

SiGe BiCMOS integrated circuits for high- speed serial communication links

D. J. Friedman
M. Meghelli
B. D. Parker
J. Yang
H. A. Ainspan
A. V. Rylyakov
Y. H. Kwark
M. B. Ritter
L. Shan
S. J. Zier
M. Sorna
M. Soyuer

Considerable progress has been made in integrating multi-Gb/s functions into silicon chips for data- and telecommunication applications. This paper reviews the key requirements for implementing such functions in monolithic form and describes their implementation in the IBM SiGe BiCMOS technology. Aspects focused on are the integration of 10–13-Gb/s serializer/deserializer chips with subpicosecond jitter performance, the realization of 40–56-Gb/s multiplexer/demultiplexer functions and clock-and-data-recovery/clock-multiplier units, and, finally, the implementation of some analog front-end building blocks such as limiting amplifiers and electro-absorption modulator drivers. Highlighted in this paper are the key challenges in mixed-signal and analog integrated circuit design at such ultrahigh data rates, and the solutions which leverage high-speed and microwave design and broadband SiGe technologies.

1. Introduction

The continuing demand for new communication services and higher user-end bandwidth has necessitated the development of Ethernet [1], Fibre Channel [2], and Synchronous Optical NETwork (SONET) [3] standards for transmission at data rates of ≥ 10 Gb/s, backed by the development of hardware capable of supporting such data rates. For example, the 10-Gb/s Ethernet standards activity which targets the broadband local area network (LAN) and wide area network (WAN) applications was completed before the end of 2002. For telecommunication applications, a well-established standard for 10-Gb/s serial links has also been defined by SONET OC-192 [3], which calls for a line rate of ~ 9.95 Gb/s. Telecommunication equipment for higher-data-rate standards [such as SONET OC-768 [3] at 40 Gb/s and its variants with forward-error correction (FEC) at 43–50 Gb/s] for high-capacity long-haul applications is currently being commercialized.

As these very-high-data-rate communication markets continue to mature, they generate a pressing need for

higher levels of integration to bring down the cost and power dissipation. However, this must be achieved while still complying with stringent serial link requirements such as minimizing jitter and bit-error rate (BER). In order to put some of these requirements into perspective, consider an example based on a realistic serial link jitter budget: A clock generator at the transmitter of a 10-Gb/s link must typically have less than 10% peak-to-peak jitter in order to enable a total link BER better than 10^{-12} . Assuming that the equivalent noise source is white, one can show that this peak-to-peak requirement translates to a sub-0.7-ps rms jitter generation specification, indicating a key challenge for circuit designers—how to design high-performance clocking circuits while keeping power dissipation low. In addition, one must still aim at high levels of integration, which may include all of the analog and digital functions related to serialization, deserialization, coding, framing, and built-in self-test (BIST). It is well known that whenever analog and digital circuits are used on the same chip, deleterious crosstalk

Table 1 IBM SiGe HBT generations.

Parameter	Units	0.5 μm	0.25 μm	0.18 μm
A_E min	μm^2	0.39	0.3	0.15
R_{BI}	$\text{k}\Omega/\square$	9	9	9
β (0.72 V)		100	100	250
V_A ($V_{BE} = 0.72$ V)	V	65	65	>120
BV_{CEO}	V	3.35	3.35	1.9
BV_{CBO}	V	10.5	10.5	6.5
r_B ($A_E = 1 \mu\text{m}^2$)	Ω	124	114	60
$NF_{\min}$	dB	0.8	0.8	0.4
f_T	GHz	47	47	120
$f_{\max}$	GHz	65	65	100

and substrate coupling may occur, especially at these speeds. These considerations clearly show that circuit designers targeting such ultra-broadband applications will need all the help they can get from process technology developers. Among the more mature silicon-based technologies with high yields, the SiGe BiCMOS integrated circuit (IC) technology fulfills this need by providing high integration capability as well as high performance levels [4]. The remainder of this paper focuses on the IC design work done at the IBM Thomas J. Watson Research Center for 10–13-Gb/s and 40–50-Gb/s serial links using two generations of the IBM SiGe BiCMOS integrated circuit technology.

Compared to III–V integrated circuit technologies such as that based on the use of InP [5], the SiGe BiCMOS IC technology is viewed as a good candidate for implementing the relatively complex digital functions involved in clock and data recovery, serialization, and deserialization at data rates to 40 Gb/s and beyond. It is appropriate for these applications because it offers high-performance devices while enabling high levels of integration density. For ultrahigh-speed analog front-end circuits such as wideband and high-voltage amplifiers and electro-absorption modulator drivers, InP implementations have produced the strongest results in the industry to date. Even for these circuits, however, SiGe implementation performance, as described in this paper, is reaching a level at which key application requirements are clearly attainable. This is likely to drive a shift to SiGe for these designs, just as is currently occurring for the high-speed clock and data recovery, serialization, and deserialization functions in the 40-Gb/s arena.

The paper is organized as follows. A brief review of the IBM 0.5- μm SiGe BiCMOS 5HP and 0.18- μm SiGe BiCMOS 7HP integrated circuit technologies (hereafter

designated simply as BiCMOS 5HP and BiCMOS 7HP, respectively) is given in Section 2. BiCMOS 5HP is the most mature IBM SiGe generation and has been used extensively in our 10–13-Gb/s serializer/deserializer (SerDes) work, as described in Section 3. The design and hardware characterization details of all of the SerDes building blocks, such as the voltage-controlled oscillators (VCOs), transmit and receive phase-locked loops (PLLs), and multiplexer/demultiplexer, as well as the fully integrated 10.3125- and 12.5-Gb/s versions of the SerDes chip, are discussed in that section. More demanding bandwidth and jitter requirements of OC-768 applications require the use of BiCMOS 7HP, as described in Sections 4 and 5. SerDes building blocks operating at 40–56 Gb/s with half-rate 20–28-GHz clocks are reviewed in detail in Section 4. Some of the challenges of packaging and testing at these speeds are also discussed. Section 5 contains an overview of two analog front-end functions and their implementation for SONET OC-768 in BiCMOS 7HP, namely the limiting amplifier and the electro-absorption modulator (EAM) driver. The hardware results from the EAM driver work show that SiGe heterojunction bipolar transistors (HBTs) can address high-voltage drive applications, in contrast to what is commonly assumed. The significance of accurate modeling of both active and passive devices including interconnection wiring is also highlighted in that section. Finally, a summary of the current status and potential directions for the future are presented in Section 6.

2. Technology

The circuits described in this paper are implemented using two generations (BiCMOS 5HP and BiCMOS 7HP) of the IBM family of 200-mm silicon–germanium (SiGe) technologies [6, 7]. Each technology generation provides a full device suite for analog/rf design, including a trench-isolated, highly planar SiGe heterojunction bipolar transistor (HBT), n-MOS and p-MOS devices, polysilicon and diffused resistors, high-density MOS and high-quality-factor metal-insulator-metal (MIM) capacitors, varactor diodes, and spiral inductors, all built on a 10–20- $\Omega\text{-cm}$ substrate. Compatibility with the CMOS ASIC library of the same generation (CMOS 5S and 7S, respectively) makes it possible to achieve full mixed-signal design and high levels of integration. Yield on all devices comparable to that achieved in the equivalent CMOS technology generation has been achieved.

Table 1 compares three SiGe generations. (BiCMOS 6HP [8], corresponding to the 0.25- μm column, is provided for reference only and was not used for designs described in this paper.) BiCMOS 5HP, the first production-level SiGe BiCMOS process, is used in applications from wireless communication integrated circuits (ICs) to wired data communications to high-speed

test equipment. This technology features an HBT with cutoff and maximum oscillation frequencies (f_T and f_{max}) of 47 and 65 GHz, respectively, with a minimum noise figure of 0.8 dB. The 0.5- μm (0.35- μm L_{eff}) CMOS devices utilize a 7.8-nm-thick gate oxide for 3.3-V operation. Three to five levels of aluminum interconnect are selectable, with spiral inductors available on the upper 2- μm -thick metal layer, located 6 to 10 μm above the substrate. The term BV_{CEO} denotes the collector-emitter breakdown voltage with the base open; the term BV_{CBO} denotes the collector-base breakdown voltage with the emitter open; r_B denotes the base resistance; and NF_{min} denotes the minimum noise figure.

BiCMOS 7HP presents a new generation of HBT to meet the challenges of high-speed wired communications at 40 Gb/s (SONET OC-768) and beyond. Aggressive vertical and lateral scaling yields an HBT with f_T and f_{max} of 120 GHz and 100 GHz, respectively, and a minimum noise figure of 0.4 dB. The CMOS devices feature 0.18- μm (0.11- μm L_{eff}) gates and 3.5-nm-thick gate oxide for 1.8-V operation, with thick oxide options for 2.5-V and 3.3-V operation. For rf applications, BiCMOS 7HP adds an accumulation-mode varactor (n-MOS in n-well) and a TaN metal resistor. It also integrates the IBM copper metallization process, reducing wiring resistance. The top layer is 4- μm -thick aluminum, 8 to 13 μm above the substrate, enabling the fabrication of high- Q inductors (e.g., 21 at 5 GHz for a 0.7-nH inductor).

3. SerDes chip at 10 to 13 Gb/s

Introduction

In this section, we describe an integrated serializer/deserializer (SerDes) chip implemented in SiGe BiCMOS 5HP. We first present PLL details in the form of early test chip hardware and results, followed by full SerDes hardware and results. Among the key challenges in such a design is building low-power, high-performance integrated circuits that support very high data rates while complying with stringent jitter specifications.

Figure 1 shows a typical block diagram of an optical serial link and associated electronics. The objective of the link is to serialize parallel data streams from one end of the link for high-speed transmission over an optical fiber, then deserialize them at the other end of the link. In the first step for the link, the serializer block takes low-speed parallel data and multiplexes it [via a multiplexer (MUX)] onto a high-speed serial line. This serial data is then provided to a laser diode driver (LDD), which, in conjunction with a laser diode, performs an electrical-to-optical conversion and sends the data over an optical fiber. On the receiving end of the serial data stream, the light is sensed by a photodetector, with the resulting signal amplified by a transimpedance amplifier (TIA) and then a




Figure 1

Typical block diagram of an optical serial link. From [9], with permission; © 2000 IEEE.

postamplifier to provide sufficient signal amplitude for the deserializer input. The deserializer includes a clock and data recovery circuit (CDR), which recovers the serial data and the clock embedded in that data, and a demultiplexer (DMUX), which converts the serial data stream to a final parallel output. The overall jitter budget for the link is very tight, particularly at high data rates, and much of this budget is typically claimed by the optical and opto-electronic conversion pieces of the link.

Within the SerDes chip, the transmit PLL (TxPLL) and receive PLL (RxPLL) are thus two critical circuits in the link; they execute key functions and must operate within a fraction of the overall link jitter budget. The TxPLL circuit (or simply TxPLL) uses a local low-frequency clock reference to synthesize a low-jitter, high-frequency clock required in the multiplexer stages and determines the random jitter present on the output serial data stream of the serializer. The RxPLL circuit (or simply RxPLL) is used to extract the clock signal from the received data stream, which is corrupted by noise due to the physical medium and to the optical and electronic components; it must effectively undo all of the distortion introduced in the data stream from its creation at the remote serializer to its arrival at the RxPLL serial input. The clock it extracts from the data stream is then used for data retiming and demultiplexing.

Standalone PLL test-site designs and results

In this section, we describe the core design for the transmit and receive PLLs used in the 10–13-Gb/s SerDes implementations described in this paper [9]. The PLLs were first implemented in standalone test sites to enable independent evaluation; results from these test sites are described below. These standalone PLLs were used, with slight modifications, in the full integrated SerDes implementations described in the next section. Initial designs were executed targeting 12.5-Gb/s data




Figure 2

TxPLL block diagram. From [9], with permission; © 2000 IEEE.




Figure 3

Schematic of LC VCO used in the TxPLL and RxPLL. From [9], with permission; © 2000 IEEE.

rates, intending to cover the 10-Gb Ethernet standard if 8-bit/10-bit (8b/10b) coding were chosen, as well as SONET rates with significant forward error correction overhead. When the 10-Gb Ethernet standards body chose a 64b/66b coding scheme, the target data rate was shifted to 10.3125 Gb/s. This new data rate was addressed primarily in the fully integrated SerDes designs but did not engender any significant architecture changes; the earlier test-site designs generally targeted 12.5-Gb/s operation.

TxPLL circuit design

Figure 2 shows a block diagram of the TxPLL. The main components of the loop, which uses a very standard architecture, are the voltage-controlled oscillator (VCO), the phase and frequency detector (PFD), the charge pump, the low-pass filter, and the divider. The TxPLL input is a low-frequency reference clock. The action of the TxPLL creates an output clock at a frequency N times that of the input, where N is the divide ratio of the divider. For 12.5-Gb/s data rates, a full-rate 12.5-GHz clock is generated. In contrast to a TxPLL based on a half-rate scheme [10],

the full-rate approach implemented here enables a final retiming of the multiplexed data stream, thus minimizing duty cycle variation. Because a low-phase-noise LC VCO is used in the TxPLL, the TxPLL bandwidth was chosen to be relatively low, ~ 300 kHz. This bandwidth choice enables the loop to suppress more in-band noise, while using a low-phase-noise VCO mitigates the effect of passing more VCO-generated phase noise to the PLL output. All circuits in the TxPLL are bipolar emitter-coupled logic (ECL) circuits or current mode logic (CML) circuits. Eliminating high-swing CMOS circuits from this block helps reduce noise injection that can degrade circuit performance.

Differential VCO

Figure 3 shows the simplified schematic of the differential LC VCO used in both the TxPLL and the RxPLL. The cross-coupled differential pair of bipolar transistors uses positive feedback to form a negative-resistance cell that restores energy lost by the oscillator in each cycle. The on-chip spiral inductor connected to the collectors of the differential pair devices resonates with the on-chip varactor and parasitic capacitance at that node, setting the base operating frequency of the oscillator. The value of the varactor capacitance is tuned by changing the V_{var} voltage, thus tuning the VCO. A top-level-metal 0.34-nH inductor with a quality factor of 11 at 12.5 GHz was used. A pair of emitter-follower stages buffers the VCO output, thus isolating the resonator from external perturbations. The emitter-followers are connected to the core of the VCO using a capacitive divider built using high- Q MIM capacitors. This divider is necessary to step down the high-amplitude signal generated in the tank (which helps reduce VCO phase noise) for use in the ECL circuits connected to the VCO. High- Q MIM capacitors are also used for feedback and bypass capacitors. Local power-supply and dc-node decoupling are achieved using on-chip MOS capacitors.

Phase and frequency detector

A conventional phase and frequency detector (PFD) circuit is used in the TxPLL. This PFD is built from two synchronous-set, asynchronous-reset flip-flops and a NOR gate. It acts linearly in response to phase errors and provides a saturating output in response to frequency errors. This fully differential circuit takes the reference clock and the frequency-divided VCO output as inputs and produces up and down pulses as needed to adjust the VCO frequency and/or phase. The two PFD flip-flops (RS-DFF) are synchronously set, one by the reference and one by the divided VCO, and are asynchronously reset by the NOR gate output. In the locked condition, narrow, matched up and down pulses are generated. In order to minimize static phase error, the NOR gate includes

dummy devices that allow the two flip-flops to drive nominally identical loads. Because of the relatively low operating speed of this circuit, bipolar CML is used throughout.

Charge pump, gain stage, and low-pass filter

Figure 4 shows the schematic of the differential charge pump. This circuit includes two differential pair switches, each one connected to the up and down differential outputs of the PFD. The charge pump current I_{cp} flows through both the capacitor C and the resistor R_Z , forming the transmission zero to stabilize the phase loop. A negative resistance cell is used to compensate the current leakage through the load resistors R_L [11]. Therefore, the loop filter acts ideally as a perfect integrator. Mismatch in V_{be} between the two transistors that form the negative resistance as well as mismatch between resistors will cause non-ideal cancellation of the filter leakage current, meaning that the real charge pump will act as an imperfect integrator. Because the charge-pump output swing is limited to a relatively narrow range to keep charge-pump transistors out of saturation, the low-pass filter output is connected through a linear gain stage to the varactor control input V_{var} . This linear gain stage extends the tuning range of the VCO when it is operating within the loop by amplifying the charge pump output. High-frequency poles are formed using $C1$, in parallel with C and R_Z , and another capacitor on the output of the gain stage. These poles act to suppress PFD ripple.

TxPLL on-wafer measurement results

Figure 5 shows the results of four phase-noise measurements and illustrates both TxPLL performance and the effect of the reference clock choice on that performance. The open-loop VCO phase noise is shown in trace A, while the phase noise of a relatively noisy input reference is shown in trace B. Trace C shows the phase noise of the TxPLL output clock when the noisy input reference is used; trace D shows the phase noise of the TxPLL output when a low-noise reference is used. As expected, when the noisy reference is used, inside the loop bandwidth the clock phase noise (trace C) follows the reference clock input phase noise multiplied by the divider ratio, and outside the loop bandwidth follows the open-loop VCO phase noise. This plot indicates that loop bandwidth is approximately 300 kHz, consistent with bandwidth prediction by linear PLL theory. In this case, jitter generation is still relatively high and is dominated by the input noise. When a low-noise source is used as the reference, in-band noise is dominated by charge-pump and other circuit element noise, not by input noise (see trace D). The resulting TxPLL clock phase noise yields jitter generation in a 10-kHz–100-MHz bandwidth of




Figure 4

Schematic of negative resistance charge pump. From [9], with permission; © 2000 IEEE.




Figure 5

Phase-noise plots: trace A, free-running VCO; trace B, reference; trace C, locked VCO, noisy reference; trace D, locked VCO, clean reference. From [9], with permission; © 2000 IEEE.

0.4 ps rms. The measured power dissipation of the circuit is 270 mW at 3.3 V (excluding the output test buffers).

The VCO oscillation frequency can be set between 11 GHz and 13 GHz by varying the tuning voltage between 3.3 V and 0. This is a tuning range of more than 16% around the center frequency. The VCO gain is approximately 300 MHz/V when oscillating at 12.5 GHz. The measured free-running VCO phase noise at 12.5 GHz is -101 dBc/Hz at 1-MHz offset from the carrier (Figure 5, trace A). This represents a phase-noise improvement of ~ 20 dB over a previously reported ring VCO design [12]. Since the TxPLL bandwidth is relatively small, a low-phase-noise VCO is mandatory to achieve low clock jitter generation. The effect of the VCO phase noise is less important in the RxPLL because of its much higher




Figure 6

RxPLL block diagram. From [9], with permission; © 2000 IEEE.

bandwidth (>10 MHz). However, the *LC* VCO has four times less gain than the previously reported ring VCO and thus is less sensitive to the noise on its input frequency control, which improves RxPLL performance.

RxPLL design

The RxPLL, shown schematically in **Figure 6**, uses a dual-loop architecture. The first of these loops is responsible for frequency acquisition, while the second executes data recovery. A frequency-acquisition aid is required because the pull-in range of a PLL in data-recovery mode is too small to guarantee that the PLL will lock to the incoming data stream independently of the initial frequency of the RxPLL VCO. Furthermore, keeping the VCO frequency close to the desired operating frequency will also prevent false locking of the data loop. In operation, the frequency-acquisition loop drives the RxPLL VCO frequency to a multiple of a stable reference clock input, where the reference clock frequency is set at $1/64$ th of the target data rate. A lock detector automatically switches the circuit to the data-recovery loop when the divided VCO clock frequency and the reference clock agree to within $\pm 0.1\%$. In this mode, the lock-in time of the data loop is very short because the VCO frequency is already very close to the incoming data rate. If lock is lost, the lock detector automatically switches the circuit back to frequency-acquisition mode. To reduce digital switching noise coupling, the CMOS PFD is automatically turned off by the lock detector while the loop is in data-acquisition mode [12]. A bang-bang-type phase detector that also acts as a decision circuit for data retiming is used in the data-recovery loop [12]. A high-dc-gain active loop filter is used to reduce the static phase error. The high-speed portion of the circuit uses bipolar emitter-coupled logic (ECL).

Frequency-acquisition loop enabling automatic RxPLL data acquisition

The frequency-acquisition loop and the data-acquisition loop share many components, with the key extra circuits

required by the frequency loop being the phase and frequency detector (PFD) and the divider. In this implementation, a conventional CMOS three-state PFD is used in the frequency-acquisition loop [13]; in later designs, this was replaced with a bipolar PFD for consistency with the TxPLL loop and to reduce CMOS switching noise in the design. The *up* and *down* signals from the PFD are fed via the selector to the charge pump/loop filter as differential signals after a CMOS-to-ECL conversion.

The CMOS lock detector used to switch the RxPLL between frequency-acquisition mode and data-recovery mode consists of two gray-counters, one driven by the reference clock (refclk) and the other by “divclk,” the VCO frequency divided by 64. The counters count for a period of 2048 refclk cycles, or 4096 divclk cycles, whichever is less. At the end of this interval, the counter values are compared to determine whether the two frequencies agree to within $\pm 0.1\%$. The counters are then reset before the next counting interval begins. Under ordinary conditions, 2048 refclk cycles are needed for counting, 9 refclk cycles for comparing and resetting the counters—a total of 2057 refclk cycles between each measurement update. If the frequencies are within tolerance, the lock detector immediately switches the RxPLL to data-recovery mode. Conversely, if the frequencies are out of tolerance for 1000 consecutive measurement intervals, the RxPLL is toggled back to frequency-acquisition mode. This bias toward data recovery is designed to allow the circuit to tolerate low-frequency, high-amplitude jitter on the data. In later, full-SerDes implementations, this feature was eliminated, so that single out-of-tolerance measurements switch the RxPLL to frequency-acquisition mode.

RxPLL on-wafer measurement results

Measurements of the standalone RxPLL circuit were made on-wafer using 12.5-Gb/s pseudorandom bit stream (PRBS) input data. The data-acquisition loop pull-in range is about $\pm 0.5\%$, which amply covers the $\pm 0.1\%$ tolerance of the lock detector. The jitter generation (in a 100-MHz bandwidth) of the recovered clock is 0.3 ps rms when the RxPLL is locked to $2^7 - 1$ PRBS input data (relevant to 8b/10b coded data, which has a maximum run length of 5 [14]). This jitter increases to 0.6 ps rms when $2^{31} - 1$ PRBS input data is used. The dc power consumption of the chip is 330 mW at 3.3 V excluding output test buffers, with half of this power consumed in its frequency-acquisition aid circuitry.

A 14-km-length single-mode optical fiber link test bench was used to evaluate jitter tolerance of the RxPLL in a real-world environment. **Figure 7** shows the RxPLL single-ended input and output data eye diagram and recovered clock when receiving $2^{31} - 1$ PRBS data. The measured




Figure 7

Input and output waveforms of the RxPLL: (a) 12.5-Gb/s input eye diagram; (b) 12.5-Gb/s recovered eye diagram; (c) 12.5-GHz recovered clock. From [9], with permission; © 2000 IEEE.

input data jitter is ~ 42 ps peak-to-peak. The measured BER of the RxPLL under test is less than 10^{-12} , with an input sensitivity of 100 mV.

PLL physical design

One of the challenges of mixed analog/digital designs such as the one presented in this section is to avoid the coupling of digital switching noise into sensitive analog nodes through the substrate and power-supply lines. In the physical design of these PLLs, several measures have been taken to reduce the digital noise coupling [15]:

- Since the high-speed part of the IC is fully differential, it has been implemented in a symmetrical way with respect to complementary signal paths.
- Local bias generators are used for each sensitive part of the IC, e.g., VCO, phase detector.
- Separate power supplies and grounds for the analog part and digital part within the two loops of the RxPLL are used to avoid coupling. Special care is also taken to reduce the parasitic inductance of supply lines and to improve ground homogeneity.
- A ring of substrate contacts surrounds the VCO in both PLLs and the CMOS logic present in the RxPLL. In




Figure 8

Chip photomicrographs of TxPLL (a) and RxPLL (b). From [9], with permission; © 2000 IEEE.

addition, this ring is surrounded by a ring of deep-trench oxide.

The chip microphotographs for the standalone PLLs are shown in **Figure 8**.

Single-chip SerDes design and results

In this section, we describe the design and evaluation of single-chip SerDes chips operating from 10 Gb/s to 13 Gb/s [16]. An initial SerDes design targeted 12.5-Gb/s operation, but later designs retuned the VCO to enable 10.3-Gb/s operation as required to meet the final 10-Gb Ethernet standard. Recent related work in this area includes a SiGe BiCMOS/SOI full transceiver implementation [17], a SiGe receiver/1:8 demultiplexer chip [18], a CMOS clock and data recovery circuit [19], and, more recently, CMOS serializers and deserializers [20, 21].

The SerDes described here is a fully monolithic, single-chip 16:1 serializer and 1:16 deserializer, including integrated transmit and receive PLLs with on-chip loop filters, featuring very good jitter and error-rate performance and low power consumption in BiCMOS 5HP; significant power reduction could be realized by




Figure 9

Transceiver block diagram. From [16], with permission; © 2001 IEEE.

porting this design to BiCMOS 7HP. At startup, an automatic trim operation sets the optimal center frequency of the VCOs used in the chip's two PLLs. This is achieved using available wrap modes and on-chip data checking.

Use is made of *LC* VCOs in the PLLs, one for the TxPLL and one for the RxPLL; both operate over multiple overlapping frequency bands. The chip automatically selects the correct band setting for each VCO at startup. Thus, trimming operations at test are not required for the circuit to accommodate process, supply-voltage, and temperature variations. An automatic phase-selection mechanism enables input parallel data to be correctly latched into the transceiver independently of the relationship between the arrival time of that data and the transceiver output clock edge that times the request for parallel data. Leakage testing and scanning features are built into the design to enhance testability. Serial and parallel wrap modes coupled with circuitry that checks a fixed pattern's successful transit of the transceiver further enable efficient on-wafer test using low-frequency test equipment.

Transceiver architecture

Figure 9 shows a block diagram of the transceiver. In the receive path, the input data stream is fed to the clock and data recovery (CDR) PLL. The recovered data output is then demultiplexed to 16 parallel recovered data streams using the recovered clock. This output is sent to the off-chip drivers and can also be wrapped to the transmit path data input when the part is operated in a parallel wrap test mode. A selectable full- or half-parallel-rate clock accompanies the parallel output data. In the transmit path, 16 parallel-input data streams are multiplexed to a single serial-output data stream. A clock multiplier unit

(CMU) PLL uses an input reference clock to generate the full-rate clock. This clock is required to create the clock frequencies and phases needed to stagger and multiplex the incoming data and also to execute the final full-rate retiming of the serial output. Finally, the transmit section outputs a full-parallel-rate clock to be used by another chip in the system to time the data sourced to the transceiver for serialization.

Test features as well as configuration and startup circuitry are controlled by the CMOS block of the transceiver. These features, implemented in a combination of CMOS and bipolar logic, include wrap modes, self-test, automatic VCO coarse tuning, leakage testability, and general scan design testability. Three different reference clock frequencies, the data rate divided by 16, 64, and 80, were supported in the initial 12.5-Gb/s design. A divide-by-66 mode replaced the divide-by-80 mode for the 10.3-Gb/s design.

Transmit and receive PLLs

As in the case of the standalone TxPLL, the transceiver transmit PLL takes a selectable low-frequency reference frequency and creates an output clock at the desired serial line rates as well as divided clock frequencies used in the multiplexer. By executing a final retiming of the multiplexed data stream, duty-cycle variation in the output is minimized. As in the case of the standalone RxPLL, the receive PLL uses a dual-loop structure: a data-recovery loop that recovers clock and data from the input data stream and a frequency-acquisition loop that moves the oscillation frequency of the Rx VCO inside the capture range of the data-recovery loop. All elements of both PLLs are implemented in bipolar emitter-coupled logic (ECL) or current-mode logic (CML), with the exception of lock to reference indicator circuits, which are implemented in CMOS logic. The detailed structure of the transmit and receive PLLs implemented here is closely related to those described above (see Figures 2 and 4) and in [9] and [12].

Each PLL uses an *LC* VCO. Such VCOs typically suffer from relatively narrow tuning ranges, making it difficult for these VCOs to run at target operating frequencies in the presence of process, supply-voltage, and/or temperature variations. Furthermore, even if an *LC* VCO with a broad tuning range were available, such a VCO would have high gain, which is generally undesirable because of noise sensitivity considerations, particularly in the transmit PLL.

The VCOs (**Figure 10**) used in the transceiver PLLs address these problems by operating in multiple overlapping frequency bands, where the VCO gain in any given band is modest. Band selection is digitally controlled. At startup, with a valid reference clock applied to the chip, the transceiver operates in a serial wrap

mode, sending a known parallel data pattern from the transmitter parallel input all the way to the receiver parallel output, where the data pattern is checked.

Combinations of band settings for the transmit and receive VCOs are automatically tried, and the transceiver searches for an optimal combination. Each combination is evaluated for a fixed number of reference clock cycles. Its quality is assessed by checking whether both PLLs successfully locked, whether the known data pattern was successfully recovered, and whether the locked VCOs are tuned near the middle of their tuning range. The output of the calibration circuit is a set of trim bits that set the coarse tuning of the Rx and Tx VCOs. As the automatic selection progresses, the output trim bits are modified each time a combination yielding a better assessment is identified, until the automatic trim process is complete. During operation, the 10.3-Gb/s implementation also detects whether the VCO control voltage is approaching a band edge and will automatically shift its coarse tuning to the appropriate adjacent band, provided this update feature is enabled. Automatic band selection and update capability thus enables the use of a VCO with low effective gain when operating within the main PLLs, yet does not require extraordinary measures such as multiple VCOs or fuses (which can be blown at wafer test) for the transceiver to be robust against process, temperature, and supply-voltage variations.

Multiplexer and demultiplexer

In both the 12.5-Gb/s and 10.3-Gb/s designs, the 16:1 multiplexer is composed of a recursive series of 2:1 multiplexer blocks. Similarly, the 1:16 demultiplexer is composed of a recursive series of 1:2 demultiplexer blocks. In both cases, data and clock paths are intertwined in the macros.

The sampling of the parallel Tx data coming into the transceiver presents a special timing problem. Since the clocking of all elements in the Tx multiplexer macro is derived from the Tx VCO, there is naturally a question regarding how to synchronize the incoming data with those clocks. Typically (as is the case in this transceiver), a clock is fed forward from the transceiver to a framer chip that is responsible for providing the parallel input data to the transceiver transmit path. Since the total delay in this loop, composed of driver delays, receiver delays, and card delays, is typically not known and is likely to be significant with respect to a baud interval, a mechanism is needed to ensure proper sampling of the incoming parallel data by the transceiver. In the initial transceiver design described here, a source synchronous clock (half parallel rate) that accompanies the parallel data is used as an input to a phase-detection and -selection circuit; in the follow-on designs, both half-rate and full-rate parallel clocks are supported. In any case, the phase-selection and -detection




Figure 10

Schematic of the LC VCO used in the PLLs. From [16], with permission; © 2001 IEEE.

circuit then selects a desirable phase of the local clock to sample the parallel input. In cases where automatic phase selection is not appropriate, phase selection can be explicitly set using CMOS control inputs. The phase-selection mechanism can operate continuously while the part is processing data to accommodate significant changes in path delay during part operation, although errors will be generated as the phase setting is updated. This implementation is more compact and power-efficient than the FIFO-based approach typically used to solve this problem.

Test environment and results

Both the 12.5-Gb/s and 10.3-Gb/s transceivers were first evaluated on-wafer using a standard chip tester connected to a probe card plus cabling with typical I/O bandwidths of 500 MHz. Such an environment does not allow generic full-rate testing, but does permit a successful demonstration of the full parallel interface, control circuitry, and test features such as leakage and logic scan. It also permitted the evaluation of the full range of frequencies at which the transmit and receive PLLs locked. Successful lock and optimal VCO trim settings were automatically generated at VCO operating frequencies from 11 GHz to 13 GHz in the 12.5-Gb/s design and from 9.6 GHz to 11.5 GHz in the 10.3-Gb/s design.

Full-rate testing was accomplished using a surface laminar circuit (SLC) carrier-packaged pair of transceiver chips. SLC packaging supports finer linewidths than typical board-level technologies. This characteristic makes it possible for the transceiver chips, which have solder bumps on their I/O pads, to be direct-flip-chip-attached to




Figure 11

Back-to-back serializer/deserializer test setup. From [16], with permission; © 2001 IEEE.




Figure 12

Sample serial output data eye diagram at 11.1 Gb/s. From [16], with permission; © 2001 IEEE.




Figure 13

Tx serial output at 10.3 Gb/s $2^{31} - 1$ PRBS.

the SLC; a ceramic or other space-transforming carrier is not required. This approach not only saves the cost of a package, but also reduces the number of transitions through which high-speed signals must pass, thus improving high-speed signal integrity.

The test configuration most like the expected operating mode of the transceiver is shown in **Figure 11**. In this case, a 12.5-Gb pulse pattern generator (PPG) feeds a $2^{31} - 1$ PRBS data stream to the serial Rx input (RxS) of the first transceiver (SerDes1). This data is demultiplexed, passed over the parallel interface from the SerDes1 parallel output (RxP) to the SerDes2 parallel input (TxP), multiplexed in SerDes2, and becomes output from the SerDes2 serial output (TxS). That output is compared with the SerDes1 input using a bit-error-rate tester (BERT2). Simultaneously, a second 12.5-Gb PPG feeds a second $2^{31} - 1$ PRBS data stream at a data rate nominally identical (but not synchronized) to the RxS input of the second transceiver, SerDes2. Again, this data is demultiplexed, passed over the parallel interface from the SerDes2 RxP parallel output to the SerDes1 TxP parallel input, multiplexed in SerDes1, and becomes output from the SerDes1 TxS port. In this configuration, then, each transceiver effectively acts as a framer for the other transceiver. Furthermore, independently controllable data streams at noncommensurate frequencies run in opposite directions through both transceivers, enabling chip performance under worst-case crosstalk conditions to be evaluated. No performance degradation and no injection locking are observed under these conditions, even when the operating frequency of one data path is dragged through that of the second. This test configuration was successfully exercised for both the 12.5-Gb/s and 10.3-Gb/s versions of the transceiver.

In the operating condition described above, the card using 12.5-Gb/s transceivers performed error-free overnight at data rates from 11 Gb/s to 11.9 Gb/s. At 12.5 Gb/s, error rates below 5×10^{-14} were measured. The degradation in error rate was due to a noise-coupling path that was eliminated in the follow-on design, resulting in error-free operation across the entire operating band of the part. Both Tx and Rx jitter generation at 12.5 Gb/s were below 0.5 ps rms, measured with a spectrum-analyzer phase-noise utility with integration limits from 10 kHz to 100 MHz. The TxS output eye diagram at 11.1 Gb/s is shown in **Figure 12**. The TxS output eye diagram at 10.3 Gb/s from the 10.3-Gb/s design is shown in **Figure 13**. Supply-voltage variations of $\pm 5\%$ about the nominal 3.3-V level did not affect performance. In both the 12.5-Gb/s and the 10.3-Gb/s designs, operation was also demonstrated on-wafer at ambient temperatures ranging from 0 to 100°C. Total power dissipation at nominal operating conditions

was 3.3 W, with significant reduction possible by porting the design to a more advanced process.

Transceiver physical design

In high-speed circuits, significant attention must be paid to physical design both to avoid speed degradation and to minimize noise coupling [15]. In this transceiver design, the target operating mode of the part, in which transmit and receive paths operate simultaneously at noncommensurate yet almost identical frequencies, poses a particularly challenging problem with respect to noise coupling. In addition, switching noise from CMOS logic and digital drivers must not significantly degrade PLL performance. Several key physical design strategies were followed in the transceiver design reported here:

- The Rx and Tx VCOs were placed far apart, with a ring of substrate contacts and deep-trench isolation surrounding each VCO. The Rx and Tx PLLs were also isolated from each other and from the rest of the circuit with substrate contact regions and deep-trench rings. The CMOS section was placed on the opposite side of the chip from the PLLs.
- Three power-supply domains were maintained on-chip, one for the PLLs, one for the multiplexer, demultiplexer, and parallel I/O, and one for the CMOS section. Low-inductance connections to power-supply pads were implemented throughout.
- Substrate contacts outside the CMOS region were not tied to the analog ground; these contacts nominally carry no current.
- Separate bias generators were used for key blocks.

A micrograph of a 12.5-Gb/s-design transceiver chip is shown in **Figure 14**. The design integrates approximately 8400 HBTs and 30000 FETs.

4. SerDes circuits at 40 Gb/s

Introduction

A 40-Gb/s serial data communication link is similar to a 10-Gb/s communication link, but with everything running four times faster. In the case of the 40-Gb/s link, the laser diode driver used in a typical 10-Gb/s data link is replaced by an electro-absorption modulator driver; for the 40-Gb/s application, the laser is operated in a CW mode, and an electro-absorption modulator is used to modulate the effective power coupled to the fiber. A generic 40-Gb/s link is shown in **Figure 15**. Beyond serial data communication links, the core multiplexing elements of serializers and core demultiplexing elements of deserializers are also necessities for test equipment, such as pattern generators and error detectors, operating at 40-Gb/s data rates.




Figure 14

Micrograph of transceiver chip. From [16], with permission; © 2001 IEEE.




Figure 15

Generic 40-Gb/s optical link.

Two of the key elements of this link are the serializer and deserializer. In this section, building blocks for serializers and deserializers that support 40-Gb/s non-return-to-zero (NRZ) serial data rates are described. Section 5 provides some details regarding front-end chips.

The main building blocks for the serializer and deserializer are the multiplexer, the demultiplexer, the clock multiplier unit PLL (CMU), and the clock and data recovery PLL (CDR). The number of signals multiplexed in the serializer and demultiplexed in the deserializer is implementation-specific, but might typically be of the




Figure 16

Simplified block diagram of the 4:1 multiplexer circuit. From [23], with permission; © 2002 IEEE.

order of 16. The most critical portions of these designs are their highest-frequency sections, namely the 4:1 multiplexer stage and the 1:4 demultiplexer stage. Implementing and demonstrating these core circuits is thus an excellent first step in developing a working 40-Gb/s serializer and deserializer, as is demonstrating the CMU and the CDR. In the remainder of this section, we present details of the core circuits needed to build serializers and deserializers that operate at 40-Gb/s data rates, as implemented in SiGe BiCMOS 7HP.

4:1 multiplexer and 1:4 demultiplexer

There are several key challenges associated with the design of the 4:1 multiplexer and 1:4 demultiplexer circuits [22]. The first of these challenges is, of course, the high data rate these circuits must handle. The base SONET OC-768 data rate is just below 40 Gb/s. Once forward-error-correction schemes are added, however, the line rate exceeds 40 Gb/s, with the actual line rate depending on the chosen scheme. Because the SiGe BiCMOS 7HP technology does not support a 40-Gb/s full-rate latch at moderate power with margin, the multiplexer, demultiplexer, CDR, and CMU PLLs use a half-rate architecture. The high operating speeds of the circuits also present significant additional challenges in layout, packaging, and test. In the following, we illustrate how

these challenges can be met through a description of 4:1 multiplexer and 1:4 demultiplexer designs that were implemented in SiGe BiCMOS 7HP, tested on-wafer to data rates above 50 Gb/s, packaged, and then demonstrated as packaged parts to data rates above 50 Gb/s.

A block diagram of the 4:1 multiplexer circuit is shown in **Figure 16**. The multiplexer was implemented using a tree architecture in which the base elements of the tree are a recursive series of 2:1 multiplexer stages. The logic family chosen for this design is emitter-coupled logic, which offers sufficient performance while enabling the use of supply voltages in the -3.3-V range. Single-ended internal signal swing levels were chosen to be 300 mV throughout the design, providing adequate signal-to-noise ratio without demanding a larger supply voltage.

The multiplexer IC takes four parallel single-ended data streams at its input and a half-rate input clock, e.g., 20 GHz for 40-Gb/s operation, to create the final multiplexed output. The parallel data are received by single-ended-to-differential-conversion buffers with on-chip $50\text{-}\Omega$ matching resistors. These buffers contain two differential pairs in order to achieve a good common-mode rejection at the input of the two first multiplexing stages. The input clock is internally divided by two with a static divider (built with a toggle flip-flop) in order to latch the four parallel data inputs and perform the first multiplexing operation. The clock input is received with a double-stage wide-bandwidth Cherry-Hooper amplifier [24] in order to speed up clock rise and fall edges, since sine-wave synthesizers are used for testing at these frequencies. These amplifiers permit the use of very-low-power-level input clocks. In each 2:1 multiplexer stage, proper data timing is achieved by using latches. These latches act to offset the two input data streams by 90° with respect to each other and must operate with a clock-to-data-output time delay of less than half the bit time of the output multiplexed data. The current consumption of the multiplexer is 410 mA from a nominal -3.6-V supply voltage; 37% of this current is consumed in the input and output buffers.

The demultiplexer IC receives a full-rate differential data stream and outputs four parallel single-ended quarter-rate data streams. The demultiplexer circuit block diagram is shown in **Figure 17**. Like the multiplexer, it uses a tree architecture, here with a recursive series of 1:2 demultiplexer stages. It uses a half-rate input clock for the first 1:2 demultiplexing stage. This clock is then divided by 2 with a static divider in order to perform the last 1:2 demultiplexing operation. Each 1:2 demultiplexer stage uses latches for demultiplexing and data alignment. The serial input data and input clock are received with double-stage wide-bandwidth Cherry-Hooper amplifiers with on-chip $50\text{-}\Omega$ matching resistors. A “bit-skip” dc signal is

XORed with the divided clock to allow bit-rotation control at the parallel outputs of the demultiplexer. The chip consumes 430 mA from a -3.6-V supply voltage, with 40% of this current consumed in the input and output buffers.

At the high data rates intended for the multiplexer and demultiplexer, attention to detail in the physical design is crucial in order to maximize circuit performance. Layout parasitics must be minimized, particularly on high-speed lines. In order to preserve signal integrity, transmission lines are used on long runs. Finally, in order to support at-speed evaluation prior to packaging, the layout used must be compatible with both wafer-level probing and packaging.

The multiplexer and demultiplexer ICs were mounted in a custom-designed package that enabled high-frequency operation of the chips [23, 25]. The package included a housing, a ceramic substrate, and the IC itself. The aluminum/graphite composite housing provided support for high-speed V-connectors, protected the ceramic and chip, and served as a heat sink. The ceramic substrate measured 1.75 in. by 1.95 in. On the ceramic, a finite-ground-plane coplanar waveguide was used for signal lines to limit the number of propagation modes and achieve better isolation between signal traces. Ribbon bonds were used for chip-level interconnection, and the bonding lengths were minimized and well controlled to minimize parasitic inductance. Feed-throughs with internal decoupling capacitors were used for power and bias supplies, and surface-mount capacitors were used on selected power traces near the chip within the package to further reduce supply noise. An assembled package and its chip-level interconnections are shown in **Figure 18**. The same package can be used for either multiplexer or demultiplexer chips.

Testing the multiplexer requires that four quarter-rate data streams and a commensurate half-rate clock be provided to the circuit. In test, the data streams and clock come from test equipment through cables to the part; delay lines are used to match the arrival times of the data and to position the clock appropriately for successful multiplexing. When the clock and data frequency are changed, some delay adjustment must generally be repeated if cable lengths do not match. The output signal is observed using a high-frequency oscilloscope, and the quality of the output eye is evaluated. Testing of the demultiplexer requires that full-rate data and a commensurate half-rate clock be provided to the circuit. Again, the data stream and the clock come from test equipment through cables to the part, so that initial delay adjustment between clock and data is necessary, as is follow-on adjustment for significant input-frequency changes. The four demultiplexed outputs can then be observed using an oscilloscope, and the quality of the




Figure 17

Simplified block diagram of the 4:1 demultiplexer circuit. From [23], with permission; © 2002 IEEE.




Figure 18

Assembled package and demultiplexer chip-level interconnections. From [23], with permission; © 2002 IEEE.

resulting four output eyes evaluated. Once packaged parts are available, testing of the multiplexer and demultiplexer in tandem is possible by connecting either the parallel ports or the serial ports of the devices. Error-rate testing of the multiplexer and demultiplexer in isolation can also be executed.

Both the multiplexer and demultiplexer circuits were first tested on-wafer to check functionality. On-wafer, the demultiplexer could not initially be evaluated at target data rates because of test equipment limitations, but basic operation was demonstrated up to 12.5 Gb/s. However, in early testing we were able to demonstrate multiplexer




Figure 19

Input and output waveforms of 4:1 MUX at 40 Gb/s: (a) 20-GHz clock input; (b) on-wafer measured 40-Gb/s output eye diagram; (c) measured 40-Gb/s output eye diagram of packaged sample. From [23], with permission; © 2002 IEEE.




Figure 20

Simplified block diagram of back-to-back MUX/DMUX test setup. From [23], with permission; © 2002 IEEE.

operation on-wafer up to 56 Gb/s at a supply voltage of -3.3 V by feeding it 4×14 -Gb/s $2^7 - 1$ PRBS inputs. At the time, 14 Gb/s constituted the maximum test equipment speed available to us. In later testing, 60-Gb/s operation of both the multiplexer and the demultiplexer was demonstrated. The input amplitude was 500 mV single-ended (the inputs are referenced to a dc level

of -250 mV). Once functionality was checked, both ICs were packaged. This enabled us to perform bit-error-rate testing on both circuits in order to confirm their functionality at full speed. **Figure 19** gives a qualitative comparison of the multiplexer (MUX) performance at 40 Gb/s when tested on-wafer and in packaged form. The IC was tested with $2^{31} - 1$ PRBS, -3.3 -V supply voltage, and 100°C chip temperature (during on-wafer testing). The modest performance degradation observed in the packaged MUX eye compared to the on-wafer MUX eye is probably due to return loss at the serial output in the packaged part.

The final test for the packaged parts was evaluation in a back-to-back configuration connected through their serial interface. When testing in this configuration, shown in **Figure 20**, error-free operation ($<10^{-15}$) at 40 Gb/s was achieved on all four parallel outputs of the demultiplexer with $2^{31} - 1$ PRBS parallel inputs at the multiplexer. At that speed, consistent error-free operation was observed when varying the supply voltage from -3.3 V to -3.9 V. The clock-phase margin of the demultiplexer (DMUX) was approximately 11 ps, or 44% of a unit interval. Error-free operation ($<10^{-15}$) was also achieved at 50 Gb/s with a supply voltage of -3.5 V and $2^7 - 1$ PRBS data at the MUX parallel inputs. At this data rate, the clock-phase margin of the DMUX was reduced to 6 ps, or 31% of a unit interval. Corresponding 50-Gb/s eye diagrams are shown in **Figure 21**. Error-free operation up to 52.2 Gb/s was also achieved under the same supply-voltage conditions.

Clock and data recovery circuit and clock multiplier unit at 40 Gb/s

Overview

In addition to the demultiplexer and multiplexer cores themselves, the clock and data recovery circuit (CDR) and the clock multiplier unit circuit (CMU) are required elements of a complete deserializer and serializer, respectively. The function of the clock and data recovery circuit is to extract clock information from the input high-rate serial data stream and then to use that clock to reconstruct the input data. Generally, the signal carrying the data is significantly degraded in the course of its transmission over the high-serial-data-rate link. The CDR must compensate for the degradation of (primarily) timing jitter and amplitude, recovering a clean copy of what was transmitted with extremely low error rates. The output data and clock from the CDR is delivered to the demultiplexer. The CMU supports the multiplexing function, taking an input reference clock and multiplying the frequency of that clock to generate the set of clocks needed by the multiplexer. In the case of a full-rate architecture, the CMU might also provide a clock that drives a final at-speed latch, enabling very uniform

data launch timing at the final output of the serializer. The clocks generated by the CMU must exhibit very little jitter, and, in the case of a half-rate architecture, as was pursued in the 40-Gb/s designs described in the following, the highest-frequency clock must have very uniform duty-cycle characteristics.

The remainder of this section focuses on the CDR for 40-Gb/s serial data. The general operation of the CMU is exactly as described in the 10-Gb/s work presented above, although a half-rate architecture was used.

The challenges faced in executing a 40-Gb/s CDR design are quite similar to those faced in executing the multiplexer and demultiplexer cores. The high data rate pushes the limits of the technology for full-rate architectures, driving the design to half-rate architectures, which unfortunately add complexity while making it possible to achieve the desirable lower operating frequency.

The 40-Gb/s CDR circuit uses a half-rate clock architecture, which means that the on-chip VCO is running at 20 GHz, in contrast to a full-rate architecture, in which the VCO runs at 40 GHz. This architecture was chosen for two main reasons. First, given npn device performance, it allows sufficient margin to ensure 40-Gb/s operation of the chip with respect to process, temperature, and power-supply variations. Second, a full-rate architecture is a more aggressive design that would have required higher power consumption and, most likely, a larger power-supply voltage to achieve the required design margin.

The CDR circuit block diagram is shown in **Figure 22**. In this first-pass design, the loop solely implements a data-recovery function. Because of the relatively narrow capture range of a typical CDR, a frequency-acquisition aid loop is usually required to bring the operating frequency of the loop close to that required for the input data rate. Once the loop frequency is close enough, the data loop takes over. A frequency-acquisition aid loop was successfully incorporated and demonstrated to be operational in a follow-on version of the chip.

The CDR itself extracts a 20-GHz clock from a 40-Gb/s serial input data stream. This clock is then used for data demultiplexing, which, in this design, includes the retiming function. The 1:2 demultiplexing function is actually part of the phase-detector operation. The VCO used in the CDR has two frequency-tuning inputs. The first of these is controlled by the output of the charge pump, forming the integral path of the loop (I-filter). The second of these is controlled directly by the output of the phase detector (P-filter), effectively implementing the zero of the loop by direct modulation of the VCO frequency. The I-filter, which has a very large time constant, is responsible for the frequency acquisition of the loop. The P-filter, meanwhile, applies high-frequency modulation pulses to effect phase




Figure 21

50-Gb/s output eye diagram of the 4:1 MUX (top) and one of the 12.5-Gb/s demultiplexed outputs (bottom). From [23], with permission; © 2002 IEEE.




Figure 22

CDR block diagram.

corrections of the VCO. The amplitude of the P-filter pulses controls the bandwidth of the loop; it is externally adjustable in the implementation described here. Applying phase corrections directly from the output of the phase detector (PD) reduces loop latency and thus jitter generation. The charge pump current of the I-filter controls the loop damping factor and has little effect




Figure 23

Input amplifier.




Figure 24

Block diagram of phase detector.

on the loop bandwidth, assuming that the damping factor is much larger than 1 (which is the case here because the PD gain is relatively high). This current is also adjustable in the implementation described here. The loop filter capacitor is integrated on chip. In addition to the main recovered data output, the chip outputs three additional

signals for design evaluation—the two 1:2 demultiplexer data streams and the VCO clock divided by 2.

Input amplifier

The input amplifier schematic is shown in **Figure 23**. It uses the well-known Cherry-Hooper architecture, which consists of a series chain of an alternating transadmittance stage (TAS) and a transimpedance stage (TIS). The input impedance of the TIS is small compared to the output impedance of the TAS. This results in a mismatch between the stages up to high frequencies, increasing the bandwidth of the amplifier. Because the amplifier fan-out is high (it drives four latches of the phase detector), a tree architecture is used in which the TIS of a first amplifier is driving the TAS of two second amplifiers. The transimpedance amplifier uses active parallel feedback in order to increase its overall gain-bandwidth product. The input stage of the amplifier consists of a pair of emitter-followers with 50-Ω input termination for impedance matching.

Phase detector

The block diagram of the PD is shown in **Figure 24**. This PD consists of three double-edge-triggered master-slave flip-flops D1, D2, and D3; two latches L1 and L2; one XOR gate X1; and one modified AND gate A1. The flip-flops D1 and D2 sample the incoming data stream on the rising and falling edges of the 20-GHz I and Q VCO clock signals, respectively. Note that the VCO Q signal is 90° out of phase with respect to the VCO I signal. The loop executes data recovery as follows: When the loop is locked, data transitions are aligned with the Q signal from the VCO clock, enabling the I clock to be used to sample the data at a favorable time. This is effected in the following way: Depending on the relative phase of the VCO clock with respect to the data, the output of D1 leads or lags the output of D2. D3 samples the output of D2 on the rising and falling edges of the output of D1 to produce a binary signal indicating whether the VCO clock is leading or lagging the incoming data. In the locked condition, the edges of the VCO Q clock signal will be aligned with the transitions of the incoming data. Assuming a 50% duty-cycle VCO clock and good I and Q matching, the VCO I clock signal then samples the data at the midpoint of each bit interval, thus acting as a decision circuit with no phase adjustment required over process, temperature, and power-supply variation. The duty cycle of the VCO clock and the I and Q matching are important to ensure effective operation of the PD, especially with regard to jitter tolerance, assuming an input data stream for which the optimum sampling point is effectively at the midpoint of each bit.

Because the SONET OC-768 standard allows the run length of random data to be large, the PD must provide a

tri-state output when no data transitions are present in order to avoid excessive VCO frequency drift which can unlock the loop. The XOR gate X1 senses the presence of transitions by comparing the outputs of L1 and L2 (these outputs are actually the demultiplexed data, D_Even and D_Odd, which are 90° out of phase). The outputs of X1 and D3 are combined in gate A1 (**Figure 25**), which generates the tri-level output of the phase detector, where the three levels are up, down, and tri-state. Gate A1 is a CML AND gate modified so that its output is tri-stated if its $\bar{Z}$ input is asserted.

Two critical constraints that must be satisfied for a successful implementation of this architecture are the delay matching between A1 and the selector of D3, and the symmetry of the XOR gate. The former constraint provides good-clock-to-data-delay matching between L1 and L2 and the latches of D3, which is easily achieved, while the latter requires matching between the two inputs to output propagation delay. Reasonable delay mismatch (a few picoseconds) will not break the PD operation but instead will slightly modify the dynamic of the loop. Though not catastrophic, this may lead to an increase in generation of jitter because of erroneous VCO frequency corrections.

Charge pump

The charge pump used in the 40-Gb/s CDR uses the same approach used in the 10-Gb/s CMU and CDR described above, in which a negative resistance cell is used to cancel filter leakage current. The basic topology for this circuit was shown in Figure 4.

20-GHz I/Q VCO

Half-rate CDR architectures require, in general, a VCO with precise quadrature phases; these phases are referred to as I and Q. The basic circuit building block in the implementation described here is a CMOS *LC* oscillator consisting of two back-to-back cross-coupled inverters with an inductor load (**Figure 26**). The inductor resonates with the varactor capacitance and the gate and drain junction capacitance of the FETs to determine the oscillation frequency. The negative resistance of the cross-coupled FETs acts to overcome the inductor loss as well as the layout wiring loss. In the I-Q oscillator implementation, two such identical oscillator building blocks, labeled A and B in Figure 26, are coupled to each other by FETs (M5–M8) of the same size as the internal negative resistance FETs (M1–M4), thus creating direct coupling in one direction and cross-coupling in the other.

In order to understand the operation of the coupled oscillators, first suppose that the two oscillators synchronize in-phase. In this case, the cross-coupled path from oscillator B to A absorbs the negative resistance current produced by (M1A, M2A) and (M3A, M4A), and oscillator A stops, which in turn stops oscillator B through




Figure 25

Modified AND gate schematic.




Figure 26

20-GHz CMOS VCO with quadrature outputs. From [26], with permission; © 2002 IEEE.

the cross-coupling path. The same process applies in reverse if the two oscillators are anti-phase. Therefore, oscillations can coexist in both coupled oscillators only when both are synchronized in quadrature. In this mode of operation, the unique phase relationship of the four outputs are 0° at M2A, 180° at M1A, 90° at M2B, and 270° at M1B. The current in the quadrature-output VCO is controlled by a large p-FET device connected between ground and the top rail of the two oscillators.

Measurement results

The CDR has been tested on-wafer using 40-GHz bandwidth probes and $2^7 - 1$ PRBS input data. Its total dc power consumption was 1.3 W at -3.6-V supply




Figure 27

Input and output waveforms of 45-Gb/s CDR.




Figure 28

VCO phase-noise performance vs. offset frequency for TxPLL VCO operating at 21 GHz under various coarse calibration settings, at operating temperatures ranging from 0 to 100°C, corresponding to the different symbols.

voltage. Without the output test buffers, the power consumption was approximately 1.1 W.

The quadrature CMOS VCO oscillation frequency was observed to be 10% higher than expected. It had a phase

noise of -97 to -100 dBc/Hz at 1-MHz offset. The VCO consumed about 16 mA from -3.6 V. The tuning range was about 600 MHz. This small tuning range ensured a small VCO gain for PLL applications. When locked to PRBS input data, the VCO oscillated between 22.1 and 22.6 GHz, corresponding to a hold range of 2.2%. In a follow-on design, the tuning range of the VCO was extended to 2 GHz using a band-switching approach similar to that used in the 10-Gb/s transceiver work described above. At 45 Gb/s PRBS input, the recovered clock phase noise was approximately 0.3 ps or 13 mUI in a 50-kHz–500-MHz bandwidth. Input and output waveforms showing 45-Gb/s CDR operation are shown in **Figure 27**.

Clock multiplier unit

As mentioned above, the architecture for the clock multiplier unit is the same as that used for the 10-Gb/s clock multiplier unit described earlier in this paper. The CMU for 40-Gb/s-data-rate applications, however, was designed using a half-rate architecture. The main consequence of this choice is that the VCO runs at half rate. The phase-noise performance of this VCO is shown in **Figure 28**; its architecture is the same as that shown in Figure 10. A system-level consequence of using a half-rate architecture on the transmit path is that careful attention must be paid to duty-cycle symmetry in the 20-GHz output clock; duty-cycle variation creates additional deterministic jitter in the output data waveform.

5. Front-end circuits at 40 Gb/s

Introduction

In addition to the serializer and the deserializer, other key elements of a 40-Gb/s link are the front-end circuits. In this section, the design and performance of two examples of these circuits implemented in SiGe BiCMOS 7HP are discussed. We first describe a limiting amplifier, which receives the output of a transimpedance amplifier and generates a signal of sufficient amplitude to satisfy the input sensitivity requirements of the serial data receiver of the CDR circuit. We then describe an electro-absorption modulator driver, which receives the output data stream from the serializer and generates sufficient swing and the correct characteristic impedance to satisfy the input requirements of an electro-absorption modulator.

40-Gb/s limiting amplifier

The 40-Gb/s limiting amplifier described here employs a CML topology to allow dc coupling of the 50- Ω I/O circuits. Because of the high impedance of the input stage, input matching is achieved using passive components. Each input signal was terminated by a 50- Ω resistor tied to ground, which also supplied the bias current for the input level-shifter stage. The layout of the input

Table 2 Comparison of simulated and measured performance for 40-Gb/s limiting amplifier in an experimental ceramic package with V-conductor signal launches.

Parameter	Extracted simulation	Measured (de-embedded)
Gain	26 dB	25 dB
Rise and fall times	7 ps	8 ps
Jitter, deterministic	<1 ps	0.9 ps

termination network employed a 50- Ω ground-signal-signal-ground (G-S-S-G) taper to lump the bond-pad capacitance into the controlled impedance of the taper line, thus avoiding a capacitive discontinuity at the bond pad.

The core of the amplifier consisted of four limiting amplifier cells, each having a gain of 7 dB (**Figure 29**). These amplifiers do not employ feedback to boost gain, since we wished to explore amplifier topologies that would be easier to extend to higher frequencies without fear of phase margin issues. Each amplifier consisted of a differential pair with two diodes as the load, followed by an emitter-follower stage to buffer and shift the output to the proper level for the next input stage. Small-signal analysis of this configuration gives a theoretical voltage gain of $g_{m\text{load}}/g_{m\text{diff pair}} = 2g_m/g_m = 2$, or a gain of 6 dB. However, large-signal analysis and simulations showed that both the gain and the bandwidth could be increased slightly by employing load diodes slightly smaller than the input differential pair device size. By so doing, the gain per cell in limiting mode was increased to 7 dB.

Though the core of the amplifier was unconditionally stable, employing no feedback, the emitter-follower buffer stage needed careful design to avoid “inductive” behavior causing overshoot in response and increasing amplifier jitter. It should be noted that if feedback amplifiers had been used, they too would have required emitter-follower level shifters between stages and thus the same care in design [27].

To support 40-Gb/s I/O, source and load terminations were assumed to minimize potential reflections. The output driver was a cascoded differential amplifier with 50- Ω loads which provided the termination of a 50- Ω G-S-S-G coplanar taper, identical to the taper used at the input.

Table 2 compares simulated and measured performance for the 40-Gb/s limiting amplifier packaged in an experimental ceramic package with V-conductor signal launches. The amplifier draws 120 mA from a single -5.2-V supply, dissipating 624 mW while running at 40-Gb/s rates. A 40-Gb/s differential pseudorandom bit stream generated by a SHF¹ pattern generator was deskewed with wide-bandwidth delay lines, attenuated,

¹ SHF Communication Technologies AG, Berlin, Germany.




Figure 29

Simplified schematic of gain stage with emitter-follower level shifter. In the figure, the term V_{cc} denotes the supply voltage; V_{i+} and V_{i-} denote the positive and negative inputs, respectively; V_{o+} and V_{o-} denote the positive and negative outputs, respectively; and I_b denotes the bias current.

then fed into the post-amplifier (or simply “post-amp”). After deskewing, attenuation, and transmission through cables, the 20%–80% rise and fall times of the signal fed to the post-amp were 14 ps (12 ps de-embedded).

Because the package and connector contribute 2 to 3 dB in attenuation at these frequencies, some de-embedding was necessary in order to compare the post-amplifier performance to simulations. The measured gain, 25 dB, agreed well with the 26-dB simulated gain for a nominal device at room temperature. De-embedding the scope response yielded 7- to 8-ps 20%–80% rise times for the amplifier operating in the limiting mode.

Figure 30(a) shows typical eye diagrams of the $2^{31} - 1$ pseudorandom data from the pattern generator. The amplifier output eye [**Figure 30(b)**] shows little added jitter and faster edge rates when compared to the input. A clean output eye is maintained for input data from 20-mV to 1-V peak-to-peak differential amplitude or over a 34-dB dynamic range.

Modulator driver

At bit rates much beyond 10 Gb/s, indirect laser modulation is required in order to prevent frequency chirp and concomitant chromatic dispersion penalties in such high-speed links. In this section, we report the results of a single-ended electro-absorption modulator (EAM) driver designed to deliver over 3 V p-p single-ended drive into 50- Ω termination while consuming 3 W from a single -7.5-V supply. The driver also supplied an adjustable dc bias voltage needed by the EAM device.

As in the case of the post-amp, the input and output terminations on this device were on-chip, low-parasitic 50- Ω polysilicon resistors used with coplanar waveguide (CPW) launches to allow dc coupling of the high-speed




Figure 30

Typical eye diagrams of the $2^{31} - 1$ pseudorandom data from (a) the 40-Gb/s pattern generator with 1.6-ps rms jitter and (b) the amplified signal with 1.9-ps rms jitter and faster edges.




Figure 31

Simplified schematic of EAM driver output stage with inductive peaking.

signals (Figure 31). Waveform symmetry adjustment was effected by applying an analog voltage to an input offset control pin. This driver circuit was tested on-wafer using 40-GHz bandwidth coplanar G-S-G probes for the high-speed input and output signals. A representative eye diagram from first-pass hardware at 20 Gb/s is




Figure 32

20-Gb/s signal source (top) and EAM driver output (bottom) showing 3.2-V peak-to-peak single-ended drive voltage. From [26], with permission; © 2002 IEEE.

shown in Figure 32. The noise in both eyes was primarily sampling head noise.

An important aspect of this application is the requirement for large voltage swings—indeed, large enough that there was some skepticism in the III–V semiconductor community that a SiGe-based design would ever succeed in generating the necessary swings with any degree of reliability. This driver was the first SiGe 7HP design reported with a single-ended peak-to-peak output swing larger than the collector–emitter breakdown voltage with the base open (BV_{CEO}); swings of 3.2 V (peak-to-peak), well in excess of BV_{CEO} , are shown in the figure. This was possible because of the low base drive impedance in the design, allowing the devices to operate with a collector-to-emitter voltage (V_{ce}) of 4 to 5 V (peak-to-peak)—well in excess of the dc-rated BV_{CEO} . The driver was further stressed by increasing the supply voltage to -10 V and running the driver for several hours. Although the waveform was distorted at this high supply voltage, the device exhibited nominal operation with no discernible degradation when the voltage was restored to -7.5 V. Device testing is underway at high dynamic V_{ce} swings to further clarify the safe operating area under low-base drive impedance conditions. These results show that SiGe bipolar technology may be able to address higher voltage drive applications than would be expected, given the low value of BV_{CEO} .

From the figure it is clear that the output rise/fall times (bottom) were faster than the 20-Gb/s input from the signal source (top). The output signal exhibited deterministic jitter (DJ) in the form of waveform

splitting on the falling edge, which was not predicted by simulations. The source of this DJ has not yet been determined, but is most likely due to an interaction between the peaking used to compensate for the additional loading presented by the dc offset devices and reflections in the CPW-to-package interface. The dc offset devices are externally controllable current references that can be used to set the operating point for the EAM. A new design that dispenses with the on-chip bias trim, which was deemed unnecessary, is underway. This redesign should enable operation at higher speeds and reduced DJ, since the peaking will be greatly reduced.

Discussion: Analog front-end circuit design for 40-Gb/s applications

Designing, packaging, and measuring high-speed analog front-end circuits targeting 40-Gb/s data rates is challenging because these efforts push the edge of what is possible while stretching modeling, circuit extraction, and simulation state of the art to the limit. Rules of thumb derived from experience as well as the results of standalone full-wave electromagnetic simulators were used in on-chip interconnect layout, but more care was taken in the design of coplanar taper launches to ensure that they provided a good match to a 50- Ω environment.

Packages for these circuits were also designed with EM simulators to provide controlled 50- Ω microstrip lines in alumina or organic substrates tapering out to 60-GHz V-connectors. The chips were placed in a cavity in the substrate to make the bonding surface of the die and package coplanar, allowing the shortest possible ribbon bonds in order to reduce parasitic inductance. At 40 GHz, for example, the magnitude of $j\omega L$ for a 100-pH inductance (a ribbon bond approximately 200 μm long) is 25 Ω —a huge discontinuity in a 50- Ω environment! We believe that such impedance discontinuities accounted for the performance-reduction results for packaged parts compared to on-wafer results.

Measurements were similarly hampered by passive and active equipment limitations. Electrical cables rated for use at a 60-GHz bandwidth and length less than 30 cm actually have 1–2-dB skin-effect losses from dc to 40 GHz which cause intersymbol interference in the signals fed into our high-speed chips. Wafer probing required even longer cables with more loss because bulky test equipment could not be placed on the prober platform. Finally, as seen in the eye diagrams, the pattern generator used to test the circuits had slower rise and fall times than the circuits themselves, and the scope sampling head, rated at 50 GHz, had 7–8-ps intrinsic rise times. For these reasons, de-embedding of measurements was required in order to estimate actual circuit performance. For example, measured 20–80% rise/fall times were 10 ps, with the

intrinsic scope response being 7–8 ps, giving 7–8-ps rise times for de-embedded estimates. As another example, the pattern generator rms jitter was 1.7 ps and the measured post-amp output was 1.9-ps rms jitter, yielding an estimate of 0.9-ps rms jitter contributed by the post-amp. Since an error detector at 40 Gb/s was not available when measurements were taken, we could not use the bathtub curve method to extract jitter information, but even this method likely would have suffered from jitter limitations of the high-speed error detector.

6. Concluding remarks

High-speed data communication networks of the 21st century will continue to require high performance, low power, and low bit-error rates from serial links at an affordable price, forcing the IC industry to move to higher levels of integration at lower power levels. This can be achieved only by using silicon-based technologies combined with cost-effective packaging and testing approaches. This paper has highlighted the challenges of such a serial transmission subsystem design environment from a circuit designer's perspective, focusing on the integration of key serial link transceiver functions. A wide range of microwave mixed-signal circuit implementations have been described, including 10–13-Gb/s serializer/deserializer chips with subpicosecond jitter performance, 56-Gb/s multiplexer/demultiplexer ICs, 40-Gb/s clock-and-data recovery/clock-multiplier phase-locked loops, limiting amplifiers, and EAM drivers. SiGe BiCMOS technologies in production at IBM have been used throughout this work. While the 10-Gb/s designs utilized full-rate architectures, the 40-Gb/s work relied on half-rate clocking schemes. However, because of the recent advances in the IBM SiGe HBT technology, which provides transistor cutoff frequencies beyond 200 GHz [28], it is now possible to implement OC-768 functions at an acceptable system-level power consumption level using a full-rate architecture.

On the basis of the leading-edge performance levels described in this paper, it can be foreseen that SiGe-based designs will continue to move into market segments previously believed to be the stronghold of GaAs- and InP-based technologies. With increased emphasis on developing microwave models for active and passive elements as well as the package and board environment, it is expected that the SiGe technology will be the choice for multi-terabit-throughput parallel links of the future. In addition to the broadband communications field, which was the focus of this paper, the future microwave test equipment markets and the emerging field of short-range wireless local area networks, which can carry Gb/s data, should also benefit from the advances described.

Acknowledgment

The authors wish to acknowledge the SiGe technology group of the IBM Microelectronics Division for chip fabrication. They also wish to thank K. Jenkins, D. Heidel, J. Trehwella, A. Joseph, G. Freeman, D. Hame, S. Subbanna, M. Oprysko, and B. Meyerson for their support.

References

1. IEEE Standard 802.3ae-2002: 10 Gb Ethernet; Institute of Electrical and Electronics Engineers, New York, 2002.
2. See ANSI/INCITS 303-1998, ANSI/INCITS 297-1997, and follow-on documents addressing 10-Gb/s Fibre Channel implementations; American National Standards Institute and International Committee for Information Technology Standards, Washington, DC.
3. "Synchronous Optical Network (SONET) Transport Systems: Common Generic Criteria," *GR-253-CORE*; available through Telcordia Technologies, Inc., 8 Corporate Place, Piscataway, NJ 08854, September 2000.
4. M. Soyuer, H. A. Ainspan, M. Meghelli, and J.-O. Plouchart, "Low-Power Multi-GHz and Multi-Gb/s SiGe BiCMOS Circuits," *Proc. IEEE* **88**, No. 10, 1572–1582 (October 2000).
5. Y. Baeyens, G. Georgiou, J. Weiner, A. Leven, V. Houtsmas, P. Paschke, Q. Lee, R. Kopf, Y. Yang, L. Chua, C. Chen, C. T. Liu, and Y.-K. Chen, "InP D-HBT ICs for 40-Gb/s and Higher Bitrate Lightwave Transceivers," *IEEE J. Solid-State Circuits* **37**, No. 9, 1152–1159 (September 2002).
6. D. C. Ahlgren, G. Freeman, S. Subbanna, R. Groves, D. Greenberg, J. Malinowski, D. Nguyen-Ngoc, S. J. Jeng, K. Stein, K. Schonenberg, D. Kiesling, B. Martin, S. Wu, D. Hame, and B. Meyerson, "A SiGe HBT BiCMOS Technology for Mixed-Signal RF Applications," *Proceedings of the IEEE Bipolar/BiCMOS Circuits and Technology Meeting (BCTM)*, 1997, pp. 195–197.
7. A. Joseph, D. Coolbaugh, M. Zierak, R. Wuthrich, P. Geiss, Z. He, X. Liu, B. Orner, J. Johnson, G. Freeman, D. Ahlgren, B. Jagannathan, L. Lanzerotti, V. Ramachandran, J. Malinowski, H. Chen, J. Chu, P. Gray, R. Johnson, J. Dunn, S. Subbanna, K. Schonenberg, D. Hame, R. Groves, K. Watson, D. Jadas, M. Meghelli, and A. Rylyakov, "A 0.18 μm BiCMOS Technology Featuring 120/100 GHz (f_T/f_{MAX}) HBT and ASIC-Compatible CMOS Using Copper Interconnect," *Proceedings of the IEEE Bipolar/BiCMOS Circuits and Technology Meeting (BCTM)*, 2001, pp. 143–146.
8. S. A. St. Onge, D. L. Hame, J. S. Dunn, S. Subbanna, D. C. Ahlgren, G. Freeman, B. Jagannathan, J. Jeng, K. Schonenberg, K. Stein, R. Groves, D. Coolbaugh, N. Feilchenfeld, P. Geiss, M. Gordon, P. Gray, D. Hershberger, and S. Kilpatrick, "A 0.24 μm SiGe BiCMOS Mixed-Signal RF Production Technology Featuring a 47-GHz f_T HBT and 0.18 μm L_{eff} CMOS," *Proceedings of the IEEE Bipolar/BiCMOS Circuits and Technology Meeting (BCTM)*, 1999, pp. 117–120.
9. D. Friedman, M. Meghelli, B. Parker, H. Ainspan, and M. Soyuer, "Sub-Picosecond Jitter SiGe BiCMOS Transmit and Receive PLLs for 12.5Gbaud Serial Data Communication," *IEEE Symposium on VLSI Circuits, Digest of Technical Papers*, June 2000, pp. 132–135.
10. J. F. Ewen, A. X. Widmer, M. Soyuer, K. R. Wrenner, B. Parker, and H. A. Ainspan, "Single-chip 1062-Mbaud CMOS Transceiver for Serial Data Communications," *IEEE ISSCC Digest of Technical Papers*, February 1995, pp. 32–33.
11. T. H. Lee and J. F. Bulzacchelli, "A 155-MHz Clock Recovery Delay- and Phase-Locked Loop," *IEEE J. Solid-State Circuits* **SC-27**, 1736–1746 (December 1992).
12. M. Meghelli, B. Parker, H. Ainspan, and M. Soyuer, "SiGe BiCMOS 3.3V Clock and Data Recovery Circuits for 10Gb/s Serial Transmission Systems," *IEEE J. Solid-State Circuits* **35**, 1992–1995 (December 2000).
13. J. F. Ewen, M. Soyuer, A. X. Widmer, K. R. Wrenner, B. D. Parker, and H. A. Ainspan, "CMOS Circuits for Gb/s Serial Data Communication," *IBM J. Res. & Dev.* **39**, No. 1/2, 73–81 (January/March 1995).
14. A. X. Widmer and P. A. Franaszek, "A DC-Balanced, Partitioned-Block, 8B/10B Transmission Code," *IBM J. Res. & Dev.* **27**, No. 5, 440–451 (September 1983).
15. M. Meghelli, M. Bouche, and A. Konczykowska, "Very High Speed Integrated Circuits: Design Methodology and Applications for Optical Communications," *Proceedings of the European Conference on Circuit Theory and Design (ECCTD)*, Budapest, Hungary, September 1997, pp. 1366–1370.
16. D. Friedman, M. Meghelli, B. Parker, J. Yang, H. Ainspan, and M. Soyuer, "A Single-Chip 12.5Gbaud Transceiver for Serial Data Communication," *IEEE Symposium on VLSI Circuits, Digest of Technical Papers*, June 2001, pp. 145–148.
17. S. Ueno, K. Watanabe, T. Kato, T. Shinohara, K. Mikami, T. Hashimoto, A. Takai, K. Washio, R. Takeyari, and T. Harada, "A Single-Bit 10Gb/s Transceiver LSI Using SiGe SOI BiCMOS," *IEEE ISSCC Digest of Technical Papers*, February 2001, pp. 82–83.
18. Y. Greshishchev, P. Schvan, J. L. Showell, Mu-Liang Xu, J. J. Ojha, and J. E. Rogers, "A Fully Integrated SiGe Receiver IC for 10-Gb/s Data Rate," *IEEE J. Solid-State Circuits* **35**, 1949–1957 (December 2000).
19. J. Savoj and B. Razavi, "A 10Gb/s CMOS Clock and Data Recovery Circuit with Frequency Detection," *IEEE ISSCC Digest of Technical Papers*, February 2001, pp. 78–79.
20. M. Green, A. Momtaz, K. Vakilian, Xin Wang, Keh-Chee Jen, D. Chung, Jun Cao, M. Caresosa, A. Hairapetian, I. Fujimori, and Yijun Cai, "OC-192 Transmitter in Standard 0.18 μm CMOS," *IEEE ISSCC Digest of Technical Papers*, February 2002, pp. 248–249.
21. J. Cao, A. Momtaz, K. Vakilian, M. Green, D. Chung, Keh-Chee Jen, M. Caresosa, B. Tan, I. Fujimori, and A. Hairapetian, "OC-192 Receiver in Standard 0.18 μm CMOS," *IEEE ISSCC Digest of Technical Papers*, February 2002, pp. 250–251.
22. M. Reinhold, C. Dorschky, E. Rose, R. Pulella, P. Mayer, F. Kunz, Y. Baeyens, T. Link, and J.-P. Mattia, "A Fully-Integrated 40 Gb/s Clock and Data Recovery/1:4 DEMUX IC in SiGe Technology," *IEEE J. Solid-State Circuits* **36**, 1937–1945 (December 2001).
23. M. Meghelli, A. V. Rylyakov, and Lei Shan, "50 Gb/s SiGe BiCMOS 4:1 Multiplexer and 1:4 Demultiplexer for Serial Communication Systems," *IEEE ISSCC Digest of Technical Papers*, February 2002, pp. 260–261.
24. W. Pöhlmann, "A Silicon-Bipolar Amplifier for 10Gb/s with 45 dB Gain," *IEEE J. Solid-State Circuits* **29**, 551–556 (May 1994).
25. L. Shan, M. Meghelli, Joong-Ho Kim, J. Trehwella, M. Taubenblatt, and M. Oprysko, "Millimeter Wave Package Design: A Comparison of Simulation and Measurement Results," *Proceedings of the 10th IEEE Topical Meeting on Electrical Performance of Electronic Packaging*, 2001, pp. 29–34.
26. G. Freeman, M. Meghelli, Y. Kwark, S. Zier, A. Rylyakov, M. A. Sorna, T. Tanji, O. M. Schreiber, K. Walter, J.-S. Rieh, B. Jagannathan, A. Joseph, and S. Subbanna, "40-Gb/s Circuit Built from a 120-GHz f_T SiGe Technology," *IEEE J. Solid-State Circuits* **37**, 1106–1114 (2002).

27. Y. Greshishchev and P. Schvan, "A 60-dB Gain, 55-dB Dynamic Range, 10-Gb/s Broad-Band SiGe HBT Limiting Amplifier," *IEEE J. Solid-State Circuits* **34**, 1914–1920 (December 1999).
28. A. Joseph, D. Coolbaugh, D. Hamee, G. Freeman, S. Subbanna, M. Doherty, J. Dunn, C. Dickey, D. Greenberg, R. Groves, M. Meghelli, A. Rylyakov, M. Sorna, O. Schreiber, D. Herman, and T. Tanji, "0.13 mm 210 GHz f_T SiGe HBTs—Expanding the Horizons of SiGe BiCMOS," *IEEE ISSCC Digest of Technical Papers*, February 2002, pp. 180–181.

Received March 12, 2002; accepted for publication November 26, 2002

Daniel J. Friedman IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (dfriedmn@us.ibm.com). Dr. Friedman received a Ph.D. degree in engineering science from Harvard University in 1992. After completing postdoctoral work in image sensor design, he joined IBM at the Thomas J. Watson Research Center in 1994. His initial work at IBM was the design of analog circuits and air interface protocols for field-powered radio frequency identification tags. Since 1999, he has focused on analog circuit design for high-speed serial data communication, publishing several papers in this field. Since June 2000, Dr. Friedman has managed a team of circuit designers working on serial data communication and wireless applications. Beyond publishing circuits papers regarding serial links, he has published articles on imagers and radio frequency identification (RFID); he holds 12 patents.

Mounir Meghelli IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (meghelli@us.ibm.com). Dr. Meghelli received an M.S. degree in electronics and automatics from the University of Paris XI in 1992 and an Engineering degree in telecommunications from the Ecole Nationale Supérieure des Télécommunications de Paris in 1994. From 1994 to 1998, he was with the France Télécom Research Center, CNET-Bagneux, as a Ph.D. student working on the design of high-speed ICs. Since 1998, Dr. Meghelli has been at the Thomas J. Watson Research Center working on SiGe BiCMOS high-speed integrated circuits for optical fiber communications.

Benjamin D. Parker IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (nj@us.ibm.com). Mr. Parker received a B.S. degree in physics from Bowdoin College, Brunswick, Maine, in 1979. He received an M.S. degree in physics from Brown University in Providence, Rhode Island, in 1981. His graduate work dealt with the optical properties of adsorbed layers on metal surfaces. In 1986 Mr. Parker joined the GaAs group at the Thomas J. Watson Research Center, where he worked on the characterization of III–V semiconductor systems. In 1991, he joined the Mixed-Signal Communications IC Design group to work on the design and verification of high-speed serial communication links.

Jungwook Yang SmaL Camera Technologies, 10 Wilson Road, Cambridge, Massachusetts 02138 (yang@smalcamera.com).

Herschel A. Ainspan IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (ainspan@us.ibm.com). Mr. Ainspan received B.S. and M.S. degrees in electrical engineering from Columbia University in 1989 and 1991, respectively. In 1989 he joined IBM at the Thomas J. Watson Research Center, where he has been involved in the design of mixed-signal and rf ICs for high-speed data communications.

Alexander V. Rylyakov IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (sasha@us.ibm.com). Dr. Rylyakov received an M.S. degree in physics from the Moscow Institute of Physics and Technology in 1989, and a Ph.D. degree in physics from the State University of New York (SUNY) at Stony Brook in 1997. From 1994 to 1999 he worked in the Department of Physics at SUNY Stony Brook on the design and testing of

high-speed (up to 770 GHz) digital integrated circuits based on a superconductor Josephson junction technology. In early 1999 Dr. Rylyakov joined IBM at the Thomas J. Watson Research Center as a Research Staff Member to work on the design and testing of full-custom digital and mixed-signal integrated circuits for serial communications (in the 1-Gb/s to 40-Gb/s range and beyond), using a broad spectrum of CMOS, SiGe bipolar, and BiCMOS technologies.

Young Hoon Kwark *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (kwark@us.ibm.com).* Dr. Kwark received degrees in electrical engineering from MIT and Stanford University. He is currently a circuit designer working in high-speed circuit design at the Thomas J. Watson Research Center.

Mark B. Ritter *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (mritter@us.ibm.com).* Dr. Ritter received a B.S. degree in physics from Montana State University in 1981 and M.S., M.Phil, and Ph.D. degrees from Yale University in 1987. In the fall of 1986, Dr. Ritter took a postdoctoral research position in the Physical Sciences Department of the Thomas J. Watson Research Center. His postdoctoral work was with Dr. David Awschalom; it centered on nonlinear optical properties of fluid-filled porous media and carrier dynamics in diluted magnetic superlattice structures. Dr. Ritter became an IBM Research Staff Member in 1989. Since 1993 he has been manager of a group designing analog and digital circuits for data communications. His group has contributed to the Fibre Channel standard for storage-area networks and to the 802.3ae 10-Gb/s Ethernet standard, as well as to the IR wireless (IrDA) standard. The group's work has also resulted in a number of products, including Fibre Channel transceivers, IrDA transceivers, and, most recently, 10-Gb/s Ethernet transceivers, for which his group contributed to the design of the analog front-end circuits at 10 Gb/s. Dr. Ritter and his group are currently exploring the use of the SiGe technology for communications links at 40 Gb/s and beyond, as well as the use of Si and other technologies for future terabit communications links within switches and servers. Dr. Ritter received the 1982 American Physical Society Apker Award. Additionally, he has received three IBM Outstanding Innovation Awards and several Research Division and Technical Group Awards. He is an author or coauthor of 30 technical publications and a co-inventor of 14 U.S. patents.

Lei Shan *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (leis@us.ibm.com).* Dr. Shan received a B.S. degree in mechanical engineering from Zhejiang University in 1986. In 2000, he received an M.S. degree in electrical engineering and a Ph.D. degree in mechanical engineering from the Georgia Institute of Technology. From 1997 to 2000, as a Graduate Assistant at Georgia Tech, he studied the tribological aspects of the chemical-mechanical polishing process used in integrated circuit fabrication. In 2000, he joined IBM at the Thomas J. Watson Research Center as a Research Staff Member to work on high-speed electronics/optoelectronics packaging design as well as electrical and thermal modeling. Dr. Shan designed and demonstrated a high-speed package for a 50-Gb/s multiplexer and demultiplexer based on IBM SiGe BiCMOS technology. He has authored more than twenty papers and patents and has received the 1999 Best Paper Award, *Journal of Tribology*, ASME.

Steven J. Zier *IBM Microelectronics Division, East Fishkill facility, Route 52, Hopewell Junction, New York 12533 (zier@us.ibm.com).* Mr. Zier received a B.E. degree in electrical engineering from the State University of New York at Stony Brook in 1983 and an M.S. degree in computer science from Rensselaer Polytechnic Institute in 1993. From the time he joined IBM in 1983 until 1988, he designed low-cost bipolar gate array libraries; from 1988 to 1994, he designed mainframe CMOS/BiCMOS SRAM chips; from 1994 to 2000, he designed low-cost Flash/DRAM chips; and from 2000 to the present he has designed SiGe CMOS/BiCMOS serial communication chips (in the 1-Gb/s to 43-Gb/s range).

Michael Sorna *IBM Microelectronics Division, East Fishkill facility, Route 52, Hopewell Junction, New York 12533 (sorna@us.ibm.com).* Mr. Sorna received a B.S. degree in electrical engineering from the University of Pittsburgh in 1984, and an M.Eng. degree in electrical engineering from Rensselaer Polytechnic Institute in 1986. He is currently working toward a Ph.D. degree in electrical engineering at Syracuse University. Mr. Sorna has been working at the IBM East Fishkill facility since 1984—since 1998 as a Senior Engineer. From 1984 to 1993, he worked in the Test Equipment Engineering Group, developing high-speed wafer-level production test equipment. From 1993 to 1998, his responsibility centered around high-speed SerDes/chip circuit design, in which a wide variety of standards were addressed: IEEE 1394, Fiber Channel, Gigabit Ethernet, Infiniband, and SONET (OC-192 and OC-768). Mr. Sorna's design experience includes both analog and digital experience in bipolar (including SiGe), CMOS (0.1–0.8- μ m), and GaAs technologies. His recent achievements include leading a team to develop the first IBM OC-768 (43-Gb/s) SONET-compliant chip set. He is currently responsible for IBM ASIC SerDes research and development, including circuit design, architecture, and laboratory characterization. He has authored numerous technical publications and presentations, primarily in the field of SerDes circuits and BIST techniques.

Mehmet Soyuer *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (soyuer@us.ibm.com).* Dr. Soyuer received a Ph.D. degree in electrical engineering from the University of California at Berkeley in 1988, subsequently joining IBM at the Thomas J. Watson Research Center as a Research Staff Member. His work has involved high-frequency mixed-signal integrated circuit (IC) designs, in particular monolithic phase-locked-loop designs for clock and data recovery, clock multiplication, and frequency synthesis using silicon and SiGe technologies. He managed the Mixed-Signal Communications IC Design group from 1997 to 2000, and since then has been the Senior Manager of the Communication Circuits and Systems Department within the IBM Communication Research and Development Center. Dr. Soyuer has authored numerous papers in the areas of analog, mixed-signal, rf, microwave, and nonlinear electronic circuit design, and he is an inventor and co-inventor of seven U.S. patents. Since 1997, he has been a technical committee member of the International Solid-State Circuits Conference (ISSCC). He was an Associate Editor of the *IEEE Journal of Solid-State Circuits* from 1998 through 2000. Dr. Soyuer also chaired the Analog, MEMS and Mixed-Signal Electronics Committee of the International Symposium on Low Power Electronics and Design (ISLPED) in 2001.