Digital-Signaling Techniques for LANs

Various approaches, like NRZ, MLT-3, and Manchester, can be used in LAN equipment design. As higher-speed and lower-quality transmission media are increasingly employed, the design of the signal encoding scheme becomes even more critical.

By William Stallings

The increasing proliferation of different types of local area networks (LANs) and the ongoing evolution of public and private wide-area networks (WANs) toward digital technology and services have led to an increased interest in designing efficient digital-signaling techniques.

Digital signaling may be illustrated by considering the case of a bit stream being generated by a source, such as a personal computer or a voice digitizer. In either case, the data is typically represented as discrete voltage pulses, using one voltage level for binary (0) and another for binary (1). This traditional encoding technique is known as Nonreturn-to-Zero (NRZ), and is common in physical interfaces such as EIA-232.

A common means of transmitting digital data is to pass it through a modem, then transmit it as analog signals. There are a number of cases, however, where this is not done:

- Baseband LANs, such as Ethernet and Token Ring
- Digital PBX connections for digital telephones and data processing devices
- Digital access to public telecommunications networks over a digital local loop.

In all of these cases, the digital data is transmitted as a series of voltage pulses — referred to as digital signaling. Although it is possible to use NRZ directly for signaling, NRZ's form is not compatible with the data rates and/or distances of LANs and digital WANs. Instead, the NRZ stream is encoded to enhance performance.

Described and compared here are the various encoding approaches used in LAN equipment design. These codes are defined in Table 1, and an example is given in Figure 1. Two of the major issues are signaling rate requirements and quality/performance.

There are two important tasks in interpreting digital signals at the receiver. First, the receiver must know the timing of each bit. That is, the receiver must know with some accuracy when a bit begins and ends, so that the receiver may sample the incoming signal each bit time to recognize the value of each bit. Second, the receiver must determine whether the signal level for each bit position is high or low.

A number of factors determine how successful the receiver will be in interpreting the incoming signal: the signal-to-noise ratio (SNR), the data rate, and the bandwidth of the signal. Other factors, such as type and length of transmission medium are held constant:

- An increase in data rate increases

bit-error rate (the probability that a bit is received in error)

- An increase in SNR decreases bit-error rate
- Increased in bandwidth of the transmission medium allows increased data rate.

Another factor that can be used to improve performance is the encoding scheme, which is simply the mapping from data bits to signal elements. A variety of approaches have been tried. All of these approaches involve mapping a stream of bits into a signal encoding format. Some of the approaches also involve a bit transformation on the bit stream prior to signal encoding, to improve signal characteristics.

To evaluate or compare the various techniques, signal spectrum, clocking, error detection, signal interference and noise immunity, and cost and complexity should be considered.

Nonreturn to zero codes

The most common, and easiest, way to transmit digital signals is to use two different voltage levels for the two binary digits. For example, the absence of voltage can be used to represent 0, with a constant positive voltage used to represent 1. More commonly, a negative voltage is used to represent one binary value and a positive voltage is used to represent the other. This latter code is known as Nonreturn-to-Zero-Level (NRZ-L). NRZ-L is generally the code used to generate or interpret digital data by terminals and other devices. If a different code is to be used for transmission, it is typically generated from an NRZ-L signal by the transmission system.

A variation of NRZ is known as NRZI (Nonreturn to Zero, Invert on Ones). As with NRZ-L, NRZI maintains a constant voltage pulse for the duration of a bit time. The data are encoded as the presence or absence of a signal transition at the beginning of the bit time. A transition (low-to-high or high-to-low) at the beginning of a bit time denotes a 1 for that bit time; no transition indicates a 0.

ILLUSTRATION BY JANE HAMBLETON

make these codes unattractive for signal transmission applications.

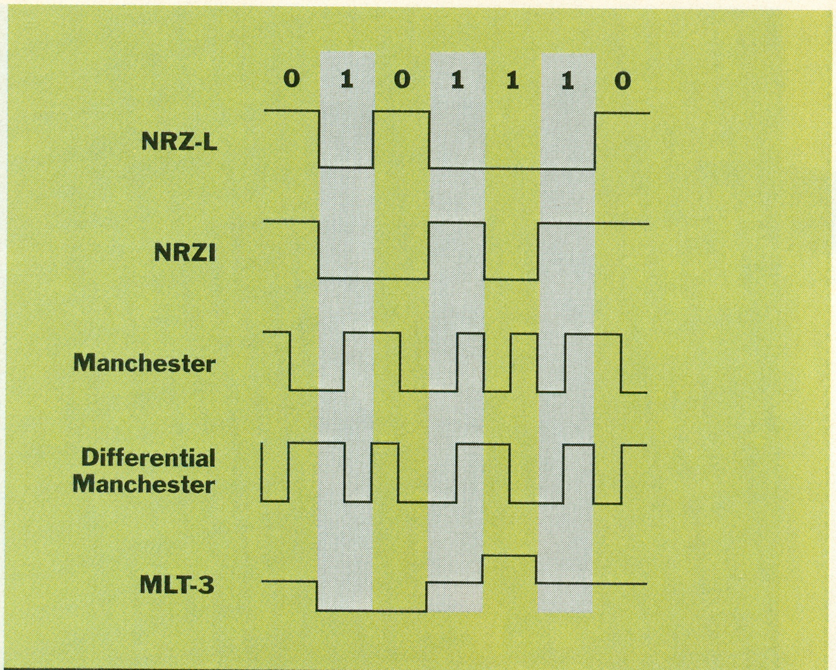

FIGURE 1: *Examples of digital-signal encoding approaches.*

NRZI is an example of differential encoding. The signal is decoded by comparing the polarity of adjacent signal elements rather than determining the absolute value of a signal element. One benefit of this scheme is that it may be more reliable to detect a transition in the presence of noise than to compare a value to a threshold. Another benefit is that with a complex transmission layout, it is easy to lose the sense of the polarity of the signal. For example, on a multidrop twisted-pair line, if the leads from an attached device to the twisted pair are accidentally inverted, all 1s and 0s for NRZ-L will be inverted. This cannot happen with differential encoding.

The NRZ codes are the easiest to engineer and, in addition, make efficient use of bandwidth. Most of the energy in NRZ and NRZI signals is between dc and at half the bit rate. For example, if an NRZ code is used to generate a signal with a data rate of 9,600 bps, most of the energy in the signal is concentrated between dc and 4,800 Hz.

The main limitations of NRZ signals are the presence of a dc component and the lack of synchronization capability. To visualize the latter problem, consider that with a long string of 1s or 0s for NRZ-L or a long string of 0s for NRZI, the output is a constant voltage over a long period of time. Under these circumstances, any drift between the timing of transmitter and receiver will result in a loss of synchronization between the two.

Because of their simplicity and relatively low-frequency response characteristics, NRZ codes are commonly used for digital magnetic recording. However, their limitations make these codes unattractive for signal transmission applications.

Biphase

There is another set of coding techniques, grouped under the term biphase, which overcomes the limitations of NRZ codes. Two of these techniques, Manchester and Differential Manchester, are commonly used in LANs.

In the Manchester code, there is a transition at the middle of each bit period. The mid-bit transition serves as a clocking mechanism and also as data: a low-to-high transition represents a 1, and a high-to-low transition represents a 0. In Differential Manchester, the mid-bit transition is used only to provide clocking. The encoding of a 0 is represented by the presence of a transition at the beginning of a bit period, and a 1 is represented by the absence of a transition at the beginning of a bit period. Differential Manchester has the added advantage of employing differential encoding.

All of the biphase techniques require at least one transition per bit time and may have as many as two transitions. Therefore, the signaling (baud) rate is as much as twice the bit rate. In contrast, the baud rate for NRZ is the same as the bit rate. Because it is the actual rate of signal

TABLE 1: Definition of signal encoding formats.

Nonreturn to Zero-Level (NRZ-L)
0 = high level
1 = low level

Nonreturn to Zero Inverted (NRZI)
0 = no transition at beginning of interval (one bit time)
1 = transition at beginning of interval

Manchester
0 = transition from high to low in middle of interval
1 = transition from low to high in middle of interval

Differential Manchester
Always a transition in middle of interval
0 = transition at beginning of interval
1 = no transition at beginning of interval

MLT-3
0 = output value same as for preceding interval
1 = if preceding value was either positive or negative, output value is 0
if preceding value was 0, output value is positive or negative, alternating for successive occurences

transitions rather than the bit rate that determines the bandwidth of a signal, the bandwidth required for the biphase schemes is considerably greater than for NRZ. On the other hand, the biphase schemes have several advantages:

- Synchronization: Because there is a predictable transition during each bit time, the receiver can synchronize on that transition. The biphase codes are therefore known as self-clocking codes.
- No dc component: Biphase codes have no dc component, yielding the benefits described earlier.
- Error detection: The absence of an expected transition can be used to detect errors. Noise on the line would have to invert both the signals before and after the expected transition to cause an undetected error.

The bulk of the energy in biphase codes is between one-half and one times the bit rate. Thus, the bandwidth is reasonably narrow and contains no dc component.

The more common Manchester code has been specified for the IEEE 802.3 standard for baseband coaxial cable and twisted-pair CSMA/CD bus LANs. Differential Manchester has been specified for the IEEE 802.5 Token Ring LAN, using shielded twisted pair.

MLT-3

MLT-3 is an encoding scheme used for the twisted-pair version of FDDI and for one of the versions of 100BASE-T. The effect of the MLT-3 scheme is to concentrate most of the energy in the transmitted signal below the frequency corresponding to one-third the bit rate. This reduces radiated emissions and therefore interference, which is a serious concern with unshielded twisted pair.

The MLT-3 encoding produces an output that has a transition for every binary 1 and that uses three levels: a positive voltage (+V), a negative voltage (-V) and no voltage (0). The encoding rules are best explained with reference to the encoder state diagram shown in Figure 2:

1. If the next input bit is 0, then the next output value is the same as the preceding value.
2. If the next input bit is 1, then the next output value involves a transition.

a. If the preceding output value was either +V or -V, then the next output value is 0.
b. If the preceding output value was 0, then the next output value is nonzero, and that output is of the opposite sign to the last nonzero output.

Bit transformation schemes

The NRZ schemes are not suitable for LAN transmission because of the lack of synchronization. For traditional LANs, such as 10-Mbps Ethernet and 4- and 16-Mbps Token Ring, biphase codes have proved adequate. However, when data rates of 100 Mbps or more are needed, the biphase codes produce an unacceptably high signaling rate. For example, at a date rate of 100 Mbps, a signaling rate of 200 Mbaud is required using a biphase code. To reduce bandwidth requirements, retain the synchronization characteristic of biphase codes, and achieve an effective signal spectrum shape, designers have turned to the use of bit transformation schemes. These schemes alter the input so that the desired transmission characteristics are achieved without resorting to biphase. The schemes that have found application in recent LAN equipment designs are 4B5B, 8B6T, 5B6B, and 8B10B.

A primary object of all of these bit mappings is to structure the input bit stream so that there are frequent transitions between the two binary digits. As long as transitions are frequent, a biphase code is not needed for synchronization, and the more efficient NRZ codes can be used. Table 2 shows which LAN systems use which bit mappings and signal encoding formats.

TABLE 2: Encoding schemes for LANs

IEEE 802.3 (CSMA/CD)	
10BASE5, 10BASE2, 10BASE-T	Manchester
100BASE-TX	4B5B/MLT-3
100BASE-FX	4B5B/NRZI
IEEE 802.5 (Token Ring)	Diffential Manchester
FDDI	
Optical fiber	4B5B/NRZI
Twisted pair	4B5B/MLT-3
100VG-AnyLAN	5B6B/NRZ
Fibre channel	8B10B/NRZ

4B5B

The 4B5B scheme, used in combination with NRZI, is a popular choice for 100-Mbps transmission over optical fiber. This scheme is used for the optical fiber versions of FDDI and Fast Ethernet.

For 4B5B, encoding is done 4 bits at a time. Each 4 bits of data are encoded into a symbol with 5 code bits, such that each code bit contains a single signal element; a block of 5 code bits is referred to as a code group. In effect, each set of 4 bits is encoded as 5 bits. There is a second stage of encoding: Each code bit of the 4B5B stream is treated as a binary value and encoded using NRZI. Because NRZI is a differential encoding format, it aids the ultimate decoding of the signal after it has been converted back from the optical to the electrical realm.

Since we are encoding 4 bits with a 5-bit pattern, only sixteen of the thirty-two possible patterns are needed for data encoding. The codes selected to represent the sixteen 4-bit data blocks are such that a transition is present at least twice for each 5-code group code. No more than three 0s in a row are allowed across one or more code groups.

The encoding scheme can be summarized as follows:

- A simple NRZ encoding is rejected because it does not provide synchronization; a string of 1s or 0s will have no transitions.

- The data to be transmitted must first be encoded to assure transitions. The 4B5B code is chosen over Manchester because it is more efficient.
- The 4B5B code is further encoded using NRZI so that the resulting differential signal will improve reception reliability.
- The specific 5-bit patterns chosen for the encoding of the sixteen 4-bit data patterns are chosen to guarantee no more than three 0s in a row to provide for adequate synchronization.

Those code groups not used to represent data are either declared invalid or assigned special meaning as control symbols. For example, two of the patterns (codes 11000 and 10001) always occur in pairs and act as start delimiters for a frame.

The combination of 4B5B and NRZI provides an efficient reliable transmission technique. At 100 Mbps, a signaling rate of 125 Mbaud is required. By contrast, when using Manchester or Differential Manchester, a signaling rate of up to 200 Mbaud is needed.

Although 4B5B/NRZI is effective over optical fiber, it is not suitable for use over twisted pair. The reason is that the signal energy is concentrated to produce undesirable radiated emissions from the wire. MLT-3, which is used on both 100BASE-TX and the twisted pair version of FDDI, is designed to overcome this problem. For these LANs, either shielded twisted pair or high-quality Category 5 unshielded twisted pair is used. In essence, the scheme used is 4B5B/MLT-3. At a data rate of 100 Mbps, the effect of using MLT-3 is to concentrate most of the energy in the transmitted signal below 30 MHz, which reduces radiated emissions.

8B6T

100BASE-T4 is designed to produce a 100-Mbps data rate over lower-quality voice grade, or Category 3 cable. The advantage is that in many existing buildings there is an abundance of voice-grade cabling and very little else. Thus, if this cabling can be used, installation costs are minimized.

With present technology, a data rate of 100 Mbps over one or two Category 3 pairs is impractical. Instead, 100BASE-T4 specifies that the data stream to be transmitted is split up into three separate data streams, each with an effective data rate of 331/3 Mbps. Four twisted pair are used. Data are transmitted using three pairs and received using three pairs. Thus, two of the pairs must be configured for bidirectional transmission.

As with 100BASE-X, a simple NRZ encoding scheme is not used for 100BASE-T4. This would require a signaling rate of 33 Mbps on each twisted pair and does not provide synchronization. Instead, a ternary signaling scheme known as 8B6T is used. With ternary signaling, each signal element can take on one of three values (positive voltage, negative voltage, or zero voltage). A pure ternary code is one in which the full information-carry capacity of the ternary signal is exploited. However, pure ternary is not attractive for the same reasons that a pure binary (NRZ) code is rejected: It lacks synchronization. 8B6T is designed to approach the efficiency of ternary and overcome this disadvantage.

With 8B6T, the data to be transmitted is handled in 8-bit blocks. Each block of 8 bits is mapped into a code group of six ternary symbols. The stream of code groups is then transmitted in round-robin fashion across the three output channels (Figure 3). Thus the ternary transmission rate on each output channel is:

$$\frac{6}{8} \times 33\frac{1}{3} = 25 \text{ Mbaud}$$

All possible 8-bit patterns are mapped into a unique code group of six ternary symbols. The mapping was chosen with two requirements in mind: synchronization and dc balance. For synchronization, the codes were chosen to maximize the average number of transitions per code group. The second requirement is to maintain dc balance so that the average voltage on the line is 0. For this purpose, all of the selected code groups either have an equal number of positive and negative symbols or an excess of one positive symbol. To maintain balance, a dc-balancing algorithm is used. In essence, this algorithm monitors the cumulative weight of all code groups transmitted on a single pair. Each code group has a weight of 0 or 1. To maintain balance, the algorithm may negate a transmitted code group (change all + symbols to - symbols and vice versa), so that the cumulative weight at the conclusion of each code group is always either 0 or 1.

5B6B

As with 100BASE-T4, a key objective of the 100VG-AnyLAN effort is to be able to achieve 100 Mbps over short distances using ordinary voice-

TABLE 3: 5B6B Encoding table.

Input quintet	Mode 2 output	Mode 4 output	Input quintet	Mode 2 output	Mode 4 output
00000	001100	110011	10000	000101	111010
00001	101100	101100	10001	100101	100101
00010	100010	101110	10010	001001	110110
00011	001101	001101	10011	010110	010110
00100	001010	110101	10100	111000	111000
00101	010101	010101	10101	011000	100111
00110	001110	001110	10110	011001	011001
00111	001011	001011	10111	100001	011110
01000	000111	000111	11000	110001	110001
01001	100011	100011	11001	101010	101010
01010	100110	100110	11010	010100	101011
01011	000110	111001	11011	110100	110100
01100	101000	010111	11100	011100	011100
01101	011010	011010	11101	010011	010011
01110	100100	100100	11110	010010	101101
01111	101001	101001	11111	110010	110010

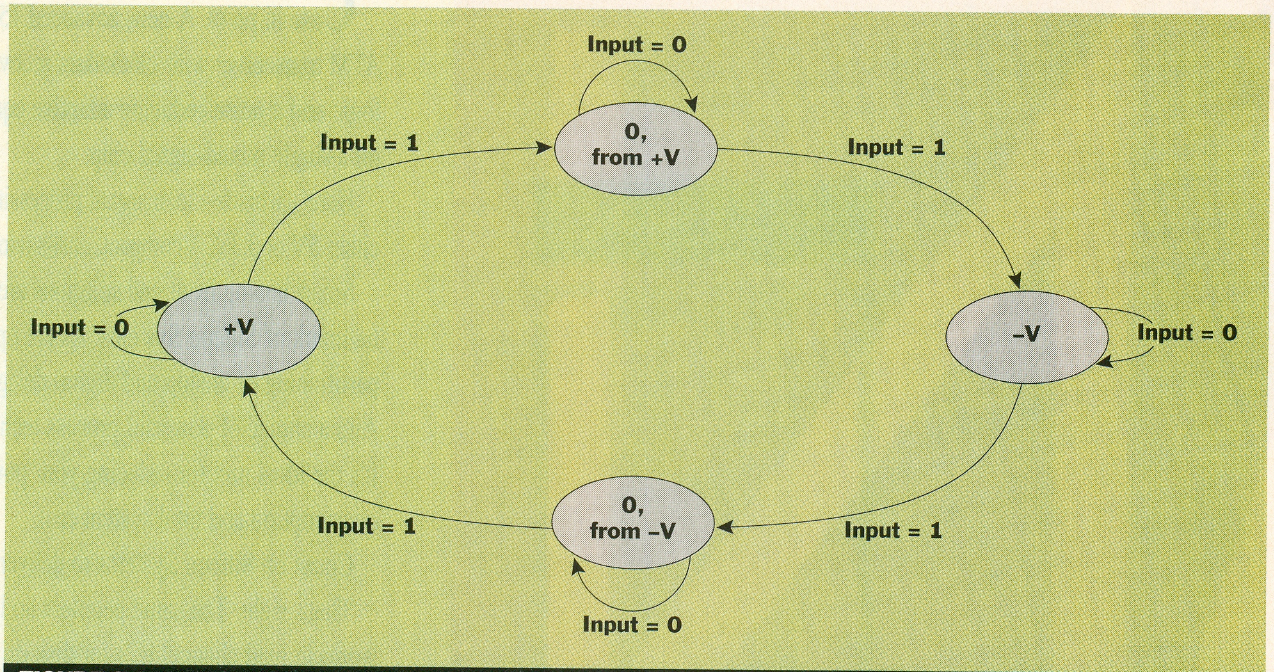

FIGURE 2: *MLT-3 encoder state diagram.*

grade (Category 3) cabling. To meet the objective, 100VG-AnyLAN specifies a novel encoding scheme that involves using four pair to transmit data in a half-duplex mode. Thus, to achieve a data rate of 100 Mbps, a data rate of only 25 Mbps is needed on each channel. An encoding scheme known as 5B6B is used.

The 5B6B encoding is used in the 100VG-AnyLAN specification in the following way. The bit stream to be transmitted is divided into 5-bit chunks (quintets), and each successive chunk is transmitted over a different channel in round-robin fashion. Thus, to achieve a data rate of 100 Mbps, a data rate of only 25 Mbps is needed on each channel.

The 5B6B scheme, which is used to ensure adequate transitions on each line for synchronization, is based on the same strategy as the 4B5B scheme described earlier. In this case, each group of 5 input bits is mapped into a set of 6 output bits. For an effective data rate of 25 Mbps, a signaling rate of 30 Mbaud is required.

With the 5B6B scheme, there are thirty-two possible 5-bit inputs. Ideally, we would like to assign to each 5-bit input a 6-bit code that has an equal number of 1s and 0s, which would maintain a dc balance of 0. However, there are only twenty 6-bit code words that have three 1s and 0s. These codes are assigned to twenty of the input patterns. For the remaining twelve input patterns, two code words are assigned: one with four 0s and two 1s (mode 2) and one with two 0s and four 1s (mode 4). Successive instances of any of these twenty-four unbalanced code words

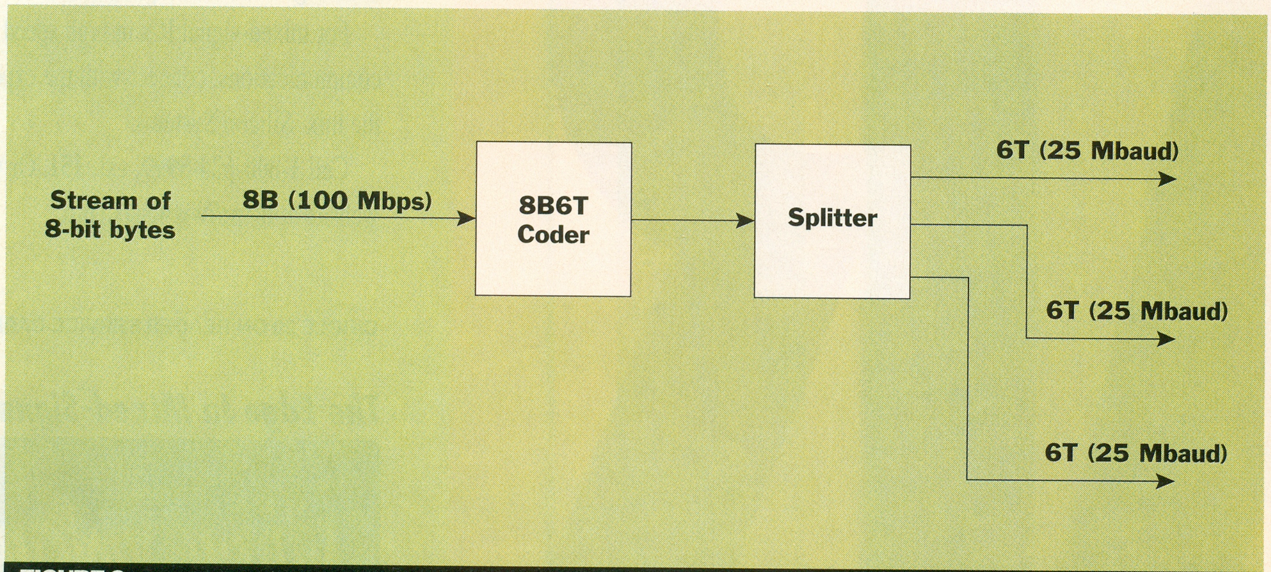

FIGURE 3: *8B6T transmission scheme.*

must alternate between mode 2 and mode 4 output to maintain balance. If, during reception, a station or repeater receives two of the same type of unbalanced words in a row (with any number of intervening balanced words), the receiver knows a transmission error has occurred and will ask for a retransmission of the data.

Table 3 shows the complete 5B6B encoding scheme. There is a unique output code word for twelve of the input patterns. For the rest, the transmitter keeps track of whether the last unbalanced transmitted word was mode 2 or mode 4, and transmits the appropriate output code word to maintain balance.

8B10B

Fibre Channel is a switched LAN specification designed for use over optical fiber at data rates up to 800 Mbps. The encoding scheme used for Fibre Channel is 8B10B. In this method, each group of 8 bits of data is converted into 10 bits for transmission. This scheme has a similar philosophy to the 4B5B scheme used for FDDI. The 8B10B scheme is more powerful than 4B5B in terms of transmission characteristics and error-detection capability.

The developers of this code list the following advantages:

- It can be implemented with relatively simple and reliable transceivers at low cost.
- It is well-balanced, with minimal deviation from the occurrence of an equal number of 1 and 0 bits across any sequence.
- It provides good transition density for easier clock recovery.
- It provides useful error-detection capability.

With 8B10B, each group of 8 input bits is mapped into a 10-bit code block. There is also a function called disparity control. In essence, this function keeps track of the excess of 0s over 1s or vice versa. If there is an excess in either direction, this is referred to as a disparity. If the current code block would add to that disparity, then the disparity control block complements the 10-bit code block. This has the effect of either eliminating the disparity or at least moving it in the opposite direction of the current disparity.

The evolution of LANs to higher speeds, sometimes over lower-quality transmission media, has sparked a corresponding evolution in digital-signal encoding techniques. The NRZ codes commonly used in EDP equipment and for magnetic recording are not efficient enough for high-speed LAN use. Even Manchester and Differential Manchester, developed for Ethernet and Token Ring equipment, fail to meet evolving needs. In recent years, a variety of schemes that combine bit mapping and signal encoding have appeared to meet these needs.

William Stallings is an independent consultant and president of Comp-Comm Consulting of Brewster, MA. He is a frequent lecturer on networking topics and the author of over a dozen books on data communications and computer networking. This article is based on material in Stallings' latest book, Local and Metropolitan Area Networks, *Fifth Edition (Prentice Hall, 1996). He can be reached at ws@shore.net.*