

How SolarFlare Communications Broke the 10 Gbit/s-on-UTP Barrier

Solarflare Communications has developed an innovative PHY solution for standard UTP cable to implement 10 Gbit/s Ethernet transmission, and is leading the industry in driving for a 10-Gbit/s standard for twisted-pair copper, known by the IEEE as 10GBASE-T.

The key to this new solution is an algorithm that mitigates the host of impairments that degrade the communication channel over UTP wiring. These consist of impairments related to the propagation of the signals themselves, such as Insertion Loss and Inter-Symbol Interference [ISI], which are caused by the limited bandwidth and real impedance of the cable itself, plus degradations due to interference, such as echo, near-end crosstalk [NEXT] and far-end crosstalk [FEXT]. In addition, background noise and other radiated signals, such as “Alien” NEXT (NEXT from other cables) can reduce the received SNR (see Fig. 1.)

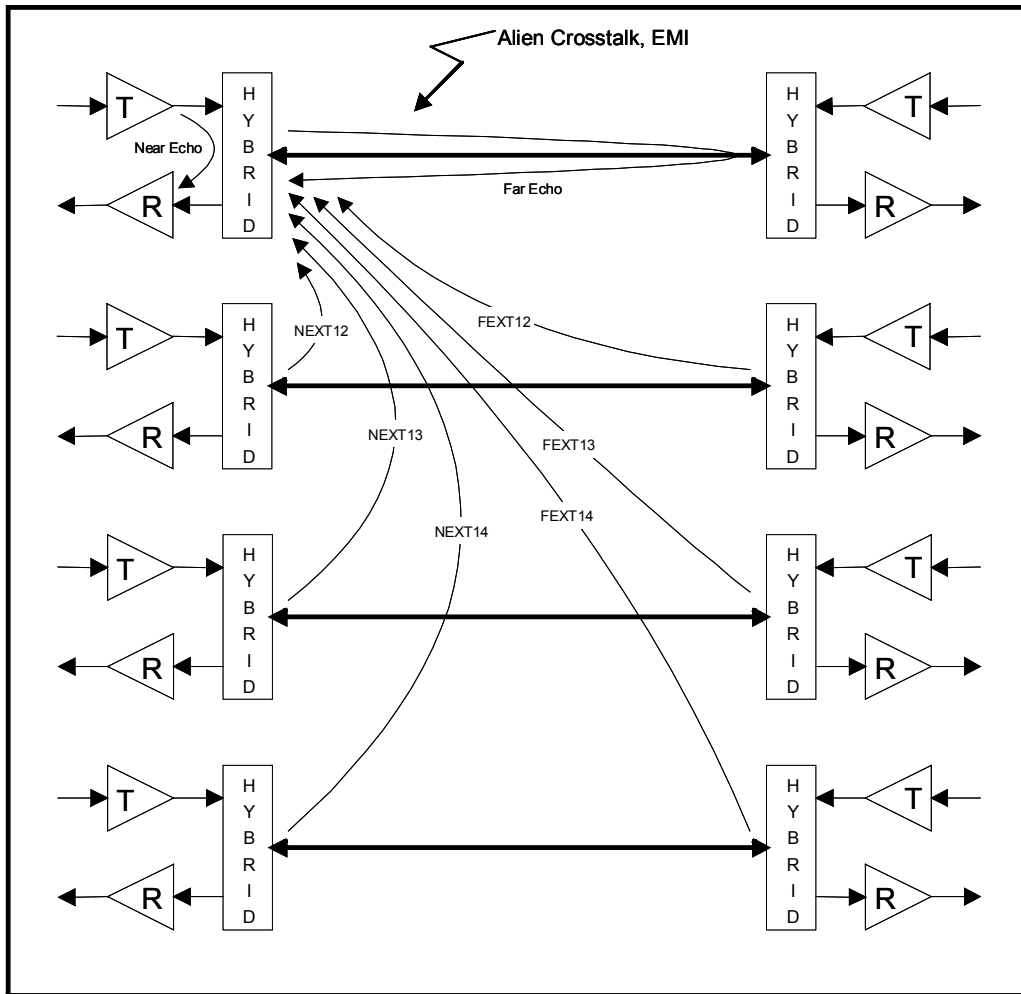


Fig. 1: Impairments to Transmission on 4-Pair Cat5e Cable

Insertion Loss, ISI, and Equalization

Insertion Loss is the measure of signal loss, over a length of cable, vs. the frequency of the signal. In the case of the 10GBASE-T, the analog bandwidth required is about 400 MHz and at this frequency 50 dB of the transmitted signal is lost by the time it reaches the end of 100 m of cable. To achieve bandwidth reduction, the 10-Gbit/s signal is encoded as 3 bits/ baud/twisted-pair, as a coded PAM-10 signal. In order to maximize the BER performance of the receiver a 4-dimensional trellis-coded Viterbi algorithm is used for symbol detection, a structure similar to the 4-D Trellis Viterbi in 1000BASE-T and providing sufficient SNR improvement to meet the BER required by network users. A PAM-10 receiver would require 26.2 dB of SNR to meet a BER of 10^{-12} .

Even with this encoding, the data channel exhibits a large amount of ISI due to the spreading of the transmitted pulse to cover nearly a hundred bit intervals. Borrowing both from previous 1000BASE-T Ethernet and DSL technologies, high-performance equalizers are used to mitigate the ISI and restore the received signal's pulse shape. However, while these are high-performance elements, such equalizers are only the beginning of the challenges facing 10GBASE-T designers.

Echo & NEXT Cancellation

Because of the limited bandwidth available over a typical wireline channel, a full-duplex transmission scheme is employed. The full-duplex transmission scheme has the advantage that both transmit and receive signals occupy the low-frequency portion of the wireline channel, thereby minimizing channel attenuation. The disadvantage of the full-duplex scheme is that transmit and receive spectra overlap, causing the transmit signal to interfere with the receive signal. This interference is known as echo.

The echo is formed by the near-end transmitted signal reflecting off the line, and acts as a linear filter on the transmitted signal. Echo cancellers are used in conjunction with a compromise hybrid to provide the required echo rejection.

Similar in many ways to echo, NEXT is interference produced by the transmitters on the other 3 pairs of the 4-pair Cat5e cable. Because receive signals are significantly attenuated by the wireline channel, and NEXT is not, interference levels due to NEXT are potentially much larger than the receive signal levels. Previously, cable manufacturers strived to keep pair-to-pair NEXT low, considering it an uncancellable impairment that limited the bandwidth of the cable. However, beginning in 1000BASE-T, and continuing to 10GBASE-T, because these transmitters are part of the same communications link, their effects can be substantially cancelled in the receiver.

The challenge of echo and NEXT cancellation in 10GBASE-T is one of the major complexity challenges to implementation. Members of the IEEE study group have

estimated that more than 40 dB of echo and NEXT cancellation is required for a 10GBASE-T system to achieve sufficient capacity to function. This is nearly twice the required cancellation (in dB) for 1000BASE-T systems. The FIR techniques used in 1000BASE-T solutions, if implemented in a straightforward way, would result in complexity increase on the order of 45x over 1000BASE-T. This has been apparent to many in the industry, and was one of the reasons given for claiming 10GBASE-T is not feasible. The high degree of cancellation required, at these speeds, also makes all-analog cancellation difficult, requiring high bandwidth of adaptive analog filters which, even if feasible, would require high power. Because of the inherent variability and randomness of the impulse responses, simple techniques to extend the length of the impulse response cancelled, such as continuous-time analog filters or IIR digital filters do not easily lend themselves to flexible solutions which will work on a variety of wiring. To make matters worse, the echo and NEXT signals so dominate the received signal that a straightforward all-digital approach would require more than 10 bits in the ADC.

To gain both the advantage of highly precise DSP and efficient analog processing, SolarFlare's approach involves both analog and digital signal processing. In order to perform the computations necessary (which, for the straightforward approach would be more than 10 Tera-Ops!) massively parallel computational methods must be brought to bear to implement these filters, allowing massive reuse of computations, performed with fast algorithmic approaches. SolarFlare's patent-pending solution to the problem allows for robust cancellation of NEXT and echo by implementing cancellers 100s of taps long, and sharing computation between all 16 echo and NEXT cancellation paths. The result requires a less than a 6x computational increase over 1000BASE-T solutions, making the echo- and NEXT-cancellers feasible. In addition, the algorithms are coupled with processing in the analog domain, which allows a reduction in NEXT and echo at the ADC, allowing an efficient ADC design, with less than 9 ENOB.

FEXT Cancellation

FEXT is interference produced by signals emanating from adjacent transmitters at the far end of the 10- Gbit/s Ethernet link (see Fig. 1, again). Because these signals are at the far end of the link, they were once thought uncancellable, similar to NEXT. Interference coupled into the receive channel at the far-end of the link is attenuated by the channel before reaching the receiver and FEXT interference levels at the receiver are much lower than the interference levels of NEXT and echo. As a result FEXT cancellers were not required for 1000BASE-T, although today some receivers have limited cancellation. For 10GBASE-T, however, FEXT interference produces a significant loss in SNR, and would prevent 10GBASE-T transmission on Category 5e and 6 cabling. For this reason, FEXT cancellers are employed, improving the capacity available.

To support 10GBASE-T, FEXT interference must be reduced by some 20 dB; while this may sound small compared to the challenges of echo and NEXT, FEXT has the

unfortunate characteristic of not only being formed by signals which do not originate at the same end of the cable, but also have non-causal arrival times (FEXT can arrive before the signal which generated it arrives on its wire).

The Result

A received Eye Diagram on a 100 m, 4 connector Cat-5e channel, clearly showing the PAM-10 signal at the input of the Viterbi detector, after all cancellation and equalization, is in Fig. 2. This indicates that the PHY has achieved the necessary interference cancellations to provide sufficient SNR to decode data in full duplex operation.




Fig. 2: Received Eye Diagram After All Cancellation and Equalization

Characterizations of existing CAT5e cable were used in the modeling of the channel in agreement with prior work done for 1000BASE-T cable modeling. A carefully-crafted system design is shown to achieve full performance data exchange by fully using the available channel bandwidth, thoroughly mitigating the impairments, using advanced signal processing to reduce the complexity and exploiting a robust Viterbi symbol detection method.

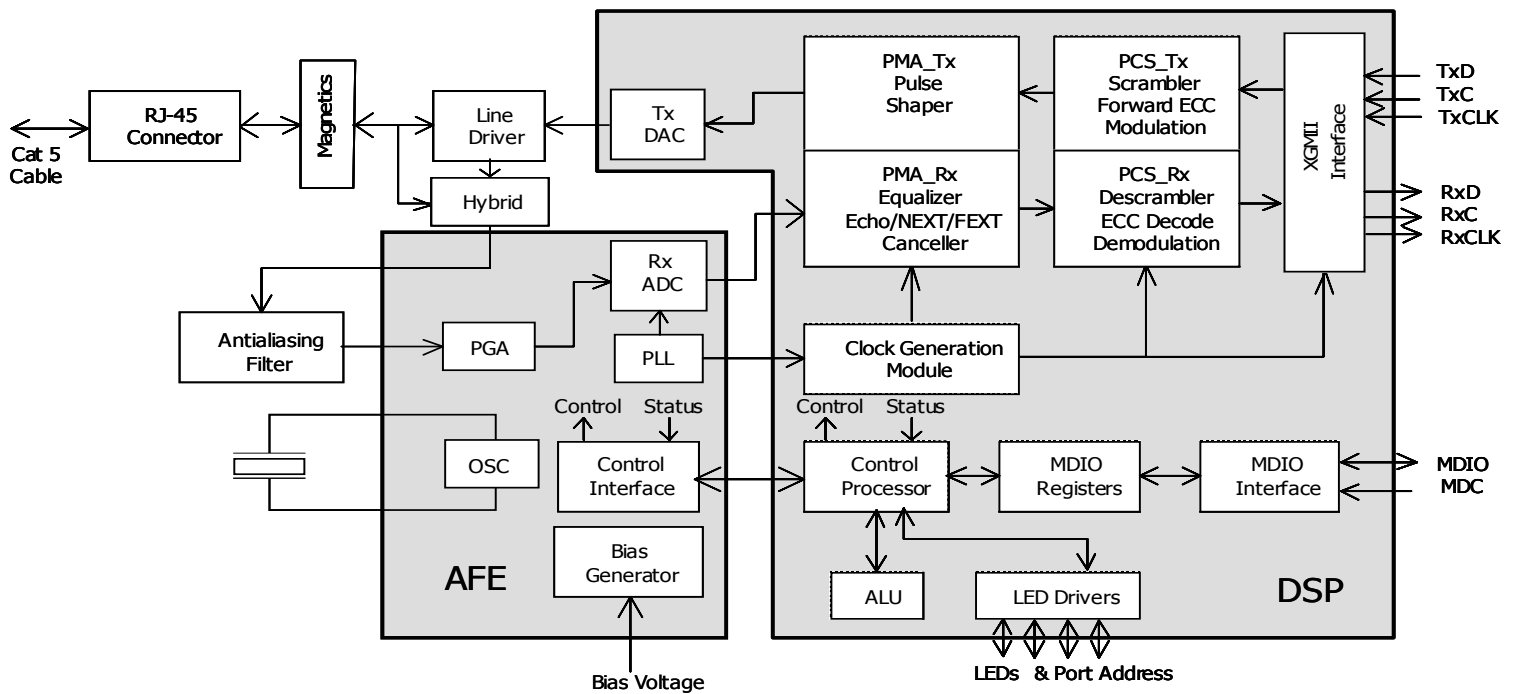


Fig. 3: Transceiver Block Diagram

A block diagram of the final transceiver (Fig. 3) consists of an analog front end (AFE), responsible for digitization, and a digital signal processor (DSP) functional blocks. The key pieces of intellectual property and their impact on the transceiver architecture are depicted in Fig. 4.

Key Intellectual Property	Solarflare Uniqueness	Impact on Design
Multilevel Coded Modulation	Better bandwidth utilization	Enhanced tolerance to cable variability
Adaptive Line Equalization	Much higher performance	Mitigates ISI
New Echo and NEXT mitigation architecture	7x circuit area reduction for function	Fabrication feasibility
Combined FEXT and equalization circuits	6x circuit area reduction for function	Economic circuit realization
Interwoven A/D and DSP architecture	Digital implementation of traditionally analog circuit attributes	Enable utilization of CMOS A/D

Fig. 4: IP and Design Impact

Summary

While the challenges associated with 10-Gbit/s transmission on UTP are formidable, SolarFlare has shown that they are not insurmountable. By cleverly using state-of-the-art DSP techniques, the issues associated with insertion loss and ISI, and degradations due to echo, near-end crosstalk and far-end crosstalk can be overcome.

