQ was written, the latest SGMII version is 1.8.

7. **How do you configure the module for 10/100/1000BASE-T operation including RxLOS functionality at all three data rates?**
   Only the 8520 version has RxLOS functionality (pin 8 of the 20-pin SFP-to-host connector). For details on how to enable tri-rate functionality please refer to the table below. The Finisar 1000BASE-T SFP can be used with a SGMII rev. 1.8 interface (without clocks). This interface supports 10, 100 and 1000 BASE-T modes of operation, as mentioned above. The table below shows how to enable SGMII and advertise all speeds and full/half-duplex using register writes to the PHY over the 2-wire serial interface (see Question 5). A simple power cycle will return the module to default operation.

PHY Address: 0xAC		
Register Address	**Write Data**	**Description**
0x1B	0x9084	Enable SGMII mode
0x09	0x0F00	Advertise 1000BASE-T Full/Half-Duplex
0x00	0x8140	Apply Software Reset
0x04	0x0DE1	Advertise 100/10BASE-T Full/Half-Duplex
0x00	0x9140	Apply Software Reset
For RxLOS pin of 8520 products to function at 10/100/1000BASE-T		
0x18	0x4108	Enable RxLOS pin at 10/100/1000BASE-T

8. **What is auto-negotiation?**
   Auto-Negotiation is the communication or handshake between two remote devices to determine if the two devices can transfer data to one another and, if so, the specifics of data transfer such as data rate, flow control, and duplex traffic.

9. **What does 1000BASE-T or 1000BASE-X mean?**
   This designation is defined in IEEE802.3:2005. The "1000" in the designation refers to the transmission speed of 1000 Mbps. The "BASE" refers to BASE band signaling, indicating that only Ethernet signals are carried on the medium. The "T" represents twisted-pair copper cable (for example Cat 5), and the "X" represents fiber optic cable.

10. **What is the difference between 1000BASE-T and 1000BASE-X auto-negotiation?**
    1000BASE-T auto-negotiation is conducted over the Cat 5 cable between the two 1000BASE-T devices (see Figure 1 below). 1000BASE-X auto-negotiation is typically conducted between two host systems over fiber. In the case where 8521's are installed in the host systems, the 1000BASE-X auto-negotiation information is used to set the configuration options the 8521 advertises during 1000BASE-T auto-negotiation (see Question 12 below).

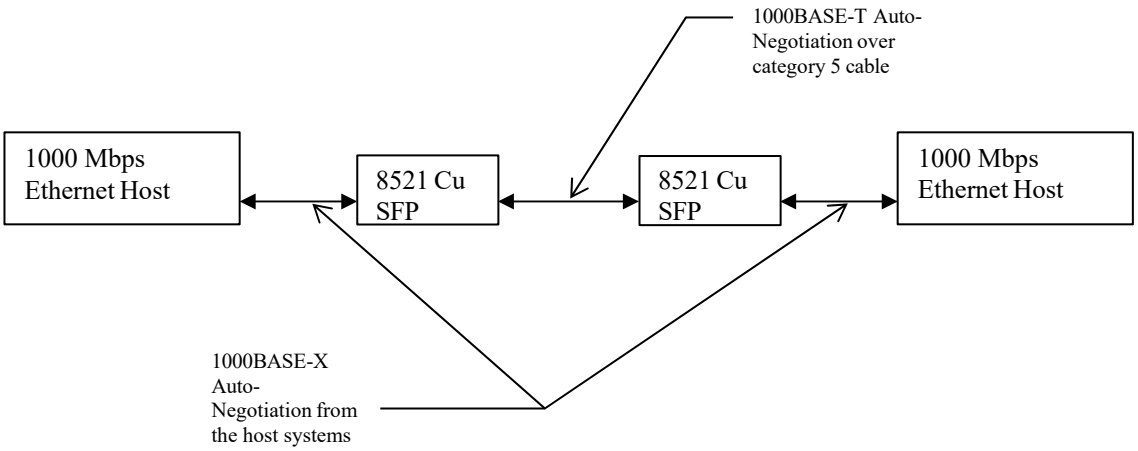

Figure 1: Link configuration of the 8521 version of the 1000BASE-T SFP.

11. What is the difference between the 8521 and the 8520 versions of the 1000BASE-T SFP?

The 8521 is designed to be a 1000BASE-X compliant device that functions properly in any slot designed for optical SFP's, with no hardware or software changes to the host system. The PHY is configured to perform 1000BASE-X auto-negotiation with the host system (see Question 12). The 8521 does not have link detection circuitry—the RxLOS pin 8 of the 20 pin SFP-to-host connector is internally grounded.

The 8520 is similar to the 8521, but the 1000BASE-X auto-negotiation is disabled. For this device to operate, auto-negotiation should also be disabled in the host system. In the 8520, the RxLOS pin functions as a link indicator—RxLOS is asserted when the 1000BASE-T link is lost. Typically, this device is used in systems where the host cannot determine the link status without using RxLOS as a link indicator.

12. Does the 8521 support 1000BASE-X auto-negotiation from the host system?

The 1000BASE-X auto-negotiation from the host systems is supported and should be enabled when using the 8521. This is not the case for the 8520, in which auto-negotiation should be disabled in the host system or a link cannot be established.

13. How do I disable 1000BASE-X auto-negotiation and 1000BASE-T auto-negotiation on the 8521?

1000Base-X Disable;

PHY Address: 0xAC		
Register Address	**Write Data**	**Description**
0x1B	0x808C	Change HWCFG_MODE to non-GBIC
0x00	0x8140	Apply Software Reset

1000Base-T Disable;

PHY Address: 0xAC		
Register Address	**Write Data**	**Description**
0x16	0x0001	Select Fiber Register Bank
0x00	0x8140	Disable Auto-negotiation
0x16	0x0000	Return to Cu Register Bank

14. How do I enable 1000BASE-X auto-negotiation and 1000BASE-T auto-negotiation on the 8520?

If the system using an 8520 is compatible with 1000BASE-X auto-negotiation, this feature can be enabled in the 8520 as shown in the below table. Please note that depending on how the RxLOS signal is used by the system, enabling this feature might cause an interface problem with the host system.

1000Base-X Enable;

PHY Address: 0xAC		
Register Address	**Write Data**	**Description**
0x1B	0x9088	Change HWCFG_MODE to GBIC
0x00	0x8140	Apply Software Reset

1000Base-T Enable;

PHY Address: 0xAC		
Register Address	**Write Data**	**Description**
0x16	0x0001	Select Fiber Register Bank
0x00	0x9140	Enable Auto-negotiation
0x16	0x0000	Return to Cu Register Bank

15. Does the host system need to support 1000BASE-T auto-negotiation?

No, the 1000BASE-T auto-negotiation is fully supported by either module (8520 or 8521). The 8521 uses the 1000BASE-X auto-negotiation information it receives from the host to adjust the configuration options that it advertises during copper auto-negotiation.

16. When using the 8521, does the host know that it is driving a copper transceiver?

No, to the host system, Finisar's 8521 will appear to be a fiber transceiver. From the host's perspective, acknowledgements received during the 1000BASE-X auto-negotiation process are coming from the remote link partner even though they are coming from the 8521.

17. When using the 8521, what keeps the 1000BASE-X auto-negotiation from finishing before the 1000BASE-T auto-negotiation?

The 8521 will hold back acknowledgement to the host until the copper auto-negotiation (1000BASE-T AN) is resolved, and then it will send the 1000BASE-X auto-negotiation advertisements and acknowledgement information. This way, the

host can complete 1000BASE-X auto-negotiation without knowing copper auto-negotiation was involved. This process will slow down 1000BASE-X auto-negotiation, because it will be delayed until the 1000BASE-T auto-negotiation is complete (<3 seconds).

18. How are abilities resolved between the host and the 8521 during auto-negotiation?

The module PHY is able to incorporate the host system's abilities in the cable-side auto-negotiation. More specifically, the module PHY first receives the abilities from the host system during 1000BASE-X auto-negotiation, but the 1000BASE-X auto-negotiation is not yet allowed to complete and no abilities are passed to the host system. The PHY then incorporates the host system's abilities into the 1000BASE-T auto-negotiation. Finally, the result of the 1000BASE-T auto-negotiation is passed on to the host system and the 1000BASE-X auto-negotiation is allowed to complete. The resolution of the 1000BASE-T auto-negotiation will therefore be compatible with the resolution of the 1000BASE-X auto-negotiation.

19. Since RxLOS is grounded on the 8521, what indication is there to determine whether the link is functioning or not?

If the host is receiving idles (/I/ ordered sets) from the 8521 then the link is good. If the host is receiving auto-negotiation coded words (/C/ ordered sets) or random data then the link is down. You can also access the PHY at register 0x11, bit 10 for real time link indication. 1=Link up, 0=Link down.

20. Your data sheet indicates that the 1000BASE-T SFP are set to "Preferred Master" mode. What does this mean?

In 1000 BASE-T, one of the link partners becomes the master, and the other becomes the slave. The master uses its local clock to transmit data on the Cat 5 cable, while the slave uses the clock recovered from the data received from its link partner. During 1000BASE-T auto-negotiation, the link partners agree to 10/100/1000 Mb/s operation, and determine who is the master and who is the slave. Forcing master or slave mode can cause conflicts, so 1000BASE-T SFPs are set to preferred master mode. If both link partners advertise the same preference, a pseudo random number generator will determine the master/slave choice.

21. The 1000BASE-T SFP data sheet indicates that TX_FAULT is not supported. What does this mean?

Pin 2 on the SFP connector is specified as the TX_FAULT output. In optical SFP's, the TX_FAULT pin is often used to indicate a possible eye-safety condition, and the pin is driven high under a set of transmitter conditions specified by the manufacturer. For electrical SFP's, this functionality is of no value. The TX_FAULT pin is permanently grounded on the 1000BASE-T SFP.

22. What is typical link up / link down response time, and can it be made shorter?

Customers are accustomed to short loss of link response times for optical transceivers (typically < 50 ms). This fast response is not possible with a 1000BASE-T transceiver. Link up should typically occur within 3 seconds; detecting link loss

typically takes less than 1 second. See the Marvell White Paper ***MV-S70000rr-00 time to link up and link down.pdf*** for more details.

23. Can you provide greater detail on PHY registers?

The Marvell datasheet for the 88E1111 is confidential, and you must register at the Marvell extranet at http://www.marvell.com to gain access.

24. What are the differences between first and second generation 1000BASE-T SFP products?

Second generation products (FCLF8520P2BTL and FCLF8521P2BTL) operate over wider case temperatures (–40C to +85C) compared with first generation products (0C to +85C). The chassis and signal grounds are tied together on first generation products, but electrically isolated from each other on second generation products. Second generation products generally do better on IEEE compliance tests. First generation products have known issues with 1000BASE-T transmitter distortion and MDI common mode output voltage tests. Otherwise first and second generation products exhibit very comparable performance, especially with respect to 2-wire serial interface programming issues.

25. Is external loopback possible for this module, and do you have any other recommendations to aid in testing and debugging?

Although the PHY does have a test mode for external loopback, the regular operation of the PHY must be severely modified to work with an external loopback cable. By design all 1000BASE-T PHYs cancel the effects of their own transmissions from the received signals, also called Near End Cross Talk (NEXT) canceling. If line A is connected to line B, the crosstalk between the lines becomes 100%, and the PHY will automatically cancel the entire signal. Therefore, to enable external loopback, the NEXT canceling must be disabled, which can then make debugging, and testing ineffective.

Finisar recommends using line loopback and internal loopback for testing and debugging purposes in 1000BASE-T mode. This mode still requires that the PHY is reconfigured, but not in a way that could disguise problems in the system.

26. How does Line Loopback mode work?

Line loopback allows a link partner to send frames into the PHY to test transmit and receive data paths. Frames sent from a link partner into the PHY, before reaching the MAC interface pins are looped back and sent out on the line side. The link partner thus receives its own frames.

To enable line loopback, the Finisar 1000BASE-T SFP must first establish copper link with another link partner. If auto-negotiation is enabled, both link partners should advertise the same speed and full duplex. If auto-negotiation is disabled, both link partners need to be forced to the same speed and full duplex. Once link is established, enable the line loopback mode by writing to register 0x14 bit 14.

0x14 bit 14 = 1 (Enable line loopback)
0x14 bit 14 = 0 (Disable line loopback)

Once the line loopback is enabled, the link partner can send data into the PHY.

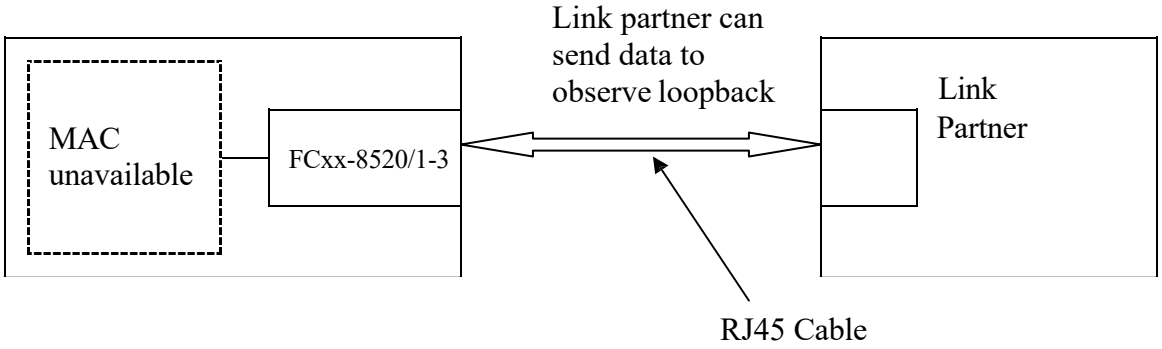

Figure 2: Line Loopback Setup

27. How do you configure the module for internal loopback testing?

Internal loopback testing is a quick way to check the integrity of the MAC to SFP connections. Data sent from the host system to the SFP is looped back internally in the SFP's PHY (data sent in to the TX pin of the SFP are looped out onto the RX pins). This requires register writes to the PHY to put it in loopback mode. For the internal loopback operation, fiber auto-negotiation should be disabled, both on the PHY and on the MAC. This is because the SFP mode of operation requires a handshake between the fiber auto-negotiation on the MAC side, and the copper auto-negotiation on the RJ45 side. Since the loopback mode automatically disables the receive functionality on the copper side, it will be impossible to carry out this handshake that is required by the SFP mode of operation.

To enable loopback, first disable 1000BASE-X auto-negotiation per Question 13. Then, set register 0x00 bit 14 = 1.

After the loopback test is completed, internal loopback can be disabled by setting register 0x00 bit 14 = 0 and 1000BASE-X auto-negotiation can be re-enabled per Question 14.

28. How are special compliance test patterns programmed with the Marvell PHY?

1000BASE-T test modes can be programmed as discussed in Clause 40 Table 40-7 of IEEE 802.3.

Test Mode	Register 0x09 Contents
Normal operation	0x 0E 00
1	0x 2E 00
2	0x 4E 00
3	0x 6E 00
4	0x 8E 00

100BASE-T compliance tests can be performed by the following 2-wire serial interface commands. In the absence of any incoming signals from the host side, the module will transmit a series of IDLES. Because of scrambling, the resulting signal on the CAT5 cable is essentially a pseudo-random bit sequence suitable for compliance tests.

Register	Contents	Comments
0x1B	0x 90 84	Enable SGMII mode
0x00	0x A1 00	Apply 100BASE-T software reset
0x10	0x 00 11	Disable auto MDI / MDIX
0x00	0x A1 00	Apply 100BASE-T software reset

10BASE-T tests require either a link pulse, a pseudo-random bit pattern, or for the harmonic distortion test a square pulse sequence. To program a 10BASE-T ***link pulse***, issue the following commands:

Register	Contents	Comments
0x1B	0x 90 8F	Enable SGMII mode; GMII to copper
0x00	0x 81 00	Apply 10BASE-T software reset
0x10	0x 04 11	Disable auto MDI / MDIX; force link good
0x00	0x 81 00	Apply 10BASE-T software reset

A 10BASE-T ***pseudo-random*** pattern is available from the Marvell PHY by issuing the following commands:

Register	Contents	Comments
0x1B	0x 90 8F	Enable SGMII mode; GMII to copper
0x00	0x 81 00	Apply 10BASE-T software reset
0x10	0x 04 11	Disable auto MDI / MDIX; force link good
0x00	0x 81 00	Apply 10BASE-T software reset
0x1D	0x 00 12	Select packet generator
0x1E	0x 89 20	Enable packet generator

The harmonic content test requires that a square pulse signal be transmitted over the CAT5 cable. A square pulse test pattern is not available from the Marvell PHY as a built-in test pattern, and must be provided externally (e.g., by a "SmartBits" tester or equivalent).

29. Are there any known problems with the Marvell PHY??

For the latest information on known problems with either Rev. B0 or Rev. B2 of the Marvell PHY, consult the Marvell website for the 88E1111 "release notes". At the time of writing, the latest version of the Rev. B0 release notes is ***MV-S300395-00N_88E1111_RevB0_RN.pdf***, and the latest version of the Rev. B2 release notes is ***MV-S300723-00E_88E1111_RevB2_RN.pdf***.

At the time of writing here is a partial list of problems that are associated with the current rev. B0 silicon of the PHY:

a. In SGMII mode, if the link partner is a 10Mb repeater hub, the Marvell PHY may potentially send frames to the MAC with an alignment error. This issue does not occur with a NIC or switch in 10Mb mode. It will only occur with link partners that are 10Mb repeater hubs.
b. Gigabit template testing of point A/B and C/D symmetry specified by the IEEE test mode may be slightly outside the limits specified in IEEE 802.3ab. Minor symmetry violations have no effect on performance, other than a possible effect in cable length performance. Finisar's 1000Base-T transceivers are tested with cables >100m in length, to make sure the parts operate error free at the max cable length.
c. The 100BASE-TX fall time may marginally exceed the IEEE802.3 limit of 5 ns (max). As a result the rise/fall time symmetry may marginally exceed the 0.5ns difference that is allowed.

At the time of writing here is a partial list of problems that are associated with the current rev. B2 silicon of the PHY:

a. The 100BASE-TX fall time may marginally exceed the IEEE802.3 limit of 5 ns (max). As a result the rise/fall time symmetry may marginally exceed the 0.5ns difference that is allowed.

30. What is the initialization sequence to optimize the pulse shape and improve BER?

PHY Address: 0xAC		
Register Address	**Write Data**	**Description**
0x1d	0x0006	Designate internal register
0x1e	0x4200	Write to internal register
0x1d	0x000A	Designate internal register
0x1e	0x0001	Write to internal register
0x00	0x9140	Software reset

31. Can the SGMII auto-negotiation be disabled?

Once the module is in SGMII mode, the SGMII auto-negotiation can be disabled using the register writes shown in the table below.

PHY Address: 0xAC		
Register Address	**Write Data**	**Description**
0x16	0x0001	Select register bank for host system side
0x00	0x8140	Turn off auto-negotiation and reset
0x16	0x0000	Return to register bank for cable side

32. When two modules are connected across the serial interface, are there any additional register writes required?

In this configuration, the data flows from the RD+/- lines of one module to the TD+/- lines of the other module, and vice versa.

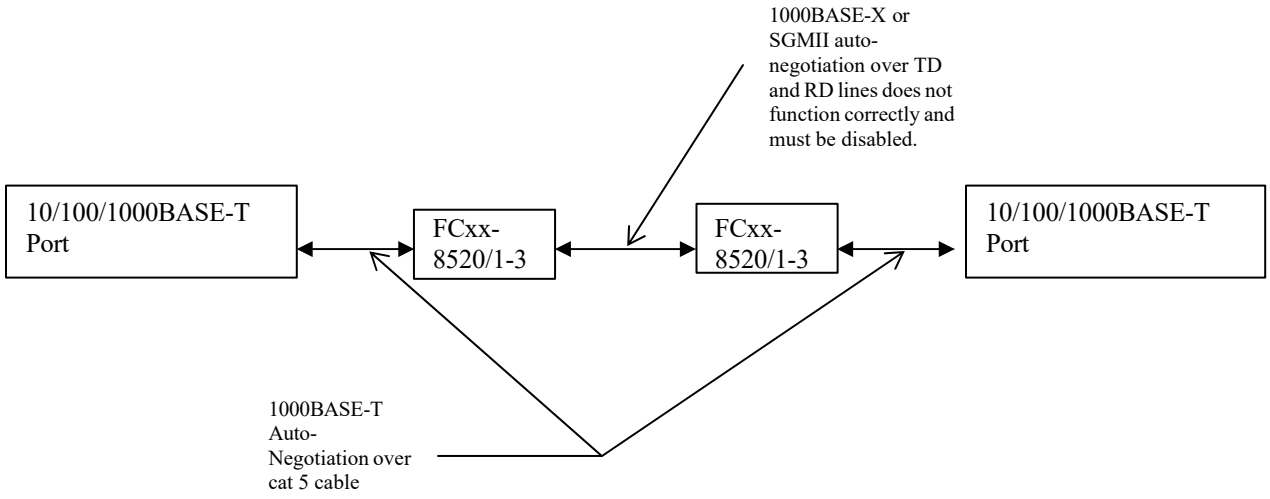

Figure 3: Connecting two modules across the serial interface

When using the FCxx-8521-3 in 1000BASE-T mode, the 1000BASE-X auto-negotiation must be disabled because of the intertwining of the 1000BASE-T and the 1000BASE-X auto-negotiations (see Question 13). Because the resolution of the 1000BASE-T auto-negotiations might result in the modules' having incompatibilities, the abilities of each module should be forced to disable flow control and only enable full-duplex. And, when the 1000BASE-T link is disconnected from one of the modules, the other 1000BASE-T has no way of knowing. The system must handle this situation by monitoring the PHY registers (see Question 19).

When using the FCxx-8520-3 in 1000BASE-T mode in this configuration, each module will link to its partner over the cat 5 cable because the 1000BASE-X auto-negotiation is already disabled (see Question 10). Data will pass between the modules without additional register writes. Again, when the 1000BASE-T link is disconnected from one of the modules, the other link has no way of knowing, and the system must handle this situation by either monitoring the RxLOS pin or the PHY registers (see Question 19).

When using the FCxx-8520/1-3 in SGMII mode, the SGMII auto-negotiation must be disabled (see Question 31). Again, the abilities of each module should be forced to disable flow control and only enable full-duplex. Again, when the link across the cat 5 cable has been disconnected from one of the modules, the other link has no way of knowing. Finally, modules can be configured for operation at 10/100/100BASE-T (see Question 7)

For more details about monitoring the PHY registers, setting the abilities and speeds of each module, please see the Marvell documentation (see Question 23).

33. Can the PHY registers be accessed when the module TX_DISABLE is asserted?
No. Asserting TX_DISABLE puts the module into its hardware reset state, and the PHY cannot be accessed until the TX_DISABLE has been negated and the reset has completed.

34. Are there any other considerations that need to be taken into account when laying out the host board interface?
The RxLOS pin on an FCxx-8520-3 is TTL, whilst on SFP modules this pin is an open collector/drain output.