

# A CMOS 10-Gb/s Power-Efficient 4-PAM Transmitter

Kamran Farzan, *Student Member, IEEE*, and David A. Johns, *Fellow, IEEE*

**Abstract**—A novel power-efficient architecture for a multilevel pulse amplitude modulation (PAM) transmitter is proposed. A data-look-ahead technique is used to pre-switch the current sources so that drive current is reduced when transmitting small voltage levels. This technique also eliminates the need for a pre-driver block, which also saves transmitter power. Based on this architecture, a 4-PAM transmitter is designed in 0.18- $\mu\text{m}$  standard digital CMOS technology. The transmitter achieves 3.5 GS/s (7 Gb/s) with a 1.7-V supply and 5 GS/s (10 Gb/s) with a 2-V supply and it occupies an area of 0.16 mm<sup>2</sup>. The output driver and the entire transmitter consume only 11.25 and 66 mW at 7 Gb/s (20 and 121 mW at 10 Gb/s), respectively, which are the lowest reported powers at this speed.

**Index Terms**—CMOS, driver, high speed, power efficient, 4-PAM.

## I. INTRODUCTION

HIGH-SPEED interconnect links are widely used in high-speed network switching, local area networks, memory buses, and multiprocessor interconnection networks. Moreover, advances in IC fabrication technology, coupled with aggressive circuit design, have led to an exponential growth in speed and integration levels. To improve the overall system performance, the communication speed between chips in a system must increase accordingly. For a given data rate, multilevel signaling can be used to reduce the channel symbol rate, intersymbol interference (ISI), and crosstalk [1], [2]. The potential benefits of four-level pulse amplitude modulation (4-PAM) signaling for increasing data rates in physical short-bus systems have been shown in [3]–[6]. However, transmitted power is often increased to compensate for the impact of multilevel signaling on bit error rate (BER). Since there are several drivers in a parallel bus signaling system, the power dissipation of each driver is extremely important, and therefore, power-efficient drivers are desirable. The reported high-speed multilevel drivers have used power-inefficient unipolar architectures [2]–[6]. The proposed driver in this work has a novel power-efficient architecture. A data-look-ahead technique is used for high-speed implementation of this architecture. Based on this architecture, a 4-PAM transmitter is implemented.

## II. DRIVER ARCHITECTURES

As shown in Fig. 1, there are generally two different architectures for drivers: unipolar and bipolar [1]. In unipolar

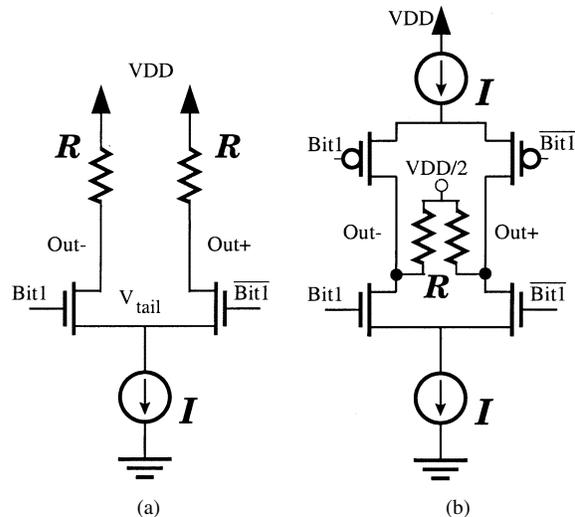


Fig. 1. Driver architectures. (a) Unipolar. (b) Bipolar.

architectures, current  $I$  is steered either in the right or left transistor in the current-steering differential pair. In this case, the single-ended output voltages would be either  $V_{DD}$  or  $V_{DD} - RI$ . Therefore, the differential swing is  $RI$ . However, in bipolar drivers [Fig. 1(b)], the output voltages would be either  $V_{DD}/2 + RI$  or  $V_{DD}/2 - RI$ , which makes the differential swing equal to  $2RI$ . Since the power dissipation in both architectures is practically the same,  $V_{DD} \cdot I$ , the bipolar architecture needs half the power of unipolar architectures for a given swing. Note that since the top and bottom tail currents are equal, the power dissipation of the power supply  $V_{DD}/2$  is ideally zero. In practice, however, there might be some mismatches between the two current sources. In this design, replica bias is used to ensure the top and the bottom tail current sources remain equal over process, voltage, and temperature (PVT) variation. This ensures negligible power dissipation in the power supply  $V_{DD}/2$ . Since the supplied current from this source is small, it is generated on-chip.

There is another source of power inefficiency in typical multilevel PAM drivers. As shown in Fig. 2, the tail current sources of a typical 4-PAM driver are always on and output differential levels ( $\pm RI$ ,  $\pm 3RI$ ) are obtained by steering current in different branches. For example, level  $+RI$  is generated when  $Bit1$  is 0 and  $Bit2$  is 1. This means that the power dissipation would be the same for all four levels. However, the right current source in Fig. 2 ( $2I$ ) can be turned off whenever levels  $\pm RI$  are to be transmitted. Since the driver power dissipation is directly proportional to the total current ( $P = V_{DD} \cdot I_{total}$ ), the driver power consumption can be reduced by a factor of 1.5 for 4-PAM. This power saving would be even more significant in PAM schemes with more than four levels (1.75 for 8-PAM).

Manuscript received July 22, 2003; revised November 5, 2003.

The authors are with the Department of Electrical and Computer Engineering, University of Toronto, Toronto, ON, M5S 3G4 Canada (e-mail: farzan@eecg.toronto.edu; johns@eecg.toronto.edu).

Digital Object Identifier 10.1109/JSSC.2003.822898

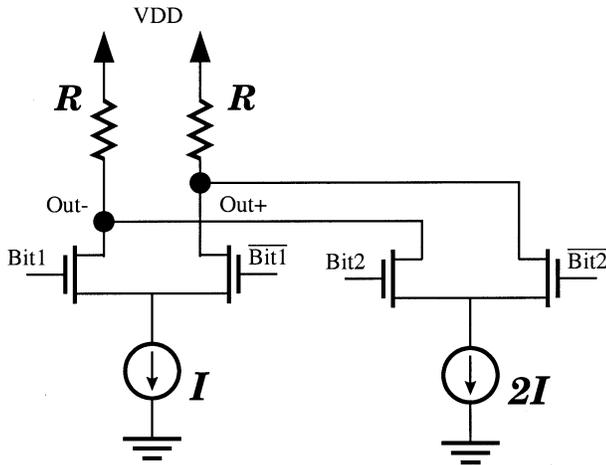


Fig. 2. Common 4-PAM driver architecture.

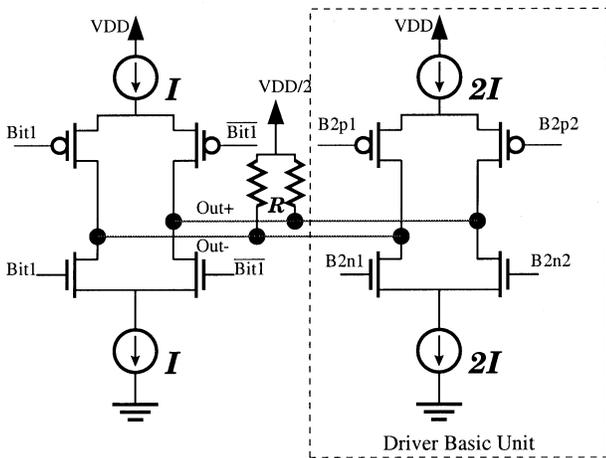


Fig. 3. Basic architecture of the power-efficient 4-PAM driver.

### III. POWER-EFFICIENT DRIVER TOPOLOGY

Fig. 3 shows the basic architecture of the proposed 4-PAM driver. As shown in this figure, this driver uses a bipolar topology to reduce the power. In this architecture, the driver is composed of two basic units. To further reduce the power, the right unit can be turned off whenever  $\pm RI$  are to be transmitted. The top current source can be turned off by pulling up both  $B2p1$  and  $B2p2$  inputs, while for turning off the bottom current source,  $B2n1$  and  $B2n2$  signals should be pulled down. Another advantage of this architecture is its modularity. This 4-PAM topology can be easily changed to 6-PAM by just adding another basic unit.

While the architecture in Fig. 3 significantly reduces the power, switching current sources reduce the maximum operating speed of the driver since current sources need some time to settle at the switching time. A data-look-ahead technique is used to overcome this problem. As shown in Fig. 4, four branches in each unit are used to pre-switch the current sources. This increases the achievable data rate of the driver. The mechanism can be described with an example. Assume both current sources in Fig. 4 are off and signal  $B2n1$  is about to go from low to high. Since signal  $B2n1$  is the output of an inverter whose input signal is  $B2n1a$ , signal  $B2n1a$  goes low before

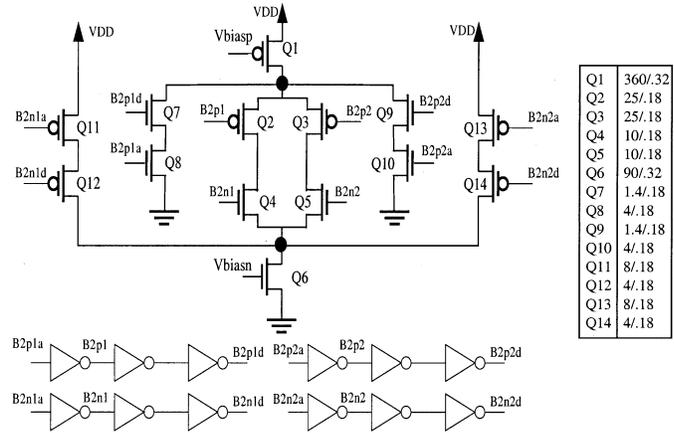


Fig. 4. Detail of each driver basic unit.

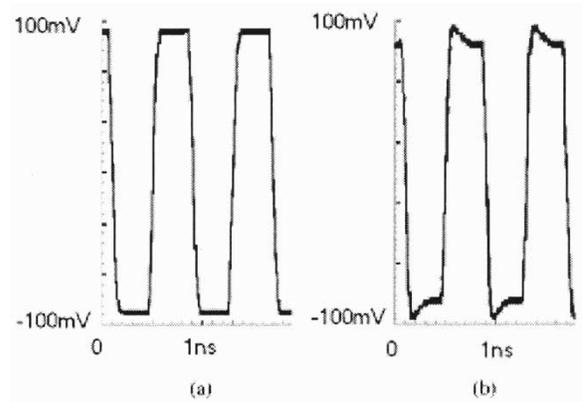


Fig. 5. Output of a two-level bipolar driver. (a) With pre-switching. (b) Without pre-switching.

the signal  $B2n1$  goes high (one inverter delay, roughly 40 ps). Since the signal  $B2n1d$  is the delayed version of signal  $B2n1$  (two inverter delays, roughly 80ps), at the transition of  $B2n1a$ ,  $B2n1d$  is still low. This turns on transistors  $Q11$  and  $Q12$ , which then turns on transistor  $Q6$ . By the time the signal  $B2n1$  goes from low to high, after one inverter delay, the current in transistor  $Q6$  has settled down. In other words, the current source is turned on slightly before the transition of signal  $B2n1$ . Two inverter delays after the transition of signal  $B2n1$ , signal  $B2n1d$  goes high and turns off the  $Q11$ – $Q12$  branch.

Another advantage of this architecture is the fact that it eliminates the need for a pre-driver. The pre-driver's function is to switch the gate voltages on the two current-steering transistors [see, for example, Fig. 1(a)] in each driver segment in such a way that current steers smoothly from one output to the other. To achieve this, tail transistors should stay in the saturation region. Thus, driver inputs should switch between  $V_{DD}$  and a voltage slightly less than  $V_{tail} + V_{th}$  (where  $V_{th}$  is the threshold voltage of the transistor) [1]. However, without a pre-driver, the crossover voltage of the inputs of the driver,  $Bit1$  and  $\overline{Bit1}$ , cannot hold both steering devices on, and thus the tail transistor will fall out of saturation. This not only reduces the speed of the driver, but also creates overshoot and undershoot in the output. Fortunately, the proposed pre-switching technique also alleviates this problem and there is no need for the pre-driver. Fig. 5

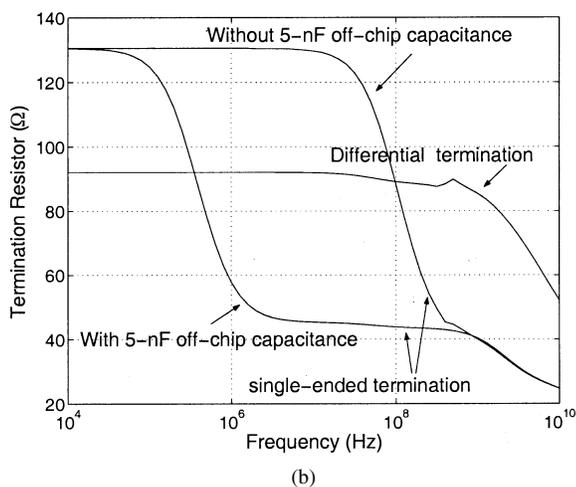
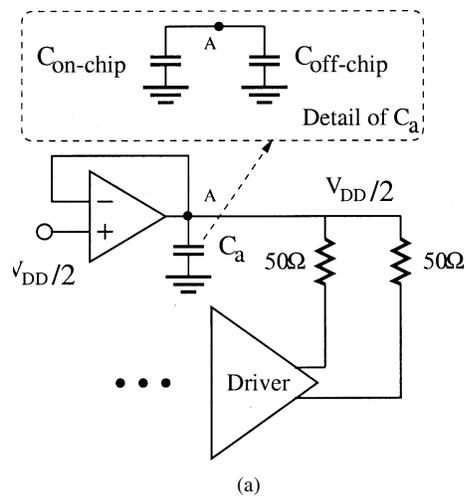


Fig. 6. Termination structure and its simulation results. (a) Termination architecture. (b) Termination resistance variation with frequency.

shows the output of a two-level bipolar driver, which is similar to the one in Fig. 1(b), in two different cases: with and without pre-switching technique. Eliminating the pre-driver from the transmitter block diagram can significantly reduce the transmitter power since a high-speed pre-driver is a power-hungry block.

#### IV. SIMULATION RESULTS

A 4-PAM transmitter is designed and implemented based on the proposed power-efficient driver architecture. A pseudo-random bit sequence (PRBS) unit consisting of two  $2^7-1$  PRBS generators produces the random data. Four encoders provide the necessary driver inputs at low speed. Six 4:1 multiplexers are used for serializing and making them ready for the drivers inputs, which all have CMOS logic values, at full speed.

The entire transmitter was simulated with HSPICE in a 0.18- $\mu\text{m}$  CMOS technology. The simulated driver power consumption is 10.5 mW and simulation results show that pre-switching circuitry consumes only 1 mW at 10 Gb/s. An important factor in transmitter design is the variation of differential and single-ended termination with frequency. Ideally, with perfect matching, only differential termination is important. However, due to the mismatch between the two

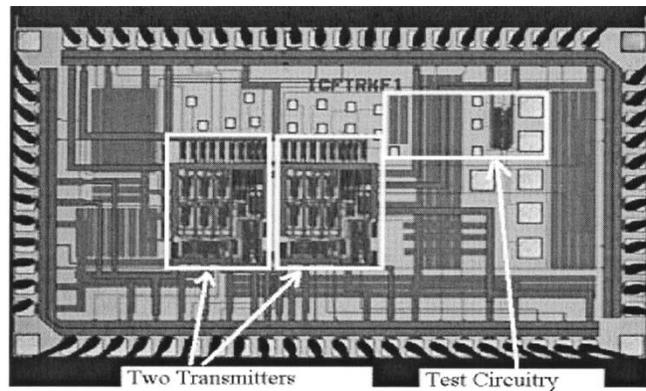


Fig. 7. Transmitter chip microphotograph.

differential outputs, single-ended termination is also important. Fig. 6(a) shows the general termination architecture. As shown in this figure, a simple on-chip buffer is used for generating the  $V_{DD}/2$  power supply. The current of this power supply is practically zero and the differential termination is always 100  $\Omega$  in parallel with the driver output impedance. Therefore, this structure is expected to show small differential termination variation with frequency. This is verified by HSPICE simulation [see Fig. 6(b)]. On the other hand, in this architecture, single-ended termination could vary significantly with frequency. At high frequencies, the capacitance  $C_a$  in Fig. 6(a) is short circuit, and therefore, the single-ended termination would be roughly 50  $\Omega$ . However, at low frequencies, there would be some impedances between node A and ground, and this would change the single-ended termination. This can be alleviated by increasing the capacitance at node A. In this design, capacitance  $C_a$  is composed of one on-chip and one off-chip capacitance. Fig. 6(b) shows the single-ended termination variation with frequency for a 10-pF on-chip capacitance with and without a 5-nF off-chip capacitor. As shown in Fig. 6(b), the single-ended termination is 50  $\Omega$  for a wide range of frequency when a 5-nF off-chip capacitor is used. Fortunately, increasing the capacitance at node A also decouples the noise of the  $V_{DD}/2$  reference voltage.

#### V. EXPERIMENTAL RESULTS

The four-level PAM transmitter is implemented and fabricated in a 0.18- $\mu\text{m}$  standard digital CMOS technology, and an 80-pin ceramic flat package (CFP80) is used for packaging. Fig. 7 shows the 2.6 mm  $\times$  1.5 mm chip micrograph. As shown in this figure, there are two transmitters and some test circuitry in the chip, and the design is pad-limited. Each transmitter occupies 0.16 mm<sup>2</sup>. The entire transmitter draws 39 mA from a 1.7-V power supply. The driver and  $V_{DD}/2$  reference generator consume roughly 12.5 mW at 7 Gb/s, which is the lowest reported power at this speed. Fig. 8 shows the eye at 7 Gb/s at the output of the transmitter. The transmitter has a random output jitter of 22 ps (peak to peak) at 7 Gb/s. A better jitter benchmark for multilevel signaling is the eye opening. This transmitter has a maximum eye height of 140 mV and an eye width of 200 ps over 0.8-m cable and 3-cm printed circuit board (PCB) traces at 7 Gb/s. An eye opening of 200 ps at the 3.5 GS/s rate

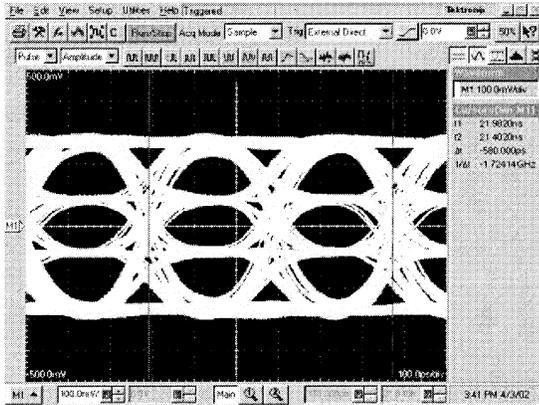
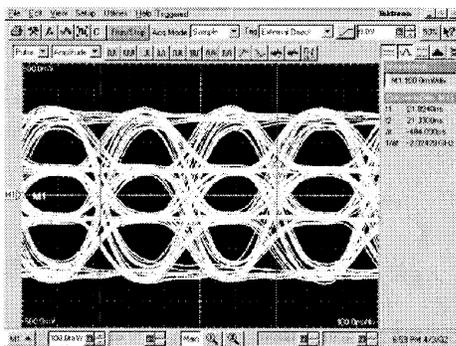
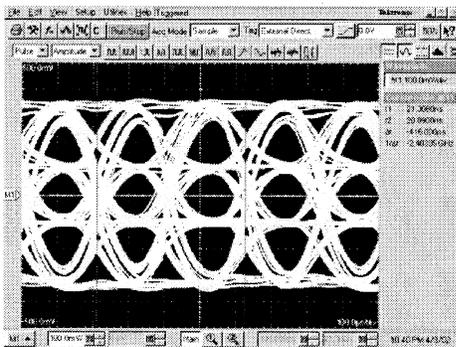


Fig. 8. Eye diagram at 7 Gb/s over 0.8-m cable and 3-cm PCB channel.



(a)



(b)

Fig. 9. Eye diagram at (a) 8 Gb/s and (b) 10 Gb/s over 0.8-m cable and 3-cm PCB channel.

corresponds to a 70% eye opening. By increasing the speed from 7 Gb/s toward 10 Gb/s with a 1.7-V power supply, the eye opening gets smaller. This problem has been resolved by increasing the power supply voltage from 1.7 to 2 V and increasing the input reference current by 25%. Fig. 9(a) shows the eye-diagram at 8 Gb/s. As shown in this figure, the eye has a maximum eye height of 140 mV and an eye width of 160 ps at 8 Gb/s. Although the eye diagram at 10 Gb/s [Fig. 9(b)] is open, the duty cycle of the clock is not 50% and this produces eyes with different widths. This shows the importance of the clock duty cycle at high speed. To solve this problem, duty-cycle-correction circuits can be used [7]. An alternative scheme is to use a delay-locked loop (DLL) to generate four different quadrature

TABLE I  
TEST RESULT SUMMARY

Specification	This design (7Gb/s)	This design (10Gb/s)	design in [6]	design in [5]	design in [2]
Driver Power	12.5mW*	20mW*	220mW	-	-
Total power	66mW	120mW	1W**	1.5W**	400mW**
Data rate	7Gb/s	10Gb/s	8Gb/s	10Gb/s	1.3Gb/s
Max. Swing (p-p)	600mV	600mV	2V	2V	1.1V
Power supply	1.7V	2V	3V	3.3V	3.3V
Technology	0.18 $\mu$ m	0.18 $\mu$ m	0.3 $\mu$ m	0.4 $\mu$ m	0.5 $\mu$ m

\*This includes the power consumption of  $V_{DD}/2$  reference generator.

\*\*This value shows the total power of the transmitter and receiver.

clocks. The small performance degradation of this transmitter due to clock duty-cycle imbalance is not related to the proposed power-efficient architecture.

Table I summarizes the test results along with the results of other state-of-the-art designs. For 10-Gb/s measurement, since power supply and the input reference current have been increased by 20% and 25%, respectively, the power dissipation is expected to increase roughly by a factor of 1.5. This is reasonably confirmed by the experimental results in Table I, bearing in mind that the dynamic power dissipation also increases with speed.

## VI. CONCLUSION

A novel power-efficient bipolar architecture for a multilevel PAM transmitter is presented. This architecture reduces the driver power by employing a bipolar architecture, reducing current when transmitting small voltage levels, and eliminating the need for pre-driver. A data-look-ahead technique is used for high-speed implementation of this architecture. Moreover, a 4-PAM transmitter based on this architecture is fabricated in 0.18- $\mu$ m digital CMOS technology. The transmitter achieves 3.5 GS/s (7 Gb/s) with a 1.7-V power supply and 5 GS/s (10 Gb/s) with a 2-V power supply. The driver draws 6.7 mA from a 1.7-V supply at 7 Gb/s and 10 mA from a 2-V power supply at 10 Gb/s.

## REFERENCES

- [1] W. J. Dally and J. Poulton, *Digital System Engineering*. London, U.K.: Cambridge Univ. Press, 1998.
- [2] D. J. Foley and M. P. Flynn, "A low-power 8-PAM serial transceiver," *IEEE J. Solid-State Circuits*, vol. 37, pp. 310–316, Mar. 2002.
- [3] J. L. Zerbe *et al.*, "1.6 Gb/s/pin 4-PAM signaling and circuits for a multi-drop bus," in *Proc. IEEE VLSI Symp. Circuits*, June 2000, pp. 128–131.
- [4] —, "A 2 Gb/s/pin 4-PAM parallel bus interface with transmit crosstalk cancellation, equalization, and integrating receiver," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2001, pp. 66–67.
- [5] R. Farjad-Rad, C.-K. K. Yang, M. A. Horowitz, and T. H. Lee, "A 0.4  $\mu$ m CMOS 10-Gb/s 4-PAM pre-emphasis serial link," *IEEE J. Solid-State Circuits*, vol. 34, pp. 580–585, May 1999.
- [6] R. Farjad-Rad, C.-K. K. Yang, and M. A. Horowitz, "A 0.3  $\mu$ m CMOS 8-Gb/s 4-PAM serial link transceiver," *IEEE J. Solid-State Circuits*, vol. 35, pp. 757–764, May 2000.
- [7] S.-J. Jang, Y.-H. Jun, J.-G. Lee, and B.-S. Kon, "ASMD with duty cycle correction scheme for high-speed DRAM," *Electron. Lett.*, vol. 37, pp. 845–847, Aug. 2001.