CMOS circuits for Gb/s serial data communication

by J. F. Ewen
M. Soyuer
A. X. Widmer
K. R. Wrenner
B. D. Parker
H. A. Ainspan

The functional characteristics and design challenges associated with a variety of communication-related circuits are presented. These include the mixed-signal design and noise issues associated with high-speed clock generation and recovery for serial data communication. Hardware results are presented on the noise properties of common integrated voltage-controlled oscillator (VCO) circuits.

Introduction

A key characteristic of ICs in many communication-related applications is the combination of analog circuits with digital logic, while maintaining maximum performance at minimum power and cost. This combination presents a number of challenges beyond basic circuit design issues, ranging from technology choice and simulation techniques to noise and crosstalk. This paper focuses on recent CMOS design work addressing these issues, with specific attention to the area of high-speed serial data communication.

Serial baseband data links, whether using fiber-optic or coaxial cables, incorporate coders and decoders, high-speed multiplexors and demultiplexors, and low-speed clock synchronization, along with phase-locked loops for high-speed clock synthesis and recovery (**Figure 1**). In addition, specialized analog circuits are typically required to interface with the particular medium (e.g., a laser driver and optical receiver circuit if a fiber-optic cable is used). The requirements and circuits for these specialized functions can vary widely depending on the transmission medium; however, the building blocks shown in Figure 1 are common to almost all serial data links and are the focus of this paper.

The basic function, while conceptually simple, presents a number of significant design challenges due to the wide range of clock speeds and the mixture of analog and digital circuits required. For example, the Fibre Channel standard [1] specifies a maximum data rate of 1063 Mbaud, or a bit interval of 940 ps for the serial data. Clock generation and recovery and data retiming must operate at this high data rate, with a timing resolution of a fraction of the bit interval. Typically, these functions are implemented using custom analog circuits in order to achieve the stringent timing requirements and high speed; however, relatively few transistors are required (~500). At the parallel end of the data chain, frame processing (address resolution, sequence generation, etc.) is done at a multiple-byte level, at clock speeds of 50 MHz or even slower. The timing requirements are modest by current standards, but moderate transistor counts ($\sim 5 \times 10^5$) are required to implement this function.

0018-8646/95/$3.00 © 1995 IBM

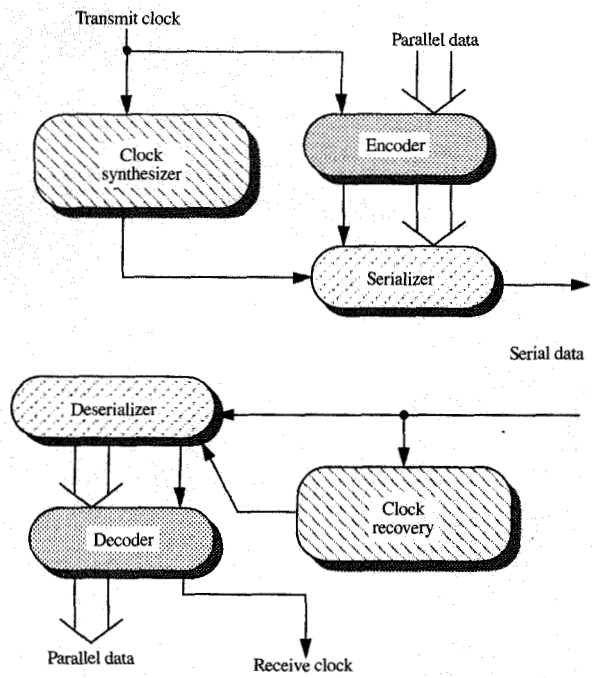

Figure 1

Serial data link block diagram.

A number of factors drive the integration of these functions to optimize cost and performance at the system level. Integration contains high-speed on-chip interfaces, eliminating the need for sophisticated high-speed packaging techniques. Integration also reduces the overall chip count, power dissipation, and board area required by the serial communication subsystem. These factors lead to lower system cost, provided that the analog circuits can be combined with the digital logic without requiring a unique technology to achieve the required performance. An integrated serial data link requires very high-speed, low-density analog functions, together with low-speed, high-density digital logic, creating a difficult set of design and technology trade-offs.

The following sections discuss the custom design approaches used to implement these functions, beginning with the analog circuits used in clock recovery and synthesis, and followed by simple, but high-speed, multiplexors and demultiplexors.

Serial data communication

The basic design issues associated with serial communication fall into two broad categories, performance and noise. The key performance goals are to customize the design, and choose logic families, to optimize the circuit speed and power while using relatively mature technologies to implement the frame processing logic (which consumes most of the chip area). The primary noise issue dealt with here is that of coupled noise between different functional blocks on the same chip. It arises from the switching activity in one section of the chip coupling to another section via the power supplies (delta-I noise) or via the common substrate. The interaction can occur between the digital and analog circuits, as well as among different analog circuits, and has the potential to seriously degrade performance. The approach taken here is to minimize the amount of analog circuitry, and to implement the remaining analog function by using circuit techniques that minimize the sensitivity to coupled noise. Random noise (e.g., thermal noise) can also be a serious concern for the analog circuits. In the context of serial communication, voltage-controlled oscillator phase noise (or jitter) is a key design parameter. Low phase-noise oscillator design in the absence of coupled noise has been considered by a number of authors [2–4] and is not discussed further here.

• *Clock generation and recovery*

The clock generation and recovery phase-locked loops (PLLs) typically place the most difficult requirements on the design and technology. These circuits typically operate at the serial data rate, with timing requirements of a fraction of a bit interval, and are very sensitive to noise within the PLL passband. This is particularly problematic for these applications, since the noise generated by the digital logic will be harmonically related to the input data to which the PLL is locked. If the clock generation (transmitter) and clock recovery (receiver) PLLs are integrated on the same chip, crosstalk between PLLs becomes another major concern. This can cause even more performance degradation than the noise from the digital logic. The key noise-sensitive circuit in the PLL is the voltage-controlled oscillator (VCO). Circuit techniques that minimize sensitivity to power-supply and substrate noise, while maintaining high performance (e.g., timing jitter, power dissipation, maximum oscillation frequency), are critical.

Integrated VCO design techniques

Designing low-phase-noise, high-frequency VCOs is a well-known art [2]; however, integrating VCOs in a digital CMOS technology presents a number of significant design constraints. Small, high-Q passive components are typically not available. Adding these components external to the chip increases the module cost, or compromises performance because of packaging parasitics. Integrated VCOs are typically based on astable multivibrator circuits (relaxation VCO) or variants of ring oscillators. These circuits can easily be integrated into a digital process, but

have very low Q, leading to larger phase noise and timing jitter. For a baseband serial communication link, the relatively large amount of phase noise is not a serious limitation. However, a low-Q VCO is also more sensitive to coupled noise, a characteristic which must be carefully examined in this mixed-signal environment.

Figure 2 illustrates three common circuit topologies for integrated VCOs. The relaxation VCO [Figure 2(a)] operates by charging and discharging a timing capacitor with a constant current. Varying the magnitude of this current varies the VCO frequency. This type of oscillator uses few transistors and small timing capacitors (at high frequencies the wiring capacitance alone may suffice), has a wide tuning range, and dissipates little power. The voltage across the timing capacitor crossing a switching threshold determines the timing jitter in this oscillator, so care must be taken in the physical design of this circuit to minimize its sensitivity to substrate noise.

A second common type of VCO is the current-controlled, or current-starved, ring oscillator [Figure 2(b)]. The basic concept is to control the current in each stage of a ring oscillator in order to vary its delay, and hence the frequency of oscillation. There are many variations on this theme, with Figure 2(b) illustrating a differential source-coupled logic (SCL) buffer whose tail current is controlled to vary the delay. Differential circuits are used to minimize pulse distortion and sensitivity to power-supply noise, which is critical for high-speed operation. Although the buffer data path is differential, the VCO control path, via the tail-current source, remains a single-ended signal susceptible to power-supply and substrate noise. Very high oscillation frequencies are possible using this design, with slightly more power dissipation than for a relaxation oscillator.

The third type of VCO shown is a variable-stage, or delay-interpolating (DI), VCO [5] [Figure 2(c)]. This VCO consists of two delay lines of different lengths. The outputs of the delay lines are combined in a multiplexor, or mixer. With the multiplexor control at one extreme, the frequency is set by the delay through the first delay line. At the other extreme, the frequency is set by the delay through the second line. When the multiplexor control is between its logical 0 and 1 states, the multiplexor acts as an analog mixer, and it is possible to continuously tune the effective delay between these two extremes. The specific design shown in Figure 2(c) is slightly more complicated in order to provide both 0° and 90° output phases, and the multiplexors select between a six-stage and a ten-stage delay line. Again, the delay elements are simple differential SCL buffers, and the multiplexors are standard two-level Gilbert-cell designs. Fully differential circuits are used to reject power-supply noise. Because of the differential multiplexor, the VCO control voltage is also differential, providing better power-supply noise rejection than the

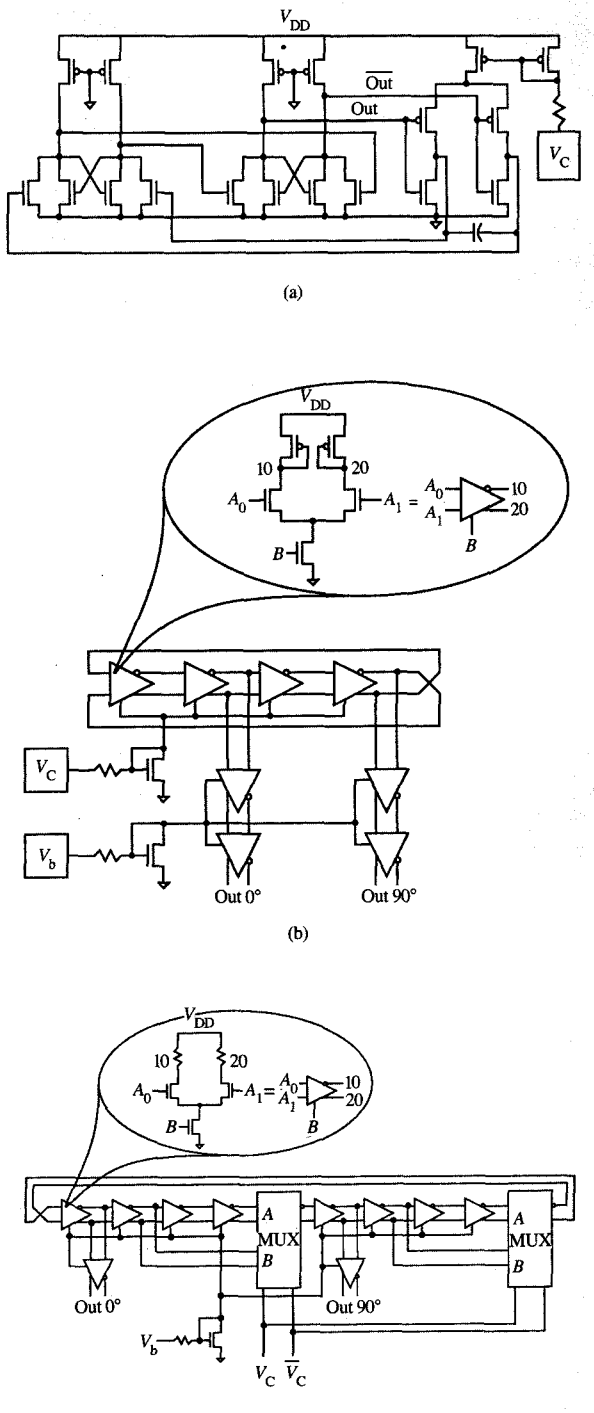

Figure 2

Typical voltage-controlled oscillator circuits: (a) relaxation VCO, (b) current-controlled ring oscillator, (c) delay-interpolating ring oscillator.

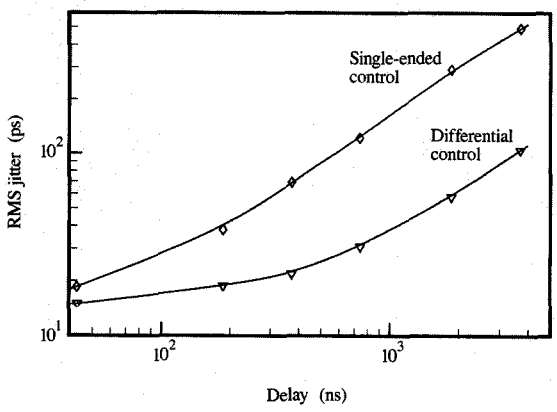

Figure 3

Typical RMS jitter characteristic.

Table 1 VCO performance.

VCO topology	*Center frequency* (MHz)	*Tuning range* (%)	*Power* (mW)	*Area* (μm^2)
Relaxation	600	±67	20	150 × 100
Ring	1000	±30	30	220 × 100
Delay-interpolating (DI)	500	±20	75	240 × 150

previous designs. The added circuit complexity and fan-out yield slightly lower operating frequencies and higher power dissipation than the other two topologies. The maximum tuning is also smaller, being limited to a 2:1 range to avoid the possibility of harmonic oscillation.

To put the preceding discussion into perspective, consider an 0.8-μm (0.45-μm effective) CMOS technology [6]. Each of these three VCO topologies has been designed and fabricated using this technology, with the resulting hardware characteristics shown in **Table 1**. The target data rate is 1 Gb/s in each case. The relaxation VCO is centered at a lower frequency in order to operate near the high end of its tuning range, resulting in lower jitter. In general, the jitter in the VCO output due to coupled noise (via the power supplies, substrate, or control input) is proportional to the VCO gain (MHz/V). Relaxation VCOs have asymmetric tuning characteristics, with VCO gain decreasing with increasing frequency. For the best noise immunity, this VCO should be operated at the high end of its tuning range—consistent with the margins required by process, supply, and temperature variations. The delay-interpolating VCO is centered at half the target frequency for use in a half-speed clock-recovery circuit. These results illustrate the previous discussion. The relaxation VCO has the largest tuning range, lowest power dissipation, and smallest size. The current-controlled ring oscillator has a smaller tuning range and dissipates more power. The delay-interpolating VCO has the smallest tuning range, dissipates the most power, and consumes the most area. Part of the extra power dissipation is due to the additional delay stages required to lower the center frequency chosen for this particular design, but even accounting for this difference it still dissipates 50% more power than the equivalent current-controlled ring oscillator.

The phase noise, or timing jitter, of the oscillator is also an important parameter to consider. The timing jitter of these oscillators can be characterized by using a digital sampling oscilloscope to measure the jitter histogram as a function of the delay between the trigger and the measurement point. Increasing this delay allows the VCO jitter to "accumulate" over multiple cycles, providing a rough estimate of the spectral characteristics of the phase noise. Also, for typical integrated PLLs, the loop Q ranges from 20 to about 200. Because the jitter of the PLL output is the intrinsic jitter of the VCO accumulated over $\sim Q$ cycles, this is an especially important parameter to evaluate for clock-recovery or clock-synthesis applications.

Figure 3 shows a typical jitter measurement for the delay-interpolating VCO operating at a center frequency of 270 MHz. These results show the characteristic increase in timing jitter as more VCO cycles are included in the measurement (i.e., as the measurement delay is increased). They also clearly show the advantage of using a fully differential input for the VCO control. Similar measurements can be obtained for the relaxation and ring VCO circuits.

The jitter can be determined by fitting a straight line to the data (on a linear scale) and extracting a "cycle-to-cycle" jitter number which is characteristic of the intrinsic performance of the VCO. This cycle-to-cycle jitter (expressed in ppm) can be used as a figure of merit for different VCO topologies. However, this figure of merit must be evaluated at equivalent VCO gains (not equivalent frequencies), in order for a valid comparison to be made. Results for the three VCO topologies of Figure 2 are shown in **Table 2**. As expected, the jitter increases with increasing VCO gain, because the oscillator is more sensitive to noise at the control input. The delay-interpolating VCO has one third of the jitter of the current-controlled ring oscillator, and only 20 to 25 percent of the jitter of the relaxation VCO, when normalized to the same VCO gain.

The circuit sensitivity to substrate noise is illustrated in **Table 3**, which summarizes a series of circuit simulations comparing the relaxation VCO with the delay-interpolating

VCO. In these simulations, a fixed-amplitude noise signal is added to the substrate node. The frequency of this noise signal is chosen to be close to the free-running VCO frequency in each case. In general, as the noise frequency approaches the free-running VCO frequency, the VCO jitter increases. Again, the relaxation VCO exhibits a much larger sensitivity to substrate noise. In the worst case, the relaxation VCO is injection-locked by the noise signal (i.e., it oscillates directly at the noise frequency). For the serial link shown in Figure 1, injection locking implies that the clocking circuits are synchronized to the coupled noise rather than the data, and leads not to a simple performance degradation, but to total failure of the link. The delay-interpolating VCO shows much less sensitivity to injected substrate noise.

These results show the importance of fully differential circuits, not only in the individual buffer stages (as for all the circuit topologies of Figure 2), but also in the control path. It should be noted that the delay-interpolating VCO still uses a single-ended input to bias the tail-current source in the buffer stages. Careful attention must be paid to generating this bias so as to reject supply and substrate noise. Because this bias point is not used to tune the DI VCO [as opposed to the current-controlled VCO of Figure 2(b)], the tuning characteristics and the noise rejection can be optimized independently. On the basis of power dissipation and size, the relaxation VCO is the design of choice for many applications; however, in highly integrated mixed-signal designs where noise performance and noise sensitivity are a concern, the current-controlled ring oscillator finds wide use [7, 8]. Those applications requiring the highest noise immunity are best served by the delay-interpolating VCO, or by an equivalent topology which provides for differential control [9].

Frequency synthesizer PLL

To illustrate the application of the VCO in more complicated functions, consider the clock generation circuit of Figure 1. The clock recovery function is similar in many respects. Clock generation is implemented using a frequency synthesizer PLL (**Figure 4**) containing a symmetric phase-frequency detector (PFD), a differential charge-pump filter (CPF), a divide by 16 (for this particular example), and a delay-interpolating VCO [10]. The fully symmetric PFD (**Figure 5**) and CPF (**Figure 6**) are designed to minimize the dead-zone problems of conventional CMOS synthesizers [11], and to improve the overall jitter performance of the generated clock. The PFD combines standard CMOS logic with custom pseudo-n-MOS logic. The pseudo-n-MOS gate of Figure 5 minimizes the delay in the critical path used to reset the SR latches at the PFD input, and ensures that the delay for each of the four inputs is closely matched. Because of the finite gate delays, the PFD outputs contain narrow pulses even when the PLL reaches steady state. Ideally, one would like to have these pulses very narrow but still with enough amplitude to provide sufficient phase-error correction through the CPF. This becomes critical especially around the zero-phase-error region, where UP and DN pulses appear almost simultaneously. One can control the energy in these pulses by making use of the pseudo-n-MOS gates as appropriate. All of these characteristics minimize the dead zone in the detector.

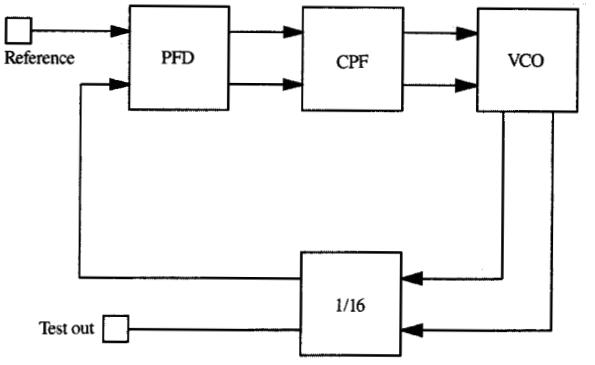

Figure 4

Frequency synthesizer PLL block diagram.

Table 2 VCO jitter performance.

VCO topology	*Center frequency* (MHz)	*VCO gain* (%/V)	*Cycle-to-cycle jitter* (ppm)
Relaxation	500	108	350
	730	51	180
Ring	800	77	156
	1000	60	100
DI	276	75	55

Table 3 RMS VCO jitter vs. substrate noise frequency.

Frequency offset (%)	*Relaxation VCO*	*Delay-interpolating VCO*
0.5	—*	1.65
1.0	—*	0.8
2.0	3.3	0.4
5.0	0.8	0.16

*VCO is injection-locked to noise signal.

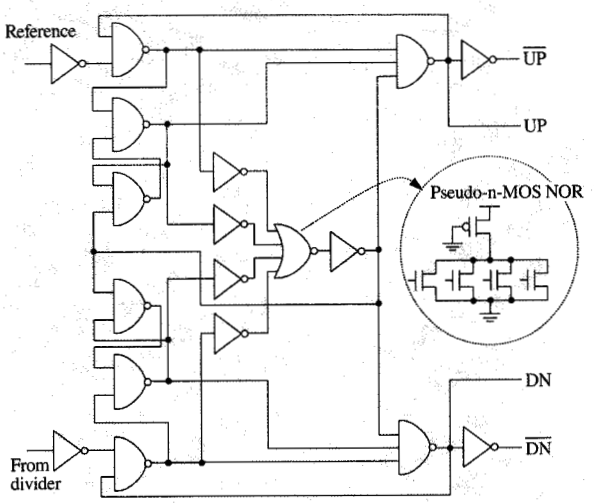

Figure 5

Phase-frequency detector circuit.

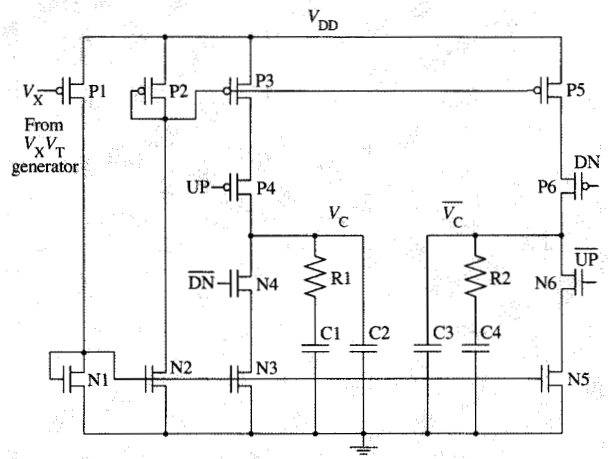

Figure 6

Simplified charge-pump filter schematic.

The CPF implements an analog loop filter in a fully differential and balanced configuration in order to minimize the sensitivity to power supply and substrate noise. Common-mode feedback circuits and/or clamp diodes can be used to limit the dynamic range of the CPF output such that the VCO common-mode input stays within a desired range. Care must be taken to minimize the effect of such additional circuitry on the static and dynamic phase-error performance of the PFD and CPF combination, since they contribute offset currents across the loop filter.

This circuit has also been implemented in 0.8-μm CMOS technology, with the design optimized for speed. The test chip consumes 270 mW at 1.25 GHz, with an active circuit area of 1000 × 750 μm^2. The RMS jitter in the generated clock is 1.4 ps (or 0.18%), and the peak-to-peak jitter is approximately 10 ps. This demonstrates the potential to push CMOS technology to very high frequencies while minimizing jitter.

- *Serializer and deserializer*

The previous section dealt with the key analog circuits in the serial link of Figure 1. This section gives a brief overview of the digital functions, concentrating specifically on the deserializer, and the basic approach used to achieve Gb/s data rates. A more detailed discussion can be found in [12]. The architecture used in the deserializer is shown in **Figure 7**. This example is specific to a Fibre Channel link using the 8B/10B code [13], but is easily extended to other environments. A complementary approach is used for the serializer in the transmitter section.

The deserialization of the serial data stream to the byte interface is done in stages so that each stage can be optimized for speed and power while using a common technology. The serial data are demultiplexed into two data streams at half the speed of the original. The half-speed data streams are clocked into two high-speed shift registers, which are latched into a five-bit register where the data are examined for special synchronization characters [13]. The data are latched into the final demultiplexor stage at the appropriate time to provide byte-aligned parallel data at the output. The ten phases of clocks generated from the bit clock provide the required timing resolution for the 8B/10B code used here. The initial 1:2 demultiplexing is implemented using custom SCL logic, which is optimized for speed and minimum pulse distortion. This circuit contains only two critical latches (so power dissipation is not a design constraint), and can actually be included as part of the clock recovery circuit. The clock generation and control for subsequent stages is based on a ten-stage ring counter. This ring counter requires relatively high-speed logic, since the critical path delay must be less than a bit interval. A combination of semicustom CMOS logic can be used to minimize the power dissipation, with custom n-MOS logic used as necessary to reduce delay in the critical path. The timing requirements are sufficiently relaxed that the final 5:10 demultiplexor (and subsequent demultiplexing if required), can be implemented in standard-cell CMOS logic. It is important to note that only the minimum function required

is implemented at high speed. In this design, only a "dumb" demultiplexor and shift registers are implemented at the bit rate and half-bit rate, respectively. The more complicated byte synchronization function is implemented at a much lower speed using parallel techniques. By staging the deserializer function and customizing the circuit techniques used to implement each stage, the wide range of performance requirements can be satisfied with relatively modest requirements on the technology. For example, the 1063Mb/s Fibre Channel deserializer can be implemented using an 0.8-μm CMOS technology while dissipating less than 80 mW.

Physical design and noise considerations

The basic building blocks needed to implement the link adapter function of Figure 1 have been discussed; however, integrating all of this function onto a single IC requires close attention to the physical design of the chip in order to control noise and crosstalk between the different sections. The transmitter and receiver sections operate asynchronously with respect to each other, creating the worst possible noise environment for the PLLs. The jitter and bit-error rate performance can be degraded significantly by crosstalk between sections, and data-pattern-dependent noise can lead to nonfunctional systems.

The first defense against noise is to design circuits that are as insensitive to noise as possible, whether coupled through the substrate [14] or via the power supplies. As discussed in detail previously, differential circuits are necessary (but not always sufficient) for those functions that are sensitive to noise.

The corollary is to use circuits that generate little noise for the bulk of the logic on chip. Unfortunately, standard CMOS logic generates significant switching noise during transitions. Some circuit solutions are available in specific cases. For example, the deserializer described in the previous section has sufficient timing margin at the parallel interface that the individual output lines are intentionally skewed across two bit times in order to reduce the simultaneous switching noise. Alternative logic families that generate little switching noise would lead to increased power dissipation, so the only practical means of minimizing the generated noise is by careful design for the on-chip power distribution and related packaging. Sufficient I/Os must be assigned to power and ground to maintain low supply inductance. Low-inductance packaging is critical, and on-chip decoupling capacitors may also be needed to decrease switching noise due to the digital logic.

Next, it is important to isolate the different sections from each other by using different power distribution nets. The analog power should be separate from the digital power, and transmitter power should be separate from

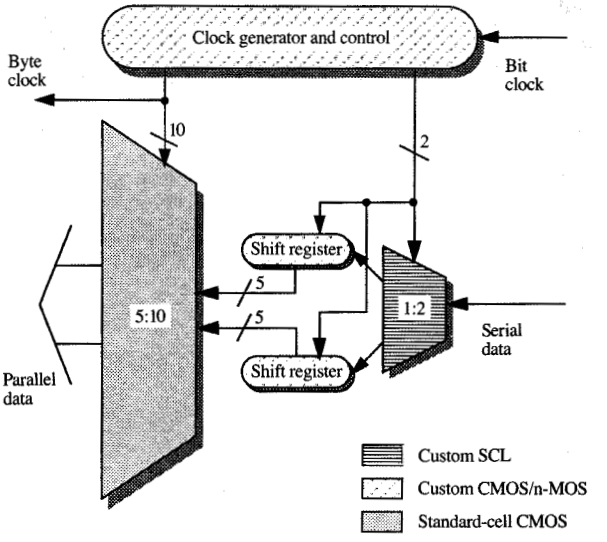

Figure 7

Deserializer architecture.

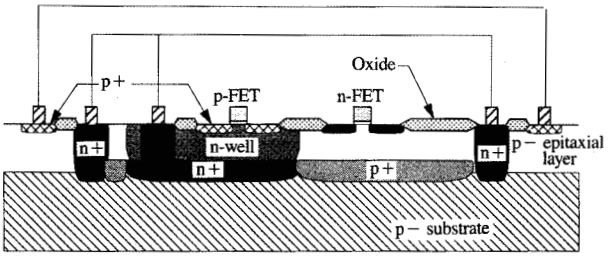

Figure 8

Typical CMOS device cross section.

the receiver power distribution nets. Only the required interface signals should be shared between sections. All necessary bias voltages and currents should be generated within their own sections, and not distributed throughout the chip. This is especially important for the analog circuits.

Finally, the different functional blocks can be isolated from each other by using careful physical design, taking advantage of a specific technology. A generic CMOS cross section is shown in **Figure 8**. The n-wells and p+ guard

rings surrounding the active circuits can be used to isolate them from one another. The p+ guard rings should be connected to non-current-carrying grounds in order to prevent noise injection from the active circuits. Guard rings for the different sections of the design should also be connected to their own ground pins, again to prevent crosstalk. If a p+ substrate is used (~0.05 Ω-cm), physical separation of the different sections has little effect on the coupled noise [15]. Once noise is coupled to the substrate, it "propagates" throughout the chip. If a high-resistivity p+ substrate is used (~10–20 Ω-cm), separating the sections can lead to lower coupled noise. For many serial data link applications, the chip area is governed by the number of I/O pads required for signals and supplies. Here, the sections can be isolated with high-resistivity regions in between, providing better immunity to noise and crosstalk without increasing chip cost. Silicon-on-insulator (SOI)-based technologies can be used to provide nearly ideal isolation between circuits at low frequencies. However, as the frequency increases, the isolation decreases dramatically. Guard-ring structures are still required, and as frequencies approach 1 GHz, SOI-based technologies offer no isolation advantage over standard CMOS [16].

Conclusion

A number of circuit building blocks for serial data communication links have been discussed. Careful partitioning of the function among analog, custom digital, and standard cell designs is required in order to achieve the required data rates without placing excessive demands on the technology. Coupled noise presents a potentially serious barrier to performance. Circuits designed specifically to reject power-supply and substrate noise are important. This has been discussed in detail with respect to voltage-controlled oscillators, where the importance of differential circuits in the control path as well as the signal path has been demonstrated. Careful choice of physical design and technology is required to isolate the digital and analog sections of the chip. A number of circuits have been described in order to illustrate these design techniques and demonstrate the levels of performance and integration that are possible using CMOS in this application environment.

References

1. *Fibre Channel Physical and Signaling Interface (FC-PH)*, Rev. 4.1, American National Standards Institute, Inc., Washington, DC, August 12, 1993.
2. W. P. Robbins, *Phase Noise in Signal Sources*, Peter Peregrinus Ltd., London, 1982.
3. A. A. Abidi and R. G. Meyer, "Noise in Relaxation Oscillators," *IEEE J. Solid-State Circuits* **18,** No. 6, 794–802 (December 1983).
4. M. H. Wakayama and A. A. Abidi, "A 30-MHz Low-Jitter High-Linearity CMOS Voltage-Controlled Oscillator," *IEEE J. Solid-State Circuits* **22,** No. 6, 1074–1081 (December 1987).
5. S. K. Enam and A. A. Abidi, "NMOS IC's for Clock and Data Regeneration in Gigabit-per-Second Optical-Fiber Receivers," *IEEE J. Solid-State Circuits* **27,** No. 12, 1763–1774 (December 1992).
6. B. J. Gross, H. Chuang, and T. Schmerbeck, "State-of-the-Art Analog BiCMOS," invited paper, IEEE GaAs IC Symposium, San Jose, 1993.
7. I. A. Young, J. K. Greason, and K. L. Wong, "A PLL Clock Generator with 5 to 110 MHz of Lock Range for Microprocessors," *IEEE J. Solid-State Circuits* **27,** No. 11, 1599–1607 (November 1992).
8. D. L. Chen and R. Waldron, "A Single-Chip 266 Mb/s CMOS Transmitter/Receiver for Serial Data Communication," *ISSCC 1993 Digest of Technical Papers*, pp. 100–101 (February 1993).
9. A. L. Mullgrav, Jr., "Integrated Differential Voltage-Controlled Ring Oscillator," U.S. Patent 5,191,301, March 2, 1993.
10. M. Soyuer, J. Ewen, and H. L. Chuang, "A Fully Monolithic 1.25GHz CMOS Frequency Synthesizer," *1994 IEEE Symposium on VLSI Circuits*, Honolulu, Paper 11.3, June 1994, pp. 127–128.
11. M. Soyuer and R. G. Meyer, "Frequency Limitations of a Conventional Phase-Frequency Detector," *IEEE J. Solid-State Circuits* **25,** No. 4, 1019–1022 (August 1990).
12. J. Ewen, A. Widmer, M. Soyuer, K. Wrenner, B. Parker, and H. Ainspan, "Single-Chip 1062Mbaud CMOS Transceiver for Serial Data Communication," to be presented at *ISSCC'95*, San Francisco, February 1995.
13. A. X. Widmer and P. A. Franaszek, "A DC-Balanced, Partitioned-Block, 8B/10B Transmission Code," *IBM J. Res. Develop.* **27,** No. 5, 440–451 (September 1983).
14. B. R. Stanisic, N. K. Berghese, R. A. Rutenbar, L. R. Carley, and D. J. Allstot, "Addressing Substrate Coupling in Mixed-Mode IC's: Simulation and Power Distribution Synthesis," *IEEE J. Solid-State Circuits* **29,** No. 3, 226–238 (March 1994).
15. D. K. Su, M. J. Loinaz, S. Masui, and B. A. Wooley, "Experimental Results and Modeling Techniques for Substrate Noise in Mixed-Signal Integrated Circuits," *IEEE J. Solid-State Circuits* **28,** No. 4, 420–430 (April 1993).
16. K. Joardar, "A Simple Approach to Modeling Cross-Talk in Integrated Circuits," *IEEE J. Solid-State Circuits* **29,** No. 10, 1212–1219 (October 1994).

Received July 25, 1994; accepted for publication October 28, 1994

John F. Ewen *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (EWEN at WATSON).* Dr. Ewen received the B.S., M.S., and Ph.D. degrees in electrical engineering from Purdue University, West Lafayette, Indiana, in 1976, 1977, and 1981, respectively. His graduate work dealt with surface acoustic wave devices and nonlinear acoustic wave propagation in lithium niobate and ZnO/Si material systems. In 1981, he joined the Fiber Optic Technologies group at the IBM Thomas J. Watson Research Center, where he has been involved in design and fabrication of high-speed optical links for computer applications, using both Si and GaAs technologies. In 1991, he became manager of the Mixed-Signal Communication IC Design group. His areas of work include high-performance fiber-optic and copper coaxial computer communication links, and high-performance wireless communication involving IR as well as RF technologies. Dr. Ewen is a member of the IEEE.

Mehmet Soyuer *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (SOYUER at YKTVMV).* Dr. Soyuer joined IBM at the IBM Thomas J. Watson Research Center as a research staff member in 1988. He received the B.S. and M.S. degrees in electrical engineering from the Middle East Technical University, Ankara, Turkey, in 1976 and 1978, respectively, and the Ph.D. degree in electrical engineering from the University of California, Berkeley, in 1988. He has received two IBM Research Division Awards and published numerous technical articles in the area of analog and nonlinear electronic circuit design. Dr. Soyuer's research interests include high-speed integrated circuits, technologies, and subsystems for communications.

Albert X. Widmer *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (WIDMER at YKTVMV).* Mr. Widmer is a member of the Communication Technology Department at the IBM Thomas J. Watson Research Center. He received the Dipl.Ing.ETH degree in electrical engineering in 1956 from the Swiss Federal Institute of Technology, Zurich. Mr. Widmer served as a teaching assistant at the Institute for Telecommunications ETH, and then joined the Advanced Systems Development Division of IBM, New York, in 1959. He has been associated with a variety of hardware projects as an analog and digital circuit designer, logic designer, and lead engineer. Most projects have been related to signal processing or data communications at the physical level. Several of his innovations have found wide use and have been recognized by IBM with corporate and divisional awards. Mr. Widmer is a senior member of the Institute of Electrical and Electronics Engineers.

Kevin R. Wrenner *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (WREN at YKTVMV).* Mr. Wrenner is an advisory engineer with the Digital Communication IC Design group. He received his B.S. and M. Eng. degrees in electrical engineering from Cornell University in 1986 and 1988, respectively. From 1988 to 1989 he worked for Digital Equipment Corporation, designing digital MOS integrated circuits. In 1989 he joined the Research Division in Yorktown Heights, New York, where he has contributed to the design of GaAs and CMOS ASICs for high-speed computer links, and custom CMOS macros for various applications including voltage conversion and signal processing. Mr. Wrenner is a member of Eta Kappa Nu and the Tau Beta Pi Association.

Benjamin D. Parker *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (BDP at WATSON).* Mr. Parker received his B.S. in physics from Bowdoin College, Brunswick, Maine, in 1979. He received his M.S. in physics from Brown University, Providence, Rhode Island, in 1981. His graduate work dealt with the optical properties of adsorbed layers on metal surfaces. In 1986 he joined the GaAs group at the IBM Thomas J. Watson Research Center, working on the characterization of III–V semiconductors. In 1991, he joined the Mixed-Signal Communications IC Design group, working on design and verification of digital circuits in high-speed serial communications.

Herschel A. Ainspan *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (AINSPAN at WATSON).* Mr. Ainspan received the B.S. and M.S. degrees in electrical engineering from Columbia University, New York, in 1989 and 1991, respectively. In 1989 he joined the IBM Thomas J. Watson Research Center, where he has been involved in designing and testing of integrated circuits for high-speed optical and infrared links for computer applications, using both Si and GaAs technologies.