

# Selecting Mixed-Signal Components for Digital Communication Systems—An Introduction

by Dave Robertson

Communications is about moving information from point A to point B, but the computer revolution is fundamentally changing the nature of communication. Information is increasingly created, manipulated, stored, and transmitted in digital form—even signals that are fundamentally analog. Audio recording/playback, wired telephony, wireless telephony, audio and video broadcast—all of these nominally analog communications media have adopted, or are adopting, digital standards. Entities responsible for providing communications networks, both wired and wireless, are faced with the staggering challenge of keeping up with the exponentially growing demand for digital communications traffic. More and more, communications is about moving *bits* from point A to point B.

Digital communications embraces an enormous variety of applications, with radically different constraints. The transmission medium can be a twisted pair of copper wire, coaxial cable, fiber-optic cable, or wireless—via any number of different frequency bands. The transmission rate can range from a few bits per second for an industrial control signal communicating across a factory floor to 32 kbits/second for compressed voice, 2 Mb/s for MPEG compressed video, 155 Mbps for a SONET data trunk, and beyond. Some transmission schemes are constrained by formal standards, others are free-lance or developmental. The richness of design and architectural alternatives produced by such variety boggles the mind. The digital communications topic is so vast as to defy a comprehensive treatment in anything less than a shelf of books.

A communications jargon and a bewildering array of acronyms have developed, making it sometimes difficult for the communications system engineer and the circuit hardware designer to communicate with one another. Components have often been selected based on voltage-oriented specifications in the time domain for systems whose specifications are expressed in frequency and power. Our purpose here, and in future articles, will be to take a fairly informal overview of some of the fundamentals, with an emphasis on tracing the sometimes complex relationship between component performance and system performance.

The “communications perspective” and analytic tool set have also contributed substantially in solving problems not commonly thought of as “communications” problems. For example, the approach has provided great insight into some of the speed/bandwidth limits inherent in disk-drive data-recovery problems, where the channel from A to B includes the writing and reading of data in a magnetic medium—and in moving data across a high speed bus on a processing board.

**Shannon’s law—the fundamental constraint:** In general, the objective of a digital communications system is:

- to move as much data as possible per second

- across the designated channel
- with as narrow a bandwidth as possible
- using the cheapest, lowest-power, smallest-space (etc.) equipment available.

System designers are concerned with each of these dimensions to different degrees. Claude Shannon, in 1948, established the theoretical limit on how rapidly data can be communicated:

This means that the maximum information that can be transmitted through a given channel in a given time increases linearly with the channel’s bandwidth, and noise reduces the amount of information that can be effectively transmitted in a given bandwidth, but with a logarithmic sensitivity (a thousandfold increase in noise may result in a tenfold reduction in maximum channel capacity). Essentially, the “bucket” of information has two dimensions: bandwidth and signal-to-noise ratio (SNR). For a given capacity requirement, one could use a wide-bandwidth channel with relatively poor SNR, or a narrowband channel with relatively good SNR (Figure 1). In situations where bandwidth is plentiful, it is common to use cheap, bandwidth-hungry communications schemes because they tend to be insensitive to noise and implementation imperfections. However, as demand for data communication capacity increases (e.g., more cellular phones) bandwidth is becoming increasingly scarce. The trend in most systems is towards greater spectral efficiency, or bits capacity per unit of bandwidth used. By Shannon’s law, this suggests moving to systems with better SNR and greater demands on the transmit and receive hardware and software.

Let’s examine the dimensions of bandwidth (time/frequency domain) and SNR (voltage/power domain) a little more closely by considering some examples.

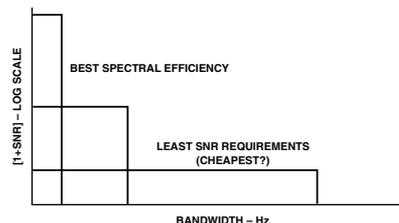


Figure 1. Shannon’s capacity limit: equal theoretical capacity.

| <b>IN THIS ISSUE</b>   |    |
|--|----|
| <b>Volume 30, Number 3, 1996, 24 Pages</b>                                       |    |
| Editor’s Notes, Authors  | 2  |
| Selecting mixed-signal components for digital communication systems—Introduction | 3  |
| TxDAC™ CMOS D/A converters are optimized for the communication                   |    |
| Transmit path  | 7  |
| Buffered multiplexers for video applications                                     | 10 |
| Single-chip direct digital synthesis (DDS) vs. the analog phase-locked loop      | 12 |
| For efficient signal processing in embedded systems, take a DSP, not a RISC      | 14 |
| New-Product Briefs:  |    |
| Operational Amplifiers and variable-compression Audio Preamplifier               | 16 |
| DSPs and Digital and Analog Audio chips  | 17 |
| ADC, 2-channel Data-Acquisition system, and Direct Digital                       |    |
| Synthesis chip   | 18 |
| Transceivers and switched-cap Regulated Inverter                                 | 19 |
| Ask The Applications Engineer—22: Current-feedback amplifiers—I                  | 20 |
| Worth Reading, More authors  | 22 |
| Potpourri  | 23 |

**PCM: A simple (but common) case:** Consider the simple case of transmitting the bit stream illustrated in Figure 2a, from a transmitter at location A to a receiver at location B (one may assume, that the transmission is via a pair of wires, though it could be any medium.) We will also assume that the transmitter and receiver have agreed upon both the voltage levels to be transmitted and the timing of the transmitted signals. The transmitter sends “high” and “low” voltages at the agreed-upon times, corresponding to 1s and 0s in its bit stream. The receiver applies a decision element (comparator) at the agreed-upon time to discriminate between a transmitted “high” and “low”, thereby recovering the transmitted bit stream. This scheme is called *pulse code modulation* (or PCM). Application of the decision element is often referred to as “slicing” the input signal stream, since a determination of what bit is being sent is based on the value of the received signal at one instant in (slice of) time. To transmit more information down this wire, the transmitter increases the rate at which it updates its output signal, with the receiver increasing its “slicing” rate correspondingly.

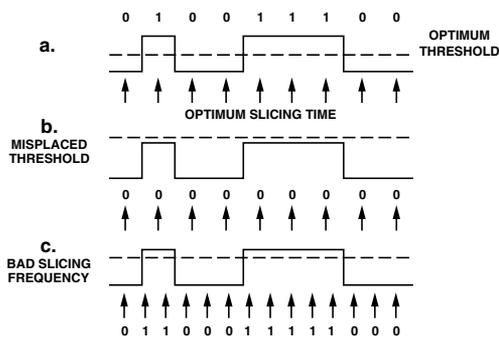


Figure 2. Simplified bit voltage transmission (PCM).

This simple case, familiar to anyone who has had an introductory course in digital circuit design, reveals several of the important elements in establishing a digital communications system. First, the transmitter and receiver must agree upon the “levels” that are to be transmitted: in this case, what voltage constitutes a transmitted “1”, and what voltage level constitutes a transmitted “0”. This allows the receiver to select the right threshold for its decision element; incorrect setting of this threshold means that the transmitted data will not be recovered (Figure 2b). Second, the transmitter and receiver must agree on the transmission frequency; if the receiver “slices” at a different rate than the bits are being transmitted, the correct bit sequence will not be recovered (2c). In fact, as we’ll see in a moment, there must be agreement on both frequency and phase of the transmitted signal.

How difficult are these needs to implement? In a simplified world, one could assume that the transmitted signal is fairly “busy”, without long strings of consecutive ones or zeros. The decision

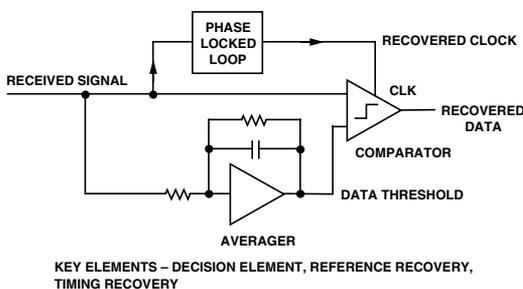


Figure 3. Idealized PCM.

threshold could then be set at the “average” value of the incoming bit stream, which should be some value between the transmitted “1” and transmitted “0” (half-way between, if the density of ones and zeros are equal.) For timing, a phase-locked loop could be used—with a center frequency somewhere near the agreed-upon transmit frequency; it would “lock on” to the transmitted signal, thereby giving us an exact frequency to slice at. This process is usually called *clock recovery*; the format requirements on the transmit signal are related to the performance characteristics of the phase-locked-loop. Figure 3 illustrates the elements of this simplified pulse receiver.

**Bandwidth Limitations:** The real world is not quite so simple. One of the first important physical limitations to consider is that the transmission channel has finite bandwidth. Sharp-edged square wave pulses sent from the transmitter will be “rounded off” by a low bandwidth channel. The severity of this effect is a function of

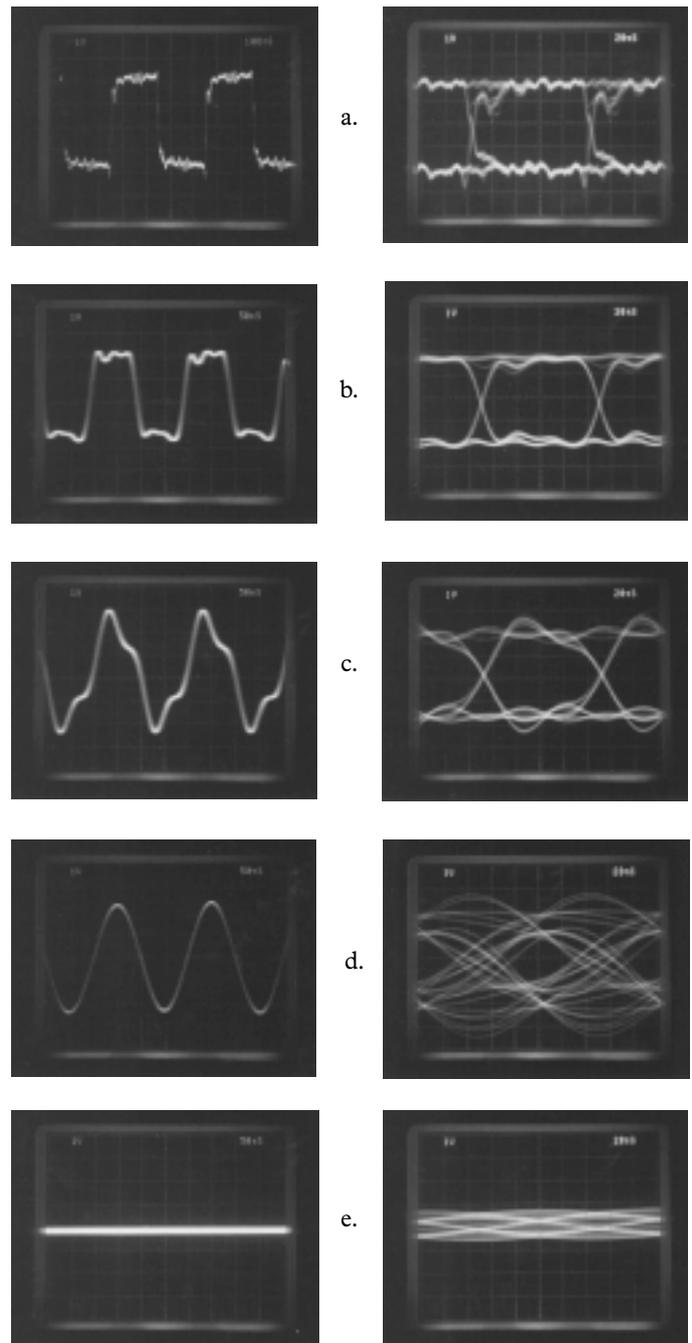


Figure 4. Scope waveforms vs. time (L) and eye diagrams (R).

the channel bandwidth. (Figure 4). In the extreme case, the transmitted signal never gets to a logical “1” or “0”, and the transmitted information is essentially lost. Another way of viewing this problem is to consider the impulse response of the channel. An infinite bandwidth channel passes an impulse undistorted (perhaps with just a pure time delay). As the bandwidth starts to decrease, the impulse response “spreads out”. If we consider the bit signal to be a stream of impulses, inter-symbol interference (ISI) starts to appear; the impulses start to interfere with one-another as the response from one pulse extends into the next pulse. The voltage seen at the Receive end of the wire is no longer a simple function of the bit sent by the transmitter at time  $t_1$ , but is also dependent on the previous bit (sent at time  $t_0$ ), and the following bit (sent at time  $t_2$ ).

Figure 4 illustrates what might be seen with an oscilloscope connected to the Receive end of the line in the simple noisy communications system described above for the case where the bandwidth restriction is a first-order lag (single R-C). Two kinds of response are shown, a portion of the actual received pulse train and a plot triggered on each cycle so that the responses are all overlaid. This latter, known as an “eye” diagram, combines information about both bandwidth and noise; if the “eye” is open sufficiently for all traces, 1s can be easily distinguished from 0s. In the adequate bandwidth case of Figure 4a, one can see unambiguous 1s, 0s, and sharp transitions from 1 to 0. As the bandwidth is progressively reduced, (4b, 4c, 4d, 4e), the 1s and 0s start to collapse towards one another, increasing both timing- and voltage uncertainty. In reduced-bandwidth and/or excessive-noise cases, the bits bleed into one another, making it difficult to distinguish 1s from 0s; the “eye” is said to be *closed* (4e).

As one would expect, it is much easier to design a circuit to recover the bits from a signal like 4a than from 4d or 4e. Any misplacement of the decision element, either in threshold level or timing, will be disastrous in the bandlimited cases (d, e), while the wideband case would be fairly tolerant of such errors. As a rule of thumb, to send a pulse stream at rate  $F_s$ , a bandwidth of at least  $F_s/2$  will be needed to maintain an open eye, and typically wider bandwidths will be used. This *excess bandwidth* is defined by the ratio of actual bandwidth to  $F_s/2$ . The bandwidth available is typically limited by the communication medium being used (whether 2000 ft. of twisted-pair wire, 10 mi of coaxial cable etc.), but it is also necessary to ensure that the signal processing circuitry in the transmitter and receiver do not limit the bandwidth.

Signal processing circuitry can often be used to help mitigate the effects of the intersymbol interference introduced by the bandlimited channel. Figure 5 shows a simplified block diagram of a bandlimited channel followed by an equalizer, followed by the bit “slicer”. The goal of the equalizer is to implement a transfer function that is effectively the inverse of the transmission channel over a portion of the band to extend the bandwidth. For example, if the transmission channel is acting as a low pass filter, the equalizer might implement a high-pass characteristic, such that a signal passing through the two elements will come out of the equalizer undistorted over a wider bandwidth.

Though straightforward in principle, this can be very difficult to implement in practice. To begin with, the transfer function of the transmission channel is not generally known with any great precision, nor is it constant from one situation to the next. (You

and your neighbor down the street have different length phone wires running back to the phone company central office, and will therefore have slightly different bandwidths.) This means that these equalizers usually must be tunable or adaptive in some way. Furthermore, considering Figure 5 further, we see that a passive equalizer may flatten out the frequency response, but will also attenuate the signal. The signal can be re-amplified, but with a probable deterioration in signal-to-noise ratio. The ramifications of that approach will be considered in the next section. While they are not an easy cure-all, equalizers are an important part of many communications systems, particularly those seeking the maximum possible bit rate over a bandwidth-constrained channel. There are extremely sophisticated equalization schemes in use today, including decision feedback equalizers which, as their name suggests, use feedback from the output of the decision element to the equalization block in an attempt to eliminate trailing-edge intersymbol interference.<sup>1</sup>

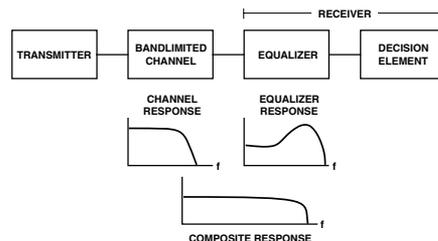


Figure 5. Channel equalization.

**Multi-level symbols—sending more than one bit at a time:**

Since the bandwidth limit sets an upper bound on the number of pulses per second that can be effectively transmitted down the line, one could decide to get more data down the channel by transmitting two bits at a time. Instead of transmitting a “0” or “1” in a binary system, one might transmit and receive 4 distinct states, corresponding to a “0” (00), “1” (01), “2” (10), or “3” (11). The transmitter could be a simple 2-bit DAC, and the receiver could be a 2-bit ADC. (Figure 6). In this kind of modulation, called pulse-amplitude modulation (PAM), additional information has been encoded in the amplitude of the bit stream.

Communication is no longer one bit at a time; multiple-bit words, or *symbols*, are being sent with each transmission event. It is then necessary to distinguish between the system’s bit rate, or number of bits transmitted per second, and its symbol rate, or baud rate, which is the number of *symbols* transmitted per second. These two rates are simply related:

$$\text{bit rate} = \text{symbol rate (baud)} \times \text{bits/symbol}$$

The bandwidth limitations and intersymbol interference discussed in the last section put a limit on the realizable symbol rate, since they limit how closely spaced the “transmission events” can be in time. However, by sending multiple bits per symbol, one can increase the effective bit rate, employing a *higher-order modulation* scheme. The transmitter and receiver become significantly more complicated. The simple switch at the transmitter has now been replaced with a DAC, and the single comparator in the receiver is

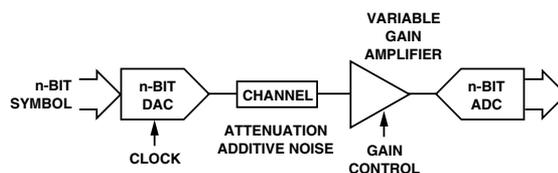


Figure 6. Simplified PAM transmitter/receiver.

<sup>1</sup>The field of disk-drive *read*-channel design is a hotbed of equalizer development in the ongoing struggle to improve access specs.

now an A/D converter. Furthermore, it is necessary to use more care to properly scale the amplitude of the received signal; more information is needed than just the sign. Making the simplifying assumption that the A/D converter, representing the receiver, is implemented as a straight flash converter, it is manifest that the receiver hardware complexity grows exponentially with the number of bits per symbol: one bit, one comparator; two bits, 3 comparators; three bits, 7 comparators, etc. Depending on the particular application, circuit cost should not quite increase exponentially with bits per symbol, but it generally will be a steeper-than-linear increase. However, hardware complexity is not the only limiting factor on the number of bits per symbol that can be transmitted.

### NOISE LIMITATIONS

Consider again the simple case of one-bit-per-symbol PCM modulation. Assuming that 1 V is used to send a "1", and -1 V to send a "0", the simple receiver (Figure 3) is a comparator with its decision threshold at 0 V. In the case where the bit being received is a "0", and the channel bandwidth is wide enough so that there is virtually no intersymbol interference, in a noiseless environment, the voltage at the receiver is expected to be -1 V. Now introduce additive noise to the received signal (this could come from any number of sources, but for simplicity and generality, assume it to be gaussian white noise that could correspond to thermal noise). At the moment the decision element is applied, the voltage at the comparator will differ from -1 V by the additive noise. The noise will not be of real concern unless it contains values that will push the voltage level above 0 V. If the noise is large enough (and in the right sign) to do this, the decision element will respond that it has received a "1", producing a bit error. *In the eye diagram of Figure 4d, the noise would produce occasional closures of the "eye".*

If the system is modified to send a 4-bit (16-level) symbol, with the same peak-to-peak voltage, -1 V corresponds to "0" (0000), and +1 V corresponds to "15" (1111). Now the incremental threshold between "0" and the next higher level, "1", is much smaller: 16 distinct states must fit into the 2-V span, so the states will be roughly 125 mV apart, center-to-center. If the decision thresholds are placed optimally, the "center" of a state will be 62.5 mV away from adjacent thresholds. In this case, >62.5 mV of noise will cause a "bit error". If the initial assumption holds and the additive noise is gaussian in nature, one can predict from the rms noise value how often the noise will exceed this critical value. Figure 7 shows the error threshold of 62.5 mV for the probability density functions of two different rms noise values. From this, one can predict the bit error rate, or how often the received data will be interpreted incorrectly for a given transmitted bit rate.

Special care must be taken as to how the data is encoded: if the code 1000 is one threshold away from the code 0111, a small noise excursion would actually cause all 4 bits to be misinterpreted. For this reason, Gray code (which changes only one bit at a time between adjacent states—e.g., 00, 01, 11, 10) is often used to minimize the bit error impact from a misinterpretation between two adjacent states.

So, despite the increase in bit rate, there are limitations to using higher-order modulation schemes employing more bits per symbol: not only will the hardware become more complex, but, for a given noise level, bit errors will be more frequent. Whether the bit error rate is tolerable depends very much on the application; a digitized

voice signal may sound reasonable with a bit error rate of  $10^{-5}$ , while a critical image transmission might require  $10^{-15}$ .

Bit errors can be detected and corrected by various coding and parity schemes, but the overhead introduced by these schemes eventually consumes the additional bit capacity gained from increasing the symbol size. One way to try to increase the signal-to-noise ratio (SNR) is to increase transmitted power; for example, increase signal amplitude from 2 V peak-to-peak to 20 V peak-to-peak, thereby increasing the "error threshold" to 625 mV. Unfortunately, increasing the transmitted power generally adds to the cost of the system. In many cases, the maximum power that can be transmitted in a given channel may be limited by regulatory authorities for safety reasons or to ensure that other services using the same or neighboring channels are not disturbed. Nevertheless, in systems that are straining to make use of all available capacity, the transmit power levels will generally be pushed to the maximum practical/legal levels.

Voltage noise is not the only kind of signal impairment that can degrade the receiver performance. If timing noise, or jitter, is introduced into the receiver "clock," the decision "slicer" will be applied at sub-optimal times, narrowing the "eye" (Figs. 4a-4d) horizontally. Depending on how close the channel is to being band-limited, this could significantly decrease the "error threshold," with increased sensitivity to voltage noise. Hence, SNR must be determined from the combination of voltage-domain and time-domain error sources.

*This is the first in a series of articles offering an introduction to topics in communications. In the next issue, we'll discuss various modulation schemes and ways of multiplexing multiple users in the same channel. A*

**For Further Reading:** This article scratches the surface of a very complex field. If your appetite for information has been whetted, here are a few suggested texts (bibliographies within these books will fan out to a wider list):

*Electronic Communication Systems—a complete course*, 2nd edition, by William Schweber. Englewood Cliffs, NJ: Prentice Hall ©1994. A good basic introduction to communications fundamentals, with an emphasis on intuitive understanding and real-world examples. No more than one equation per page.

*Digital Communication* (2nd edition), by Edward Lee and David Messerschmitt. Norwell, MA: Kluwer Publishing, ©1994. A more comprehensive and analytical treatment of digital communications.

*Wireless Digital Communications: Modulation and Spread-Spectrum Applications*, by Dr. Kamilo Feher. Englewood Cliffs, NJ: Prentice Hall, ©1995. A fairly rigorous analysis of different wireless modulation schemes, with insights into particular strengths and weaknesses of each, and discussion of why particular schemes were chosen for certain standards.

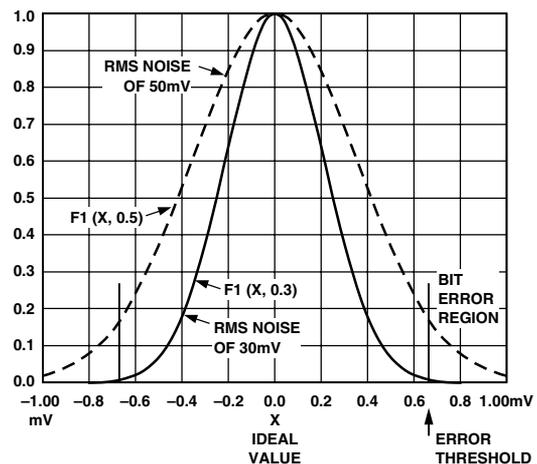


Figure 7. Ideal signal plus noise vs. error threshold: Threshold at  $2\sigma$ , and threshold at  $1\sigma$ .