High-Speed Digital Design

Key similarities and differences between high-speed digital and microwave hardware

First published in IEEE Microwave Magazine, August 2011.

> High-speed digital design, in contrast to digital design at low speeds, emphasizes the behavior of passive circuit elements. These passive elements may include the wires, circuit boards, and integrated-circuit packages that make up a digital product. At low speeds, passive circuit elements are just part of a product's packaging. At higher speeds they directly affect electrical performance.—*High-Speed Digital Design, a Handbook of Black Magic* [1]

Introduction

High-speed digital design and microwave design share a rich heritage of electrical engineering theory, yet, key differences separate these two fields and the people that work in them.

For example, most digital engineers are not equipped either by training or natural inclination to comprehend the intricacies of traveling-wave physics [2]. They may learn necessary bits and pieces of this theory through on-the-job training, but rarely attain the full level of physical understanding common among professionals in the microwave industry.

If digital technology is to advance, with speeds rising to levels previously reserved for only the highest-technology microwave systems, this author believes that more cross-pollination of the fields must occur.

Towards the goal of developing better communication between the two disparate fields, the following survey of 29 selected topics addresses design factors common to both digital and microwave engineering, with an eye towards explaining the terminology and conventions prevalent in digital engineering. The targets of discussion include digital transmitters, transmission pathways and receivers, all in the context of a high-speed serial link.

Such a collection of topics can never be all-encompassing. I can only hope that, at least, you find in this list some topics that seem unusual or pique your interest in further reading.

About the author: Dr. Howard Johnson gained his PhD from Rice University in 1982 in the field of electrical engineering, specializing in communications and system theory. He has played a key role developing groundbreaking digital backplane traveling-wave structures (ROLM), transceivers for the world's first gigabit local area network (Ultra Network Technologies), transceivers for the first 10G backplane applications (Accelerant Networks), and international standards for computer networking (IEEE 802.3 Fast Ethernet and Gigabit Ethernet). His books, *High-Speed Digital Design: a Handbook of Black Magic* (1993) and *High-Speed Signal Propagation: More Black Magic* (2003), set the pace for on-the-job engineering education in his field. He is the "Signal Integrity" columnist for EDN magazine, and a frequent guest lecturer at Oxford University. Dr. Johnson teaches public and private seminars around the globe. www.sigcon.com.

Digital Signal Characterization

This paper contemplates high-speed digital signals of an electrical nature ranging in amplitude from a few hundred millivolts to two and one-half volts, synchronously clocked at rates of 1-20 Gb/s, as typically found in consumer, communications, industrial, medical and military computing devices. Such signals are characterized by a level, a maximum bandwidth, a low-frequency cutoff, and jitter.

Levels

The level of a digital signal is characterized by its peak-to-peak voltage swing into a stated load, not its average signal power as is common in the RF world.

Most (but not all) high-speed serial links use binary coding, that is, they transmit one of only two discrete signal levels during each digital baud interval. For a symmetric binary signal centered about zero, the peak signal voltage and the root-mean-square (RMS) voltage are the same. Manufacturers of digital drivers universally specify the full peak-to-peak output, also called the "signal swing", rather than the one-sided peak (equal to rms) signal level because the former number makes a larger, more impressive-sounding specification.

A high-speed signal may be single-ended (one wire plus return path) or differential (two wires, with or without an additional reference path). When differential signaling is employed, the peak-to-peak differential signal level equals twice the peak-to-peak level on either wire alone.

In cases where multi-level signaling is employed, such as in the 1000 Mb/s, unshielded twisted-pair, physical layer transceiver specified as part of Gigabit Ethernet [3], the peak-to-peak voltage swing shown in Figure 1 still serves as the primary means of signal level specification.

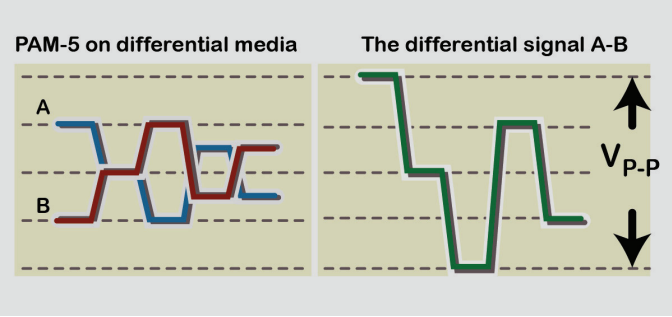

Figure 1—*The peak-to-peak swing V_{P-P} for a differential signal comprising two antipodal signal elements equals twice the swing of either element alone.*

Over time, the typical output levels for popular serial link drivers have decreased from a nominal value of +/-12V in 1969 (EIA RS-232-C [4]) to +/- 0.6V in 2007 (PCI-SIG PCIe 2.1 [5]). This decrease brings with it the advantage of a natural reduction in unwanted signal emissions, but the concomitant disadvantage of an increased susceptibility to external interference.

Maximum Bandwidth

RF communication channels are often designed to accommodate a particular frequency range from some low-frequency cutoff to an upper high-frequency limit. The high-frequency limit required for adequate transmission of digital signals can be quite high, in some cases orders of magnitude higher than the data transmission rate. The old rule that the spectrum of a digital signal contains significant power only up to the 5th harmonic of the data rate (or 7th, or 9th) is simply not true. As indicated in Figure 2, the actual upper limit to the bandwidth of a digital signal depends on the rise and fall times of the individual signal edges, not the data transmission rate, according to the following approximate rules which apply for Gaussian-shaped signal edges:

$$F_{3dB} \approx 0.338/T_{10\text{-}90\%} \qquad \text{[EQN 1]}$$

$$F_{6dB} \approx 0.5/T_{10\text{-}90\%} \qquad \text{[EQN 2]}$$

These numbers indicate the frequency at which the actual signal spectrum falls, due to the low pass filtering effect of its finite-speed rising and falling edges, either 3dB or 6dB, respectively, below the spectrum of an theoretical signal with the same data pattern, but infinitely-fast rising and falling edges. In Figure 2, the spectrum of a theoretical infinitely fast-edged waveform would project forward along the -20dB/decade slope forever. The high-frequency cutoff marked in the figure as "H.F. cutoff" marks the -6dB rolloff point.

Many digital signals are specified in terms of their 20-80% rise/fall time because that standard yields a faster-looking, more attractive specification. The relation between the 20-80% rise or fall time and the 10-90% rise or fall time depends heavily on the exact shape of the rising and falling edges. Signals subject to large amounts of skin-effect loss, for example, exhibit an exaggerated ratio between the two measures of signal edge speed and do not conform well to the bandwidth approximations listed.

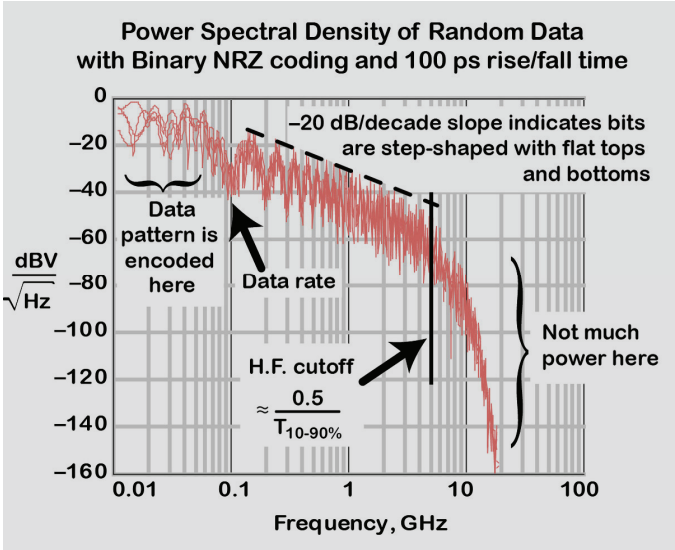

Figure 2—*Negligible amounts of signal power extend beyond the high frequency cutoff (four example data sequences shown).*

Low-Frequency Cutoff

The majority of digital signals contain a significant DC bias and are DC-coupled. The low-frequency content of a digital signal includes both a fixed DC bias and a time-varying component. The value of the time-varying component depends on the particular data being transmitted. An infinitely long datastream of digital one's, for example, concentrates all its power into a single spectral peak at DC. Such a signal could not pass through a transformer, or an AC-coupled amplifier, without requiring DC level restoration. While specialized DC-level restoration circuits are common in the RF world (the clamping circuit used to restore the blanking pulse amplitude within National Television System Committee (NTSC) analog television receivers comes to mind), digital designers typically use *data coding* to perform this function. A data coding algorithm works on blocks of data bits. It modifies the data pattern within each block and adds additional bits in a way that guarantees two functionally important properties: a minimum edge transition density and DC balance. For example, the popular 8B10B code [6] popularized in Fiber Channel and other local area network standards operates on blocks of 8 bits, producing 10-bit code words (Figure 3). Each 10-bit code word contains, at the expense of the two additional bits that must be transmitted, a minimum of three edge transitions and a very nearly balanced number of one's and zero's. The high transition density simplifies phase-locked loop clock recovery in the receiver. The DC balance property excludes long strings of one's or zeros, making it possible to send this signal through AC-coupled circuits without requiring DC level restoration [7], [8].

Jitter (Phase noise)

Carrier-based RF communication systems always exhibit small sidebands of power on either side of the main carrier frequency. The total sideband power, divided by the main carrier power, indicates the amount of carrier phase noise. Digital systems have no central carrier. Their spectrum is splattered from near-DC up to the baud rate and beyond. Imperfections in the exact baud rate of a digital transmission system are measured in the time domain, using tools that examine the precise location of signal zero crossings and compare those times against an ideal template. The measure of digital timing imperfection is called *jitter*. In 1998 the ANSI Fibre Channel committee T11.2 documented a general methodology now widely accepted for defining and measuring random and deterministic sources of jitter. [9]

Transmitters

Transmitting circuits for high-speed digital applications are fully integrated with no external discrete components. They may occur in numbers ranging from 1 to several thousand on a single integrated circuit (IC). Each circuit, operating from a single power supply of 1 to 2.5 Volts, creates an output signal with a AC rms power level of approximately 5-50 mW. The circuits often, but not

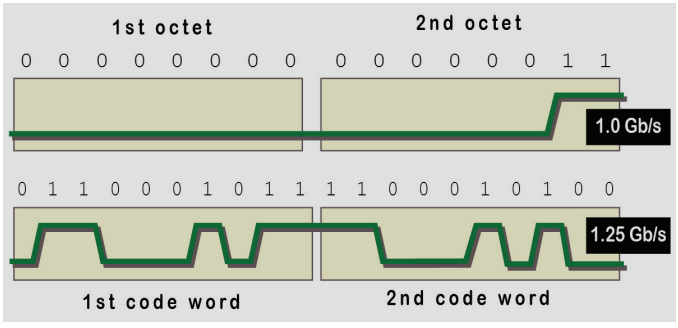

Figure **3**—*The 8B10B data code adds two bits of overhead to every 8 data bits.*

always, operate in a Class-D mode [31], exhibit poor control over their output impedance, and suffer greatly from inadequate packaging.

Class-D operation

The measure of success for a binary driver is the quality of the waveform it creates when driving its designated nominal load in just two discrete states. Since no intermediate output states are required and there is no requirement for linear superposition of signals, the driver may employ highly non-linear output stages. Figure 4 illustrates a typical *totem-pole* output stage. A zero-volt input turns on the P-channel field-effect transistor (FET) at the top of the circuit and cuts off the N-channel FET at the bottom, creating a logic-high output. An input of Vcc does the opposite, creating a logic-low output [10]. A totem-pole circuit behaves much like the output stage of a class-D push-pull amplifier.

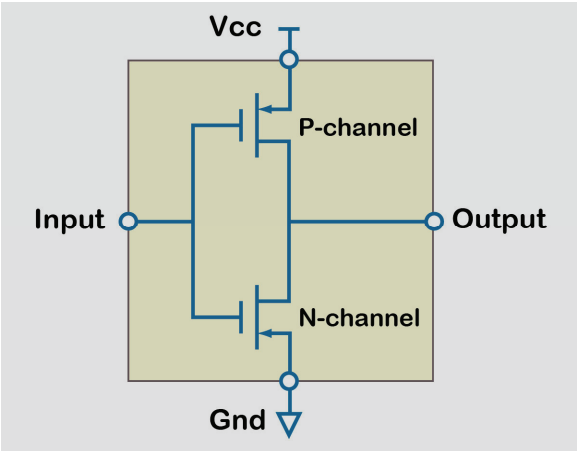

Figure 4—*A simple totem-pole inverter incorporates two switching transistors.*

Variation in output impedance

A digital totem-pole driver operates on a single power supply set to a very low voltage in the range of 1 to 2.5 volts. At very low voltages the pre-driver "Input" in Figure 4 may lack sufficient drive voltage to turn either FET completely on or off. As a result, both halves of the totem-pole forever swing back and forth between states of "more on" and "less on".

The circuit output impedance in that condition varies wildly with temperature, age, power supply voltage, and from part to part. One strategy for using such devices is to simply manufacture rather large FET outputs, with extraordinarily low output impedances, and place in series with each output an accurate resistor. The variations in FET output impedance, while still large in relative terms, will thereby appear small in comparison to the resistor, rendering the overall circuit performance and output levels more predictable.

Variety of output structures

Many other output structures, including pull-up only circuits (bipolar Emitter-Coupled Logic , ECL, [11]), pull-down-only circuits (Gunning Transistor Logic, GTL, [12]), and current-steering output stages (Low-Voltage Differential Signaling, LVDS, [13], [14]) are possible. Regardless the marketing claims made in favor of these alternative structures, in every case what matter most to the user of a digital device are the quality of the waveform created in the two necessary states and the cost of the component not how the driver was built internally.

Multiple output levels

The driver circuits employed in high-speed serial links at 2.5 Gb/s and beyond often produce not just two, but four or more discrete output levels. The driver uses these additional levels to temporarily boost the output at each data transition while holding the output steady during consecutive runs of identical bits. The little boosts on each signal edge emphasize the spectral content at higher frequencies. The boosting technique, called *pre-emphasis*, helps overcome dielectric and skin-effect power losses that can occur at higher frequencies [15], and section 8.2 of [27]. A related term, *de-emphasis*, employs an almost identical strategy except that the overall boosted waveform amplitude is scaled back so that the peak excursions appear no larger than the original signal. Of the many possible forms of equalization, pre-emphasis has gained favor in digital applications because it is easy to understand, it can be fully integrated within the transmitter silicon, and its output can be tested against a specification to verify compliance. Equalization employed within a receiver is more difficult to test (see "Receivers"). Multiple output levels can also be used to implement advanced data-coding techniques that pack more than one bit of information into each transmitted baud interval [29], [30].

Maximum speed

The maximum attainable operating bit rate for a transmitter is limited by two main factors: the driver's natural rise/fall time, for which a good figure of merit is the product of the driver's output impedance times its own parasitic capacitance, and the frequency response of the driver's physical packaging [16]. The

natural rise/fall time improves with each successive generation of silicon technology. The packaging performance tends to lag behind, generally being improved only when an IC manufacturer encounters a critical failure forcing re-evaluation of packaging technology. The IEEE journal, *Transactions on Advanced Packaging*, is a good place to begin your researches about the upper limits of package performance [17].

Signal Conveyance

This paper contemplates digital signals operating at speeds of 1-20 Gb/s, in a single-ended or differential format, spanning distances from 0.01 to 1 meters. The primary means of conveyance for such signals is electrical, using the Transverse Electric and Magnetic (TEM) mode on printed circuit board (PCB) traces or coaxial cables. Digital signals occur in systems at astonishing density, and must endure terrifying amounts of crosstalk as a result. In addition, the signals undergo significant degradation due to connectors, vias, and PCB materials.

TEM mode

High-speed digital signals within IC packages, on printed circuit boards, and in cables are carried almost exclusively in a TEM mode of propagation. The TEM mode requires the flow of physical current in a signal conductor, and an equal and opposite flow of current in a return conductor. The signal and return currents remain in-phase at each point along the transmission medium.

Many of the difficulties associated with reflections, crosstalk and radiation that plague digital architectures result from places like wirebond wires, vias, and connectors where the signal conductors are removed from their otherwise continuous and uniform relation to the return wires.

Microstrip traces

In the case of a microstrip trace, the signal propagates in a *quasi-TEM mode*, a hybrid mode very much like the TEM mode, but slightly different [18]. The quasi-TEM mode is useful to digital designers up to a much higher frequency than it might be to a microwave designer owing to the fact that most high-speed digital designs use very thin layers, a spacing of 4 mils between the reference plane and the traces being typical. Digital signals at bit rates up to 40 Gb/s are routinely conveyed through microstrip traces without incident.

Solid reference plane layers

Most high-speed digital boards contain one or more solid reference planes used as return conductors for the purpose of enabling TEM mode propagation on the signal wires. The impedance of each trace to its reference plane (or combination of reference planes) typically ranges from 25-85 ohms. There is no optimum value of PCB trace impedance, as there is for coaxial cables [19], [20].

Broadband nature of digital signals

Tricks involving Smith Charts, stub tuning, and reactive component neutralization only work on narrowband signals. There are only two "cute tricks" I know that work in the digital world. First, the constant-resistance network in Figure 5 can counteract reflections from the parasitic capacitance of a receiver, at the expense of an degradation in the received signal rise/fall time [21]. Second, a narrowing of the PCB trace width on either side of a capacitive load, as illustrated in Figure 6, helps mitigate the impact of reflections that might otherwise occur [22].

Layout density

A typical digital system laid out using 4-mil traces, 8-mil on center, using 24 total layers with 8 horizontally-oriented routing layers achieves an astonishing theoretical layout density of 1000 signal lines per linear inch. A Ball-Grid Array (BGA) package with balls on a square grid with 0.8mm centers, and an 4:1 ratio of signal pins to power and ground pins achieves a connection density of 800 signals per square inch. Such numbers are common for ordinary, synchronous digital signals. High speed serial links, on the other hand, cannot be laid out with traces so fine, or on such a dense pitch. Such fine traces contribute excessively to skin effect loss at high

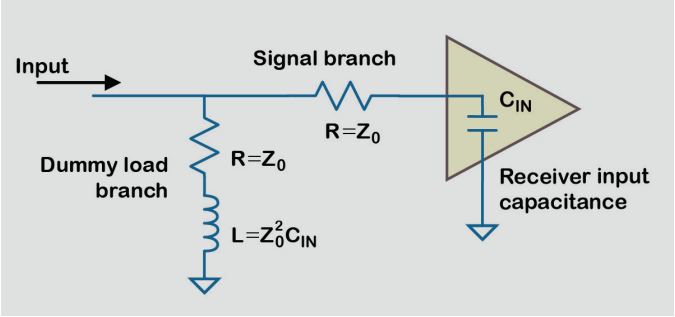

Figure 5—*The input impedance of this constant-resistance network, if perfectly implemented, equals Z_0 at all frequencies.*

frequencies, and the dense pitch contributes too much crosstalk. [23]

SSO Noise

Signals passing through an array of wirebond wires, BGA balls, or connector pins, pick up a type of crosstalk called simultaneously switching output noise (SSO Noise) [24], the exact modeling of which is best tackled using a fullwave electromagnetic field solver in conjunction with a lumped-element circuit model of the driving circuit. SSO noise is the result of many signals activating at once, all affecting a single receiver or set of receivers. SSO noise tends to affect all the receivers within a group in the same way; it may thus be mitigated by the use of differential signaling.

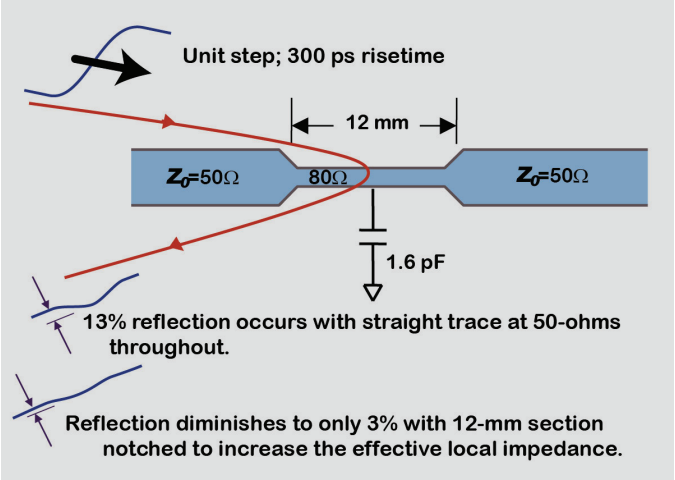

Figure 6—*A skinny section of PCB trace mitigates the reflections from a small capacitive load.*

Differential signaling

Digital signals carried in antipodal form on two wires are called *differential signals*. The two element s of a differential pair are generally routed close together on a PCB. Differential connections can help mitigate the otherwise destructive potential of simultaneously switching output noise. The practical result of differential signaling is a reduction by a factor of two to four in the number of power and ground pins required to control SSO, at the expense of a doubling of the number of signal pins. Differential connections do relatively little to mitigate crosstalk picked up by parallel traces within a PCB. Crosstalk with the PCB itself is controlled mostly by the spacing between adjacent traces.

Common-mode balance

Differential connections as implemented in typical digital systems are balanced only to within about one part in ten (at best). This contrasts with the standard implementation of unshielded-twisted pair connections which typically use transformer coupling and common-mode chokes to attain a very fine degree of balance. Because the two traces of an ordinary digital differential pair are not well balanced, and thus convey a substantial common-mode signal, they must be routed using the same general rules for single-ended signals, namely, both signals should maintain continuous proximity to their respective return conductors. [25] When placed close together, the fields surrounding the two elements of a differential pair interact, lowering the effective differential impedance of the pair. The widths of the PCB traces comprising a differential pair are commonly adjusted to account for that interaction. Even when placed exceptionally close together, PCB traces still couple more strongly to the underlying reference plane than to each other.

Connectors

In comparison to good-quality RF connectors, digital connectors are often produced at much greater density, and much lower cost, with a concomitant loss of signal fidelity, A large backplane connector, for example, may comprise some 400 signal connections each with ±3 dB passband performance and 15% aggregate crosstalk. The best connectors for high-speed digital applications maintain continuous proximity of the return conductors to the signals all the way through the connector from end to end. This requires either a vast number of ground pins almost equal to, or in some cases exceeding, the number of signal pins, or solid reference plane layers isolating rows of pins.

Vias

After connectors, the next most deleterious structures in a typical PCB are signal vias. As Figure 7 makes clear, the larger the via, the greater its parasitic capacitance, C[p]. Large vias on thick backplanes are absurdly difficult to manage at high speeds.

The primary measure of performance for a signal via is its transmission gain, however, since the via is a passive, lossless device the reflection coefficient provides a more sensitive indicator of its imperfections. A healthy focus on reflection coefficients also keeps one constantly reminded that double-reflections in a complex system create

forward-moving waves which ultimately affect the overall system transmission gain.

Low-loss dielectric materials

Alumina, while popular for some microwave applications, rarely appears in high-speed digital designs. It has a high dielectric constant that, while good for shrinking the size of microwave resonant structures, imposes too much latency for digital applications. An advanced digital design would more likely use a low-loss material with a dielectric constant between 2.6 and 3.5 (Table 1).

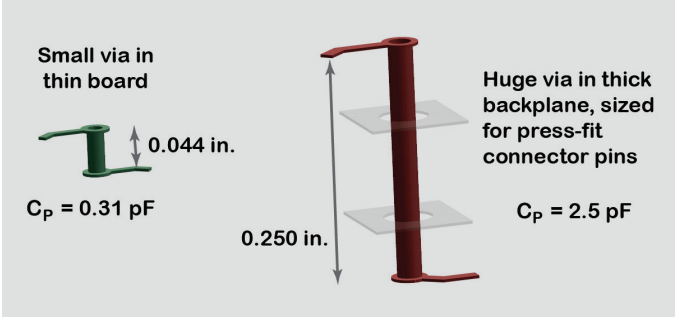

Figure 7—*The quasi-static parasitic shunt capacitance to ground of a via structure grows with its physical scale.*

Table 1—*Dielectric Properties of Selected Low-Loss PCB Materials*

Brand	Type	at 10GHz	tan θ	T_g (°C)
Nelco®	N4000-13-EP-SI®	3.2	0.008	210
Arlon	25N	3.25	0.0024	225
Gore®	Speedboard® C	2.6	0.003	220
Rogers	RO4000® series	3.5	0.004	280
Taconic	RF 35	3.5	0.0025	315

Crosstalk

Digital systems at the PCB level suffer crosstalk primarily due to near-field mutual inductance between circuits. This happens because digital PCB traces have impedances on the order of 25-85 ohms, well below the impedance of free space, a value defined by basic physics at 377 ohms. The low impedance suggests that a greater proportion of magnetic-field, as opposed to electric-field, energy surrounds the circuit. The best mechanism for controlling mutual-inductance interactions is to maintain a continuous, solid return-current path in close proximity to every signal conductor. Given a choice, the returning-signal current associated with signal **X** will be naturally attracted to the least inductive return current path. That path will be the path which lies in closest proximity to **X**, because in that configuration the magnetic fields from the outbound and return currents very nearly cancel each other, thus minimizing the overall circuit inductance. A solid reference plane provides an infinite continuum of possible paths for returning signal current, thus minimizing crosstalk simultaneously for all circuits in the entire PCB. That is the main purpose of having a solid reference plane.

Receiver

Digital receivers for high-speed applications are fully integrated with no external discretes. They may occur in numbers ranging from 1 to several thousand on a single integrated circuit. In some cases the circuitry may be quite sophisticated.

Return loss

The concept of standing wave ratio (SWR) commonly used to characterize input impedance in a carrier-based system is replaced by the more general concept of return loss, which, when expressed in the frequency domain, is the same as the S-parameter reflection term, s_{11}, expressed as an attenuation in units of dB. A return loss of 20 dB, for example, implies a reflection coefficient of plus or minus one part in ten. The return loss function is often measured in the time domain using a time-domain reflectometer (TDR). The TDR instrument injects a fast-edged step from a 50-ohm coaxial cable into the equipment under test and records the reflected signal. It then computes the s_{11} parameter using a Fourier Transformation. [26]

Reflection budget

A good reflection budget adds together all worst-case possible reflections from every pair of facing interfaces and ensures that the sum of the absolute values of all possible combinations amounts to a value substantially less than the received signal swing. For example, one reflected pattern might comprise a return loss of 15 dB at the receiver, another 15 dB at the transmitter, plus a round-trip cable loss of 5 dB. That combination suggests a worst-case round-trip residual reflection due to this one particular combination of interfaces no greater than -35 dB, or 1.8 percent of the transmitted signal level. When summing reflections in

a digital system, never assume signal cancellation unless the reflections emanate from interfaces positioned so close that the reflections occur within the interval of a single baud, in which case you may assume some degree of correlation between the main transmitted signal and its reflected artifacts. Otherwise, add the absolute values of all reflected products. [27]

Quantization of voltage

Input levels in a serial link are resolved to one of a few (typically two) possible discrete output levels within each baud interval. The quantizing process leaves room for some types of noise and interference that, if sufficiently small, have zero effect on the receiver's ability to perfectly reconstruct the input data sequence. The design of a serial link typically contemplates some form of noise margin budget stipulating the maximum degree of interference that may be present at input terminals of the receiver.

Quantization of time

Digital signal quality matters only at those specific points in time when the signal is sampled; noise at other times generally has no effect on system operation. Synchronous systems routinely tolerate substantial surges of noise that are carefully timed to occur between sampling instants.

Equalization

Binary serial links can typically tolerate 3 dB of signal loss at high frequencies without compensation, and 6-9 dB of high-frequency loss with the addition of a simple one-pole, 6-dB-per-octave equalizer. More elaborate compensation for known or measurable linear channel distortion is theoretically possible, but finds limited application at the present time in single-chip transceivers at high speeds. Good examples of equalization circuits that have become popular are the pre-emphasis circuit employed in PCIe 2.1 and other serial links at 2.5 Gb/s and 5.0 Gb/s, and the decision-feedback equalization (DFE) used on Gigabit Ethernet links at 250 Mb/s per channel. A good DFE requires analog design skills far beyond the capabilities of a typical digital designer. [28]

Multiplicity of channels

A digital system, in contrast to an analog RF system, can often tolerate 10-15% aggregate crosstalk with no degradation in its operational characteristics. That characteristic makes possible the implementation of multiple I/O channels on a highly-dense fully-integrated silicon die. Large IC packages commonly incorporate 32 or more channels of multi-gigabit transceivers on one chip. Such a feat is not trivial. This author is aware of many design teams that failed to assess the package and chip-level interference that results from such an architecture and were forced to cancel their projects late in the development phase as a result. Grouping the transmitters together, and the receivers together, in the IC pinout is a common strategy that helps prevent the big, loud local transmissions from interfering with reception of weak, distant signals.

Designers

If there is a *typical* digital designer, she brings a solid mix of computer science and practical experience to her job, but very little formal understanding of electromagnetic field theory. Management presses her to complete each job in the minimum time, so she resorts to rules of thumb and copy-cat engineering to make do. There is little time in a digital design lab for study or experimentation. Digital designers this author has known (and he has met literally thousands) share two characteristics, one good and one bad.

Mindset

In digital applications the *length of a rising or falling signal edge*, defined as the signal rise or fall time divided by its propagation velocity, replaces the carrier-based concept of signal wavelength. The successful digital signal integrity engineer thinks little about frequency, but much about the properties of his transmission media as averaged under the length of each rising edge as that edge propagates through the medium. If he can maintain a uniform ratio of series inductance to shunt capacitance everywhere, signals will propagate flawlessly through the medium.

Widely held misconception

Many digital engineers believe that a TEM-mode signal is conveyed by solid conductors when, in reality, the signal is conveyed as the result of a field that exists in the space between the conductors. A grasp of this basic point opens the way to an better understanding of the meaning of inductance, capacitance, and the importance of continuous proximity of every signal to its associated signal return conductor.

Open Problems

Despite continuing claims made over the last thirty years that optical interconnections will soon replace

metallic conductors, this author believes that further reduction in the physical scale of connectors and IC packages will satisfy our need for increased speed for many years to come.

Rich areas for further development include:

- Good-quality, low-cost dielectric materials that don't melt at high temperatures,
- Thinner dielectric layers and finer-pitched PCB etching technology,
- Thinner IC packages with more flexible, durable attachments,
- Simple, cheap ways to implement blind vias and coaxial vias,
- Connectors with substantially less crosstalk and lower emissions, and
- Design tools that tell you not just what's broken, but what to do about it.

Conclusion

Your author appreciates this chance to address the microwave design community and share with you, on behalf of the many thousands of digital engineers with which I have personally worked, some sense of the similarities, and key differences, between your field of study and mine.

In the event that you are called upon to assist in the development of high-speed digital hardware, I hope that this brief summary helps orient your thinking about the nature of typical digital design, and stimulates you to create something well above and beyond the typical.

References

[1] H. W. Johnson and M. Graham, High-Speed Digital Design: A Handbook of Black Magic, Englewood Cliffs, NJ: PTR Prentice Hall, 1993, Chapter 1, p. 1.

[2] H. W. Johnson, "Why digital engineers don't believe in EMC", *EMC Soc. newsletter*, Mar. 2, 1998, www.sigcon.com/Pubs/news/noEMC.htm

[3] IEEE Standard for Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications, IEEE 802.3, 1998, Clause 40 (a.k.a., Gigabit Ethernet 1000BASE-T).

[4] Recommended Standard 232 Version C, Interface Between Data Terminal Equipment and Data Communications Equipment Using Serial Interface, EIA RS-232-C, 1969.

[5] *PCI-Express Base Specification Version 2.1*, PCI-SIG PCIe 2.1, 2007. http://www.pcisig.com/specifications/pciexpress/

[6] A. X. Widmer, P. A. Franaszek, "A DC-Balanced, Partitioned-Block, 8B/10B Transmission Code," *IBM J. Res. Devl.*, vol. 27, no. 5, pp. 440-451, 1983; also see U.S. pat. 4,486,739.

[7] H. Johnson, "DC Blocking Capacitor Value", High-Speed Digital Design online newsletter, vol. 7, no. 9, 2004 www.sigcon.com/Pubs/news/7_09.htm .

[8] H. Johnson, "SONET Data Coding", High-Speed Digital Design online newsletter, vol. 5, no. 5, 2002 www.sigcon.com/Pubs/news/5_5.htm .

[9] Working Draft, Fibre Channel - Methodologies for Jitter Specification, ANSI T11.2, Project 1230, rev. 6, Sep. 23, 1998.

[10] J. F. Wakerly, *Digital Design Principles and Practices*, Englewood Cliffs, NJ: Prentice-Hall, 1990, Chapter 2, Digital Circuits, pp. 123-133.

[11] W. R. Blood, Jr., *MECL System Design Handbook*, 4th ed., Phoenix, AZ: Motorola Semiconductor Products, Inc., 1988.

[12] *Pentium® Pro Processor GTL+ Guidelines*, Intel Corp., Santa Clara, CA, AP-524, ord. no. 242765-001, 1996

[13] Electrical Characteristics of Low Voltage Differential Signaling (LVDS) Interface Circuits, ANSI/TIA/EIA-644, 1996.

[14] LVDS Data Outputs for High-Speed Analog-to-Digital Converters, Analog Devices AN-586, 2002.

[15] H. Johnson, "Extend Your Reach," *Xilinx Xcell Journal*, #53, 2nd quarter 2005 (or Spring 2005) www.xilinx.com/publications/archives/xcell/Xcell53.pdf

[16] H. B. Bakoglu, *Circuits, Interconnections and Packaging for VLSI*, Reading, MA: Addison-Wesley, 1990.

[17] D G. Kam and J. Kim, "40-Gb/s package Design Using Wire-Bonded Plastic Ball Grid Array," *IEEE Trans. Advanced Packaging*, vol. 31, No. 2, May. 2008, pp. 258-266.

[18] K. C. Gupta, R. Garg, I. Bahl, and P. Bhartia, *Microstrip Lines and Slotlines*, Boston: Artech House, 2nd ed., 1996.

[19] H. Johnson, "Why 50 Ohms?", *EDN magazine*, Sept. 14, 2000. www.sigcon.com/Pubs/edn/why50.htm

[20] H. Johnson, "50-Ohm Mailbag", *EDN magazine*, Jan. 4, 2001. www.sigcon.com/Pubs/edn/why50mail.htm

[21] H. Johnson, "Constant-resistance Termination", EDN magazine, Jun. 12, 2003. www.sigcon.com/Pubs/edn/ConstantRTermination.htm

[22] H. Johnson, "Potholes", EDN magazine, Nov. 11, 1999. www.sigcon.com/Pubs/edn/potholes.htm

[23] C. Coombs, Jr. (ed.), *Printed Circuits Handbook*, 6th ed., NY: McGraw Hill, 2007.

[24] R. Senthinathan and J. L. Prince, *Simultaneous Switching Noise of CMOS Devices and Systems*, Boston: Kluwer Academic Publishers, 1994.

[25] B. Archambeault and J. Drewniak, *PCB Design for Real-World EMI Control*, Norwell, MA: Kluwer Academic Publishers, 2002.

[26] D. Derickson and M. Muller (ed.), *Digital Communications Test and Measurement*, Upper Saddle River, NJ: Prentice Hall, 2008.

[27] H. Johnson and M. Graham, *High-Speed Signal Propagation: Advanced Black Magic*, Upper Saddle River, NJ: Prentice Hall, 2003, Chapter 8, "100-Ohm Balanced Twisted-Pair Cable," pp. 471-481.

[28] J. A. C. Bingham, *The Theory and Practice of Modem Design*, NY: John Wiley & Sons, 1988, Chapter 9, "Non-Linear Equalizers."

[29] H. Johnson, "Shannon Says", EDN magazine, Nov. 13, 2003. www.sigcon.com/Pubs/edn/ShannonSays.htm

[30] Sheikholeslami, A., "Multi-level Signaling for Chip-to-Chip and Backplane Communication (A Tutorial)", *Proc. IEEE 39th International Symposium on Multiple-Valued Logic*, pp. 203-207, May 2009.

[31] Frederick H. Raab, et. al., "Power Amplifiers and Transmitters for RF and Microwave", IEEE *Trans. Microwave Theory and Techniques*, March 2002.