PMC Gigabit Ethernet UTP-5 PHY Proposal

Frequently Asked Questions

Version: 07 Nov. 1996

1. PMC's UTP-5 proposal uses 4 pairs of UTP-5 cable. How does this work with standard UTP-5 cable used in Fast Ethernet?

Standard UTP-5 cable bundles consist of 4 twisted pairs. In Fast Ethernet, two pairs are used for transmit and receive, and two are terminated to reduce EMI. PMC's scheme uses all four pairs (for 50 meters), two pairs *each* for transmit and receive.

2. Is this a Jumper Interconnect, a 50 m., or a 100 m. Solution?

The low cost inherent with the PMC proposal makes it a good candidate for short jumper interconnects within an equipment room. A distance of 50 meters[1] is feasible over 4 pairs of UTP-5 cable. 100 meters can be reached with 8 pairs of UTP-5 for critical near-term applications that require a longer distance.

3. Why 50 meters?

PMC believes that widespread market acceptance of Gigabit Ethernet depends on the timely availability of low-cost, multiple-vendor solutions. Time-to-market and complexity are important factors to consider. Proposed solutions seen up to now for reaching 100 meters involve an inordinate increase in complexity out of proportion for a twofold increase in distance. PMC's approach emphasizes time-to-market, re-use of existing technology, low cost and complexity over distance reach. We believe this is the sensible approach to take until a satisfactory, cost-effective solution for 100 meters over 4 pairs of UTP-5 can be found.

[1] Corresponding to 70% of installed backbone-server links according to a market study conducted by the Dell'Oro Group for the Gigabit Ethernet Alliance.

4. Can 100 meters be reached?

Yes, using 8 pairs (2 cable bundles) of UTP-5.

5. Is this proposal intended for structured horizontal wiring?

Yes, within the given distance limit for the number of cable bundles. To reiterate, the horizontal reach is 50 meters or less for 70% of existing backbone-server links.

6. Does the PMC scheme work across punchdown blocks?

Yes. In addition to the length of UTP-5 cable, our proposal has been tested with two 3-meter patch cords that include two punchdown blocks.

7. Will the PMC scheme be forward compatible to a 100 meter 4-pair solution?

Yes. That is the hope and intention. The market need for a 100 meter solution is still two to three years out; it is expected that this will be sufficient time for the state of technology and the standards process to deliver an acceptable solution by then. In the meantime, the PMC proposal provides a cost-effective solution for 100 meters (in the near term) and shorter reaches. We believe that a 50 meter proposal provides a very solid base from which 100 meters can be attacked. Furthermore, without this solid base, a 100 meter solution will likely be more than two to three years away.

8. What is 4LZS?

4LZS (4 Levels with Zero State) is a baseband modulation scheme that uses 4 amplitude levels to encode 2 data bits per transmitted symbol (as for example, in 2B1Q, used in ISDN), and one amplitude level (the zero state / amplitude level) for marking an escape sequence that precedes a control code. The bit mappings are as follows:

Bit Pattern	4LZS Level
00	-3
01	-1
11	+1
10	+3
(escape)	0

The data bits are mapped in such a way that symbols erroneously transformed to corresponding adjacent symbols (symbol errors most likely to occur) result in only one bit error.

9. Is 4LZS different from a general 5-level code?

There are 5 amplitude levels, but with 4LZS, the zero level is used only for control codes. Based on our current design, there is some gain in receiver sensitivity margin with 4LZS over a true 5-level code (all five levels used for data). A 5-level coding scheme merits further investigation. For example, 8B4P (8 bits encoded into four 5-level symbols) is being considered.

10. Has the PMC proposal been simulated?

Yes. A model was constructed, based on PMC's experience with 155 Mb/s implementations over UTP-5. Effects of increased data rate and changing from two to four amplitude levels (NRZ to 2B1Q) were factored in using known theoretical results to predict their impact on distance reach.

11. Has the PMC proposal been implemented?

Yes. Circuit board prototypes have been constructed for 2B1Q to verify the theoretical model and to characterize BER performance and FCC emissions. A distance of 50 meters has been reached in the lab, with negligible BER, while meeting FCC radiated emissions requirements. It appears that the BER objective of 10^{-10} can be attained with margin. Error-free runs at least 5 hours long have been observed (bounding the error rate to less than 10^{-13}).

12. Why not use 8B10B coding?

8B10B coding increases the transmission overhead by 25%. At gigabit rates, the resulting bandwidth requirements are likely beyond the limits of current standard CMOS technology. Moreover, the use of 8B10B coding would result in a shorter distance reach, all other parameters equal. Control codes *functionally* equivalent to their 8B10B counterparts are used in the PMC scheme.

13. In 100Base-T4, there were problems with transmitting over multiple pairs. Is PMC's proposal vulnerable in this regard as well?

The PMC scheme is full duplex, and hence, does not have the attendant problems of operating in half duplex (as in T4).

14. Does the PMC proposal meet FCC limits?

PMC recently brought the lab prototype to an FCC-certified field testing center to evaluate EM radiation characteristics over UTP-5 cable. Test results indicate that FCC Class B and CISPR 22 limits can be satisfied with margin.

15. PMC's proposal operates at frequencies above the 100 MHz limit to which UTP-5 cable is specified. Why should we believe that your proposal can work in an acceptable manner?

It is not uncommon for cables to be used beyond the frequency range to which they were originally intended. In ISDN and xDSL for example, the range exceeds far beyond the original 4 kHz limit used for analog voice (POTS) twisted pair telephone lines. With the PMC proposal, most of the spectrum is contained within 250 MHz. The question then is whether cable characteristics in the unspecified range of roughly 100 through 250 MHz can be so detrimental as to compromise data transmission characteristics. PMC has characterized two cables from AT&T and Belden out to 250 MHz. We have found that the cable characteristics behaved consistently with frequency into the unspecified range. Moreover, we have tested the prototype board with some cables and thus far have not encountered problems. Nevertheless, further work is needed to characterize more cable samples.

16. What is a "killer packet"?

Killer packet refers to a data packet that can bring down a link. The principal mechanism in which this can happen is when the bit pattern on the line is such that a long sequence of unchanging symbols are transmitted. In the case of a frame synchronous scrambler for example, a packet could be constructed to have a bit pattern coincident with the scrambling pattern. A long absence of transitions at the receiver can cause the clock recovery circuit to lose synchronization, causing subsequent packets to be lost until synchronization is restored.

17. Is there a "killer packet" problem with the PMC proposal?

No. Substitution control codes are used after scrambling in the PMC proposal to preclude the existence of long constant symbol runs.

18. What are the cost implications of the PMC proposal?

It is reasonable to expect that the PMC solution should cost approximately 2-3X more than a corresponding implementation for 100Base-T. This estimate is derived by noting that an implementation of the PMC proposal involves twice the number of line drivers and receivers from 100Base-T. Taken individually, each component for gigabit ethernet is only slightly more complex compared to an implementation for 100Base-T.

19. When does PMC believe they can have silicon samples?

We believe that it is possible in 1997.