A 10-Gb/s Compact Low-Power Serial I/O With DFE-IIR Equalization in 65-nm CMOS

Byungsub Kim, *Student Member, IEEE*, Yong Liu, *Member, IEEE*, Timothy O. Dickson, *Member, IEEE*, John F. Bulzacchelli, *Member, IEEE*, and Daniel J. Friedman, *Member, IEEE*

Abstract—A compact and power-efficient serial I/O targeting dense silicon carrier interconnects is reported. Based on expected channel characteristics, the proposed I/O features low-impedance transmitter termination, high-impedance receiver termination, and a receiver with modified DFE with IIR filter feedback (DFE-IIR). The DFE-IIR receiver uses a single additional IIR filter feedback tap to compensate many postcursors without paying the power and area penalty that would be incurred with a conventional high tap-count DFE. Equalization capabilities of the compact I/O at 10 Gb/s are demonstrated over various channels including conventional chip-to-chip and backplane links with half-baud losses of up to 27 dB. Finally, a transmitter-receiver pair operating over a 40-mm on-chip emulated silicon carrier channel was demonstrated to 8.9 Gb/s, at a link power efficiency of 1.9 mW/Gb/s.

Index Terms—Backplane channel communication, chip-to-chip communication, compact I/O, continuous-time IIR filter, decision feedback equalizer, serial link, silicon carrier links.

I. INTRODUCTION

T HE design of compact low-power I/O transceivers continues to be a challenge for both chip-to-chip and backplane communication applications. The introduction of dense, fine-pitch silicon packaging technologies, in principle capable of supporting tens of thousands of high data rate I/O for local chip-to-chip interconnect, will make I/O area and power requirements even more stringent. An example of such a dense packaging technology is silicon carrier, in which chips to be packaged are mounted on a silicon substrate and inter-chip signaling is supported with dense back end of line (BEOL) copper wires. Because silicon carrier technology offers not only fine pitch wiring but also good thermal coefficient of expansion matching with active chips, microbumps can be used to join active chips with a silicon carrier [1]. Fig. 1 depicts two chips mounted on a silicon carrier, connected to each other using dense BEOL wires. Signals are transmitted and received by the I/O circuits on the two chips; the two chips communicate to the
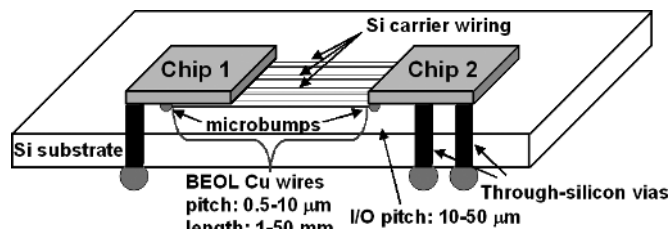

Fig. 1.   Conceptual silicon carrier link.

outside world via through-silicon vias. To capitalize on the potential advantages of the silicon carrier packaging technology, however, compact and low-power I/O circuitry suited to the silicon carrier's channel characteristics must be developed.

Note that the BEOL wires used in fine-pitch silicon carrier interconnect are very narrow and thin. Silicon carrier channels therefore have significant channel losses at high data rates ($>$5 Gb/s), leading to significant inter-symbol interference (ISI). Therefore, channel equalization is needed to compensate the channel loss for these channels. Among the various channel equalization methods that have been developed, nonlinear decision feedback equalizer (DFE) [2]–[7] has been widely recognized as an effective approach because it reduces signal distortion without amplifying noise or crosstalk. In order to apply DFE techniques to silicon carrier channels, the DFE design must address the fact that long silicon carrier channels exhibit a very long pulse response tail, suggesting that a high tap-count DFE receiver is required. However, the power and area usage of such a high tap-count DFE receiver would be prohibitive in the context of compact and low-power silicon carrier links. Therefore, the challenge of compact I/O in silicon carrier links is achieving the performance of a high tap-count DFE without paying an unacceptable area and power penalty. To meet this challenge, this paper reports a modified DFE with IIR filter feedback (DFE-IIR), which takes advantage of silicon carrier channel characteristics to enable efficient equalization of carrier channel loss [8].

The paper is organized as follows. Section II introduces the silicon carrier channel characteristics and the challenges of compact low-power I/O for such channels. The compact I/O architecture of this work is presented in Section III, including termination and equalization strategies. Section IV presents the circuit design details of the transmitter and receiver in the proposed compact I/O. To verify the proposed I/O scheme and evaluate its performance over silicon carrier channels, prototype test sites were implemented. These test sites and associated measurement results are presented in Section V and
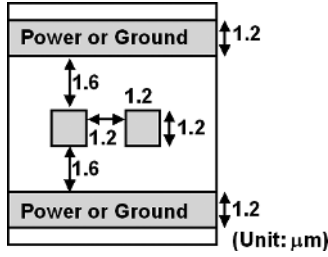
Fig. 2. Cross section of an example silicon carrier channel.

Fig. 3. Simulated eye diagrams of 5-Gb/s data transmitted over various silicon carrier channels. (a) 2.5-mm channel. (b) 7.5-mm channel. (c) 10-mm channel. (d) 20-mm channel.
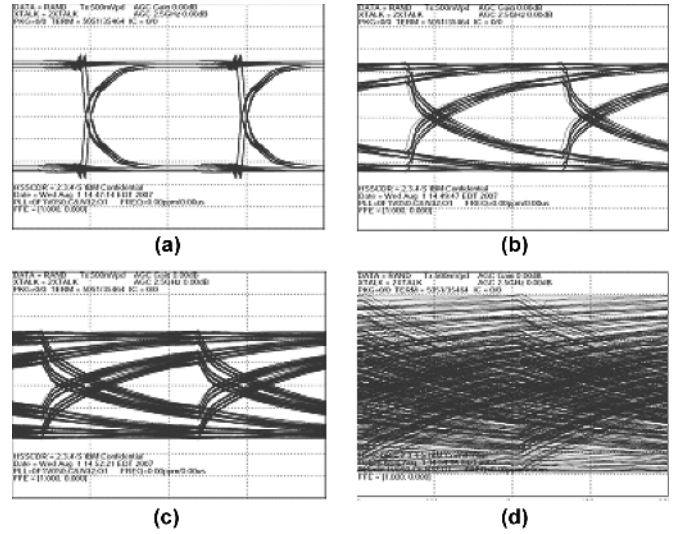
Section VI, respectively. Finally, the conclusions are given in Section VII.

II. SILICON CARRIER CHANNEL

A. Characteristics of Silicon Carrier Channel

A key goal of the I/O circuits in the silicon carrier links is achieving low power and area costs. To achieve this goal, the properties of the silicon carrier channel must be understood, enabling optimization of the I/O circuits for this packaging medium.

Fig. 2 shows the cross section of a representative silicon carrier channel, here realized as a differential stripline structure. To support ultrahigh aggregate bandwidth between chips, silicon carrier channels must have very high wiring density. Consequently, fine-pitch carrier wiring is required. In the depicted channel, the two signal wires are 1.2 μm wide, 1.2 μm high, and spaced apart by 1.2 μm. These wiring dimensions result in significant resistive losses when the on-carrier channel is long, e.g., over 5 mm. Furthermore, just as is the case for traditional backplane channels, silicon carrier channels have limited channel bandwidth, so the silicon carrier channel demonstrates higher losses at higher frequencies. Therefore, when high-frequency data is transmitted over long carrier channels, channel losses and limited channel bandwidth make the recovery of received data difficult. Fig. 3 shows simulated eye diagrams (observed after the channel) of data transmitted over silicon carrier channels with lengths ranging from 2.5 mm to 20 mm. As can be seen in the figure, the data eye opening is clearly reduced when the channel length increases from 2.5 mm to 10 mm, and the eye is completely closed at 20 mm. This result shows that equalization is generally required for transmission over long channels (>10 mm) at high data rates (>5 Gb/s). The typical range of the single-ended characteristic impedance and the effective relative permittivity of the silicon channels are 50 Ω to 110 Ω and 4 respectively.

Simulated characteristics of a 20-mm silicon carrier channel built using the wire dimensions of Fig. 2 are shown in Fig. 4. The measured frequency response for this line shows 6 dB of *dc* attenuation (Fig. 4(a)), with loss rising to 17 dB at 5 GHz. Note, however, that since the silicon carrier channel is built from microbumps and on-carrier transmission lines, signal reflections due to impedance discontinuities are not expected. The S21 transfer function of the 20-mm carrier channel therefore varies smoothly with frequency, in contrast to the equivalent transfer function of a typical backplane link, where elements such as via stubs produce a more structured S21 response. In the time domain, the channel transfer function corresponds to a long decaying pulse response tail [Fig. 4(b), derived from the measured S21 transfer function in Fig. 4(a) and normalized to make the peak amplitude to 1]. This long decaying tail creates a challenge and an opportunity for the equalizer design, as discussed below.

B. Challenges of Compact Low-Power I/O for Silicon Carrier Channel

The normalized pulse response curve of Fig. 4(b) shows that an isolated '1' sent over the silicon carrier channel results in received data with significant postcursor content. As discussed in Section I, DFE is preferred to equalize the ISI from the postcursors since it does not amplify noise or crosstalk [2]–[7]. To equalize a silicon carrier channel with such a long pulse response tail, a conventional DFE would require many DFE taps to be effective. However, in the context of compact and low-power silicon carrier links, the power and area penalty would be prohibitive. The challenge of equalizers for silicon carrier channels is to achieve the performance of a high tap-count DFE without paying an unacceptable area and power penalty.

Another question associated with compact low-power I/O for the silicon carrier channel is the selection of the proper termination strategy. Conventionally, the best link performance is achieved when both transmitter and receiver are terminated with 50 Ω, and matched with the channel impedance. However, 50-Ω termination in the receiver introduces significant power dissipation and thus may not be the optimal choice for compact I/O designed for silicon carrier links. Analyzing the power dissipation of a simple termination network provides quantitative insight regarding this issue. Fig. 5 illustrates a conceptual differential signaling termination network. To achieve a required differential voltage amplitude V_{RX} on the receiver differential termination resistor $2R_{RX}$ to meet certain signal to noise ratio (SNR), the static current flowing through the transmitter, channel, and receiver is $I = 0.5V_{RX}/R_{RX}$. Note that the current direction depends on the transmitter output $TXOUT$. Then for 50-Ω
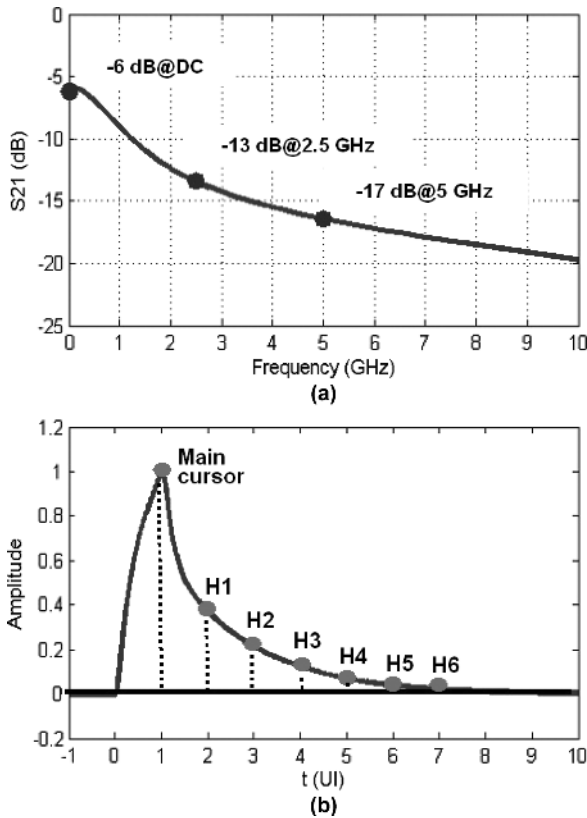
Fig. 4. Simulated characteristics of a 20-mm silicon carrier channel. (a) Transfer function. (b) Normalized pulse response sampled at 100-ps UI.
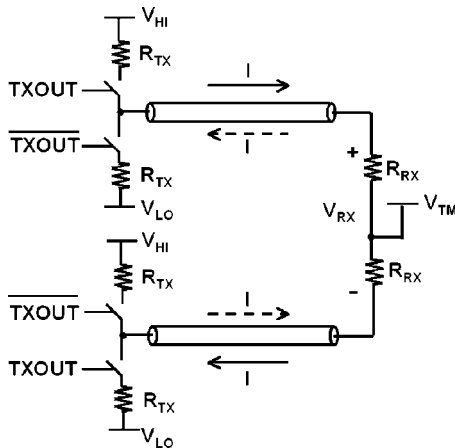

Fig. 5. Conceptual differential signaling termination network.

Fig. 6. Comparison of simulated output eye diagrams of transmitters driving silicon carrier channels. (a) Transmitter terminated with 50 Ω driving a 10-mm channel at 5 Gb/s. (b) Transmitter terminated with 10 Ω driving a 40-mm channel at 10 Gb/s.

III. COMPACT I/O ARCHITECTURE

A successful compact I/O architecture operating over silicon carrier channels must tackle the termination and equalization challenges discussed in Section II. The approach taken to meet these challenges is to leverage the characteristics of the silicon carrier channel itself, as follows [8].

A. Termination Strategy: Low-Impedance Transmitter and High-Impedance Receiver

While the focus in the previous section was on the receiver termination, it is also important to consider the transmitter termination strategy as part of the overall I/O design. The transmitter output impedance determines the drive strength and slew rate. To achieve high drive capability for to increrase the system SNR, a low impedance termination was chosen for the compact, carrier-optimized I/O transmitter. Fig. 6 shows the simulated differential transmitter output eye diagrams (observed at the transmitter output) with different source impedances driving silicon carrier channels of different lengths. When the transmitter terminated with 50 Ω drives a 10-mm silicon carrier channel at 5 Gb/s, the output signal eye diagram has very small opening [Fig. 6(a)]. In contrast, when the transmitter output impedance is 10 Ω, it can drive a 40-mm channel at 10 Gb/s with a much wider data eye opening [Fig. 6(b)]. The better performance of the low-impedance transmitter as compared to the 50-Ω transmitter, even driving a longer channel at a higher data rate, drove the selection of a low-impedance transmitter.

On the receiver side, as discussed in Section II, 50-Ω termination results in significant, undesirable signal power dissipation. If, instead, the receiver is terminated with high impedance, significant signal power can be saved. For example, a 1-kΩ receiver termination results in only 0.2 mA of static signal current for 0.8-V peak-to-peak differential signal swing, much less than the 4 mA drawn if a 50-Ω termination network is chosen. As an additional benefit, the high receiver impedance gives large signal swing.

From the discussion presented to this point, it appears that to achieve large signal swing, high slew rate and low power

single-ended receiver termination, to get a 0.8-V peak-to-peak differential signal swing $(V_{RX} = 0.4 \text{ V})$ at the receiver, it results in 4 mA of static current. In the total power budget of the silicon carrier compact I/O application, this represents excessive overhead. Reducing the signal static power dissipation demands the use of a high impedance termination of the receiver. This approach risks introducing signal reflections due to the mismatch of receiver termination impedance and channel impedance, adversely affecting signal integrity. Joint optimization of power and performance for the receiver termination is therefore necessary.
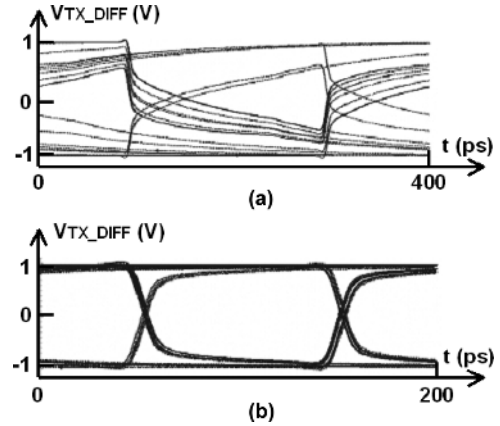
dissipation from the termination network, the combination of a low-impedance transmitter and a high-impedance receiver is desired. Because silicon carrier channels have typical single-ended characteristic impedances ranging from 50 Ω to 110 Ω, such a termination strategy leads to impedance mismatches at both the transmitter and the receiver. Therefore, the signal integrity impact due to this impedance mismatch must also be evaluated to assess whether the performance penalty these mismatches will create is acceptable.

Practical lengths of silicon carrier channels range from a couple of millimeters to tens of millimeters. Long carrier channels show very large signal loss, as was shown in the transfer function of the 20-mm channel [Fig. 4(a)]. When considering the impact of signal reflections, this large channel loss is beneficial as it attenuates the reflected signal. While this impedance mismatch scheme does not create significant performance penalties for long links, it could be problematic for shorter channels, where the channel loss is too low to effectively suppress reflections. Therefore, to study the short channel case, detailed channel simulations are needed. Fig. 7 shows simulated receiver eye diagrams of 10-Gb/s data transmitted over silicon carrier channels at various lengths. The transmitter is terminated with less than 50 Ω and the receiver is terminated with 1 kΩ.[1] Three short carrier channels, with lengths 2.5 mm, 5 mm, and 7.5 mm, were simulated and found to introduce propagation delays of 17 ps, 33 ps, and 50 ps, respectively. At a 10-Gb/s data rate, these delays correspond to 0.17 UI, 0.33 UI, and 0.5 UI. Note that 0.5-UI channel delay is the worst case for reflections; since the reflected signal has a round-trip delay of 1 UI, it has the greatest impact on the next received bit. For all three short channels, despite clear signal integrity impact from the termination choice, in the center of the received data eyes, voltage margin is more than 450 mV.

B. Equalization Strategy: DFE-IIR Receiver

In order to understand the DFE-IIR design, it is first important to understand a typical DFE-based wireline receiver. Fig. 8 shows a conceptual wireline communication system with a bandwidth-limited channel and its response to a single "1" bit transmitted over the channel. The received bit is spread over multiple unit intervals in the time domain due to the dispersive channel. Consequently, intersymbol interference (ISI) from this bit will affect subsequent bit decisions by the receiver. One approach used to overcome this postcursor ISI is to add a DFE in the receiver. In such a receiver, previous bit decisions are fed back with appropriately weighted tap coefficients to subtract ISI from the current bit. In Fig. 8, a 2-tap DFE is shown, here enabling cancellation of the inter-symbol interference from the preceding two bits. Note that because DFE relies on decisions made about previous data bits, it can only cancel the interference from postcursors. To cancel the interference from precursors (ISI affecting preceding bits), feed-forward equalization (FFE), typically realized in the transmitter, is commonly used.
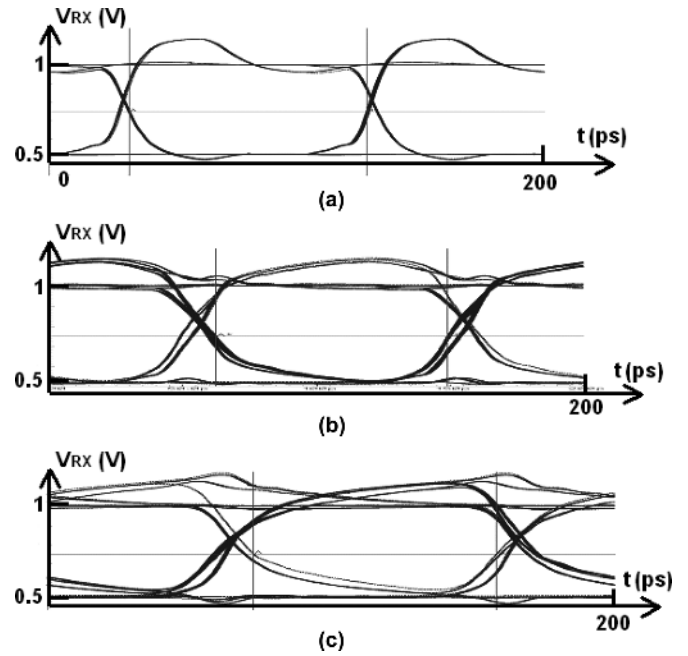
Fig. 7. Simulated receiver eye diagrams of 10 Gb/s data transmitted over channels at various lengths with <50-Ω terminated transmitter and 1-kΩ terminated receiver. (a) 2.5-mm channel. (b) 5-mm channel. (c) 7.5-mm channel.
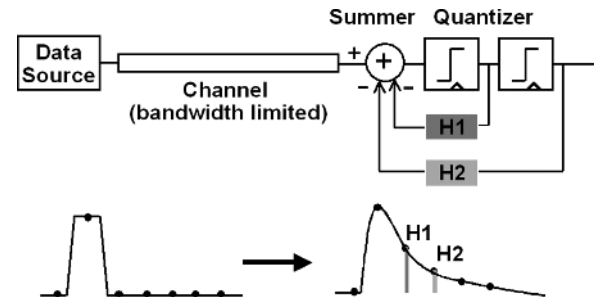

Fig. 8. Conceptual wireline communication system with a conventional 2-tap DFE receiver.

The pulse response curve of the channel to be equalized determines the specification of DFE and FFE. The normalized pulse response curve of the 20-mm silicon carrier channel is shown again on the left side of Fig. 9. It is observed that there is minimal precursor content, implying that no FFE is required. However, the postcursor content is significant. With conventional DFE architecture, a high tap-count DFE with substantial power and area overhead would be needed in the receiver.

A more efficient equalization solution can be found by studying the silicon carrier channel characteristics. Specifically, as shown in Fig. 9, a significant portion of the pulse response tail can be mathematically represented by an exponentially decaying function, enabling the overall channel pulse response curve to be well modeled by the combination of one discrete tap and one continuous-time filter. Based on this observation, we propose a DFE with one discrete tap and one adjustable continuous-time IIR filter-driven tap [8]–[10]. In this architecture, the pulse response of the continuous-time IIR filter models the behavior of the exponentially decaying tail in the channel response. The first discrete tap is kept because it does not fall
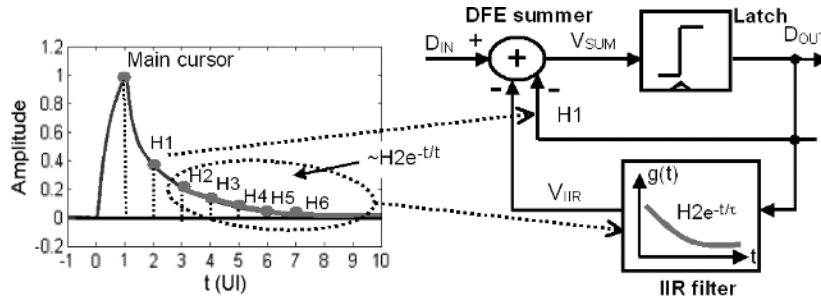
Fig. 9. Modified DFE receiver with continuous-time IIR filter feedback tap for silicon carrier links.
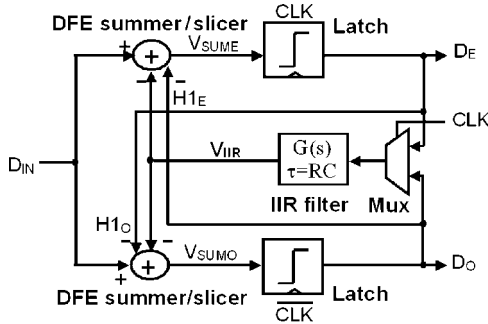

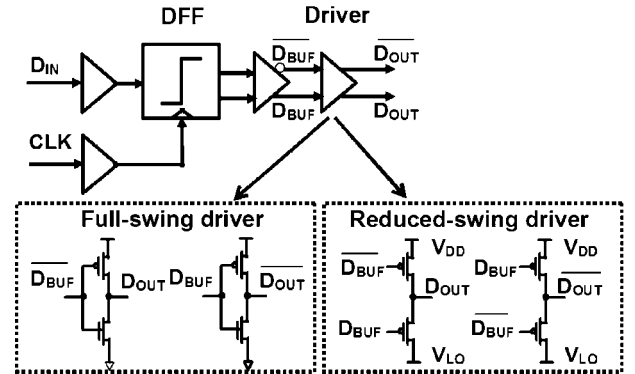Fig. 10. Half-rate DFE-IIR receiver.

Fig. 11. Transmitter with full-swing and reduced-swing drivers.

Fig. 12. Receiver termination network. (a) Reduced-swing input. (b) Full-swing input.

in the exponentially decaying tail. This approach enables one single additional feedback tap in the DFE to compensate the effects of multiple postcursors while consuming far less power and area than realizing a conventional DFE with many taps.

To implement this concept at the target data rate of 10 Gb/s, a half rate DFE implementation is more practical than a full-rate one [5]–[7]. Fig. 10 shows the half-rate DFE-IIR receiver architecture [8]. Full-rate input data are sampled by half-rate clock and then half-rate data are equalized in even and odd data paths. Since the channel acts on the full-rate data, a full-rate IIR filter is needed; this requirement implies that both even and odd data be fed to the IIR filter.

The implementation of this proposed DFE-IIR receiver features three key circuit elements. The first of these is a DFE summer with integrated signal slicing function; this summer supports high speed operation at low power. The second key circuit element is a full-rate continuous-time adjustable IIR filter with multiplexer, where the chosen full-rate approach greatly simplifies filter design as compared to alternatives based on half-rate structures. The third key circuit element is a double-regenerating latch enabling high-speed operation. The details of these three circuit blocks will be discussed in the following section.

IV. COMPACT I/O CIRCUITS

A. Transmitter

Fig. 11 depicts the transmitters adopted in the compact I/O for silicon carrier links. Based upon the discussion in Section III, the output impedance of the transmitters is designed to be less than 50 Ω to improve the driving strength. Full-rate data are retimed by a full-rate clock and sent to the channel by CMOS

drivers. Based on silicon carrier channel characteristics, two driver options were designed. For long channels, a full-swing driver with rail-to-rail output (0 to V_{DD}) is used to maximize signal strength. For short channels, a reduced-swing driver with output ranging from V_{LO} to V_{DD} is used to save output switching power [11]. In this design, V_{LO} is set to be $0.5V_{DD}$, although its value can be increased closer to V_{DD} for more aggressive power savings.

B. Receiver: Termination Network

As discussed in Section III, a high-impedance termination in the receiver is desired to save termination power dissipation. Fig. 12 shows the adopted termination networks in the compact I/O. To accommodate the two types of transmitters (Fig. 11), two types of receiver termination network structures were designed as drawn in Fig. 12. In the receiver termination network for the reduced-swing transmitter [Fig. 12(a)], two resistors, $R1$ and $R2$, set an appropriate common-mode voltage for the input stage of the receiver (DFE summer), which is set to be $0.75V_{DD}$ to match the transmitter output common-mode
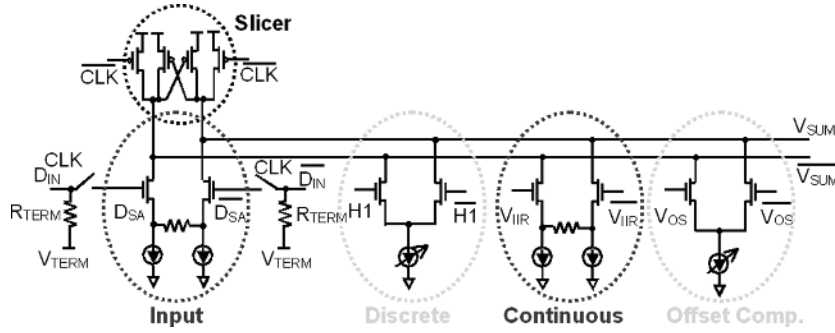
Fig. 13. Modified current-integrating DFE summer with slicer.

voltage. The peak-to-peak single-ended amplitude of the reduced-swing transmitter output is $0.5V_{\mathrm{DD}}$. To keep the receiver input stage in its linear operation range, the receiver input signal amplitude is reduced by adding a resistor, $R3$. In the resulting design, with $R1 = 1\,\mathrm{k\Omega}$, $R2 = 3\,\mathrm{k\Omega}$, $R3 = 250\,\Omega$, the input impedance is $1\,\mathrm{k\Omega}$ and the amplitude of the input signal is attenuated by 2.5 dB. Turning to the receiver termination network for the full-swing transmitter [Fig. 12(b)], since the output common-mode voltage of the full-swing transmitter is $0.5V_{\mathrm{DD}}$, which is different from that for the receiver input, dc voltage level shifting is required. This is accomplished by adding a resistor $R4$. In the resulting design, $R1 = R4 = 750\,\Omega$, $R2 = 1.5\,\mathrm{k\Omega}$, $R3 = 250\,\Omega$. With these values, the effective receiver input impedance is $1\,\mathrm{k\Omega}$ and the amplitude of the input signal is attenuated by 8.5 dB.

C. Receiver: Summer With Slicer

As part of the development of the overall compact I/O, a new DFE summer circuit was designed. This circuit, shown in Fig. 13, is modified from a previously developed current-integrating DFE summer [7], [12]. In the earlier current-integrating DFE summer, the resettable integrator mitigates the settling time constraint typical in resistively loaded summers to realize high speed and low power operation. The new feature added in this design is that the summer output current is injected into a current-comparator pMOS load; this feature is realized by the addition of two cross-coupled pMOS transistors to the summing node reset subcircuit of the previous design. This current-comparator pMOS load helps amplify the differential input signals, essentially acting as a slicer to accelerate signal regeneration. As in previous current-integrating DFE summers [7], [12], both H1 tap and offset compensation signals are fed to differential current switches. The H1 tap weight is determined by the effective charge injected into the DFE summer output nodes during the current integration operation. In practice, the tap weights are adjusted by the corresponding tail currents. In this design, H1 tap and offset compensation tap are adjusted by 6-bit and 5-bit current digital-to-analog converters (DACs), respectively. In contrast, the continuous-time IIR filter output is fed to a linear transconductor whose transconductance matches that of the summer input differential pair. The IIR filter feedback signal amplitude and time constant are adjusted to compensate the exponentially decaying channel response tail of the target channel.
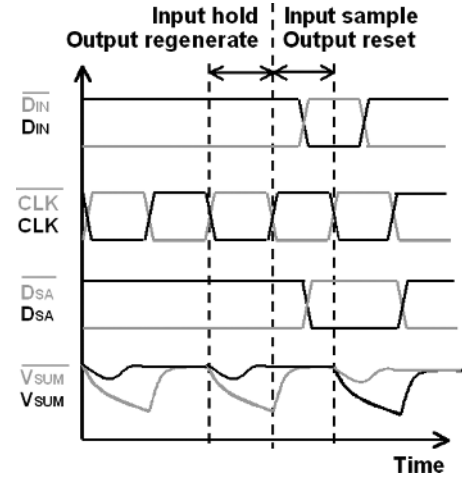

Fig. 14. Timing diagram illustrating the operation of the DFE summer with slicer.

In this prototype-level design, all the tap weights in the DFE summer and the IIR filter were manually adjusted and no adaptation control loop was implemented. Note, however, that adaptation approaches for DFE-IIR receivers have been reported in [10], and further that on-chip incorporation of this or of alternate adaptation approaches would be feasible for future DFE-IIR designs.

The operation of the DFE summer with slicer is shown in Fig. 14. When the clock signal CLK is high, the input switch is on and the input data are sampled. At the same time, the DFE summer outputs are reset to V_{DD} by the pull-up pMOS transistors. When CLK turns low, the input switch is turned off and the sampled input data are stored on the parasitic capacitance of the input transistors. The reset pMOS transistors are turned off and based on the differential input voltages to the differential pairs in the DFE summer, differential currents are drawn from the two output nodes, effectively summing all the taps. Once the common-mode voltage of the output signals falls to activate the cross-coupled pMOS load, the cross-coupled pMOS load starts to amplify the two output signals, regenerating the output data.

D. Receiver: IIR Filter

Fig. 15 depicts the adjustable continuous-time IIR filter. As a full-rate IIR filter is needed, both the even and odd data are fed to the IIR filter through multiplexer controlled by the clock signal CLK. When CLK is high, the odd data are fed into the IIR filter
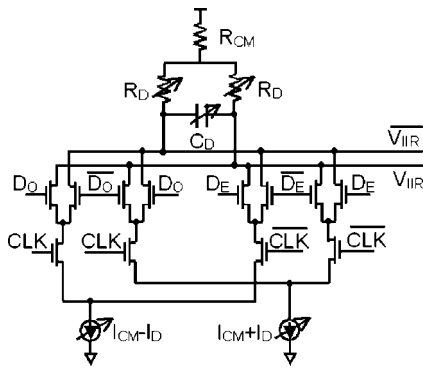
Fig. 15. Continuous-time IIR filter with adjustable time constant and signal amplitude.

Fig. 16. Double-regenerating latch.

Fig. 17. Timing diagram illustrating the operation of the double-regenerating latch.

and when CLK is low, the even data are fed into the IIR filter. To match the IIR filter response function with the exponentially decaying channel response tail (Fig. 9), both the output amplitude and the filter time constant of the IIR filter need to be adjusted. This tunability is implemented using adjustable resistor R_D, adjustable capacitor C_D and adjustable current I_D. More specifically, the filter's time constant is adjusted using R_D and C_D. Filter differential output amplitude is adjusted using R_D and I_D. When the differential output amplitude is changed, the filter output common-mode voltage remains constant. In this design, R_D and C_D are implemented by on-chip components and have 1-bit and 5-bit resolution respectively. The 3 most significant bits (MSBs) and 2 least significant bits (LSBs) of C_D are implemented with thick-oxide decoupling capacitors and higher-resolution vertical metal-insulator-metal capacitors, respectively. With the parasitic effects, the time constant $R_D C_D$ covers a tuning range of 0.5 UI to 10 UI. The adjustable current I_D is achieved with off-chip current supply, but of course could be implemented with on-chip DAC in a future design. The output common-mode voltage is maintained at the same level as the input stage of the DFE summer, namely $0.75 V_{DD}$.

E. Receiver: Double-Regenerating Latch

The output from the DFE summer with slicer is further regenerated by a high-speed double-regenerating latch shown in Fig. 16. This latch has two cascaded differential regenerating stages that achieve higher speed and sensitivity than a static CMOS latch. Fig. 17 depicts its operation. When the clock signal CLK is high, the output signals of the first stage are precharged to V_{DD}. In the second stage, the pMOS transistors are turned off, so this stage just retains its output at levels indicative of the previous bit decision. When CLK goes low, the first stage is turned on and begins to regenerate the input signal by the cross-coupled pMOS load. At the same time, the output common-mode voltage of the first stage falls, which then turns on the input transistors of the second stage. When the output of the first stage is regenerated to a sufficiently high level, the logical state of second stage is switched. Since the second stage has cross-coupled nMOS transistors, its output is amplified with additional regeneration. This latch is particularly useful when receiving a weakly regenerating signal from the summer with slicer. In such a case, the first stage of the latch regenerates at
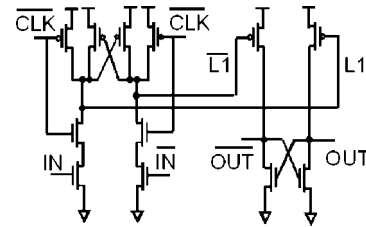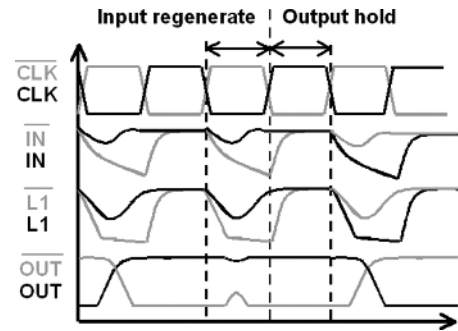
the same time as the summer with slicer so that the regeneration speed is improved. The transistors in the second stage are sized such that its rail-to-rail CMOS output has a crossing voltage level appropriate for driving nMOS differential current switches in the current-mode logic (CML) DFE feedback circuits, which eliminates the need for a CMOS-to-CML converter.

V. IMPLEMENTATIONS

The compact I/O with the DFE-IIR receiver described in Section IV was designed and fabricated in IBM's 65-nm bulk CMOS technology. To evaluate the performance, stand-alone transmitters and receivers were implemented, including full-swing and reduced-swing transmitters, 5-Gb/s and 10-Gb/s DFE-IIR receivers, as well as 5-Gb/s and 10-Gb/s conventional 2-tap DFE receivers. To improve testability, 50-Ω input termination was used for all receiver stand-alone sites instead of the high-impedance termination discussed in Section III. In addition to stand-alone receivers and transmitters, a large integrated transmitter-channel-receiver test site was also implemented. The goal of the integrated test site was to evaluate the compact I/O design in the context of channels emulating those of the target silicon carrier application. The die micrograph of this integrated test site is shown in Fig. 18. In this site, the channels connecting the transmitters and receivers were realized using on-chip back-end wire, and were designed to match the measured channel characteristics of silicon carrier transmission lines. Due to lack of thick metals in the CMOS back-end process, the emulated channels are differential microstrip structures with parameters designed to match the frequency rolloff characteristics of the silicon carrier channels. In total, the large test site included 12 individual transmitters (some full-swing and some reduced-swing) and 12 individual receivers (split between DFE-IIR and conventional 2-tap DFE instances, with some of each targeting 5 Gb/s and some 10 Gb/s operation),
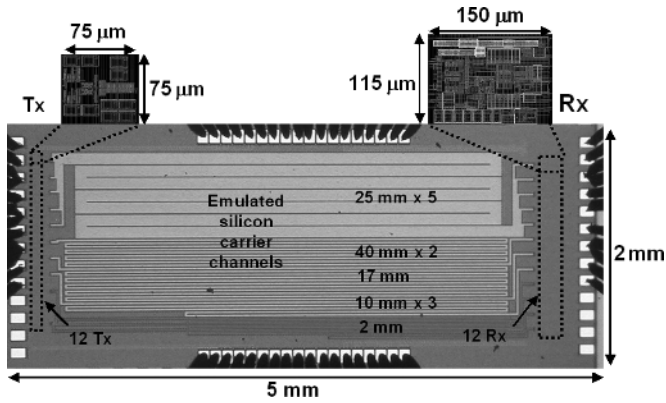
Fig. 18. Die micrograph of the integrated transmitter-channel-receiver test site.

connected by a variety of channels with lengths ranging from 2 mm to 40 mm. In this site, the transmitters and receivers are terminated with low impedance and high impedance, respectively, as presented in Section IV. In this prototype, half-rate receiver clock signals are from off chip and no on-chip eye monitor was implemented. Five 25-mm channels are grouped together so that crosstalk effects can be measured and studied. The full-swing transmitter layout and the DFE-IIR receiver layout are shown as insets in Fig. 18. The transmitter core, consisting of buffers, latches, and drivers, occupies an area of 75 μm $\times$ 75 μm. The DFE-IIR receiver core, consisting of termination network, DFE summers, IIR filter, latches, current DACs, and buffers, occupies an area of 150 μm $\times$ 115 μm. The inclusion of 2-tap DFE receiver blocks enables direct comparison between a conventional DFE approach and the DFE-IIR approach. Finally, an on-chip test structure was also added to allow characterization of the BEOL channel characteristics to enable assessment of the match between the characteristics of the on-chip channels and those of the target silicon carrier channels.

VI. EXPERIMENTAL RESULTS

A. Measurements of Stand-Alone Receivers

The measurement of the stand-alone DFE-IIR receiver shows that it has an input sensitivity of 64 mVppd at 10 Gb/s for a bit-error rate (BER) of better than 10^{-9}. It operates error-free up to 13 Gb/s with clean input data.

The equalization capabilities of the DFE-IIR receiver were first tested over 30″, 40″, and 50″ traces on a high-quality board made from Nelco4000-13SI material. Fig. 19(a) and (b) shows the frequency responses and normalized pulse responses of the channels, respectively. The channels have smooth frequency rolloff characteristics and exponentially decaying long response tails similar to those expected in silicon carrier channels (Fig. 4). Including ~2 dB of loss from the test setup, the channels losses are 15.5 dB, 19.6 dB and 23.2 dB at 5 GHz, respectively. Fig. 20 shows the measured receiver bathtub curves when equalizing 10-Gb/s PRBS7 data transmitted with an amplitude of 600 mVppd over these carrier-like channels [8]. For the 30″, 40″, and 50″ channels, at a BER = 10^{-9}, the DFE-IIR receiver produced 71%, 57%, and 45% horizontal eye openings, respectively. Note that these measurements do
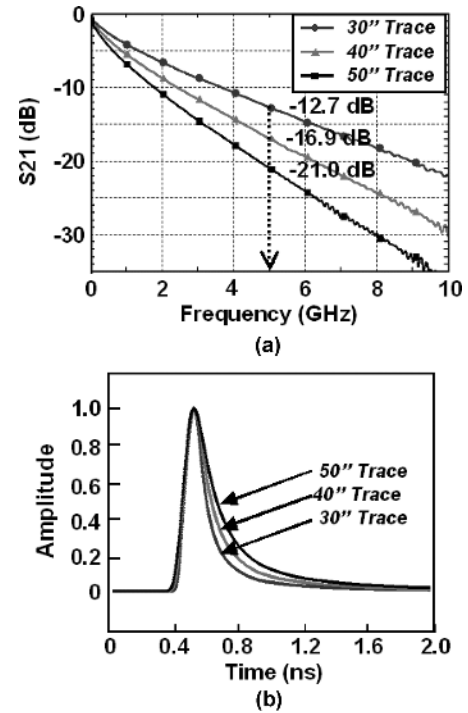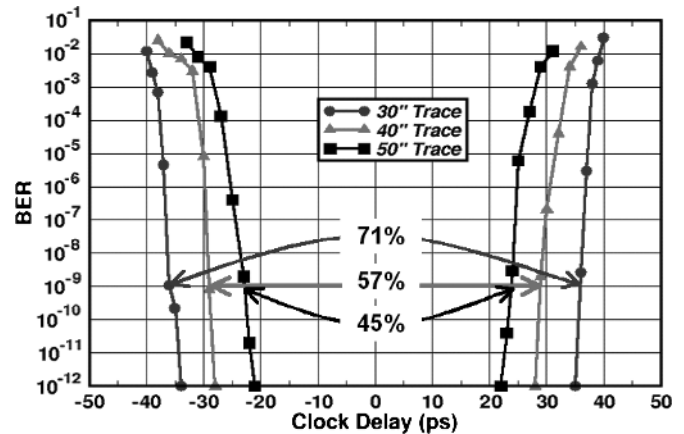

Fig. 19. Channel characteristics of 30″, 40″, and 50″ traces on a Nelco4000-13SI board. (a) Transfer functions. (b) Normalized pulse responses.

Fig. 20. Measured BER bathtub curve after equalizing 10-Gb/s PRBS7 data passed over the 30″, 40″, and 50″ traces characterized in Fig. 19. Horizontal eye openings at BER = 10^{-9} are reported.

not include recovered clock jitter from a CDR, which was not implemented in this work. In the center of the eye, data error rate much less than 10^{-12} is achieved. The power consumption of the DFE-IIR receiver was 6.8 mA from a 1-V supply (0.68 mW/Gb/s). The power breaks down as follows: 20% from the DFE summers with slicers, 25% from the latches, 50% from the clock buffers, and 5% from the IIR filter, current DACs and biasing circuitry.

In a second equalization experiment, 700 mVppd data was transmitted over a 16″ Tyco legacy backplane channel with two daughter cards. Fig. 21 illustrates the channel characteristics, including frequency response and normalized pulse response. Compared to silicon carrier channels, this 16″ Tyco legacy backplane channel has much less smooth frequency rolloff and
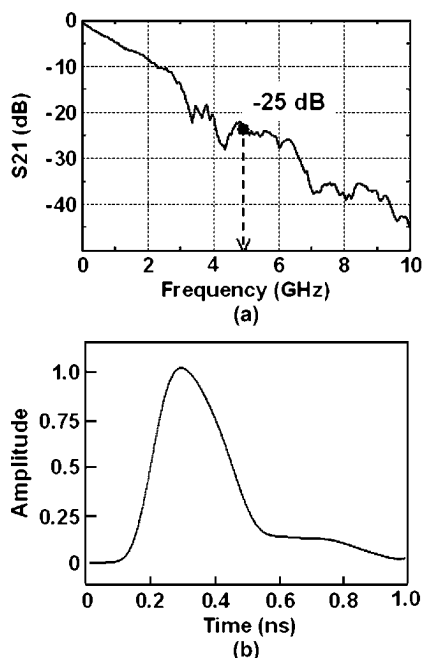
Fig. 21. Channel characteristics of a 16″ Tyco legacy backplane with two daughter cards. (a) Transfer function. (b) Normalized pulse response.

TABLE I
SUMMARY OF STAND-ALONE RECEIVER TEST RESULTS

Channel		Horizontal eye opening (BER<1e-9)	
		DFE-IIR	2-tap DFE
30" trace	PRBS7	71%	47%
	PRBS31	57%	24%
40" trace	PRBS7	57%	28%
	PRBS31	41%	Closed eye
50" trace	PRBS7	45%	Closed eye
16" Tyco	PRBS7	28%	Closed eye

channels were also performed with a conventional 2-tap DFE receiver. To allow clear comparisons to be made, the conventional 2-tap DFE was implemented using the same base components and power consumption level as the DFE-IIR receiver. Table I summarizes 10-Gb/s equalization results of the DFE-IIR receiver and 2-tap DFE receiver over various channels [8]. The addition of the IIR filter improved performance over all tested channels, highlighting the effectiveness of this equalization scheme.

The stand-alone full-swing and reduced-swing transmitters were demonstrated to be fully functional up to 10.8 Gb/s.

B. Measurements of Transmitter-Channel-Receiver Test Site

Before equalization experiments on the integrated transmitter-channel-receiver test site were performed, the on-chip channel measurement test structure, whose die micrograph is shown on the left side of Fig. 23, was measured. This test structure, which includes de-embedding structures and BEOL transmission lines, yielded data enabling extrapolation of S21 characteristics for the 25-mm and 40-mm BEOL channels realized in the transmitter-channel-receiver test site. These extrapolated values are compared with their measured silicon carrier counterparts in Fig. 23. It can be seen that the emulated and silicon carrier channels exhibit similar frequency rolloff behavior, though the carrier channels exhibit a few more dB insertion loss than the BEOL channels.

Fig. 24 shows the experimental setup. All measurements were performed on-wafer using high-speed probes. Full-rate data were generated using a pulse pattern generator and passed to the transmitters on the test site chip. The same full-rate clock used to time the pulse pattern generator was divided to produce the half-rate clock required for the receivers. Adjustable delay lines were used to compensate for skew in the data and clock paths to optimize the receiver performance. On the integrated test site chip, the transmitters send full-rate data through the on-chip channels, and the received data are equalized by the half-rate receivers. Both half-rate outputs from the receiver were monitored using error detectors. Serial data provided from a computer and controlled through a GUI, are used to program the DFE tap weights and offset compensation and

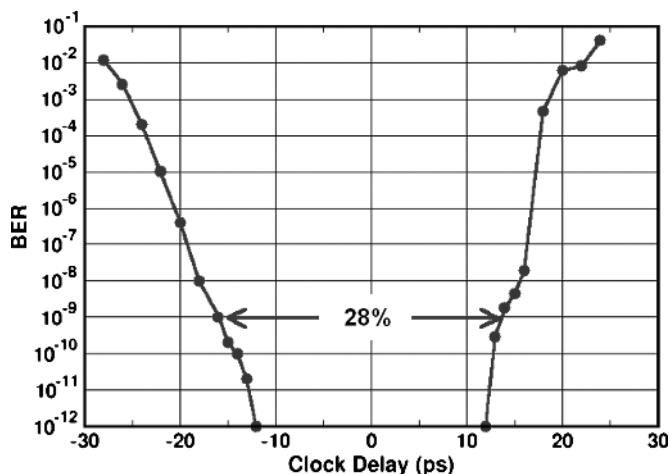
Fig. 22. Measured BER bathtub curve after equalizing 10-Gb/s PRBS7 data passed over the 16″ Tyco backplane characterized in Fig. 21.

pulse response characteristics, mainly due to significant signal reflections from via stubs in the channel. Again including ~2 dB of loss from the test setup, this channel shows ~27 dB of loss at 5 GHz. The measured bathtub curve of Fig. 22 shows that the DFE-IIR receiver recovers 10-Gb/s PRBS7 data transmitted over the Tyco backplane channel with a horizontal eye opening of 28% at a BER of 10^{-9}. In this experiment, the power consumption of the DFE-IIR receiver was approximately 7 mW from a 1-V supply. This demonstrates the potential of the DFE-IIR receiver for low-power backplane equalization in addition to silicon carrier applications.

Besides equalizing the aforementioned channels with the DFE-IIR receiver, equalization experiments over the same
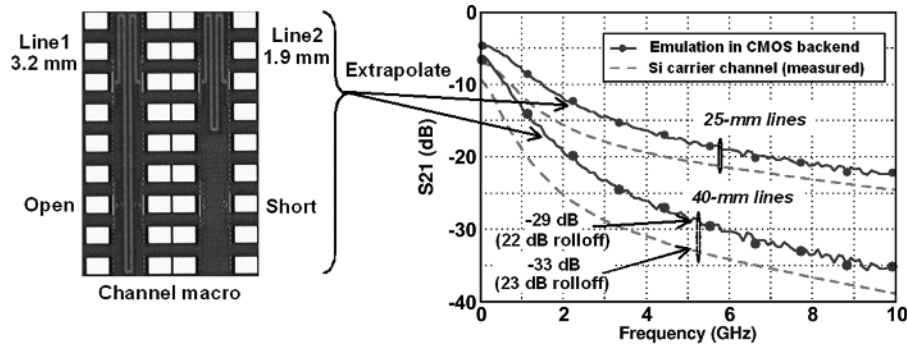
Fig. 23. Die micrograph of on-chip channel characterization macro and extrapolated channel data compared with silicon carrier data.
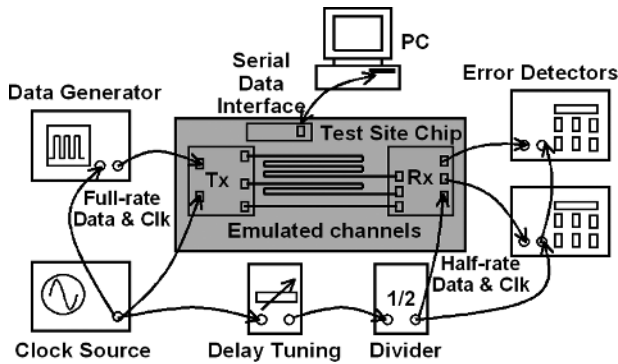

Fig. 24. Test setup for the integrated transmitter-channel-receiver test site.

to select transmitter and receiver from the twelve individual transmitter-channel-receiver sites.

The equalization results of a PRBS7 data pattern for all the individual sites are summarized in Fig. 25, which shows measured maximum data rates for 15% horizontal eye opening at a BER $= 10^{-9}$ and the I/O power efficiency (including transmitter and receiver) at maximum data rate. This integrated test site has various combinations of transmitter, channel, ESD and receiver. Successful equalization has been demonstrated across wide range of channels. Over a 40-mm link with 19-dB rolloff at 4.45 GHz, the 1-V transmitter-receiver pair operates at 8.9 Gb/s with a 15% eye opening while consuming 17 mW (5.4 mW from the receiver and 11.6 mW from the transmitter) yielding a 1.9-mW/Gb/s power efficiency. When equalizing data transmitted over a 25-mm link, the DFE-IIR achieves 9.5-Gb/s maximum data rate, higher than the 6.7-Gb/s maximum data rate achieved by a 2-tap DFE, validating DFE-IIR's superior performance as compared to conventional DFE. Additionally, a DFE-IIR receiver with a 550-fF input ESD protection device is shown to equalize a 17-mm channel with a 35% eye opening at 10 Gb/s. These results demonstrate the feasibility of the DFE-IIR for use in low-power compact silicon carrier link receivers.

C. Crosstalk Measurements

Crosstalk experiments were performed on the group of five 25-mm channels shown on the top of the test site micrograph in Fig. 18. The crosstalk experiment setup is depicted in Fig. 26. The channel under test is surrounded by two aggressor channels.

The aggressor channels can be individually turned on and off by the enable signals, $ETxi$ ($i = 1, 2, 3$). The two aggressor channels share clock and data inputs while the channel under test has independent clock and data. For crosstalk experiments, to remove signal correlations between aggressor channels and the channel under test, two aggressor channels and the channel under test are fed with a PRBS31 data pattern and a PRBS7 data pattern, respectively. The relative clock phase between the two clock signals to aggressor channels and channel under test can be tuned so that the crosstalk effect under various relative clock skews can be studied.

By measuring and comparing the maximum data rates of the channel under test with aggressor channels turned on and off, the effect of crosstalk can be evaluated and studied. However, when aggressor channels are turned on, effects like clock coupling, data coupling, and supply ripple can also influence the maximum data rate achievable for the channel under test, and thus must be evaluated to determine whether or not measurement results reflect true crosstalk-driven impairment. Control experiments were performed on the test site to explore this question, with the results showing only slight impact from these effects on the link performance. The wire-to-wire crosstalk from adjacent channels is therefore the primary source of the performance degradation observed in the crosstalk experiments.

Table II summarizes measured maximum data rates of four channels with various aggressor channels. With aggressor channels turned on, the measured channel maximum data rates decrease, with the amount of the decrease depending on the number of aggressor channels and the amplitudes of data transmitted on the channels. For instance, without any aggressor channels turned on, Channel I, which transmits full-swing data and is equalized by a DFE-IIR receiver, can operate up to 9 Gb/s. The maximum data rate drops to 8.8 Gb/s with one reduced-swing aggressor channel turned on and further drops to 8.6 Gb/s with two reduced-swing aggressor channels turned on. In another crosstalk experiment, Channel II, which transmits reduced-swing data and is equalized by a DFE-IIR receiver, can operate up to 6.9 Gb/s without any aggressor channels turned on. With one full-swing aggressor channel turned on, the maximum data rate drops to 6 Gb/s and further drops to 4.8 Gb/s with two full-swing aggressor channels turned on, demonstrating more crosstalk effects than the first experiments. Note that in this test site, no extra shielding between adjacent channels was added so that the crosstalk effects measured here
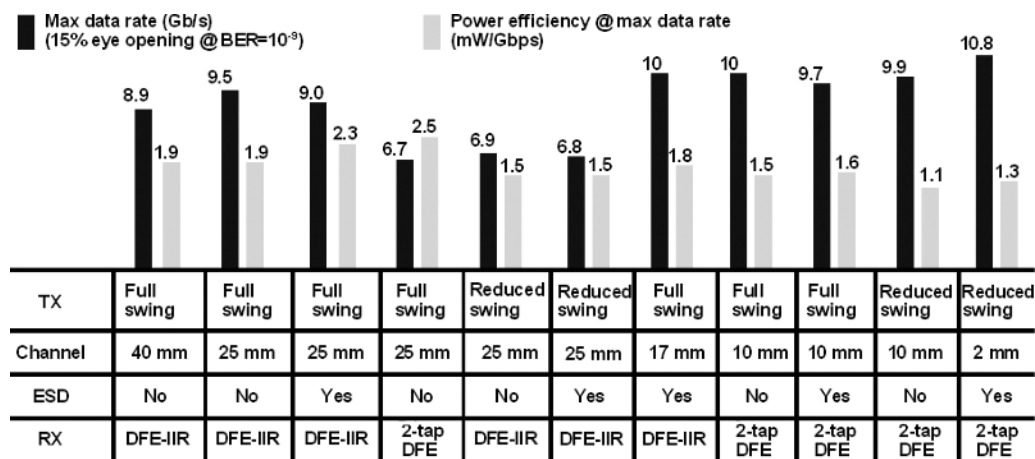
Fig. 25. Summary of integrated transmitter-channel-receiver test site results.
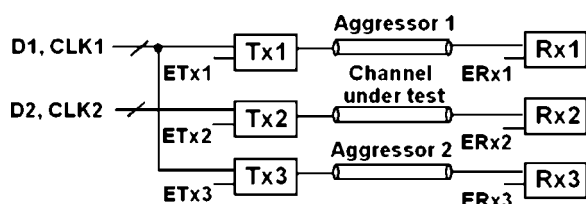

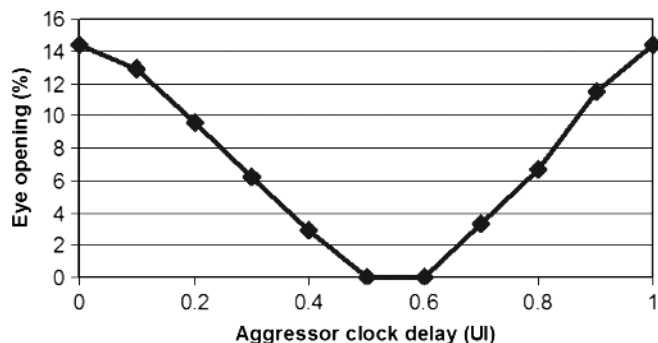
Fig. 26. Crosstalk experiment setup.

Fig. 27. Measured horizontal eye opening of the channel under test with various aggressor clock delays.

reflect worst-case carrier routing; clearly, such shielding in the carrier channels could be introduced to ameliorate crosstalk-induced degradation.

Crosstalk effects depend not only on aggressor channel signal amplitude but also on relative phase difference between aggressor and test channel data. Fig. 27 shows the horizontal eye openings of 4.8-Gb/s data transmitted over channel III with two aggressor channels turned on at various aggressor clock delays. With a 0.5-UI aggressor clock delay, the data eye is closed and the data eye is increased to 14.5% when the aggressor clock is delayed by 1 UI. Since the skew between channels is expected to be small in the target application, the data presented in Table II were measured with zero aggressor clock delay.

TABLE II
SUMMARY OF MEASURED CHANNEL MAXIMUM DATA RATES
WITH VARIOUS AGGRESSOR CHANNELS

| Channel	No aggressor (Gb/s)	With aggressor 1 (Gb/s)	With aggressor 2 (Gb/s)	With aggressors 1 & 2 (Gb/s)
DFE-IIR Full swing (I)	9	8.8 (reduced swing)	8.8 (reduced swing)	8.6 (both: reduced swing)
DFE-IIR Reduced swing (II)	6.9	6 (full swing)	6 (full swing)	4.8 (both: full swing)
2-tap DFE Full swing (III)	6.7	6.6 (reduced swing)	6.1 (full swing)	5.7 (reduced and full swing)
DFE-IIR Full swing (IV)	9.5	9.0 (full swing)	n/a	n/a

VII. CONCLUSION

A compact I/O with a power-efficient DFE-IIR receiver targeting high-speed serial I/O for dense silicon carrier links is reported. Based on silicon carrier channel characteristics, low-impedance transmitter and high-impedance receiver terminations and DFE-IIR based receiver equalization approach are adopted. In the receiver, a modified DFE with a continuous-time IIR feedback was developed for lossy channels with long pulse response tail requiring multi-tap DFE for ISI cancellation. Because it does not include a high tap-count DFE, this equalizer consumes less area and power than would be required with a conventional approach. Equalization capabilities of this DFE-IIR receiver were successfully demonstrated over a range of channels, including smoothly varying PCB channels with more than 20-dB losses and a legacy backplane channel with $\sim$27-dB loss. Furthermore, the compact I/O operates over various on-chip emulated silicon carrier channels ranging from 2 mm to 40 mm; for example, over a 40-mm on-chip emulated channel, a transmitter-receiver pair operates at 8.9 Gb/s with 1.9 mW/Gb/s power efficiency. These results validate the use of the DFE-IIR for serial I/O in advanced silicon packaging technologies.

REFERENCES

[1] J. Knickerbocker, C. Patel, P. Andry, C. Tsang, B. Paivikki, E. Sprogis, G. Hua, R. Horton, R. Polastre, S. Wright, and J. Cotte, "3-D silicon integration and silicon packaging technology using silicon through-vias," *IEEE J. Solid-State Circuits*, vol. 41, no. 8, pp. 1718–1725, Aug. 2006.

[2] T. Beukema, M. Sorna, K. Selander, S. Zier, B. L. Ji, P. Murfet, J. Mason, W. Rhee, H. Ainspan, B. Parker, and M. Beakes, "A 6.4-Gb/s CMOS SerDes core with feed-forward and decision-feedback equalization," *IEEE J. Solid-State Circuits*, vol. 40, no. 12, pp. 2633–2645, Dec. 2005.

[3] R. Payne, P. Landman, B. Bhakta, S. Ramaswamy, S. Wu, J. D. Powers, M. U. Erdogan, A. L. Yee, R. Gu, L. Wu, Y. Xie, B. Parthasarathy, K. Brouse, W. Mohammed, K. Heragu, V. Gupta, L. Dyson, and W. Lee, "A 6.25-Gb/s binary transceiver in 0.13-μm CMOS for serial data transmission across high loss legacy backplane channels," *IEEE J. Solid-State Circuits*, vol. 40, no. 12, pp. 2646–2657, Dec. 2005.

[4] K. Krishna, D. A. Yokoyama-Martin, A. Caffee, C. Jones, M. Loikkanen, J. Parker, R. Segelken, J. L. Sonntag, J. Stonick, S. Titus, D. Weinlader, and S. Wolfer, "A multigigabit backplane transceiver core in 0.13-μm CMOS with a power-efficient equalization architecture," *IEEE J. Solid-State Circuits*, vol. 40, no. 12, pp. 2658–2666, Dec. 2005.

[5] J. Bulzacchelli, M. Meghelli, S. Rylov, W. Rhee, A. Rylyakov, H. Ainspan, B. Parker, M. Beakes, A. Chung, T. Beukema, P. Pepeljugoski, L. Shan, Y. Kwark, S. Gowda, and D. Friedman, "A 10-Gb/s 5-tap DFE/4-tap FFE transceiver in 90-nm CMOS technology," *IEEE J. Solid-State Circuits*, vol. 41, no. 12, pp. 2885–2900, Dec. 2006.

[6] A. Rylyakov, "An 11 Gb/s 2.4 mW half-rate sampling 2-tap DFE receiver in 65 nm CMOS," in *Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 2007, pp. 272–273.

[7] T. Dickson, J. Bulzacchelli, and D. Friedman, "A 12-Gb/s 11-mW half-rate sampled 5-tap decision feedback equalizer with current-integrating summers in 45-nm SOI CMOS technology," *IEEE J. Solid-State Circuits*, vol. 44, no. 4, pp. 1298–1305, Apr. 2009.

[8] Y. Liu, B. Kim, T. Dickson, J. Bulzacchelli, and D. Friedman, "A 10 Gb/s compact low-power serial I/O with DFE-IIR equalization in 65 nm CMOS," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2009, pp. 182–183.

[9] E. Mensink, D. Schinkel, E. Klumperink, E. van Tuijl, and B. Nauta, "A 0.28 pJ/b 2 Gb/s/ch transceiver in 90 nm CMOS for 10 mm on-chip interconnects," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2007, pp. 414–415.

[10] P. Crespo and M. Honig, "Pole-zero decision feedback equalization with a rapidly converging adaptive IIR algorithm," *IEEE J. Sel. Areas Commun.*, vol. 9, pp. 817–829, Aug. 1991.

[11] K. Wong, H. Hatamkhani, M. Mansuri, and C. Yang, "A 27-mW 3.6 Gb/s I/O transceiver," *IEEE J. Solid-State Circuits*, vol. 39, no. 4, pp. 602–612, Apr. 2004.

[12] M. Park, J. Bulzacchelli, M. Beakes, and D. Friedman, "A 7 Gb/s 9.3 mW 2-tap current-integrating DFE receiver," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2007, pp. 230–231.
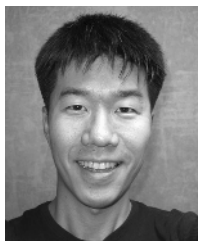
Byungsub Kim (S'06) was born in Busan, Korea. He received the B.S. degree in electronic and electrical engineering from Pohang University of Science and Technology, and the M.S. degree in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, where he is currently pursuing the Ph.D. degree.

In the summers of 2006 and 2007, he was an intern at the IBM T. J. Watson Research Center, Yorktown Heights, NY, where he developed DFE-IIR architectures and circuits for a compact I/O.

Mr. Kim was a co-recipient of the Beatrice Winner Award for Editorial Excellence at the 2009 IEEE International Solid-State Circuits Conference.

Yong Liu (S'03–M'07) received the B.S. and M.S. degrees in electrical engineering from Tsinghua University, Beijing, China, in 2000 and 2003, respectively, and the Ph.D. degree from Harvard University, Cambridge, MA, in 2007. His Ph.D. work examined applications of CMOS ICs in medicine and biotechnology by directly interfacing CMOS ICs with biological systems. Specifically, he developed CMOS ICs in conjunction with microfluidic systems to magnetically manipulate individual biological cells for cell sorting applications and a CMOS RF biosensor utilizing NMR for medical diagnostics in a portable platform.

He is currently with the IBM T. J. Watson Research Center, Yorktown Heights, NY. In 2001, he was with Tsinghua Tongfang Microelectronics Co., working on the second-generation Chinese RF ID card. In 2005 and 2006, he was with the Mixed-Signal Communications IC Group at IBM T. J. Watson Research Center as a summer intern. In November 2007, he joined the IBM T. J. Watson Research Center, where he is involved with the development of high-speed and low-power data links for conventional wireline communication systems as well as emerging silicon packaging technologies (silicon carrier links and 3-D silicon integration). He has 16 publications, three U.S. patents, and three patents filed.

Dr. Liu was the recipient of the Seagate Scholarship, Motorola Scholarship, Second Prize in the China National Graduate EDA Competition, and Analog Devices Outstanding Student Designer Award. He was a co-recipient of the Beatrice Editorial Excellence Award in the 2009 IEEE International Solid-State Circuits Conference.

Timothy O. Dickson (S'01–M'06) received dual B.Sc. degrees in electrical and computer engineering with highest honors from the University of Florida in 1999. He completed the M.Eng. degree at the University of Florida in 2002 and the Ph.D. degree at the University of Toronto in 2006, both in electrical engineering. His Ph.D. work was in the area of serial transceivers operating up to 80 Gb/s in SiGe BiCMOS technologies, focusing on the development of low-noise and low-power design methodologies.

In September 2006, he joined the IBM T. J. Watson Research Center in Yorktown Heights, NY, where he is involved with the development of low-power serial transceivers. He is also an Adjunct Assistant Professor at Columbia University where he teaches graduate level courses in analog and RF circuit design. His research interests include circuits and semiconductor technologies for high-speed wireline and mm-wave applications.

Dr. Dickson was named an undergraduate University Scholar by the University of Florida in 1999, and was an Edward S. Rogers, Sr Scholar during his doctoral work at the University of Toronto. He received the Best Student Paper Award at the 2004 VLSI Circuits Symposium in Honolulu, HI, and was a co-recipient of the Beatrice Winner Award for Editorial Excellence at the 2009 IEEE International Solid-State Circuits Conference. He is currently a member of the Technical Programming Committee for the IEEE Compound Semiconductor Integrated Circuit Symposium.

John F. Bulzacchelli (S'92–M'02) was born in New York, NY, in 1966. He received the S.B., S.M., and Ph.D. degrees in electrical engineering from the Massachusetts Institute of Technology (MIT), Cambridge, in 1990, 1990, and 2003, respectively.

From 1988 to 1990 he was a co-op student at Analog Devices, Wilmington, MA, where he invented a new type of delay-and-phase-locked loop for high-speed clock recovery. From 1992 to 2002 he conducted his doctoral research at the IBM T. J. Watson Research Center, Yorktown Heights, NY, in a joint study program between IBM and MIT. In his doctoral work, he designed and demonstrated a superconducting bandpass delta-sigma modulator for direct A/D conversion of multi-GHz RF signals. In 2003 he became a Research Staff Member at this same IBM location, where his primary job is the design of mixed-signal CMOS circuits for high-speed data communications. He also maintains strong interest in the design of circuits in more exploratory technologies. He holds two U.S. patents.

Dr. Bulzacchelli received the Jack Kilby Award for Outstanding Student Paper at the 2002 IEEE International Solid-State Circuits Conference (ISSCC) and was a co-recipient of the Beatrice Winner Award for Editorial Excellence at the 2009 ISSCC.

Daniel J. Friedman (S'91–M'92) received the Ph.D. degree in engineering science from Harvard University, Cambridge, MA, in 1992.

After completing consulting work at MIT Lincoln Labs and postdoctoral work at Harvard in image sensor design, he joined the IBM Thomas J. Watson Research Center, Yorktown Heights, NY, in 1994. His initial work at IBM was the design of analog circuits and air interface protocols for field-powered RFID tags. In 1999, he turned his focus to analog circuit design for high-speed serializer/deserializer macros. Since June 2000, he has managed a team of mixed-signal circuit designers, with recent focus on serial data communication and clock synthesis applications. In addition to circuits papers on serial links and PLLs, he has published articles on imagers and RFID, and he holds more than 20 patents. His current research interests include high-speed I/O design, PLL design, and circuit/system approaches for variability compensation.