By David Marsh, Contributing Technical Editor

IT'S HARD TO IMAGINE A TECHNOLOGY THAT CAN ALMOST SEAMLESSLY EMBRACE A THREE-ORDER-OF-MAGNITUDE BANDWIDTH INCREASE USING COPPER WIRING THAT WAS ORIGINALLY DESIGNED FOR VOICE SERVICES. THIS FIRST ARTICLE OF A TWO-PART SERIES DESCRIBES HOW ETHERNET TODAY, AS 20 YEARS AGO, CONTINUES TO AMAZE.

# ETHERNET KEEPS PUMPING THE DATA

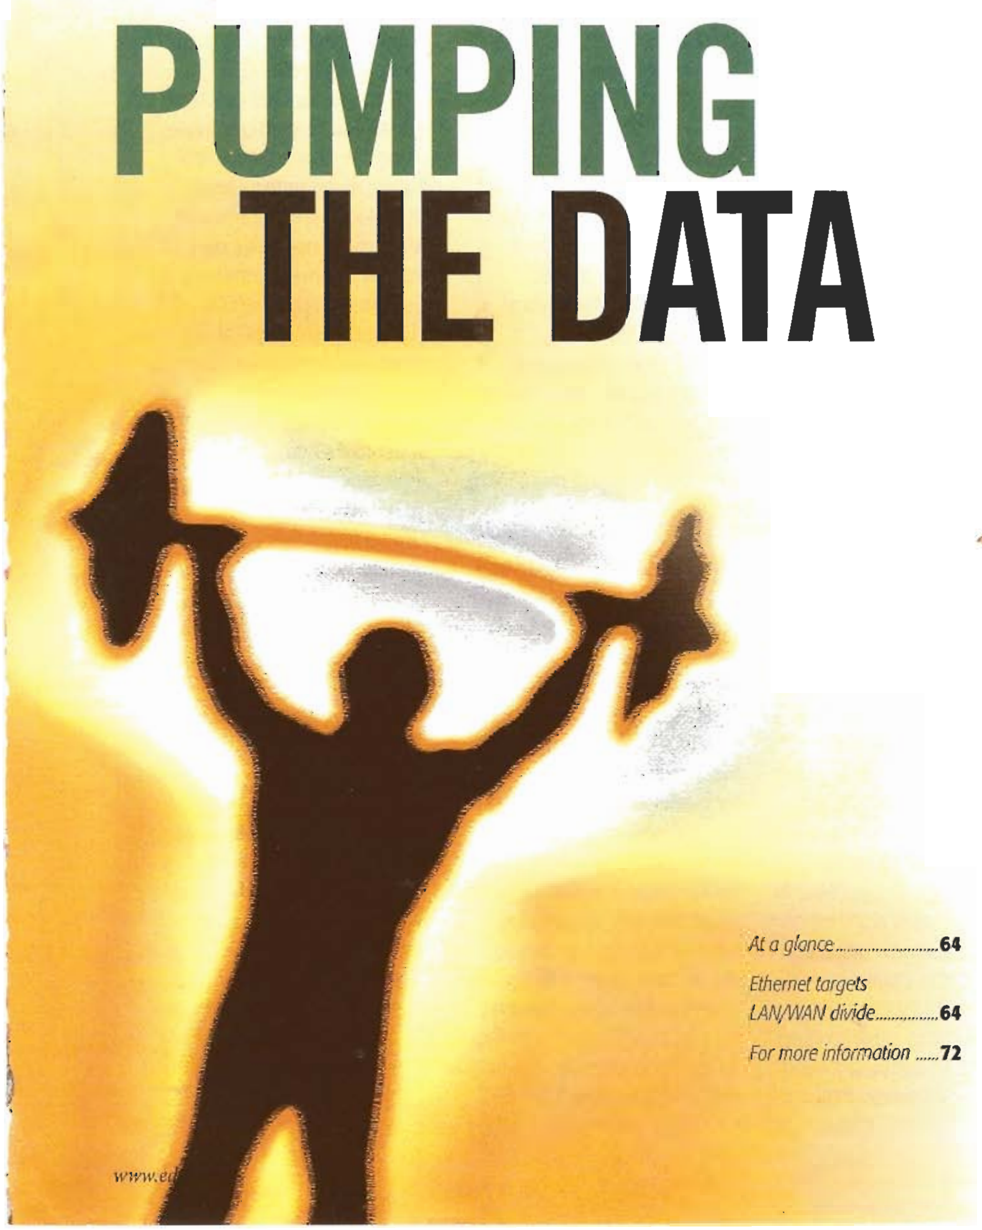

Back in the 1970s, connecting computers to peripherals or other computers was a nightmare. A typical data-processing centre then comprised a central minicomputer or mainframe and a number of dumb terminals and line printers. Connections between the central processor and its peripherals were most often RS-232 serial at speeds of no more than 9600 baud. As a result, terminal communications were slow, and the wiring closet necessary to support even a small-scale operation was virtually impenetrable. Anticipating the explosive growth of pervasive computing, Bob Metcalfe and David Boggs, then researchers at Xerox Corp's now-legendary Palo Alto Research Center, were working on a system to interconnect the company's computers and printers at a data-transmission rate of 2.94 Mbps. Their 1976 paper describes the broadcast-communication system for transporting packets of data among locally distributed computer systems that they dubbed Ethernet (**Reference 1**). Collaboration with Digital Equipment Corp and Intel began the vendor-independent set of standards that the IEEE subsequently embraced under its 802 umbrella—under which development thrives to this day.

With vendors annually shipping an estimated total of 300 million ports, Ethernet pervades business and local-area-networking in general. Automation engineers will quickly acknowledge its long struggle for acceptance in the face of field-bus technologies, principally due to intrinsic-safety issues. This situation remains a bone of contention within the process-control industry, in which many deployments restrict Ethernet to data-collection and back-office duties. But a host of silicon vendors include Ethernet ports on their microcontrollers, on which the technology continues to make inroads into the deeply embedded-system space. Elsewhere, Ethernet continues its challenge for metropolitan-area data transport, with many vendors offering bridge chips to facilitate connections with wide-area backbone technologies, such as SONET/SDH (synchronous optical network/synchronous digital hierarchy). And in consumer and professional-user markets alike, the recent wireless variants are sharply eroding wired-equipment sales.

If you're starting out with Ethernet design or simply want to appreciate contemporary developments as diverse as Gigabit-over-copper, power-over-Ethernet, and ZigBee, it's helpful to reflect on the technological basics that account for continuing development two decades after the first commercial implementations. You'll almost certainly be surprised by the flexibility that even the earliest Ethernet systems made available to implementers and the ingenuity that its current-generation developers continue. This first article in a two-part series will look at the most popular wired flavors that now span 10 Mbps to 10 Gbps. Part 2 will examine power-over-Ethernet, wireless implementations and issues such as the IEEE's efforts to bring quality of service to the wireless domain.

**AT A GLANCE**

▷ Vendors are shipping around 300 million Ethernet-port units/year.

▷ Backward compatibility and scalability are key factors in Ethernet's success.

▷ Modulation techniques provide order-of-magnitude bandwidth increases.

### CSMA/CD CHARACTERIZES ETHERNET

The original DIX (DEC/Intel/Xerox) Ethernet standard that the IEEE in 1983 adjusted before ratification describes a 10-Mbps, half-duplex communication medium that uses a CSMA/CD (carrier-sense-multiple-access-with-collision-detection) protocol. This networking component fits in at the lower two levels of the ISO's (International Organization for

## ETHERNET TARGETS LAN/WAN DIVIDE

Ethernet has long enjoyed a raft of fibre physical-layer options, but cheaper, more flexible copper alternatives have overshadowed them. There comes a time, however, when it's difficult for copper speeds to keep up with fibre over any distance. Initially at least, debate about the viability of a 10-Gbps copper Ethernet standard gave way to a fibre-based project driven by the 10GbE (10-Gigabit Ethernet) Alliance. This 10GbE family retains Ethernet's original frame format and size, as well as its media-access-control protocol to provide a step forward in speed and reach. But with its target markets set toward the metropolitan-area-, storage-area-, and wide-area-networking arenas, IEEE-802.3ae 10GbE is a fibre-only, full-duplex implementation that's compatible with SONET/SDH OC-192 links.

Bob Brand, senior solutions architect at Applied Micro Circuits Corp, notes that Ethernet's cost and familiarity relative to other wide-area technologies now makes it especially attractive for first-mile, point-to-point connections. Brand estimates that only about 7 to 9% of all enterprises in North America have fibre access, so Ethernet's fibre-to-copper transparency will become ever more important. Accordingly, the company now offers various 10GbE-compatible devices, such as the S19235 transceiver and the nP7510 network processor, to complement its wide-area silicon for primary-transport and multi-service-provision applications.

Other companies to recently introduce silicon include Galazar Networks, whose MDX250 mapper includes virtual-concatenation and generic-framing-procedure capabilities. These abilities help map Ethernet traffic onto SONET/SDH links, supporting as many as four Gigabit Ethernet channels and as many as 16 Fast Ethernet channels. The chip also includes a 16-channel SPI-3 (system-packet-interface) for connecting network processors. This interface operates in 8- or 32-bit mode at speeds as high as 104 MHz for transfer rates reaching 3.2 Gbps.

Marc Kimpe, senior staff scientist at network-access provider Adtran and an 802.3 committee member, says that laying fibre in the access network is still expensive for the telecom companies. As a result, fibre uptake is gradually growing and getting closer to the subscriber, either to street cabinets for onward distribution over copper or directly into green-field industrial developments. Cost issues with fibre—and the fact that copper still completes the end user's connection—drive the service providers' desire to offer Ethernet over copper, but they then encounter many of the same issues that haunt Broadband technologies (www.reed-electronics.com/ednmag/CA339737).

Within the access arena, Ethernet is challenging technologies such as SONET and ATM (asynchronous-transfer-mode) to become a transport mechanism that bridges the LAN/WAN divide. Kimpe notes that telecom companies in the United States have for about 10 years offered point-to-point Ethernet, but access costs are now plummeting due to reduced equipment costs. He says that it's relatively easy to define a LAN topology, but access networks provide a far greater challenge, because you can neither predetermine where subscribers lie nor control the length and type of cable the signal runs on. You can examine Ethernet's attempt to become the de facto protocol that runs over first-mile access networks, such as the xDSL family, in the latest EFM (Ethernet-in-the-first-mile) standard, which the IEEE was due to ratify as 802.3ah-2004 at press time (see www.ieee802.org/3/efm/).

The new standard describes end-to-end connections made using three basic component blocks. At the primary level, point-to-point fibre carries 100-Mbps to 1-Gbps traffic over 10 to 20 km between repeaters. A point-to-multipoint EPON (Ethernet passive-optical-network) can then deliver fibre access to residential blocks and businesses with sufficient bandwidth to distribute the so-called triple play of video, data, and voice services via multiport Gigabit transceivers. The third element is delivery over twisted-pair copper, which trades off data rate as a function of reach. A short-reach option allows symmetrical 10-Mbps communications over distances to about 750m, and a long-reach option allows symmetrical 2 Mbps over about 2.7 km.

Kimpe observes that growth in mobile services is freeing up the copper infrastructure, making available surplus capacity. The EFM standard exploits this situation by allowing the possibility of combining as many as 32 twisted pairs to form a scalable data pipe.

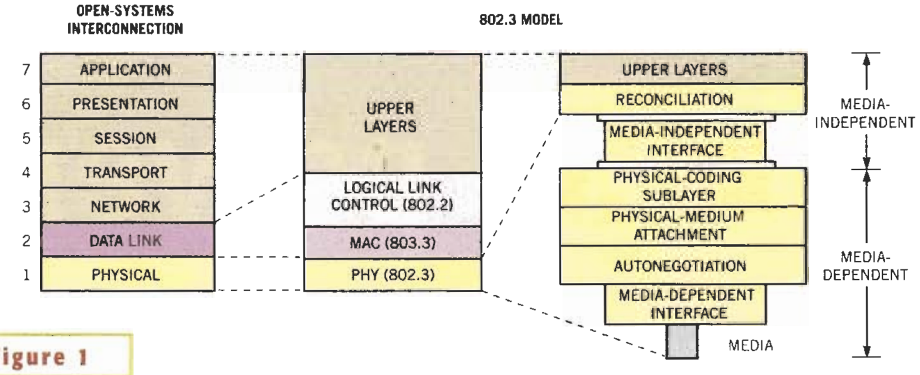

**Figure 1**

Ethernet's services map into the lower two layers of the standard ISO/OSI model to provide media-independent communications.

Standardisation's) OSI (open-systems-interconnection) model (ISO/IEC 7498-1:1994) to furnish physical-layer and data link-layer services (**Figure 1**). Colloquially known as the network level, these two layers comprise the network-connection medium, such as cable; the communications hardware, such as transceiver ICs; and the driver software that translates and transports physical-layer data streams to higher level layers for further processing. In this sense, traditional Ethernet is a connectionless and unreliable, or "best-effort-delivery," system. Therefore, the protocol packages data and endeavors to transport it to its intended destination without first checking whether the destination is available or whether the data arrives safely; if the application demands it, higher level software performs these functions.

Today's most popular Ethernet flavors still use the original frame format (**Figure 2**) and CSMA/CD mechanism, with relatively few tweaks to accommodate successively higher speeds or different application profiles. Although full-duplex variants exist that dispense with the need for arbitration between two nodes, most engineers immediately associate Ethernet with CSMA/CD. Unlike later protocols, such as CAN (controller-area network), that require each node to have a predefined network-access priority to facilitate deterministic response, all Ethernet nodes have equal access rights; hence, Ethernet doesn't suit real-time-control use. When a traditional Ethernet node wishes to transmit data, it first listens to the network to ascertain whether the medium is in use—the carrier-sense phase—and, if it is not, the node starts to transmit. If another node simultaneously attempts transmission, a data collision occurs that creates an excessively high voltage on a coaxial network, or simultaneous transmitting/receiving activity on a half-duplex twisted-pair network. Sensing the anomaly, each node ceases data transmission and transmits a jam pattern that instructs all other nodes to temporarily hold off any pending transmissions. This jam pattern is 32 bits long and comprises any sequence (typically, all ones) that doesn't match the CRC value in the frame-check-sequence field. Hence, all receiving nodes discard the corrupt packet.

Meanwhile, the colliding nodes delay their retransmission attempts by a random time. The back-off delay algorithm uses the 48-bit MAC (media-access-control) address that's unique to every Ethernet device as part of its randomization routine; if the devices still collide, the process repeats as many as 16 attempts with ever-increasing delay values—after which senders give up and flag an error. Because a transmitter can detect a collision only while it's transmitting, it's essential for the first bit of any transmission to reach every node before the sender ceases transmission. Accordingly, the protocol specifies maximum cable lengths to ensure that the first 64 bits of every packet propagate the entire cable length. If a collision occurs after the sender transmits its last bit, the protocol can't detect the error; hence, higher layer functions may be necessary to detect data losses. This "late-collision" condition normally indicates a hardware failure or cabling misconfiguration. But collisions are a normal feature of Ethernet traffic that increase with network usage and inevitably degrade throughput. In general, you can expect significant degradation above about 80% network usage, but most systems run at about 30 to 40%.

Using a thinner and much less expensive RG-58 coaxial cable, 10Base2 quickly overtook 10Base5, or "thick Ethernet," technology, which uses industrial-strength RG-8 coaxial cabling. (In Ethernet jargon, the leading "10" signifies a speed of 10 Mbps, "base" means baseband, and the final number is the maximum transmission length in hundreds of meters.) These early Ethernets employ bus architectures with 50Ω termination resistors at either extremity and a ground connection at one end. The 10Base5 technology achieves its 500m reach by using a single cable segment that's about 10 mm in diameter with takeoff points no closer than 2.5m. Having no breaks in the cable avoids impedance discontinuities due to connectors or differing cable lots and requires an insulation-displacement connector called a "vampire tap." This connector immediately links into a transceiver that interfaces with end-user equipment via an AUI (attachment-unit-interface) cable. You now see echoes of 10Base5's separate transceiver arrangement in Fast Ethernet's 40-pin MII (media-independent interface) option and some higher speed versions, too.

By contrast, 10Base2 daisy-chains each device on the bus using BNC (Bayonet Neill Concelman, after the designer of the coaxial-connector series) T-connectors, and the transceiver—or NIC (net-

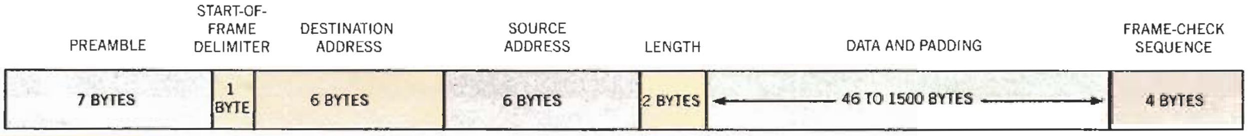

**Figure 2**

Ethernet's original frame format remains unchanged across the 10-Mbps to 10-Gbps divide.

work-interface card), in the case of a PC—resides inside the device. Simplicity apart, the bus architecture has some disadvantages, such as 10Base2's danger of corruption when adding or removing a node from a running system—a common application requirement. The requirement to daisy-chain nodes is difficult to accommodate in a system of any size, and coaxial cable isn't the easiest medium to work with. Also, the bit-timing characteristics limit transmission speed to 10 Mbps for both bus topologies and maximum segment length to 185m for the dominant RG-58 "Cheapernet" version.

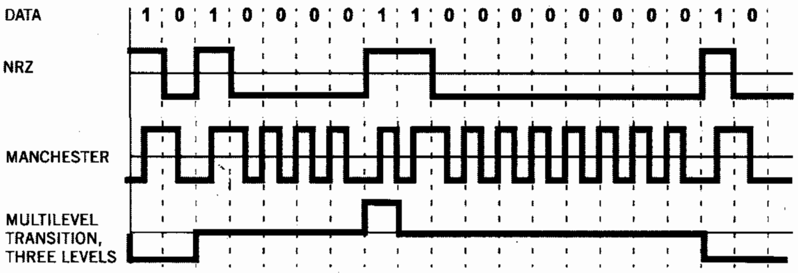

**Figure 3** **Fast Ethernet's MLT-3 coding is Ethernet's first attempt at bandwidth reduction.**

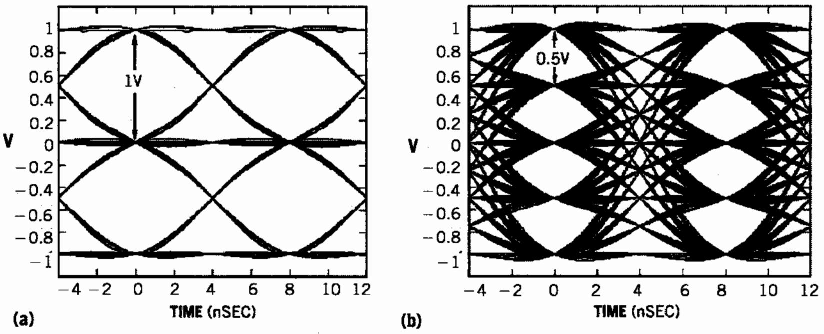

**Figure 4** **Compared with MLT-3 (a), Gigabit Ethernet's coding halves the 1-Gbps technique's SNR (b).**

As a result, the industry embraced a star topology that employs 100Ω UTP (unshielded-twisted-pair) wiring to connect devices via hubs that act as both wiring centers and signal repeaters. The respective IEEE-802.3i-1990 standard that's known as 10BaseT—in this case, the "T" signifies twisted-pair—forms the basis of most of today's wired-Ethernet deployments, with an estimated 800 million UTP ports in use. Much of this popularity comes from Ethernet's scalability, which hugely eases a network administrator's life. With the advent of structured cabling, buildings almost invariably have Ethernet sockets distributed around work areas, along with other wiring services. Typically, groups of RJ-45 Ethernet sockets connect to one of a number of multiport hubs in a local wiring closet, in which each hub connects to its neighbor. Within constraints that the so-called 5-4-3 rule dictates, adding more hubs or devices to a typical LAN simply requires cabling the new devices to the infrastructure. When you exceed these limits, you can create another Ethernet system and join the two using a network bridge that may also access another network technology.

Ethernet's 5-4-3 rule states that, in one collision domain, the permissible route between any two nodes mustn't exceed five cable segments that interconnect a maximum of four hubs, and only three of these segments can be mixing segments. Here, "collision domain" refers to a network in which simultaneous transmissions by two nodes cause a collision, and "mixing segments" are segments that connect more than two nodes. This rule of thumb attempts to account for round-trip signal-propagation times—that is, the worst-case time for a bit to travel the length of the network and back again. Assuming that this value is within limits, every node's CSMA/CD mechanism can determine whether another node is attempting a simultaneous transmission. Notice that the precise delay value depends on the media, but another rule of thumb states that the average velocity of propagation over copper cable is about 0.7 times the speed of light, or around 5 nsec/m. Interframe gap shrinkage, which refers to the amount that network conditions shorten the normal 96-bit delay between packets, is also a potentially significant value. This delay ensures that Ethernet's half-duplex interfaces have enough time to switch between transmitting and receiving modes. If you're building a complex topology and need a more accurate assessment than 5-4-3, use a protocol analyzer to find the delay and shrinkage parameters for your worst-case path.

The 5-4-3 rule has some interesting architectural implications. For example, in a coaxial-cable bus network, you can interconnect as many as five cable segments via four signal-amplification repeaters. Because each repeater has only two ports, the cabling between each can be a mixing segment that serves multiple nodes or a link segment that joins only two nodes. This setup infers that the maximum collision-domain diameter for a 10Base2 network is five times 185m—which isn't far short of a kilometer; 10Base5 manages 2500m. These bus technologies can therefore address applications such as linking separate buildings. By contrast, a UTP network contains only link segments, because its repeaters are multiport hubs. With 10BaseT specifying a 100m maximum segment length, you can configure four hubs with as much as 100m between each, with each hub serving its nodes via further link segments of as much as 100m. This arrangement especially suits buildings, with a hub serving each successive floor. Further flexibility comes from being able to mix and match coaxial and UTP segments within 5-4-3 limits. If you take into account maximum cable runs for the largest permissible number of nodes, 10BaseT's collision domain limit is about 2800m.

## MODULATION ENABLES GBPS OVER COPPER

Although there's a raft of fibre-based Ethernet standards, cost considerations continue to promote copper media. Here, advances in physical-layer implementations allow order-of-magnitude speed increases across compatible cabling systems, such as Category 5e and Category 6 (**Reference 2**). As a result, virtually all of today's PC networking hardware is dual 10/100-Mbps-capable. Providing that your cabling is up to the job, 100BaseTX Fast Ethernet hardware can

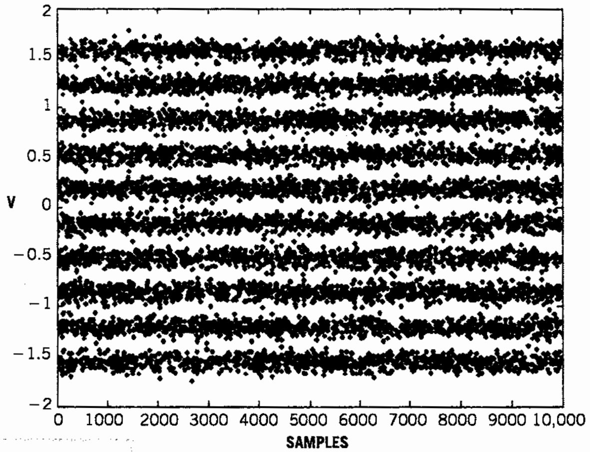

**Figure 5**

**An eye diagram shows that it's now feasible to run 10GBaseT over 100m of Category 5e cable.**

run without modification alongside 10BaseT. This specification remains today's most popular implementation in every IC format, including intellectual-property blocks from programmable-logic specialists such as Actel, Altera, and Xilinx. The frame format remains the same, and there are no changes to signalling values, such as the 512 bit-time for a slot or the 96 bit-time interframe gap. But 100BaseTX's 10-times greater signalling speed and the subsequent timing restrictions constrain its collision-domain diameter to a relatively meagre 205m. Like 10BaseT, Fast Ethernet nominates separate transmitting and receiving wire pairs to enable an optional full-duplex capability. In this case, each segment is a link segment between a device and an Ethernet switch. Using a switch rather than a hub/repeater isolates the segment from other segments, so each segment becomes in effect its own collision domain. Accordingly, this strategy is key to increasing the length of a 100BaseTX network. It also renders superfluous the CSMA/CD mechanism, which a switch that's configured for full-duplex operation ignores.

The IEEE-802.3u specification enables Fast Ethernet's compatibility by supplanting 10BaseT's binary-level Manchester encoding with a three-level symbol-transmission system. This scheme uses a 4B/5B code that was originally developed for the FDDI (fibre-distributed-data-interface) system. The coding translates 4-bit data nibbles into a 5-bit code that enables error detection and adds control codes, such as start- and end-of-stream delimiters. Upping the symbol rate to 125 Mbps compensates for the 4B/5B's inherent 20% data-transmission inefficiency, but this bandwidth increase creates a spectrum that Manchester encoding would extend into hundreds of megahertz. Attenuation losses and EMC issues prohibit this approach, so 100BaseTX uses a MLT-3 (multiple-level-transition, three-level) carrier.

Like Manchester coding, MLT-3 encodes a bit according to transitions, but it packs its output into a three-level waveform that crudely simulates sine-wave-energy distribution. These three levels continuously alternate from +1 to 0 to −1 and back again, with a logic zero effectively halting the sequence, and a one restarting it (**Figure 3**). A long sequence of ones creates the output signal's highest frequency, when the signal repeats the 1, 0, −1, 0 pattern. This pattern's cycle length is one-fourth the clock rate, which reduces the worst-case energy component to 32.5 MHz. Because a long stream of zeros generates a constant output that makes it difficult to synchronize the receiver's read-data clock with the data stream, the protocol limits such sequences and continuously transmits an "idle" signal in the absence of data. These steps also preserve the neutral dc balance on the cable that's an intrinsic feature of Manchester encoding. Such neutral balance avoids dc offsets on the cable that would otherwise accumulate to distort the signal and saturate signal-coupling transformers.

The transparent 10/100-Mbps capability relies on Fast Ethernet's optional autonegotiation sequence, which takes advantage of 10BaseT's normal-link-pulse signals. These signals verify link integrity between two nodes and have 2-nsec period and 16.8-nsec interval-timing characteristics. Modifying the protocol to include a 16-bit data packet that includes a link-code-word within a burst of fast-link-pulses with the same timing as normal-link-pulse signals allows 10BaseT devices to ignore the link-code word's selector and technology-ability fields. During autonegotiation, devices configure themselves to provide the best performance level that their design allows, from full-duplex 1000BaseT down to 10BaseT.

It's also possible to run 1000BaseT over high-spec cabling within a UTP environment, as IEEE-802.3ab describes. For backward compatibility, this Gigabit Ethernet standard supports CSMA/CD in half-duplex mode, and it is subject to the same topological constraints as 100-BaseTX. Autonegotiation is now mandatory. But Gigabit Ethernet's developers conceived it primarily as a full-duplex medium to support enterprise backbones, in which it enjoys widespread acceptance. Here, immediate compatibility with legacy environments is less crucial and is better made via bridges to lower bandwidth routes. To achieve gigabit speeds, 1000BaseT uses all four of the cable's wire pairs and can simultaneously send and receive over each wire pair. In this approach, transmission and reception signals simultaneously occupy the same low-frequency portion of the channel to minimize attenuation problems. But their spectra then overlap to create interference in the form of echoes, when the near-end transmission signal reflects off the line and acts as a linear filter on the transmitted signal. Because each end's receiver "knows" what it's just sent, it can subtract its transmission patterns from the composite signal to recover the opposite end's sent data. This technique, full-duplex echo-cancelling transmission, implemented using DSP FIR (finite-impulse-response)-filter techniques.

Deriving from RF communications practice, 1000BaseT uses a PAM (pulse-amplitude-modulation) scheme, PAM-5, that features a 5×5 constellation with 2 bits per symbol and a 125-MHz symbol rate. Combining five-level coding across four wire pairs permits 1000BaseT to send 1 byte per signal pulse (four wire pairs times 125 million symbols/sec times 2 bits/symbol equals 1 Gbps). One effect of increasing transmission levels from Fast Ethernet's three to Gigabit Ethernet's five and maintaining the same overall voltage swing is to increase sensitivity to noise disturbances by some 50%, or 6 dB (**Figure 4**). Because only four levels are necessary to transmit the symbol's 2 bits, Gigabit Ethernet uses the fifth level for FEC (forward-error-correction) to compensate for this SNR loss.

Its four-dimensional, eight-state trellis coding and Viterbi decoding combination provides as much as 6 dB of SNR improvement to theoretically match Fast Ethernet's performance. The IEEE approved 100BaseT and 1000BaseT in 1995 and 1999, respectively.

### 10 GBPS OVER COPPER DEBUTS

The most recent evolution of the traditional wired-Ethernet model appears as 802.3ae, the fibre-only, 10-Gbps derivative that the IEEE approved in 2002 (see **sidebar** "Ethernet targets LAN/WAN divide"). But with an installed base of at least 800 million UTP ports, there's considerable interest in taking the 10BaseT/1000BaseT progression one step further. Drivers include cost sensitivity in small to midsized enterprises, as well as niche applications, such as digital-video-production houses that route Gigabit Ethernet to the desktop. Although 62.5-micron multimode fibre is cheap, the connectors certainly aren't, and the skill level that's required to successfully mate them is way beyond the capabilities of personnel who are unfamiliar with polishing regimes and microscopic inspection. The hybrid optical circuits are expensive, too, because they combine technologies that are mechanically difficult to align, and impossible to functionally test until the entire assembly is complete; by comparison, most Ethernet-over-copper circuits are built using conventional CMOS processes. Although fibre easily wins over copper as a data-transmission medium, in quality terms, these points help explain why manufacturers first implemented previous Ethernet bandwidth jumps on fibre. Interestingly, it's generally newcomers, such as Broadcom and Marvell, rather than Ethernet giants, such as Advanced Micro Devices, Analog Devices, Intel, Motorola (now Freescale), National Semiconductor, and Texas Instruments, that have first migrated fibre to copper.

This year, the IEEE embraced the 10-Gbps-over-copper challenge with a study group that's looking at developing a 10GBaseT standard for UTP cable, which current estimates suggest may this year see a first draft, with ratification scheduled for July 2006. As for previous incarnations, cabling quality is critical to meet speed and reach demands. Following extensive simulation, the IEEE study group concludes that it should be possible to run 10GBaseT for distances of 55 to 100m on Category 6 cable, with a Category 7 installation ensuring 100m reach. Extending Gigabit Ethernet's bandwidth-reduction strategy, 10GBaseT's developers propose a coded PAM-10 signal that transmits 3 bits over each twisted pair for every clock pulse at an 833.33M-baud rate (**Reference 3**). The new standard will take advantage of Gigabit Ethernet's multidimensional trellis-coder/Viterbi-decoder combination, which together with PAM-10 requires about 26 dB of SNR to meet a $10^{-12}$ bit-error-rate target. But the data channel still exhibits significant intersymbol interference that's due to transmit pulses spreading over something approaching 100 bit-time intervals. This problem is common to 1000BaseT and various DSL (digital-subscriber-line) technologies, leading to the development of high-performance equalizers to recover the signal's shape.

Some critical measurements that also

## FOR MORE INFORMATION...

For more information on products such as those discussed in this article, contact any of the following manufacturers or organizations directly, and please let them know you read about their products in *EDN*.

| | |
|---|---|
| **Actel** | www.actel.com |
| **Adtran** | www.adtran.com |
| **Advanced Micro Devices** | www.amd.com |
| **Altera** | www.altera.com |
| **Applied Micro Circuits Corp** | www.amcc.com |
| **Analog Devices** | www.analog.com |
| **Broadcom** | www.broadcom.com |
| **Freescale** | www.freescale.com |
| **Galazar Networks** | www.galazar.com |
| **IEEE** | www.ieee.org |
| **Intel** | www.intel.com |
| **ISO** | www.iso.org |
| **KeyEye** | www.keyeye.net |
| **Marvell Technology** | www.marvell.com |
| **National Semiconductor** | www.national.com |
| **SolarFlare** | www.solarflare.com |
| **10 Gigabit Ethernet Alliance** | www.10gea.org |
| **Teranetics** | www.teranetics.com |
| **TIA** | www.tiaonline.org |
| **Texas Instruments** | www.ti.com |
| **Xerox PARC** | www.parc.xerox.com |
| **Xilinx** | www.xilinx.com |

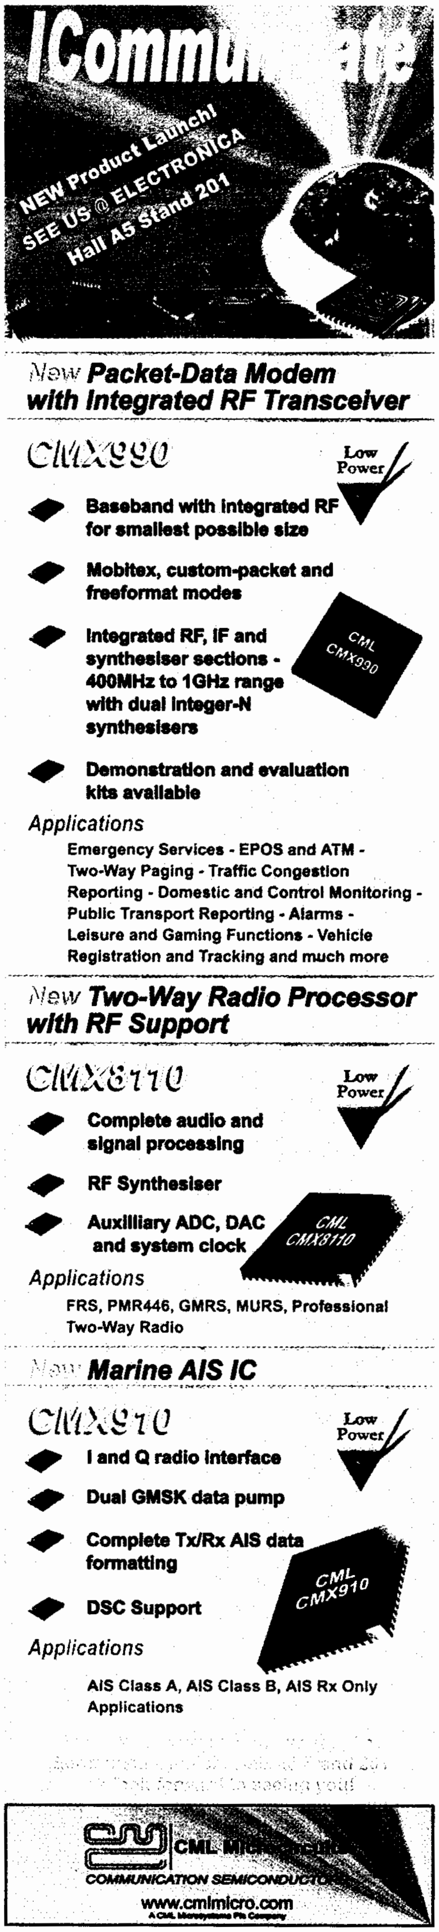

apply to Fast and Gigabit Ethernet to lesser extents include NEXT (near-end crosstalk), FEXT (far-end crosstalk), insertion and return loss, and intersymbol interference. Because the copper link significantly attenuates transmission signals, NEXT interference that aggregates from the other three adjacent wire-pair transmitters can be much higher than the reception signal. This problem increases with increasing transmission frequency, but because all the transmitters are part of the same communications link, it's possible to implement cancellation within each receiver. Some 1000BaseT transceivers also employ a limited amount of FEXT cancellation to attenuate interference that emanates from adjacent transmitters at the far end of the cable. But this capability is mandatory for 10GBaseT and requires some 20 dB of cancellation, because the resulting SNR impairment would otherwise prevent operation over Category 5e/6 cabling. This goal is difficult to achieve because, unlike near-end anomalies, FEXT's nature means that wide variations exist in the interference signals' arrival times; interference may well arrive before the signal of interest.

The 10GBaseT technology adds ANEXT (alien crosstalk) to the signal-impairment list, meaning that it's susceptible to noise from adjacent cabling. Because DSP can't sufficiently cancel ANEXT or other interference due to common sources, such as TV and radio broadcasts, fluorescent lights, and motors, installations may have to observe wider separation between UTP runs, employ better-shielded cable, or both. Projects are under way at the Telecommunications Industry Association's Category 6 Forum (www.category6.org) to produce enhanced cabling that will provide 100m reach. In the meantime, the cabling committees are using the current Category 6 characteristics for their modeling, but extend the frequency-response range to 625 MHz rather than the baseline 250 MHz. Their calculations indicate that 10GBaseT requires about 400 MHz of analogue bandwidth rather than 1000BaseT's 80 MHz, and that a 100m cable attenuates signals at this higher frequency by as much as 50 dB.

You can reach Contributing Editor David Marsh at forncett@btinternet.com.

Continuing Ethernet's tradition of small companies' driving innovation, KeyEye Communications, SolarFlare Communications, and Teranetics feature heavily in today's 10GBaseT arena. Ron Cates, vice president of marketing at Intel-backed start-up SolarFlare, likens 10-Gbps-over-copper signal-recovery challenges to receiving satellite signals from deep space. He explains that echo- and NEXT-cancellation issues are among the most difficult implementation challenges, requiring more than 40-dB attenuation to permit 10GBaseT to function—or about twice the value that 1000BaseT requires. He estimates that simply extending 1000BaseT's FIR approach would increase computational complexity by as much as 45 times to require some 10 tera-operations/sec. Also, the variability of the impulse responses excludes pure-analogue filters or IIR (infinite-impulse-response) digital designs. Worse still, echoes and NEXT interference dominate the reception signal to the extent that an all-digital approach requires more than 10 bits of ADC resolution.

Accordingly, SolarFlare's approach combines analogue processing with a digital scheme to reduce ADC-resolution needs to less than 9 bits. The digital processing employs cancelers several hundred taps long and shares computations among all 16 echo-and-NEXT cancellation paths to yield a feasible times-six complexity increase over 1000BaseT. As a feasibility demonstrator, the company last March demonstrated the first working 10GBaseT silicon. Measurements prove that 100m reach is now possible over Category 5e cable, which together with Category 6 accounts for about three-quarters of UTP's installed base (**Figure 5**).□

---

REFERENCES

1. Metcalfe, Robert M and David R Boggs, "Ethernet: distributed packet switching for local computer networks," Communications of the Association for Computing Machinery Inc, 1976, www.acm.org/classics/apr96.

2. Marsh, David, "Category 6 cable: Gigabit Ethernet over copper," *EDN*, Dec 9, 1999, pg 71, www.edn.com.

3. Zimmerman, G and Brett McClellan, "PHY proposal for 10GBase-T: Encoding, mapping and framing," SolarFlare Communications, www.ieee802.org/3/an/public/may04/mcclellan_1_0504.pdf.

---

TALK TO US

*Post comments via TalkBack at the online version of this article at www.edn.com.*