

A 24 Gb/s Software Programmable Analog Multi-Tone Transmitter

Amir Amirkhany, *Member, IEEE*, Aliazam Abbasfar, *Senior Member, IEEE*, Jafar Savoj, *Senior Member, IEEE*, Metha Jeeradit, *Student Member, IEEE*, Bruno Garlepp, *Member, IEEE*, Ravi T. Kollipara, *Senior Member, IEEE*, Vladimir Stojanovic, *Member, IEEE*, and Mark Horowitz, *Fellow, IEEE*

Abstract—A 24 Gb/s transmitter employs a digital linear equalizer and a 12 GS/s 8-bit digital-to-analog converter (DAC). Implemented in a 90 nm CMOS technology, the transmitter can be programmed to support a variety of communication modes including 4-channel and 2-channel analog multi-tone (AMT), as well as various baseband (BB) modes ranging from 2 to 256 PAM. Selection of the transmission mode is enabled through software programming of the appropriate tap coefficients into the equalizer. The transmitter dissipates 510 mW of power and is fabricated over an area of 0.8 mm². Experimental results confirm clear eye diagrams at 28 Gb/s.

Index Terms—Decision feedback equalizers, equalizers, frequency division multiplexing, high-speed circuits, MIMO systems.

I. INTRODUCTION

MODERN high-speed electrical links enable transmission of data at multi-gigabits per second between integrated circuits. Examples of such links include the network routers in the backbone of the internet, the interface between a controller and multiple memory modules inside a personal computer (PC), or the interface between the central and the graphics processing units (CPU and GPU) inside a PC. In all these applications the communication media are copper traces on printed circuit boards (PCB). Signal quality of such channels is distorted by various loss mechanisms as well as reflections from the impedance discontinuities in the signal path. The latter are caused by the vias, stubs, and connectors necessary to route the signal traces across PCB layers and multiple boards. Signal reflections caused by these discontinuities create resonance frequencies and result in notches in the frequency response of the link channels. Fig. 1(a) shows the frequency response of three different channels used for backplane [see Fig. 1(b)], multi-drop [see Fig. 1(c)], and chip-to-chip [see Fig. 1(d)] communications.

Manuscript received August 24, 2007; revised October 29, 2007.

A. Amirkhany is with Rambus Inc., Los Altos, CA 94022 USA and also with the Department of Electrical Engineering, Stanford University, Stanford, CA 94305 USA (e-mail: amirkhany@stanford.edu).

A. Abbasfar, M. Jeeradit, and R. T. Kollipara are with Rambus Inc., Los Altos, CA 94022 USA.

J. Savoj was with Rambus Inc, Los Altos, CA 94022 USA. He is currently with Qualcomm, Campbell, CA 95008 USA.

B. Garlepp was with Rambus Inc., Los Altos, CA. He is currently with SiTime Corp., Sunnyvale, CA 94085 USA.

V. Stojanovic is with the Department of Electrical Engineering, Massachusetts Institute of Technology, Cambridge MA 02139 USA.

M. Horowitz is with the Department of Electrical Engineering, Stanford University, Stanford, CA 94305 USA, and also with Rambus Inc., Los Altos, CA 94022 USA.

Digital Object Identifier 10.1109/JSSC.2008.917520

Due to the tight constraints on the power consumption of high-speed links and their extremely high operating rates, the majority of the state-of-the-art links employ 2 PAM BB signaling with relatively simple signal processing to compensate for the dispersive nature of the channel [1], [2]. A state-of-the-art BB link generally includes a discrete linear equalizer at the transmitter to cancel precursor inter-symbol interference (ISI), a linear peaking amplifier at the receiver front-end to increase sensitivity and compensate for the magnitude distortion of the channel, and a decision feedback equalizer (DFE) at the receiver to cancel post-cursor ISI. Recently, links that additionally utilize a low-resolution analog-to-digital converter (ADC) and a digital feed-forward equalizer (FFE) at the receiver have also been demonstrated [3]. Commercial links today generally operate at data transfer rates of 6 to 12.5 Gb/s and power-efficiencies of less than 30 mW/Gb/s, while prototype systems have achieved data rates of 20 Gb/s [4] or power efficiencies of 2.2 mW/Gb/s [2].

A study of the limits of signaling in backplane links [5], however, reveals that there is a large gap between the fundamental limits of signaling, commonly known as the Shannon capacity, and the limits achievable with BB signaling techniques. In particular, inspection of the channel characteristics, shown in Fig. 1, reveals that, in a class of applications, notches in the frequency domain are part of the frequency response of the link channels. It is well known in communication theory that multi-tone (MT) signaling has the potential to achieve superior performance over such channels compared to BB signaling by better allocation of the transmit energy and avoiding wasting energy over the notch frequencies. MT signaling can therefore be employed to potentially reduce the gap between current link performances and the Shannon capacity in these applications. However, as it is usually the case in link environments, the real challenge lies in the implementation. Implementing conventional MT techniques, like discrete multi-tone (DMT), which is widely used in digital subscriber line (DSL) and wireless systems, is not energy-efficient at the multi-gigabit per second operating rates necessary for high-speed links [6]. This is because such MT techniques require high-speed moderate-resolution ADCs at the receiver and relatively sophisticated digital signal processing, which significantly add to the total system power consumption.

We recently proposed an MT technique, called analog multi-tone (AMT), which is customized to the link characteristics and has the potential to achieve superior performance compared to BB systems while being energy efficient in a link environment [7]. Similar to a BB link, an AMT link employs linear transmit equalization and receive DFE to compensate for the

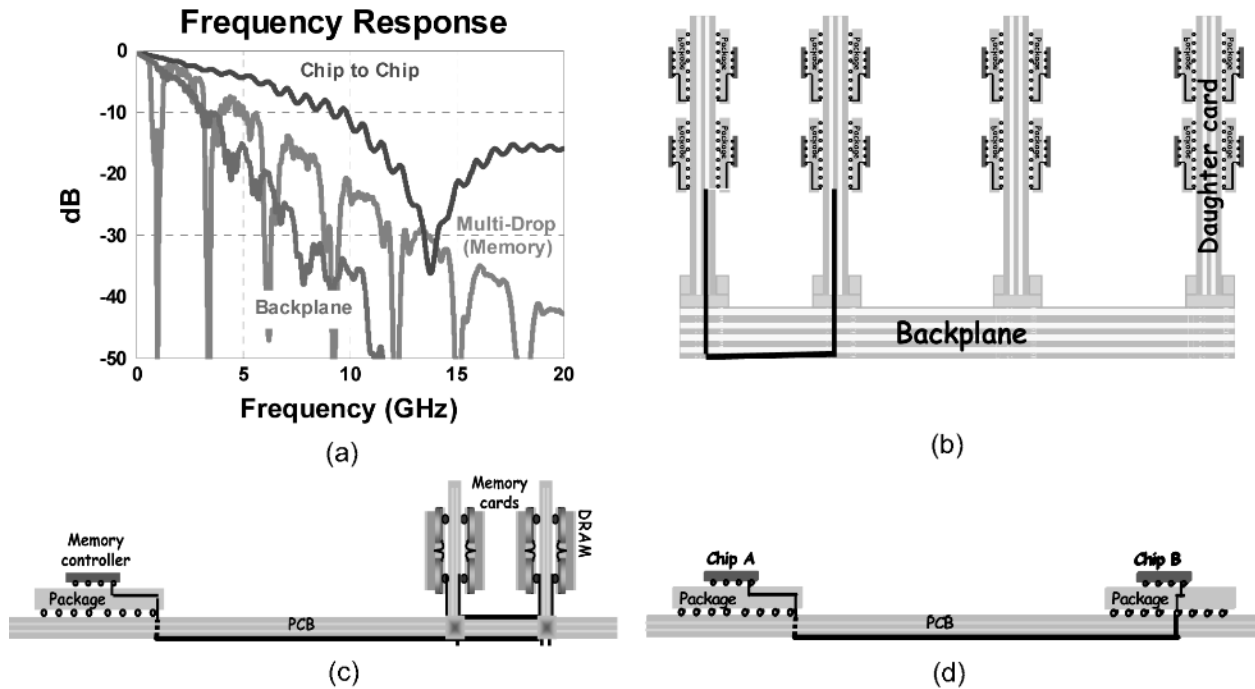


Fig. 1. (a) Channel characteristics in frequency domain. (b) An electrical link in a network router. (c) A multi-drop memory interface. (d) A CPU-GPU link.

effects of the band-limited channel. This paper describes the design of a 24 Gb/s software programmable transmitter, implemented in a 90 nm CMOS technology, which supports 4-channel and 2-channel AMT as well as a variety of BB transmission modes including 2 and 4 PAM. With 16 effective FFE taps, 10-bit tap coefficients, and no constraints on the dynamic range of the taps, the transmitter has sufficient equalization capabilities to enable the study of both AMT and BB transmission algorithms over a wide range of environments and applications. In other words, the transmitter is an extremely flexible test instrument for high-speed link applications. The transmitter architecture was originally described in [8].

Section II reviews the architecture of an AMT system and its operation from both time domain and frequency domain perspectives. With this understanding of the overall system, Section III describes the architecture of the transmitter, and measurement results follow in Section IV. The architecture of the equalizer also enables compensation of analog distortions caused by the mismatch in the paths of the on-chip time-interleaved DAC through cyclic time-variant equalization. Section V explains how this correction is done and also provides results for the multi-PAM BB operation modes supported by the transmitter.

II. ANALOG MULTI-TONE

In order to have an energy-efficient MT architecture at the extremely high operation rates necessary for high-speed links, a BB link can be extended into a bank of parallel links operating at different carrier frequencies. Fig. 2(a) shows a conceptual MT system based on this idea. Each sub-channel (the path

from an input at the transmitter to the corresponding output at the receiver) in this figure can potentially have a bandwidth of a few gigahertz and only a small number of sub-channels may exist. This architecture is, however, difficult to implement since fully integrated filters with sharp roll-off are not available. Therefore, energy from one sub-channel will inevitably spill over to the neighboring sub-channels, causing inter-channel interference (ICI). ICI is similar to ISI in nature, with the exception that it represents interference from symbols from another sub-channel. Therefore, as long as the transfer function from the input of the interfering sub-channel to the output of the target sub-channel does not change from sample to sample, i.e., as long as the system is linear time invariant (LTI) in the discrete domain, ICI can be cancelled in the same way as ISI; through equalization.

For the system to be considered LTI in the discrete domain, the following two constraints should be fulfilled:

- 1) all sub-channel symbol rates should be the same;
- 2) all carrier frequencies should be integer multiples of the sub-channel symbol rate.

These two constraints will cause all carrier frequencies to complete full cycles from one symbol to another. This means that from the input sequences' perspective the system does not change from one cycle to another, and is therefore time-invariant in the discrete domain. However, ISI and ICI cancellation in this multi-input-multi-output (MIMO) system requires MIMO linear transmit equalization and MIMO receive DFE [9]. Alternatively, it can be shown that as a consequence of the two constraints imposed on the system, for an N -sub-channel system, the MIMO transmit equalizer, the low-pass filters, and the mixers in the transmitter can be

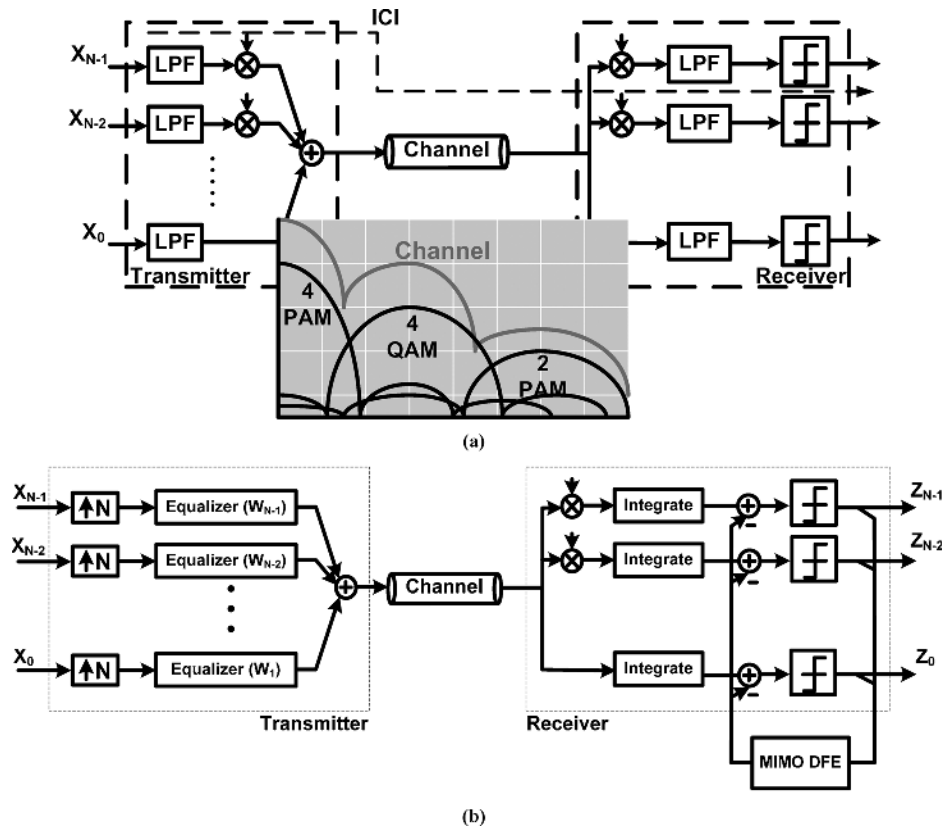


Fig. 2. (a) Conceptual multi-tone system with low-pass filters and mixers at the transmitter and receiver to create band-limited sub-channels. X_0, \dots, X_N are input sequences. (b) AMT architecture with per-sub-channel linear N -times over-sampled equalizers at the transmitter, and mixer and integrate-and-dump at the receiver.

replaced by N -times¹ over-sampled equalizers per sub-channel [7]. An appropriate choice for the low-pass filter in the receiver that leads to superior performance and can be implemented reliably on chip is an integrate-and-dump circuit that integrates over one sub-channel symbol period [10]. Fig. 2(b) shows the finalized AMT architecture. In the architecture of Fig. 2(b), each N -times over-sampled equalizer can shape the transmission bandwidth (from dc to the Nyquist frequency) of the entire system. Therefore, when all the sub-channel equalizers are optimized simultaneously, they work together to cancel both ISI and ICI at the same time.

From a time-domain perspective, the AMT system is a transmultiplexer operating based on the principle of perfect reconstruction [11]. As an example, the time-domain waveforms of a 2-channel AMT system are plotted in Fig. 3(a) to illustrate this point further. In this example, the transmit equalizers are 2-times over-sampled 2-tap filters, a zero-order hold filter performs discrete-to-analog conversion, the channel is assumed to be ideal, and there is no MIMO DFE in the receiver. Continuous time ISI and ICI patterns at the input to the samplers are also shown along with the sampled sequences Z_1 and Z_2 . It can be seen that even though the transmit equalizers are very short, and consequently, significant energy overlap exists between the

two sub-channels at the transmitter output, both the ISI and the ICI are forced to zero at the sampling points at the receiver. As a result, X_1 and X_2 are fully recovered at the sampler outputs without any interference. However, the ISI and the ICI are not necessarily zero at points other than the sampling point. This is very similar to the operation of an equalized BB system where a linear equalizer forces the ISI to zero (only) at the sampling points.

Ignoring the MIMO DFE, for large number of sub-channels, the AMT system shown in Fig. 2(b) is a very inefficient implementation of a DMT system in which the IFFT operation at the transmitter is performed using N -times over-sampled equalizers. However, for a small number of sub-channels, the AMT system is highly efficient and its performance is not limited by cyclic prefix overhead [12]—a limiting factor in small-block-size DMT systems. In fact, it can be shown that the equalization complexity (and consequently equalization power) of the AMT system is comparable to a BB system operating at the same overall throughput. To demonstrate this point further, Fig. 3(b) shows a 2-way parallelized 4-tap BB transmit equalizer and a 4-tap-per-channel 2-channel AMT equalizer. The BB and AMT equalizers look structurally identical except that the taps in the lower branch of the 2-way parallelized BB equalizer are constrained to be a shifted version of the upper branch. The taps in the two branches of the AMT equalizer, on the other hand, have more degrees of freedom and can take values indepen-

¹ N -times with respect to sub-channel symbol rate. Assuming all sub-channels are 2 PAM, each sub-channel equalizer is operating at the overall data rate of the link.




Fig. 3. (a) Signal waveforms in an example 2-channel AMT system. Two-tap equalizers per sub-channel at the transmitter, ideal channel, and no DFE at the receiver. Continuous-time ISI and ICI patterns are shown at the sampler inputs. (b) A 2-way parallelized 4-tap linear equalizer and a 2-channel 4-tap per channel AMT equalizer.

dently.² In other words, the AMT equalizer is a general form of a BB equalizer. The additional degrees of freedom in the AMT equalizer enable the AMT system to shape the transmit spectrum better than a BB system. As a result, an AMT system performs considerably better than a BB system over channels with notches in their frequency response or in the presence of frequency selective interference, where optimal shaping of the transmit spectrum is crucial. Over smooth channels, however, where both AMT and BB can achieve close to optimum transmit power allocation, for the same transmit peak voltage, the AMT system can put less signal energy on the line compared to a BB system, because a MT signal in general has a higher peak-to-average power ratio (PAPR) than a single-tone signal. Therefore, BB signaling may be optimal for signal transmission over these kinds of channels.

Fig. 3(b) also indicates that an AMT transmitter can be built to support BB signaling with very small overhead. Such a transmitter can achieve the best performance over a very wide range of channel characteristics using either AMT or BB transmission depending on which one is optimal. Section VII describes this type of transmitter designed to support 24 Gb/s in both AMT and BB modes.

III. TRANSMITTER ARCHITECTURE

The 24 Gb/s transmitter, fabricated in a 90 nm CMOS technology, includes an on-chip pattern generator (PG), a linear

²Same complexity (power) argument applies to the MIMO DFE in the receiver of the AMT system as well, and the argument is independent of whether FFE and DFE are implemented in digital domain or as pseudo-DAC current mode equalizers.

equalizer, and a DAC. The design of the equalizer in the digital domain, with 16 10-bit full-range taps, makes the transmitter a platform with sufficient flexibility to enable evaluation of different transmission algorithms in different environments. Fig. 4(a) shows the top level block diagram of the transmitter. The PG creates 24 Gb/s of pseudo-random (PN) binary data for the system in the form of four (2-way parallelized) 3 GS/s 4 PAM (or 2 PAM) sequences. The 16-tap linear equalizer is parallelized four ways with each parallel branch allowed to take independent tap values. This way the transmitter can be configured to support 4-channel and 2-channel AMT, as well as multi-PAM BB signal transmission ranging from 2 to 256 PAM. Each parallel branch of the equalizer is 2-way parallelized to operate with a 1.5 GHz clock and receives a 4 PAM (or 2 PAM) sequence as input. A thermometer encoder, included at the output stage of the equalizer, converts the three most significant bits of the output to unary code. An on-chip 8-bit 12 GS/s 2-way output multiplexed current-mode DAC converts the digital outputs to an analog signal and directly drives the 50- Ω chip output.

Fig. 4(b) shows the clock network of the transmitter. The system uses an external 12-GHz reference signal along with CML frequency dividers to minimize the impact of clock jitter and achieve optimum duty cycle for the synchronizing half-rate clock. The 6 GHz half-rate clock is further divided on-chip to create the 3 GHz clock required for the DAC, and the 1.5 GHz clock required for the digital equalizer and the PG. A phase interpolator and a phase detector are placed at the interface between the DAC and the equalizer to ensure the two circuits have independent clock distribution networks without potential setup




Fig. 4. (a) Transmitter block diagram. (b) Transmitter clock network. Black dots in the equalizer indicate the placement of the clock drivers. A phase interpolator (PI) shifts the phase of the input clock to the transmitter based on information from a phase detector (PD). The phase detector samples a 1.5 GHz clock that branches off from a leaf of the equalizer's clock grid with the 3 GHz input clock to the DAC.

and hold violations at their common interface. The phase detector samples a 1.5 GHz clock that branches off from the equalizer network with the 3 GHz in-phase clock and provides information on the phase alignment of the two clocks. The phase interpolator is programmed offline based on this information. The 1.5 GHz clock distribution inside the equalizer is in the form of an 8-mesh with the clock being routed from the sides toward the center, in the direction of the data flow [see Fig. 5(a)]. The PG block, implemented with standard library cells following a full ASIC flow, operates with a 375 MHz clock, with the exception of a small block of serializers at the output stage which operate with a 1.5 GHz clock. The 1.5 GHz clock for the PG also branches off from the equalizer network, and the 1.5 GHz clock distribution latency in the PG is included in the critical path of the data interface between the equalizer and the PG.

The equivalent functionality of the equalizer is a 16-tap FIR filter with 10-bit tap coefficients and 2-bit (4 PAM) inputs operating at 12 GS/s. Fig. 5(a) shows the datapath of one phase of the 4-phase equalizer. The relative placement of the blocks is chosen to minimize power dissipation in the long wires. In one phase, 16 10-bit by 2-bit multiplications are performed and the results are summed in one 3 GHz cycle. For multiplications, the 10-bit values of W and $3W$ (where W is an equalizer tap coefficient) are stored in flip-flops and a 4:1 multiplexer selects the correct multiplication output ($\pm W$ or $\pm 3W$) based on the 2-bit data input. The flip-flops storing the tap coefficients do not dis-

sipate any active power and only impose a small area overhead. Additions are performed with three stages of 4:2 compressor units and a final pseudo Kogge–Stone adder. In this design, the compressor architecture proposed in [14] was used, except that parts of the logic were duplicated to create true and complementary outputs in parallel to reduce the total number of logic stages in the critical path (see Fig. 6). Optimizing the compressors to meet a 3.0 GHz cycle time would lead to larger area and higher power consumption than a 2-way parallelized architecture to meet a 1.5 GHz cycle. Consequently, the design was 2-way parallelized to operate with a 1.5 GHz clock and each compression was placed in a separate pipeline stage. Overall, the equalizer has five stages of pipelining, with flip-flops placed before every compressor, before the final adder and before the thermometer encoder. A final 2:1 multiplexer that converts the 1.5 GHz odd and even outputs to a 3.0 GHz output is placed after the thermometer encoder and precedes the long output wires. The thermometer encoder was sized to settle in a 3.0 GHz cycle to reduce the transient power dissipated on the long output wires.

The equalizer is implemented in static logic and transistor sizing of the building components is fully custom. However, these building components are further characterized as standard cells to enable the automation of the design using a combination of several commercial ASIC design tools, including a synthesis and a place and route (P&R) tool, in addition to an in-house hierarchical MATLAB placement tool. The MATLAB




Fig. 5. (a) Datapath of one phase of the equalizer consisting of three stages of 4:2 compression, a pseudo Kogge–Stone adder and a thermometer encoder. Each block is 2-way parallelized and a 2:1 multiplexer (serializers) is included in the thermometer encoder. (b) Equalizer floorplan in the Matlab placement tool. Large rectangles on the sides are areas where low-speed flip-flops holding equalizer tap coefficients are placed by the P&R tool. Large rectangles in the middle column are areas where the adder and the thermometer encoder are placed by the P&R tool. The small isolated blocks in the middle row are latches, implementing the shift registers for the input data sequence. X - and Y -axes show actual sizes in microns.

tool interacts with the P&R tool to place high-speed components at exact desired locations. Fig. 5(b) shows the complete floorplan of the equalizer. Routing is fully handled by the P&R tool and verification is performed with commercial Static Timing Analysis tools. The overall flow has a precision close to a full custom design while utilizing the vast automation and verification capabilities of the commercial ASIC design tools [17]. The equalizer was designed to meet the 24 Gb/s specification in the SS/125°/0.95 V corner. Correct operation was observed in the lab at room temperature and nominal supply at 29 Gb/s before the clocking circuitry start to break.

In order to achieve superior linearity and higher operation speeds, data multiplexing and current switching are combined within the same circuit inside the DAC cell. The 6-GHz synchronizing clock signal utilizes CML swings to avoid device stress in presence of a 1.8-V termination voltage. Calibration is performed so that the multiplexing 6-GHz clock is centered with respect to the data bits. The calibration process consists of two

steps. In the first step, two binary phase detectors identify the zero crossings of the 6 GHz clock with respect to the two 6 Gb/s LSB signals. Consequently, a fixed phase offset is added to the data bits. The circuit utilizes inductive techniques to increase the clock driver and the output termination bandwidths. The design of the DAC is described in [13]. The circuit can achieve higher operation speeds if the clock dividers and the 6-GHz driver are redesigned to achieve an even higher bandwidth.

IV. MEASURED TRANSMITTER PERFORMANCE

Fig. 7 shows the measured eye diagrams in different operation modes. The three figures to the left show the eye diagrams captured on an equivalent-time scope when the transmitter is operating in BB mode. In AMT mode, the eye diagrams can only be drawn after mixing and integration at the receiver. Since an AMT receiver was not available, the analog signal at the output of the transmitter was sampled with a scope and mixing and integration was performed in MATLAB to generate the eyes.




Fig. 6. 4:2 compressor schematic. The critical path is from X1 to CO.

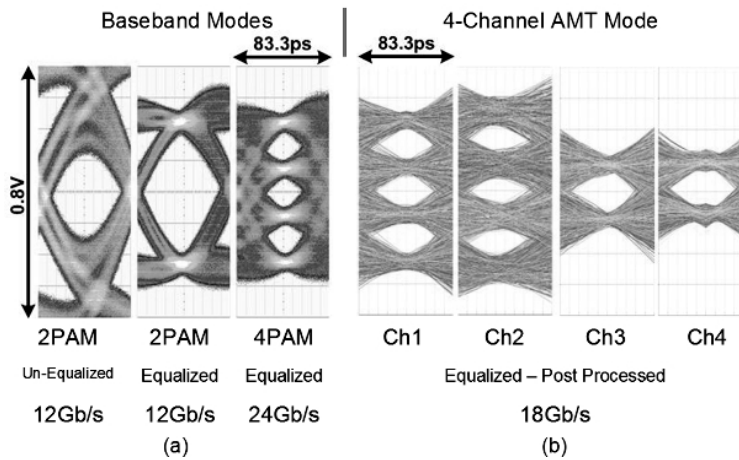


Fig. 7. (a) Eye diagrams measured on a scope when the transmitter is operating in baseband mode: unequalized 2 PAM at 12 Gb/s (left), equalized 2 PAM at 12 Gb/s (middle), and equalized 4 PAM at 24 Gb/s. (b) Eye diagrams measured at the transmitter output with a scope and post-processed in Matlab (mixing and integration per channel—no DFE) when the transmitter is operating in 4-channel 18 Gb/s AMT mode.

Fig. 7(b) shows the four eye diagrams corresponding to the four sub-channels when the transmitter operates in 4-channel AMT mode. Two of the sub-channels are operating in 4 PAM mode and the other two in 2 PAM mode for an aggregate data rate of 18 Gb/s. The two 2 PAM sub-channels can also be programmed to operate in 4 PAM mode.

For the eye diagrams shown in Fig. 7, the communication channel consists of the chip plastic package, 2 inches of differential traces on the PCB, 3 ft of cable, and the front-end of the scope. The overall pulse response of the channel exhibits 14 dB of attenuation at 6 GHz. Optimum equalizer tap coefficients are calculated offline using the analytical optimization framework described in [7] such that the receiver achieves BER of 10^{-15} with minimum required transmit swing. The taps are subsequently scaled such that transmit voltage swing is 1.6 V_{pp}.

Measured 3 dB bandwidth of the DAC with respect to an ideal zero-order-hold pulse is at 7.1 GHz. Maximum measured Integral Nonlinearity (INL) and Differential Nonlinearity (DNL) are 0.31 and 0.28 LSB, respectively. Measured signal to noise and distortion ratio (SNDR) and spurious free dynamic range (SFDR) using single tones are 41 and 51 dB, respectively, at

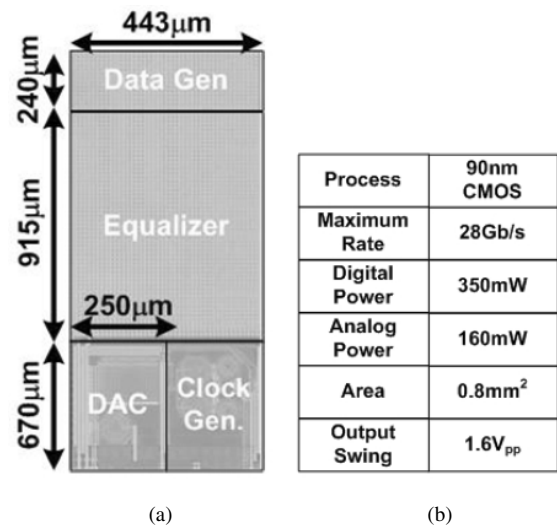


Fig. 8. (a) Chip micrograph. (b) Performance summary.

750 MHz and 32.5 and 35 dB, respectively, at 1.5 GHz. Measured wideband linear signal to distortion ratio (SDR) [16] of the DAC is 32 dB.

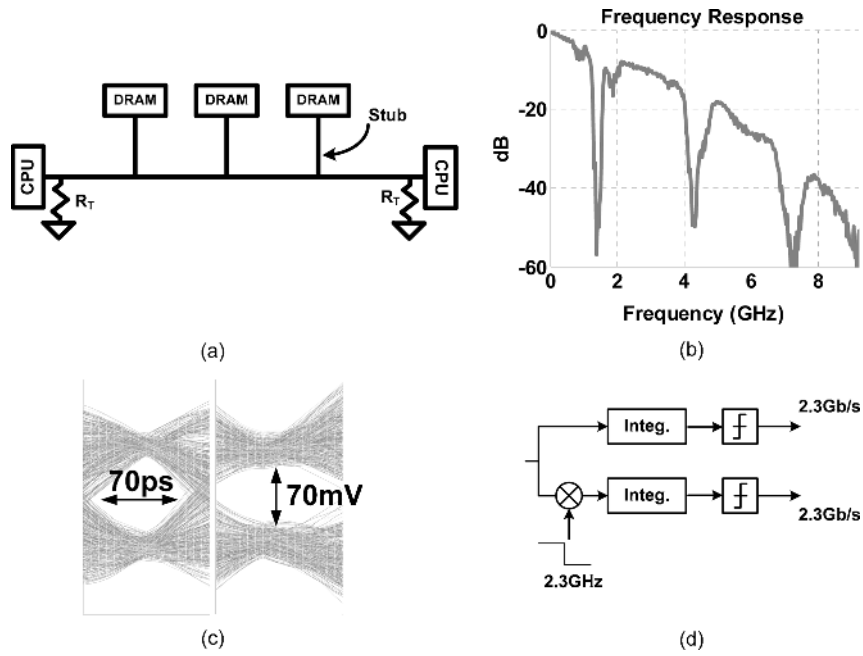


Fig. 9. (a) Multi-drop configuration. (b) Measured frequency response of a three-drop 16 in FR4 trace. (c) Eye diagrams based on measured data when only two 2 PAM, 2.3 GS/s sub-channels are used. Total throughput is 4.6 Gb/s. (d) Block diagram of the receiver simulated in MATLAB to generate the eyes.

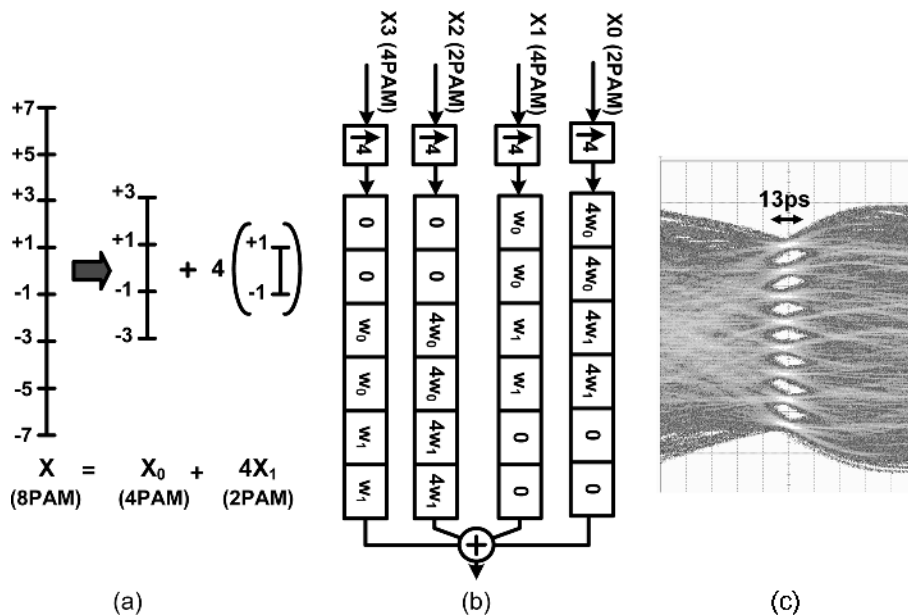


Fig. 10. (a) 8 PAM symbol decomposition to a 4 PAM and a 2 PAM symbol. (b) Transmitter configured in 8 PAM mode (only six taps per phase shown). (c) Eye diagram on a scope in 8 PAM, 36 Gb/s mode.

Fig. 8(a) shows the chip micrograph. The main measured characteristics are included in Fig. 8(b). At nominal operating rate, the entire transmitter has an energy efficiency of 21 mW/Gb/s. Clear eyes can be seen in the lab at 28 Gb/s at nominal supply voltage and room temperature.

AMT can particularly improve the performance of high-speed links where stubs dominate the channel characteristic. Fig. 9(a), for example, shows a configuration where multiple devices (for example memory modules) are connected to two CPUs. In this application, when the two CPUs communicate, other modules connected to the trace act as reflection generators and

create notches in the frequency response of the communication channel. However, notch frequencies are directly related to the length and the impedance of the stubs, and it is possible to tune these parameters to ensure that all reflections from different stubs resonate at the same frequencies [15]. Stubs are waveguides and waveguides have periodic frequency responses. As a result, in these applications, when a notch occurs at a certain frequency, other notches appear at odd harmonics of the first notch frequency. Equally-spaced notch frequencies are ideal for AMT, because in an AMT system all sub-channels have equal symbol-rates, and consequently, the same bandwidths.

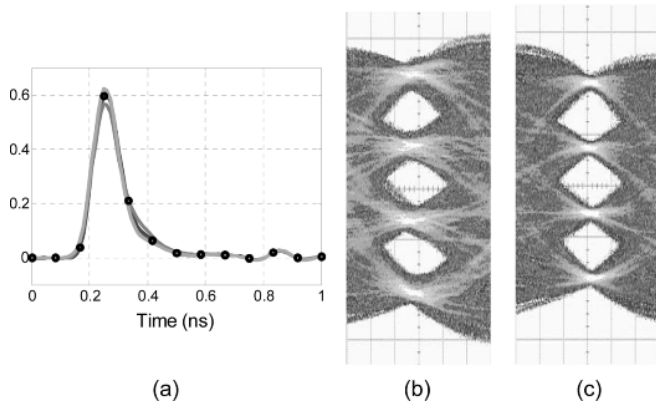


Fig. 11. (a) Measured pulse responses of the 4 phases of the DAC. (b) LTI equalized BB 4 PAM 28 Gb/s eyes on a scope. (c) Cyclically time-variant equalized BB 4 PAM 28 Gb/s eyes on a scope.

Fig. 9(b) shows the measured frequency response of a 3-stub 16 in FR4 trace. The first notch frequency is slightly above 1 GHz. A property of AMT, and the implemented transmitter in particular, is that the entire frequency response of the system can be scaled by changing a single clock frequency. In this case in order to place the 3 dB bandwidth of the BB sub-channel around 1.1 GHz, the input clock frequency to the transmitter is reduced from 12 to 9.2 GHz. Fig. 9(c) shows the measured eye diagrams when only two of the sub-channels are used in 2 PAM mode for an aggregate data-rate of 4.6 Gb/s. Again the receiver [see Fig. 9(d)] is implemented in MATLAB to post process measured data at the output of the channel. If the transmitter is configured in 2 PAM BB mode over the same channel, no open eyes can be observed on the scope beyond 2.6 Gb/s.

V. OTHER TRANSMITTER OPERATION MODES

The degrees of freedom available in the transmitter enable supporting more complex multi-PAM (for example 8 PAM) BB modulations. In addition the freedom in choosing the tap coefficients for different branches of the parallelized equalizer enables digital compensation of the analog distortions through cyclic time-variant equalization. These modes are briefly described in this section.

A. Multi-PAM Operation

By programming the correct tap coefficients to the equalizer, the transmitter can support signal transmission in 2^M -PAM ($1 \leq M \leq 8$) mode. Operation in 8 PAM, for example, is enabled based on the observation that an 8 PAM sequence can be decomposed to the summation of a 4 PAM sequence and a 2 PAM sequence as shown in Fig. 10(a). Fig. 10(b) shows the transmitter configured to operate in 8 PAM mode, as a 2-way parallelized system where each of the parallel branches consists of two branches which together generate an equalized 8 PAM output. In this mode the transmitter can only support up to 18 Gb/s of uncorrelated data since two of the branches are operating in 2 PAM mode instead of 4 PAM. However, in order to show the operating limits of the analog circuits in the transmitter, some of the equalizer taps were used to delay the input sequence for a few cycles and add it back to itself. This effectively doubles the throughput by creating a virtual source with a

weakly correlated sequence. As long as the delay is larger than the delay spread of the pulse response of the channel, such correlation has negligible effect on the eye diagrams. Fig. 10(c) shows that discernible, although perhaps too small to be practical, eyes observed on the scope at 36 Gb/s in 8 PAM mode. Similar ideas can be applied to program the transmitter in higher multi-PAM BB modes as well as 2-channel AMT with 6 GS/s (12 Gb/s) sub-channels.

B. Linear Cyclic Time-Variant Equalization

The transmitter also supports cyclic time-variant equalization. This mode is useful for applications utilizing wide-band interleaved data converters, where the response from the input to the output may be different from time to time due to the mismatch between the interleaved paths. For example, in the transmitter described in this paper, four different paths can be identified from the input of the DAC to its output. Fig. 11(a) shows the estimated pulse responses for these four paths. It can be seen that one path is visibly different from the other three. This difference in the responses makes the system cyclically time-variant rather than time-invariant. Therefore, as long as the system is treated like an LTI system, the time-variant nature manifests itself as analog distortion. Fig. 11(b) shows the measured eye diagrams when the transmitter is operating in BB 4 PAM 28 Gb/s mode, and LTI equalization is performed. Measured signal to interference and distortion ratio (SIDR) at the middle of the eye is 26 dB. However, if the four phases of the 4-way parallelized equalizer in the transmitter are programmed independently to perform cyclic time-variant equalization [16], measured SIDR improves to 31 dB [see Fig. 11(c)], which indicates at least 5 dB of the distortion is related to the mismatch.³

VI. CONCLUSION

A software programmable transmitter can be implemented by parallelizing a conventional FIR filter and letting each of the parallel branches to be programmed independently. The architecture of the transmitter enables supporting bandwidth scalable AMT and 2^M -PAM BB transmission. In addition, the degrees of freedom available in an AMT system and in the transmitter in particular, enable cyclic time-variant equalization to compensate for analog distortions of the system. The additional degrees of freedom come at a small cost, since the additional equalizer tap values which are needed only add a small additional area to the chip, but do not increase the active power.

REFERENCES

- [1] M. Meghelli *et al.*, "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.
- [2] R. Palmer *et al.*, "A 14 mW 6.25 Gb/s transceiver in 90 nm CMOS for serial chip-to-chip communication," in *IEEE ISSCC Dig. Tech. Papers*, 2006, pp. 440–441.
- [3] M. Harwood *et al.*, "A 12.5 Gb/s SerDes in 65 nm CMOS using a baud-rate ADC with digital receiver equalization and clock recovery," in *IEEE ISSCC Dig. Tech. Papers*, 2006, pp. 436–437.
- [4] B. Casper *et al.*, "A 20 Gb/s forwarded clock transceiver in 90 nm CMOS," in *IEEE ISSCC Dig. Tech. Papers*, 2006, pp. 90–91.

³The actual SDR is 32 dB which corresponds to 6 dB improvement in SDR. The 1 dB difference between SDR and SIDR is due to the residual interference which cannot be suppressed to better than 34 dB with only transmit equalization.

- [5] V. Stojanović, A. Amirkhany, and M. A. Horowitz, "Optimal linear precoding with theoretical and practical data rates in high-speed serial-link backplane communication," in *Proc. IEEE Int. Conf. Communications*, 2004, pp. 2799–2806.
- [6] A. Amirkhany, A. Abbasfar, V. Stojanović, and M. A. Horowitz, "Practical limits of multi-tone signaling over high-speed backplane electrical links," in *Proc. IEEE Int. Conf. Communications*, 2007, pp. 2693–2698.
- [7] A. Amirkhany, A. Abbasfar, V. Stojanović, and M. A. Horowitz, "SPC03-5: Analog multi-tone signaling for high-speed backplane electrical links," in *Proc. IEEE GLOBECOM'06*, Nov. 2006, pp. 1–6.
- [8] A. Amirkhany *et al.*, "A 24 Gb/s software programmable multi-channel transmitter," in *Symp. VLSI Circuits Dig. Tech. Papers*, 2007, pp. 38–39.
- [9] A. Amirkhany, V. Stojanović, and M. A. Horowitz, "Multi-tone signaling for high-speed backplane electrical links," in *Proc. IEEE GLOBECOM'04*, Nov.–Dec. 2004, vol. 2, pp. 1111–1117.
- [10] S. Sidiropoulos and M. A. Horowitz, "A 700-Mb/s/pin CMOS signaling interface using current integrating receivers," *IEEE J. Solid-State Circuits*, vol. 35, no. 5, pp. 681–690, May 1997.
- [11] P. P. Vaidyanathan, *Multirate Systems and Filter Banks*. Englewood Cliffs, NJ: Prentice-Hall, 1993.
- [12] S. B. Weinstein and P. M. Ebert, "Data transmission by frequency-division multiplexing using the discrete Fourier transform," *IEEE Trans. Commun. Technol.*, vol. COM-19, no. 10, pp. 628–634, Oct. 1971.
- [13] J. Savoj, A. Abbasfar, A. Amirkhany, M. Jeeradit, and B. Garlepp, "A 12-GS/s phase-calibrated CMOS digital-to-analog converter," in *Symp. VLSI Circuits Dig. Tech. Papers*, 2007, pp. 68–69.
- [14] M. Nagamatsu, S. Tanaka, J. Mori, K. Hirano, T. Noguchi, and K. Hatanaka, "A 15-ns 32×32 -b CMOS multiplier with an improved parallel structure," *IEEE J. Solid-State Circuits*, vol. 25, no. 2, pp. 494–497, Apr. 1990.
- [15] W. Beyene, "Controlled inter-symbol interference design techniques of conventional interconnect systems for data rates beyond 20 Gbps," *Electr. Perform. Electron. Packag.*, pp. 159–162, Oct. 2006.
- [16] A. Amirkhany, A. Abbasfar, J. Savoj, and M. Horowitz, "Time-variant characterization and compensation of wideband circuits," in *Proc. IEEE Custom Integrated Circuits Conf. (CICC)*, Sep. 2007, pp. 487–490.
- [17] A. Amirkhany, "Multi-carrier signaling for high-speed electrical links," Ph.D. dissertation, Stanford Univ., Stanford, CA, Mar. 2007 [Online]. Available: <http://mos.stanford.edu/group/people.html>



Amir Amirkhany (S'04–M'08) received the M.Sc. degree from the University of California, Los Angeles, in 2002, and the B.Sc. degree from Sharif University of Technology, Tehran, Iran, in 1999, both in electrical engineering. He is currently pursuing the Ph.D. degree in electrical engineering from Stanford University, Stanford, CA.

Since July 2007, he has been with Rambus Inc., Los Altos, CA. Since 2003, he has been a Research Assistant with the VLSI Group, Stanford University, and in close collaboration with Rambus Inc., where

he has been involved with the design of chip-to-chip electrical links. Prior to Stanford, he was with Sequoia Communications, working on the ASIC design of WCDMA systems. His main research interests include the design and implementation of communication systems, VLSI circuit design, and application of communication and signal processing techniques to the design of low power circuits.

Mr. Amirkhany was a recipient of a Best Student Paper Award at the IEEE Global Communications Conference in 2006 for his work on the design and analysis of an analog multi-tone system for chip-to-chip interconnects.



Aliazam Abbasfar (S'01–M'05–SM'07) received the B.Sc. and M.Sc. degrees from the University of Tehran, Tehran, Iran, in 1992 and 1995, respectively, and the Ph.D. degree from the University of California, Los Angeles (UCLA), in 2005, all in electrical engineering.

Upon graduation from UCLA, he joined Rambus Inc., Los Altos, CA, where he is involved with high-speed data communications on wireline backplane links. From 1992 to 1994, he was with the Iran Telecommunication Research Center (ITRC), where he was involved with switching networks for data communications. Between 2001 and 2004, he held positions as a Senior Design Engineer in the areas of communication system design and digital VLSI ASIC design with Innovics Inc., Sequoia Communications, and Jaalaa Inc. His main research interests include wireless communications, equalization, error correcting codes, and VLSI ASICs for digital data communications.



Jafar Savoj (S'98–M'02–SM'07) received the B.Sc. degree in electrical engineering from Sharif University of Technology, Tehran, Iran, in 1996, and the M.Sc. and Ph.D. degrees in electrical engineering from the University of California, Los Angeles, in 1998 and 2001, respectively.

He is currently with the Core Technology Group, Rambus, Los Altos, CA. Prior to that, he held positions with Transpectrum and Marvell.

Dr. Savoj was a recipient of a IEEE Solid-State Circuits Society Predoctoral Fellowship for 2000–2001, the Beatrice Winner Award for Editorial Excellence at the 2001 ISSCC, and the Design Contest Award of the 2001 Design Automation Conference. He is with the technical program committee of the IEEE Symposium on VLSI Circuits. He served as a technical program committee member and most recently as the chair of the wired committee of the IEEE Custom Integrated Circuits Conference (CICC) until 2007. He was a Guest Editor for the IEEE JOURNAL OF SOLID-STATE CIRCUITS in 2005 and 2006.



Metha Jeeradit (S'02) received the B.S. and M.Eng. degrees in electrical and computer engineering degrees from Cornell University, Ithaca, NY, in 2001 and 2002, respectively. In 2002, he started his Ph.D. program in electrical engineering with Stanford University, Stanford, CA.

In 2004, he joined Rambus Inc., Los Altos, CA. His main research interests include PLLs and circuit optimizations.



Bruno W. Garlepp (M'97) was born in Bahia, Brazil, in 1970. He received the B.S.E.E. degree from the University of California, Los Angeles, in 1993, and the M.S.E.E. degree from Stanford University, Stanford, CA, in 1995.

In 2007, he joined SiTime Corp., Sunnyvale, CA, as Director of Circuit Engineering to lead the design of synthesizer and timing ICs based on silicon MEMS resonators. In 1993, he joined the Hughes Aircraft Advanced Circuits Technology Center, Torrance, CA, where he designed high-precision

analog ICs for A/D applications and RF circuits for wide-band communication applications. In 1996, he joined Rambus Inc., Mountain View, CA, where he designed high-speed CMOS clocking and I/O circuits for synchronous chip-to-chip interfaces. In 2000, he joined Silicon Laboratories, Austin, TX, where he designed high-performance CDR and clock synthesis ICs for SONET applications. In 2003, he returned to Rambus Inc., Los Altos, CA, where he designed multi-gigahertz signaling interfaces for serial data communications and led a team investigating multi-tone techniques for multi-gigahertz serial links.



Ravi T. Kollipara (M'88–SM'07) is a Senior Principal Engineer with Rambus Inc., Los Altos, CA, responsible for the signal integrity of the high-speed serial link channels. His responsibilities include design and development of models for packages, line cards, backplanes, connectors, traces and vias, and performing simulations for system level voltage and timing budgets and jitter characterization.



Vladimir Stojanovic (S'96–M'05) received the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 2000 and 2005, respectively, and the Dipl. Ing. degree from the University of Belgrade, Belgrade, Serbia, in 1998.

He is currently an Assistant Professor with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Boston. He was also with Rambus, Inc., Los Altos, CA, from 2001 to 2004. He was a Visiting Scholar with the Advanced Computer Systems Engineering

Laboratory, Department of Electrical and Computer Engineering, University

of California, Davis, during 1997–1998. His current research interests include design, modeling, and optimization of integrated systems, from standard VLSI blocks to CMOS-based electrical and optical interfaces. He is also interested in design and implementation of digital communication techniques in high-speed interfaces and high-speed mixed-signal IC design.



Mark Horowitz (S'77–M'78–SM'95–F'00) received the B.S. and M.S. degrees in electrical engineering from Massachusetts Institute of Technology, Cambridge, in 1978, and the Ph.D. degree from Stanford University, Stanford, CA, in 1984.

He is currently the Yahoo Founders Professor of the School of Engineering, Stanford University. In 1990, he took leave from Stanford to help start Rambus Inc., Los Altos, CA, a company designing high-bandwidth memory interface technology. His current research includes multiprocessor design, low

power circuits, high-speed links and new graphical interfaces.

Dr. Horowitz was a recipient of a 1985 Presidential Young Investigator Award, the 1993 ISSCC Best Paper Award, the ISCA 2004 Most Influential Paper of 1989, and the 2006 winner of the IEEE Donald Pederson Award in Solid State Circuits. He is a Fellow of ACM and a member of the NAE.