

A 3.125 Gb/s Limit Amplifier in CMOS With 42 dB Gain and 1 μ s Offset Compensation

Ethan A. Crain, *Student Member, IEEE*, and Michael H. Perrott, *Member, IEEE*

Abstract—A fast offset compensation method for high-gain amplifiers is presented that leverages a novel peak detector design and a dynamic, multi-tap feedback system to achieve roughly three orders of magnitude improvement in settling time over traditional compensation methods. Design tradeoffs between gain, bandwidth, power dissipation, and noise performance of the limit amplifier are discussed. Measured results of a custom 3.125 Gb/s limit amplifier in 0.18 μ m CMOS employing the proposed compensation technique demonstrate a sub-1-ms settling time while still achieving less than 4 ps rms output jitter with a 2.5 mV peak-to-peak input at 2.5 Gb/s.

Index Terms—Compensation, high speed, offset, peak detector.

I. INTRODUCTION

HIGH-GAIN amplifiers require DC offset compensation to achieve high input sensitivity. Without offset-compensation, the output of the amplifier could saturate and block the desired signal due to the offset voltage, as shown in Fig. 1(a). The addition of offset compensation prevents the output from saturating, as shown in Fig. 1(b).

Long compensation times are acceptable for many point-to-point links because offset compensation occurs once and the long settling times do not significantly degrade the efficiency of the data-link. However, long settling times pose a severe obstacle for future many-to-one links where the input to a single receiver is switched between multiple channels. In particular, since switching of the channel may introduce a different offset, the speed of the compensation loop may determine how quickly the receiver can switch between channels.

Classic offset compensation in wide bandwidth applications with non-return to zero (NRZ) data streams, as encountered in SONET applications, utilizes a feedback path from the output of the amplifier back to its input through an RC low-pass filter as shown in Fig. 2(a) [1]–[3]. Unfortunately, this approach leads to an undesirable tradeoff between offset compensation settling time and output jitter—a higher offset compensation bandwidth has the benefit of achieving faster settling time at the expense of




Fig. 1. (a) Output of amplifier without offset compensation saturated due to offset. (b) Adding offset compensation prevents the output from saturating.

higher data-dependent jitter. As a result of this limitation, current approaches suffer from long compensation times (typically > 1 ms for SONET OC-48 applications) and often require an off-chip capacitor to achieve an acceptably low compensation bandwidth.

We propose an offset compensation scheme that leverages a novel CMOS peak detector structure and a variable tap feedback system to dramatically improve the tradeoff relationship between offset compensation settling time and data-dependent jitter due to the offset correction loop. Fig. 2(b) illustrates a simple block diagram of the proposed system—the low-pass filter of the traditional compensation method is replaced with a peak detector and integrator. The peak detector measures the output-referred offset of the limit amplifier and the integrator filters the instantaneous peak detector output and forces the steady-state, output-referred offset voltage to zero. The proposed approach supports continuous operation, does not require a clock, and allows seamless integration into data links transporting NRZ data.

We will demonstrate that the approach enables the settling time of a 3.125 Gb/s limit amplifier to be reduced from approximately 1 ms, using traditional compensation methods, to less than 1 μ s while simultaneously meeting SONET OC-48 jitter

Manuscript received March 23, 2005; revised August 11, 2005. This work was supported in part by the MARCO Focus Center for Circuit and System Solutions (C2S2), under MARCO Contract 2003-CT-888, and in part by National Semiconductor.

E. A. Crain was with the MTL High Speed Circuits and Systems Group, Massachusetts Institute of Technology, Cambridge, MA 02139 USA. He is now with Silicon Laboratories, Austin, TX 78735 USA (e-mail: eacrain@silabs.com).

M. H. Perrott is with the MTL High Speed Circuits and Systems Group, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: perrott@mit.edu).

Digital Object Identifier 10.1109/JSSC.2005.862352

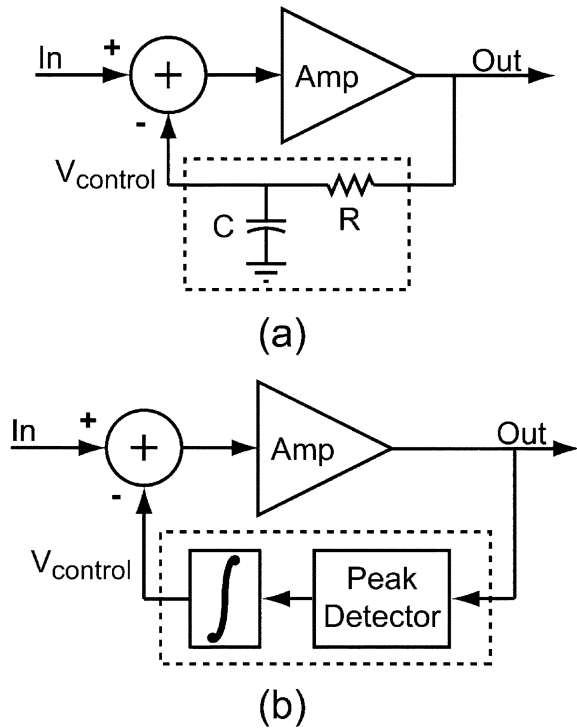


Fig. 2. (a) Traditional offset compensation using a low-pass filter. (b) Proposed offset compensation using the peak detector and an integrator.

specifications at 2.5 Gb/s. Thus, the proposed system enables a three orders of magnitude improvement in offset compensation time over the classical approach while maintaining very low data-dependent jitter levels.

The remainder of this paper is divided into four sections. Section II describes the key system component that enables the dramatic improvement in settling time—the proposed peak detector. We then present a variable tap feedback architecture that leverages the peak detector in a manner that avoids corruption of the offset information when the output of the limit amplifier saturates. Section III covers the design of each of the system blocks and presents an analysis of design tradeoffs for the limit amplifier with respect to gain, bandwidth, power, and noise. Section IV briefly discusses the behavioral modeling of the overall system. Section V presents measured results, and we summarize the benefits of the approach in the conclusion.

II. PROPOSED METHOD

We will show that an undesirable tradeoff between fast tracking performance, required for fast compensation, and data-dependent droop limit the use of traditional CMOS peak detectors. We then introduce the proposed peak detector design and demonstrate how it dramatically improves this tradeoff. We also present the proposed variable-tap feedback system that employs the new peak detector to achieve fast and reliable offset compensation despite possible saturation of the limit amplifier output.

A. Peak Detection

Rather than using a low-pass filter to determine the output-referred offset of the main amplifier as the difference between the

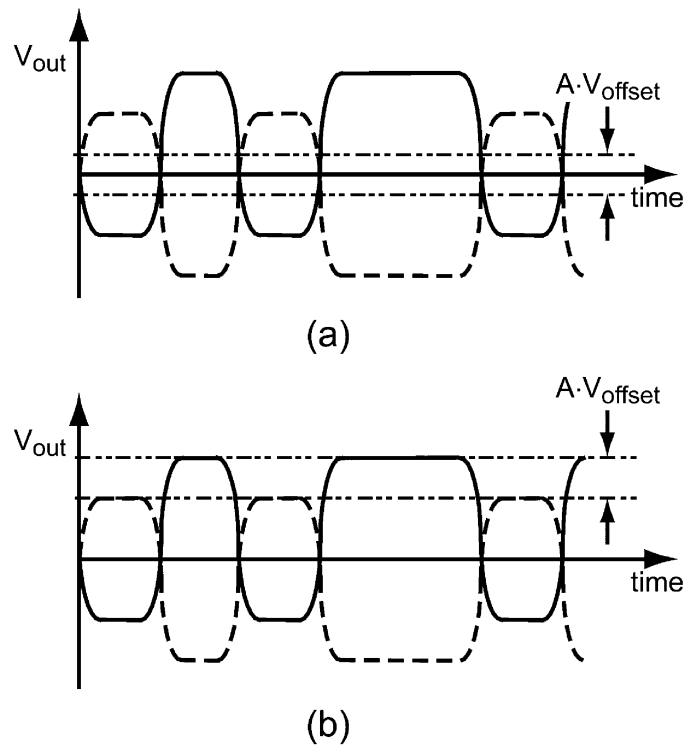


Fig. 3. (a) Typical view of measuring output-referred offset voltage. (b) Proposed method of measuring output-referred offset voltage.

two differential common-mode voltages in Fig. 3(a), we assume that the offset of a differential amplifier can instead be estimated from the difference between the peak values of the two outputs as shown in Fig. 3(b). In doing so, we assume approximately equal gains through each side of the differential limit amplifier, which is reasonable for integrated amplifier designs. We also assume that the amplifier is processing NRZ data such that the output alternates between two different levels.

While the use of peak detection for offset cancellation is not a new concept [4], its implementation in pure CMOS has previously proven challenging due to the lack of diode structures capable of high-frequency operation. While CMOS source-follower circuits can be used to perform the diode function of signal rectification, the resulting circuit lacks the ability to simultaneously achieve fast tracking of the signal peaks and long holding times of the sampled peak values. In particular, source followers require a considerable amount of channel current to achieve fast tracking of the input peak values and the presence of such current leads to severe droop problems at the peak detector output.

To illustrate the limitation of traditional CMOS peak detector circuits, consider Fig. 4(a), in which the peak detector input is high and the output is attempting to track the input signal. In this case, the peak detector input device must have a large g_m to achieve a fast tracking time, which implies the need for a large bias current through the device. Now consider Fig. 4(b), where the same peak detector attempts to hold its output value when the input signal is low. In this case, the input device is cut off so that, ideally, the last sampled input is stored on the load capacitor. Unfortunately, the tail bias current continues to strip charge from the load capacitor, thereby causing droop at the output that is proportional to both the bias current and the

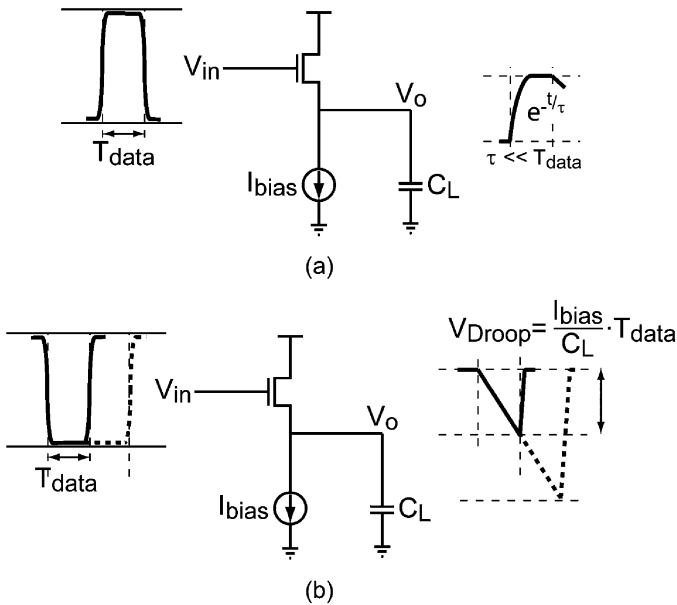


Fig. 4. Traditional CMOS peak detector (source follower) (a) requires large channel current to achieve fast compensation settling times, and (b) suffers from excess droop that is proportional to both the bias current and the symbol period.

symbol period of the input signal: $V_{\text{Droop}} = I_{\text{bias}}/C_L \cdot T_{\text{data}}$. This data-dependent droop directly translates to data-dependent jitter, or inter-symbol interference (ISI), at the output of the limit amplifier.

If we attempt to use the traditional CMOS peak detector within the proposed compensation loop in Fig. 2(b), then we are prevented from simultaneously achieving fast compensation time (which requires the peak detector to have fast tracking performance) and low jitter at the output of the limit amplifier (which requires the peak detector to have low droop). The classical use of CMOS source followers as peak detectors therefore presents an unacceptable tradeoff between settling time and droop. We need to improve the peak detector design in order to achieve both fast compensation time and low jitter.

We propose a CMOS peak detector structure that substantially mitigates the track and hold tradeoff discussed above by utilizing a simple switch that appropriately directs the bias current according to the input as shown in Fig. 5. Implementation of the proposed series switch is straightforward if the limit amplifier is processing NRZ data, as is commonly done [5]. When the input is high [Fig. 5(a)], the switch is closed and the output tracks the input—the peak detector operates exactly as the traditional peak detector in Fig. 4(a). When input is low [Fig. 5(b)], the input device is off, the last sampled input is stored on the load capacitor and, since the switch is open, the output is isolated from the bias current. Therefore, we simultaneously achieve high speed operation during the tracking phase and minimal data-dependent droop during the hold phase.

Ideally, the droop at the output of the proposed peak detector is zero since no charge is stripped from the load capacitor by the bias current source in the hold mode (when the input is low). In reality, incomplete switching due to limited input amplitude causes a finite amount of current to flow out of the load capacitor during the hold mode and results in nonzero droop. However, the proposed current switching scheme dramatically alters the

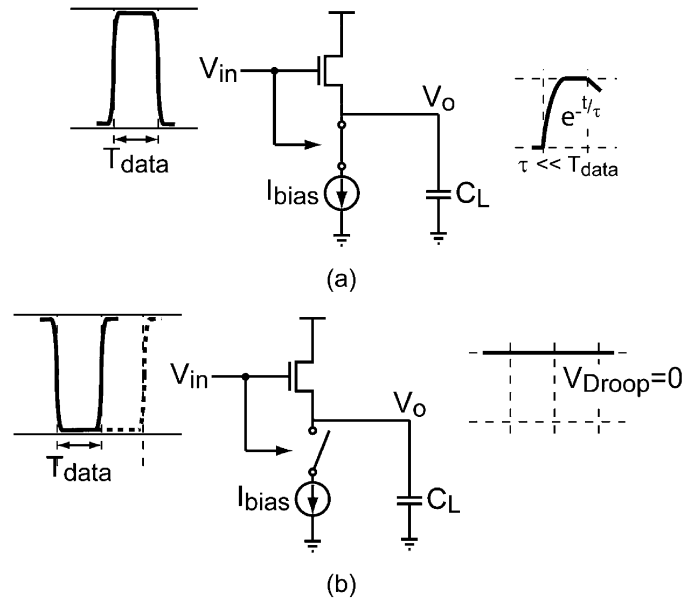


Fig. 5. Proposed CMOS peak detector (a) operates like a traditional peak detector in the track mode, and (b) ideally has zero droop at the output during the hold phase.




Fig. 6. Final differential design of the proposed CMOS peak detector.

tradeoff between the settling time and droop compared to the classical peak detector design such that the magnitude of droop is dramatically reduced. By leveraging the new peak detector design within the proposed compensation loop, we can achieve a significant improvement in settling time without compromising jitter performance of the limit amplifier [5].

A differential implementation of the proposed peak detector design leveraging the switched current technique is shown in Fig. 6. The top transistors are the source-follower devices and the bottom transistors are the current-steering switch devices. As such, a very efficient implementation of the switch is achieved—the input signal simply commutates the bias current between the two source follower circuits of the differential peak detector. Fig. 6 shows the left side of the peak detector tracking the positive input and the right side of the circuit holding the last sampled value of the negative input.

B. System Design

Since the proposed offset compensation technique estimates the output referred offset of the main amplifier as the difference between the peak values of the differential outputs, it is important that feedback occur from an unsaturated amplifier

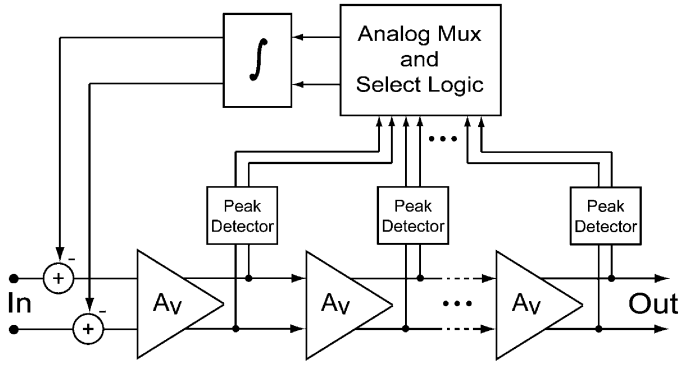


Fig. 7. Full system design showing dynamic, multi-tap offset compensation control loops.

output. Otherwise, offset information is lost. The multi-tap feedback system architecture shown in Fig. 7 is proposed and implemented to avoid corruption of the offset information. The proposed peak detector circuits are placed at the output of each amplifier stage—digital selection logic then selects the last *unsaturated* peak detector output to feed back to the input through a continuous-time integrator. The digital selection logic compares the output of each peak detector to a fixed reference; we assume that if an output is unsaturated then all preceding outputs are also unsaturated. Based on the digital selection logic outputs, the control logic dynamically adjusts the tap location of the control loop. The integrator filters any discontinuities at its input and generates a continuously adjustable output. This approach allows the offset to be continuously compensated in the presence of a time-varying input referred offset or input signal amplitude.

As the tap location changes, so does the number of limit amplifier stages in the compensation loop and, therefore, the overall gain in the feedback loop. To keep consistent closed-loop dynamics, the overall loop gain is kept relatively constant by adjusting the integrator gain as a function of which feedback tap is selected. Since the peak detectors and select logic are small and consume little power, multiple inclusions of the cells do not pose a significant power or area penalty.

III. CIRCUIT DESIGN

To demonstrate the proposed offset compensation technique, a seven-stage resistor-loaded limit amplifier utilizing the proposed approach was designed and implemented in a 0.18 μm CMOS process. We examine the design tradeoffs in achieving high gain, high bandwidth, low power dissipation, and low noise of the overall limit amplifier, and then present our design choices.

A. Limit Amplifier

The main amplifier is implemented as a cascaded multi-stage limit amplifier, as shown in Fig. 8, to achieve a high gain-bandwidth product [7]. Selection of the number of stages is a compromise between achieving a high gain-bandwidth product, low power dissipation, and high input sensitivity. Specifically, we will show that a large gain-per-stage is desirable for low-power

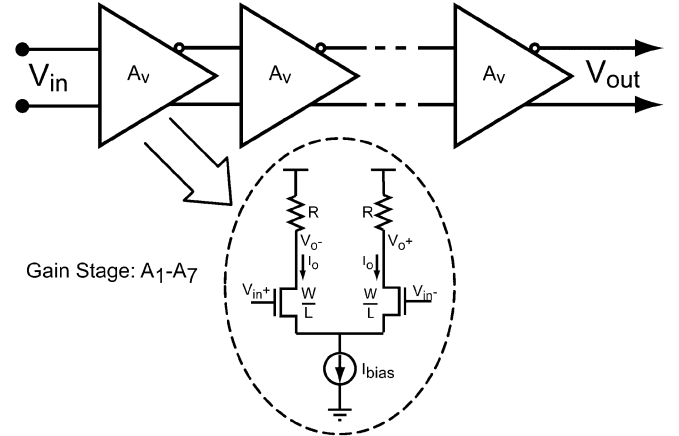


Fig. 8. Limit amplifier design and individual amplifier stage design.

operation and high input sensitivity at the expense of bandwidth. The gain-per-stage required for minimum power dissipation approaches the value required for maximum bandwidth for high-speed designs with the penalty of low input sensitivity. High input sensitivity generally necessitates high power dissipation.

Each stage in the limit amplifier was implemented as a resistor-loaded differential pair since this topology is very amenable to high-speed operation without requiring excessive area as is often encountered in inductor-enhanced designs. The specific device and bias conditions were determined using the design methodology described in [8] implemented in a Matlab script [13]. The script implements a simple numerical design procedure for designing high speed, differential amplifiers with resistor loads without relying on square-law assumptions for CMOS devices. By combining hand analysis with SPICE generated data, highly accurate designs can be achieved.

1) *Maximum Gain-Bandwidth Product*: Following the derivation in [9, ch. 8], we can determine the optimal gain per stage, or, conversely, the optimal number of stages, of the limit amplifier to maximize gain-bandwidth product. The aggregate transfer function of the n -stage limit amplifier is

$$H(s)_{\text{LA,TOT}} = \left(\frac{A_V}{(1 + s/p_1)} \right)^n. \quad (1)$$

The -3 dB frequency of the cascaded amplifier is

$$\omega_0 = p_1 \cdot \sqrt{2^{1/n} - 1}. \quad (2)$$

Examining these two equations, we see that the gain increases faster than the bandwidth decreases as the number of stages, n , increases. This means that we can increase the gain-bandwidth product as the number of stages increases, to a limit [7].

If the overall gain of the limit amplifier is G , then the gain-per-stage is $A_V = G^{1/n}$. The solid lines in Fig. 9 show the normalized total bandwidth of the limit amplifier for G equal to 10, 100, and 1000. The optimal number of stages occurs at the maximum of each curve. Note that the slopes of the curves are very shallow near the optimal values for n , providing some degree of freedom in the choice of n . After a significant amount of algebra, the optimal gain per stage required to maximize the total bandwidth is calculated as $A_V = e^{1/2} = 1.65$ [9]. The




Fig. 9. Normalized total bandwidth of multi-stage amplifier versus number of stages to determine optimal gain-bandwidth product.

normalized total amplifier bandwidth is plotted against n for A_V equal to 1.65, 2.0, and 3.0 in the dashed curves of Fig. 9. Notice that the curve for $A_V = 1.65$ intersects each of the solid curves at their peak, confirming that a gain of 1.65 per stage provides the maximum gain-bandwidth product.

For a gain of 42 dB, the optimal number of stages is nine when striving for maximum bandwidth. However, we should also consider power dissipation and input sensitivity (i.e., rms noise voltage) in the design. We will look at these issues next.

2) *Power Dissipation and Noise:* Rather than simply trying to achieve maximum bandwidth, let us consider seeking the optimal value of n to achieve an acceptable tradeoff between the lowest power dissipation and highest input sensitivity for a specified overall limit amplifier gain and bandwidth. Utilizing the Matlab design script [13], we determined the total power dissipation and input referred rms noise of a limit amplifier designed for an overall gain of 100 with overall amplifier bandwidth ranging from 1 to 4 GHz as a function of n within the range of 3 to 15. Amplifier designs with fewer than three stages were not possible because of the large gain and bandwidth requirements for each amplifier stage. The power dissipation results are plotted as solid lines in Fig. 10 and the noise results are plotted in Fig. 11 in terms of input-referred rms noise.

For small values of n , each stage of the limit amplifier must support both very high gain and bandwidth, which leads to very high power dissipation. As n increases toward the value that provides minimum overall power dissipation, the required gain per stage decreases according to $A_V = G^{1/n}$, the required bandwidth per stage increases according to (2), and the power dissipated per stage decreases sharply due to the diminishing gain-bandwidth requirement of each amplifier stage. Below the optimal value of n , the power dissipation per stage decreases faster than the number of stages increases so that the overall power dissipation of the limit amplifier decreases. However, above the optimal value of n , the total power dissipation increases due to the escalating bandwidth requirement of each amplifier stage.

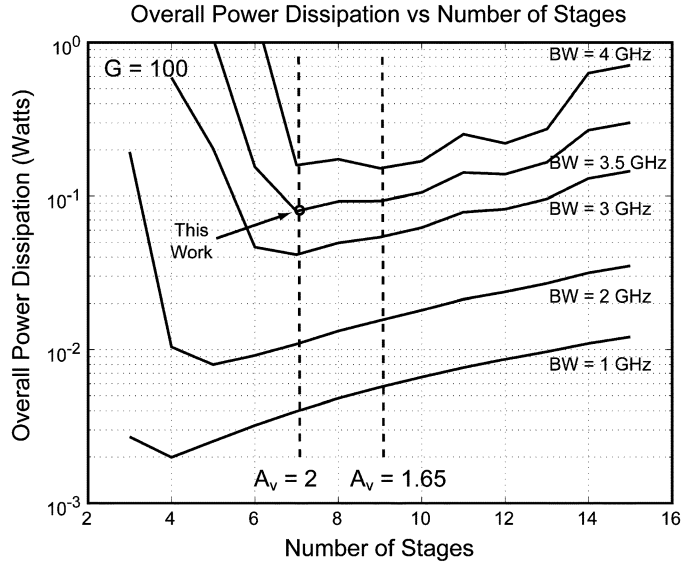


Fig. 10. Total power dissipation versus number of stages of a multi-stage limit amplifier with overall gain of 100 and bandwidths in the range 1–4 GHz.

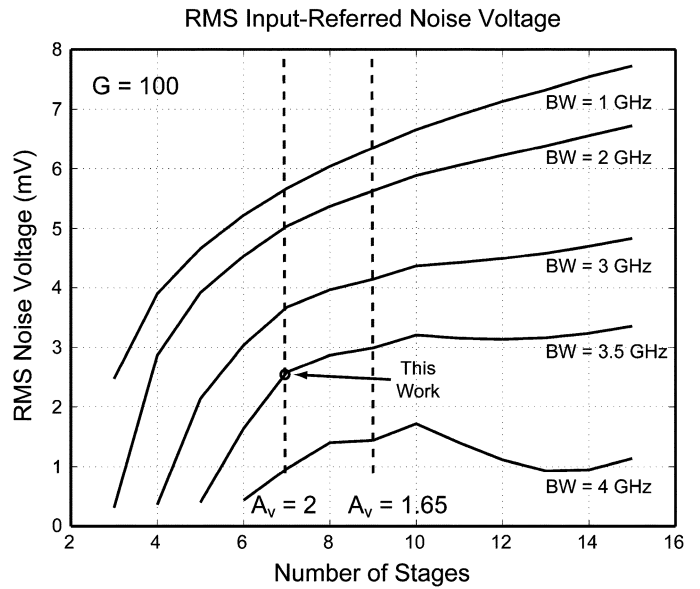


Fig. 11. Input-referred rms noise voltage versus number of stages of a multi-stage limit amplifier with overall gain of 100 and bandwidths in the range 1–4 GHz.

The overall power dissipation of a limit amplifier designed with $A_V = e^{1/2} \approx 1.65$ and overall amplifier bandwidth ranging from 1 to 4 GHz corresponds to the intersection of the right-most vertical dashed line with each solid line in Fig. 10. For small overall amplifier bandwidths, the overall power dissipation with $A_V = 1.65$ is much larger than the minimum achievable power dissipation at the specified bandwidth. As the overall bandwidth of the amplifier increases toward the process limitations, the value of A_V required to achieve minimum power dissipation approaches 1.65. This demonstrates that the number of stages required for maximum bandwidth approaches the number of stages required for minimum power dissipation for high speed amplifier designs, but that the number of stages should generally be set lower to minimize power.

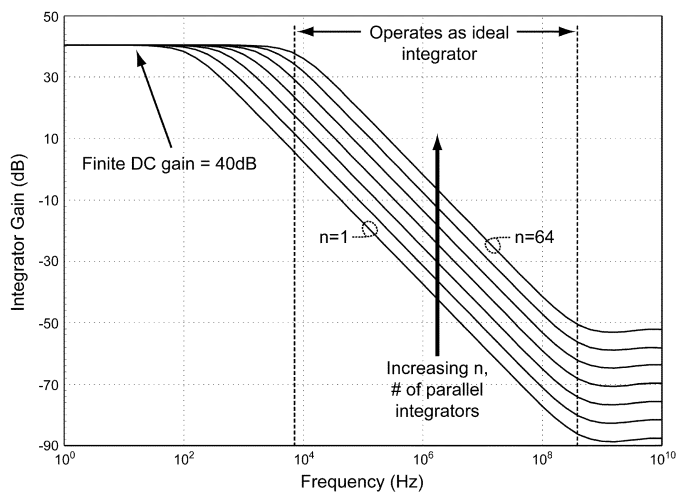


Fig. 13. Frequency response of integrator versus number of stages in parallel.

of the integrator is shown in Fig. 13, and verifies that the effective integrator gain is scaled as a function of the number of parallel g_m stages.

The common-mode feedback (CMFB) of the unit g_m cell was implemented using devices M4a and M4b, as shown in Fig. 12, [10], which operate in the linear region. This approach has the advantages of requiring no extra power dissipation beyond that of the basic g_m cell and of consuming negligible area. In order to obtain wider dynamic range and improve current matching, replica biasing was employed using transistors M1R–M4R and M5. The design is based on [11] and is common to all of the unit g_m cells.

D. Control Logic and Comparator

The control logic and comparators in the feedback path are used to determine the last unsaturated output, based on the peak detector outputs at each stage. The comparator shown in Fig. 14 was designed to be un-clocked, so devices M5/M6 must overdrive the cross-coupled devices M7/M8 to switch the output. The amount of hysteresis of the comparator can be designed by appropriate scaling of the relative over-drive of M5/M6 compared to M7/M8.

IV. BEHAVIORAL MODELING

A model for the overall system was created using linearized system blocks and was simulated using a custom behavioral simulator called CppSim [12]. The behavioral model was used to ensure that the system was stable and that desired response shape and compensation loop settling time was achieved. The system model and a brief tutorial on the system simulation are available at the web site referenced in [12].

V. MEASURED RESULTS

A seven-stage resistor loaded limit amplifier utilizing the proposed offset compensation method was fabricated in National Semiconductor's 0.18 μ m CMOS process [5]. The final die size is 1 mm² and the total active area is 0.5 mm². A die micrograph is shown in Fig. 15. We present measured eye diagrams of the limit amplifier output, which reveal <4 ps rms jitter at 2.5 Gb/s.




Fig. 14. Un-clocked comparator design for feedback control logic.




Fig. 15. Micrograph of the final chip highlighting the major system blocks.

We also show measured step responses of the limit amplifier feedback control voltage for changes in its input referred offset, which reveal <1 μ s settling time for the offset compensation loop.

A. Test Setup

The die was wired bonded in a standard ceramic package which was then soldered to a five-layer FR4 PC board containing low-noise supply voltages and digital interface circuitry. The prototype was tested up to 3.125 Gb/s with a 2³¹-1 PRBS input pattern that had input amplitudes ranging from 2.5 to 50 mV_{pp}. The PRBS input was generated with an HP 70843B 12.5 GB/s error performance analyzer driven by an HP 70340A signal generator. The chip is programmed with a PC connected to the board with a National Instruments DAQCard-6533 digital I/O card via a serial I/O port. On chip, 4-bit binary codes are converted to 15-bit thermometer codes to adjust the input-referred offset of the limit amplifier by ± 15 mV and independently adjust the offset compensation loop-gain and bandwidth to test settling times from 1 μ s down to 100 ns. The

TABLE I
POWER DISSIPATION SUMMARY

	LIMIT AMPLIFIER AND COMPENSATION	OUTPUT BUFFER	TOTAL
POWER	113 mW	225 mW	338 mW

TABLE II
PERFORMANCE OF LIMIT AMPLIFIER (IGNORING THE OUPUT BUFFER)

PARAMETER	VALUE
CMOS TECHNOLOGY	1.8V, 0.18 μ M
AREA	0.02MM ²
BANDWIDTH (-3DB)	5GHZ (SIMULATED)
DC GAIN (DIFFERENTIAL)	42DB
INPUT DYNAMIC RANGE (P-P)	2.5MV-1V
POWER	113 mW

TABLE III
MEASURED JITTER PERFORMANCE WITH 1 MHz COMPENSATION BW

INPUT AMPLITUDE	MEASURED RMS JITTER (PS)		
	1.0 GB/s	2.5 GB/s	3.125 GB/s
2.5 mV _{pp}	3.94	3.71	5.90
10 mV _{pp}	2.65	2.86	6.30
50 mV _{pp}	1.13	2.52	7.98

output jitter is measured from eye diagrams of the limit amplifier output using an Agilent 81600A 50 GHz oscilloscope. The offset compensation step-response settling time is measured by plotting the feedback control voltage with an Agilent 54832D mixed-signal oscilloscope triggered from the serial I/O.

B. Measured Results

The total differential gain was measured to be 42 dB and a summary of the power dissipation for each block is shown in Table I. A summary of the limit amplifier performance (ignoring the output buffer and package parasitics) is listed in Table II. Eye diagrams of the limit amplifier outputs using a 2.5 mV_{pp} PRBS 2³¹-1 input at 2.5 and 3.125 Gb/s are shown in Fig. 16. The measured rms jitter is 3.71 and 5.90 ps, respectively, with the offset compensation bandwidth set to 1 MHz. The increased jitter in the 3.125 Gb/s case is attributed to the limited bandwidth of the packaged amplifier—package parasitics (such as bondwires) reduced the overall amplifier bandwidth to less than the 5 GHz target value.

A summary of rms jitter versus input amplitude and data rate is shown in Table III—here the offset compensation loop was set to a 1 MHz bandwidth and the average measured settling time was 650 ns. The 2.5 Gb/s results demonstrate offset settling times less than 1 μ s while still maintaining SONET OC-48 jitter levels. Note that the 3.125 Gb/s results counterintuitively show increasing jitter levels with increasing input amplitude. The cause of this behavior was not fully determined, but is likely an artifact of board reflections in the test setup.

Plots of the control voltage step response with a 5 mV_{pp} input at 3.125 Gb/s for offset compensation bandwidths of 1 and 3 MHz are shown in Fig. 17. Settling times are 635 and 353 ns, respectively, and the measured rms jitter is 5.75 and 5.90 ps, respectively.

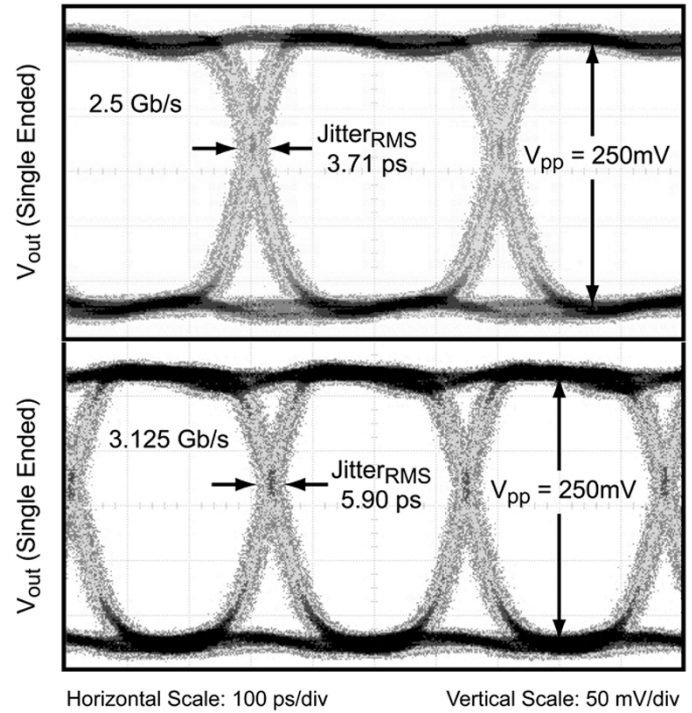


Fig. 16. Eye diagrams of single-ended limit amplifier output with 2.5 mV_{pp} PRBS 2³¹-1 input at 2.5 Gb/s (top) and 3.125 Gb/s (bottom) with the offset compensation loop set to 1 MHz bandwidth.




Fig. 17. Step response of offset compensation control voltage for offset compensation bandwidths of (a) 1 MHz and (b) 3 MHz with 5.0 mV_{pp} PRBS 2³¹-1 input at 3.125 Gb/s.

VI. CONCLUSION

A fast offset compensation method for amplifiers was presented that leverages a novel peak detector and multiple tap

feedback to achieve roughly three orders of magnitude improvement in settling time over classical offset compensation methods. An analysis of high-speed, resistor-loaded CMOS differential amplifier design tradeoffs revealed that maximum bandwidth is achieved at the expense of both higher power and lower input sensitivity. Designing for minimum power dissipation yields an acceptable compromise between large gain-bandwidth product, low power dissipation, and high input sensitivity. Measured results of a custom 3.125 Gb/s limit amplifier in 0.18 μ m CMOS, which employs the proposed technique, indicate that less than 4 ps rms jitter is achievable with $< 1 \mu$ s settling time of the offset compensation loop at a 2.5 Gb/s data rate.

REFERENCES

[1] S. Galal *et al.*, "10 Gb/s limiting amplifier and laser/modulator driver in 0.18 μ m CMOS technology," *IEEE J. Solid-State Circuits*, vol. 38, no. 12, pp. 2138–2146, Dec. 2003.

[2] E. Sackinger *et al.*, "A 3-GHz 32-dB CMOS limiting amplifier for SONET OC-48 receivers," *IEEE J. Solid-State Circuits*, vol. 35, no. 12, pp. 1884–1888, Dec. 2000.

[3] H. Rein, "Multi-gigabit-per-second silicon bipolar ICs for future optical-fiber transmission systems," *IEEE J. Solid-State Circuits*, vol. SC-23, no. 3, pp. 664–674, Jun. 1988.

[4] M. Aiki *et al.*, "446 Mbit/s integrated optical repeater," *J. Lightw. Technol.*, vol. 3, no. 2, pp. 392–399, Apr. 1985.

[5] E. A. Crain and M. H. Perrott, "A 3.125 Gb/s limit amplifier with 42 dB gain and 1 μ s offset compensation in 0.18 μ m CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2005, p. 232.

[6] C. Y. Lau and M. H. Perrott, "Phase locked loop design at the transfer function level based on a direct closed loop realization algorithm," in *Proc. Design Automation Conf.*, vol. 11, Jun. 2003, pp. 526–531.

[7] M. H. Perrott. MIT Open Course Ware site, Course 6.976: High Speed Communication Circuits and Systems, Lecture 6. [Online]. Available: <http://www-mtl.mit.edu/~perrott/>

[8] E. A. Crain and M. H. Perrott, "A numerical design approach for high speed, differential, resistor-loaded, CMOS amplifiers," in *Proc. Int. Symp. Circuits Syst. (ISCAS)*, May 2004, pp. V-508–V-511.

[9] T. H. Lee, *The Design of Radio-Frequency Integrated Circuits*. New York: Cambridge Univ. Press, 1998.

[10] P. R. Gray, P. J. Hurst, S. H. Lewis, and R. G. Meyer, *Analysis and Design of Analog Integrated Circuits*. New York: Wiley, 2001.

[11] K. Gulati and H.-S. Lee, "A high-swing CMOS telescopic operational amplifier," *IEEE J. Solid-State Circuits*, vol. 33, no. 12, pp. 2010–2019, Dec. 1998.

[12] M. H. Perrott. CppSim Behavioral Simulator. [Online]. Available: <http://www-mtl.mit.edu/researchgroups/perrottgroup/tools.html>

[13] E. A. Crain. Amplifier Design Script. [Online]. Available: <http://www-mtl.mit.edu/researchgroups/perrottgroup/tools.html>

[14] J. Gowar, *Optical Communication Systems*. London, U.K.: Prentice-Hall, 1984.



Ethan A. Crain received the B.S. and M.Eng. degrees in electrical engineering and computer science from the Massachusetts Institute of Technology (MIT), Cambridge, in 1995 and 2004, respectively.

From 1996 to 1998, he was a Device Engineer for National Semiconductor. From 1998 to 2003, he was a Staff Designer in Fairchild Semiconductor's Interface group working on the LVDS backplane product family. He joined Silicon Laboratories, Austin, TX, in May 2005 and now works as a Design Engineer in the Wireline Division.



Michael H. Perrott received the B.S. degree in electrical engineering from New Mexico State University, Las Cruces, in 1988, and the M.S. and Ph.D. degrees in electrical engineering and computer science from Massachusetts Institute of Technology (MIT), Cambridge, in 1992 and 1997, respectively.

From 1997 to 1998, he was with Hewlett-Packard Laboratories, Palo Alto, CA, working on high-speed circuit techniques for Sigma-Delta frequency synthesizers. In 1999, he was a Visiting Assistant Professor at the Hong Kong University of Science and Technology, where he taught a course on the theory and implementation of frequency synthesizers. From 1999 to 2001, he worked at Silicon Laboratories, Austin, TX, and developed circuit and signal processing techniques to achieve high-performance clock and data recovery circuits. Currently, he is an Assistant Professor in the Department of Electrical Engineering and Computer Science, MIT, where he focuses on high-speed circuit and signal processing techniques for data-links and wireless applications.