

A 10-mW 3.6-Gbps I/O Transmitter

Hamid Hatamkhani¹, Koon-Lun Jackie Wong¹, Robert Drost², Chih-Kong Ken Yang¹

¹University of California, Los Angeles, ²Sun Microsystems

9 Marisol, Newport Coast, CA, 92657, USA

Tel: 1-949-678-3993, E-mail: hhatam@icsl.ucla.edu

Abstract

This paper describes a low-power self-terminated transmitter. A novel architecture is proposed to perform impedance matching and channel equalization with low power consumption. The test chip is fabricated using 0.18- μm digital CMOS process with 1.8-V supply. The transmitter operates at 3.6Gbps and consumes 9.66mW. The total transmitter area is 0.072 mm².

Introduction

Power dissipation by off-chip I/O is an increasing concern; bandwidth of ICs is anticipated to exceed 1Tb/s within the next 10 years. In order to maintain reasonable power consumption, the signaling and the transmitter architecture must be adapted toward low power dissipation and lower power supplies. With higher data rates per I/O port, the design must also provide good signal integrity. Recent designs have demonstrated multi-Gb/s transmitters with minimum power dissipation of 7.5mW/Gbps [1], [2].

This paper describes a very low-power design for a 3.6-Gb/s transmitter for a widely parallel I/O targeting <10mW (2.7mW/Gbps) per port. The design uses low signal swing and low common mode voltage. For good signal integrity, the source termination is built-in along with the slew-rate control and signal pre-distortion to equalize the channel.

Transmitter Design

The transmitter architecture is shown in Fig. 1. The input data is at 225 Mbps. A 16:1 multiplexer serializes the input data to achieve an output at 3.6Gbps. The 16:1 multiplexer is a binary tree of 2:1 multiplexers. The pre-driver buffers the multiplexer output for the output driver. The clock signals that drive the multiplexers are 1.8GHz, 900MHz, 450MHz and 225MHz. Each clock is divided down from a 1.8-GHz PLL clock.

To minimize power, the design uses low peak-to-peak swing of 250mV and a low common-mode voltage of 250mV. A separate signal supply voltage, V_s , is used and the channel is terminated to $V_s/2$. With impedance matching of 50 Ω to the source and load, the signal supply is chosen to be 500mV.

With such low voltages, the design opts for voltage-mode signaling. The driver is push-pull with the transistors in triode region to provide the source impedance. The signaling scheme dissipates only 0.75mW.

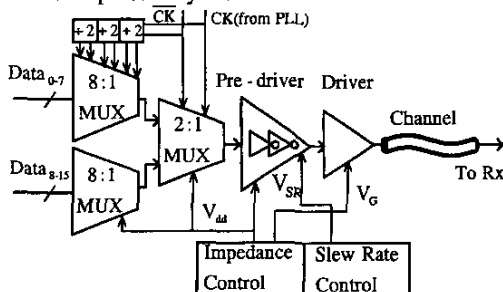


Fig. 1: Transmitter block diagram

The swing at the output of the driver is determined by the attenuation of the channel and the sensitivity of the receiver. In this design the target sensitivity of the receiver is <35mV. The pre-distortion design handles 14dB loss of 6.5 meters of RG55 or 80 centimeters of FR-4 20 mils PCB at 1.8GHz.

NMOS devices are used for both pull-up and pull-down to minimize device size and capacitance. The source impedance is adjusted by adapting the pre-driver supply, PV_{dd} (Fig. 2(a)). Using the same PV_{dd} for both up and down paths yields different impedances. The design degenerates the pull-down path with an extra transistor, M_1 . The gate of M_1 is connected to a voltage V_G that independently adapts the pull-down resistance. Figure 2(b) shows the impedance controller, which consists of 2 loops for the up and down impedances. Note that only the loop generating PV_{dd} needs to source current. A linear regulator is currently used but a switching regulator would normally be used to minimize power dissipation. The control loop uses a known impedance (50 Ω) as reference. The control voltages are shared among all I/O ports.

Since PV_{dd} tracks the speed of the process, it is conveniently used as the supply voltage of the remaining digital logic. Thus the logic is regulated to operate at a fixed speed regardless of the process corner. At a fast corner, significant power can be saved.

This work introduces a novel 2-tap pre-distortion filter for a voltage-mode driver without sacrificing the output impedance matching. The high-pass filter is:

$$Y[n] = X[n] - \alpha X[n-1] \quad (1)$$

Increasing the number of taps has diminishing returns for power dissipation at short channel lengths [1]. To implement $X[n-1]$, the data is delayed by a half-cycle.

The output driver devices are split into binary weighted transistors (Fig. 3). Depending on the coefficient of the high-pass filter, the inputs to each of the pull-up device are either $X[n]$ or $\bar{X}[n-1]$. Four bits, b_{0-3} , determine the filter coefficient. A series of switches assign the appropriate input to each branch.

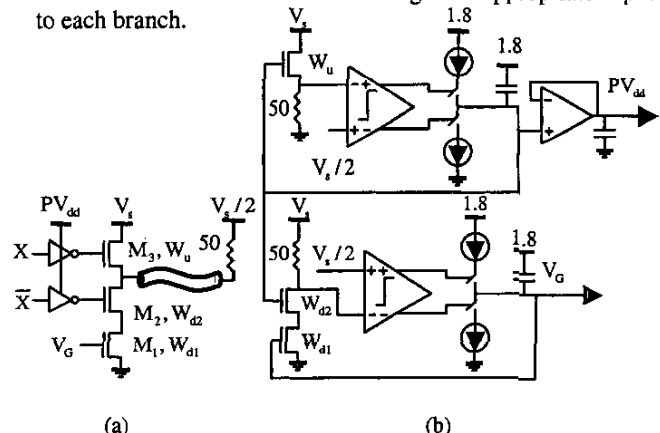


Fig. 2: (a) Output driver (b) Dummy circuit for impedance matching




Fig. 3: Output driver to be used for matching and equalization

For instance, with $\alpha=3/15$, then $b_3b_2b_1b_0=0011$ and therefore $X_3=X_2=X[n]$ and $X_1=X_0=\overline{X[n-1]}$. Assuming $X[n]=X[n-1]=1$, the output voltage is

$$V_{out} = \frac{15R/3}{15R/3 + 15R/12} V_s = \frac{4 \times V_s}{5} = (1-\alpha) \times V_s \quad (2)$$

Where R is the impedance of all branches in parallel (50Ω). The output impedance is: $(\frac{15}{12}R) \parallel (\frac{15}{3}R) = R = 50$

Note that the output impedance does not change with α . Intuitively, the design always connects each driver branch to either V_s or Gnd thus the total parallel resistance is always 50Ω. The additional power is only the cost of a half-cycle delay and the selection switches.

To improve signal integrity and minimize ISI and SSI, the design limits the output slew rate to roughly $1/3 T_{bit}$. The output slew rate is controlled by limiting the pre-driver's slew-rate (Fig. 4(a)). Since the gate capacitance of the driver is quite process independent, the pre-driver's slew rate is controlled by the drive resistance. Since PV_{dd} is from the impedance control, the NMOS resistance of the pre-driver is fixed. Simply sizing the pull-down adjusts the pre-driver's falling slew-rate. A separate loop controls the pull-up slew rate. The PMOS has 2 devices in series: the top device has a control voltage that limits the pull-up resistance. The sizing significantly favors the top device to minimize the capacitance in the signal path. The slew-rate control circuit (Fig. 4(b)) uses a replica of the pre-driver ($M_1 - M_3$). The loop turns on both NMOS and PMOS of the dummy pre-driver and adjusts V_{SR} until the output is $PV_{dd}/2$ (equal up and down resistance).

Measurement Results

A transmitter with area 0.6×0.12 mm is fabricated in a 0.18- μ m CMOS process. There are eight transmitters in the chip, which shares only one impedance and slew rate controller. The core transmitter power consumption is 9.66mW. The power breakdown is shown in Table I. Fig. 6 shows the output data-eye after passing through a 6.5-m cable, with and without equalization. Fig. 5 shows the output impedance for various high-pass filter coefficients. It can be inferred from the figure that the output impedance in this design is constant with different filter coefficients.

The authors thank National Semiconductor for fabrication and funding by UC-Micro:01-102.




Fig. 4: (a) Changing the last pre-driver inverter for Slew rate control (b) dummy circuit that equalizes the NMOS and PMOS speed

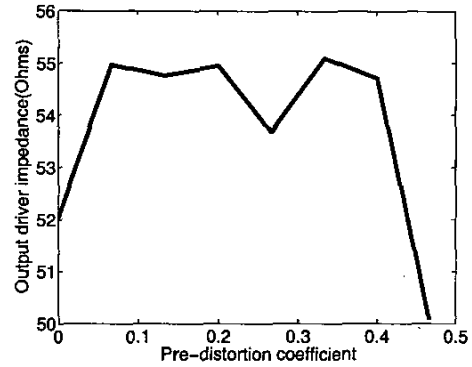


Fig. 5: Static output impedance for the driver with various pre-distortion coefficients.




Fig. 6: Data eye after passing through the cable at 3.6 Gb/s (a) without Pre-distortion (b) with Pre-distortion. $\alpha=1/3$




Fig. 7: Chip microphotograph

Table I
Measurement results at 3.6Gbps ($T_{bit}=278$ ps)

| Signaling Power(mW) | 0.75 |
|-----------------------------------------------|------|
| Impedance controller(per I/O) | 1.1 |
| Slew rate control(per I/O) | 0.1 |
| 8:1 Multiplexers | 0.4 |
| 2:1 MUX and Pre-driver | 5.91 |
| Frequency dividers and on-chip PRBS generator | 1.4 |
| Total(mW) | 9.66 |
| Output data rise time(ps) | 112 |
| Output data fall time(ps) | 107 |

References

- [1] M.-J. E. Lee, W. J. Dally, and P. Chiang, "Low-power, area efficient, high speed I/O circuit techniques," *IEEE J. Solid-State Circuits*, vol. 35, pp. 1591-1599, Nov. 2000.
- [2] F. Yang, J. H. O'Neill, D. Inglis, and J. Othmer, "A CMOS low-power multiple 2.5-3.125-Gb/s serial link macrocell for high IO bandwidth network" *IEEE J. Solid-State Circuits*, vol. 37, pp. 1813-1821, Dec. 2002.

A 27-mW 3.6-Gb/s I/O Transceiver

Koon Lun Jackie Wong, Mozghan Mansuri, Hamid Hatamkhani, and Chih-Kong Ken Yang

University of California, Los Angeles

56-147A Engineering IV Building, Box 951594

Los Angeles, CA 90095-1594, USA

Tel: 310.206.3665 Fax: 310.206.8495 E-mail: jwong@icsl.ucla.edu

Abstract

This paper describes a 3.6-Gbps 27-mW transceiver for chip-to-chip applications. A novel data receiving and timing recovery technique are presented with very low power penalties while maintaining high signal integrity. The input comparator filters noise with built-in bandwidth control and digital offset compensation while consuming 300uW. Static phase offset introduced onto the charge-pump permits phase recovery with no additional power. The entire design occupies 0.2mm² in a 0.18-um 1.8-V CMOS technology.

Keywords: Transceiver, I/O, Receiver, Data recovery, Comparator, and Low-power.

I: Introduction

Aggregate bandwidth from a large IC is anticipated to exceed 1Tb/s within the next 10 years. Widely parallel, multi-Gb/s chip-to-chip I/O links are an integral part of these systems. Power consumption of these links is an increasing concern. Recently demonstrated links have shown power dissipation of 20-40mW/Gbps [1-3]. This work demonstrates a scalable design capable of 7.5mW/Gbps in a 0.18-um CMOS technology. The transceiver operates at 3.6Gbps per port.

Section II describes the system and signaling architecture of the transceiver. The details of the transmitter are described in [4]. Section III and IV describe the design of a low power receiver and a novel timing recovery technique, respectively. Section V summarizes the measurement results from an 8 channel test chip.

II: Transceiver Architecture

The transceiver shown in Fig 1 is designed for widely parallelized half-duplex I/O. Each I/O cell can achieve 3.6-Gbps data rate by using low swing push-pull voltage-mode signaling. A separate supply is used for the low-swing driver. For transmitting through a medium with up to 12dB of loss, the output swing is chosen to be 0.25V to accommodate a receiver resolution of 35mV. The driver is series terminated at the transmitter and parallel terminated at the receiver. With this signaling scheme, the driver supply is chosen to be 0.5V.

To reduce power and area, transmitter (Tx) and receiver (Rx) share the same phase-locked loop (PLL). A 450-MHz reference clock (CKref) is distributed into each I/O cell by using a low-jitter clock distribution technique. The clock frequency is multiplied by 4 with a low-power low-jitter PLL [7] and the multi-phased outputs of the PLL are used for data recovery and data transmission. The data input are mesochronous in that each of the ports has the same data rate but can significantly vary in phase.

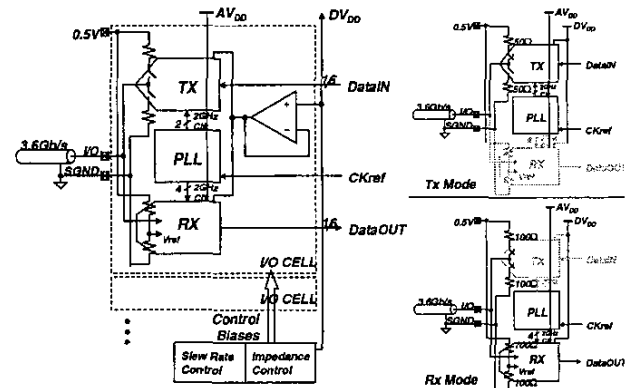


Fig. 1 Transceiver Architecture

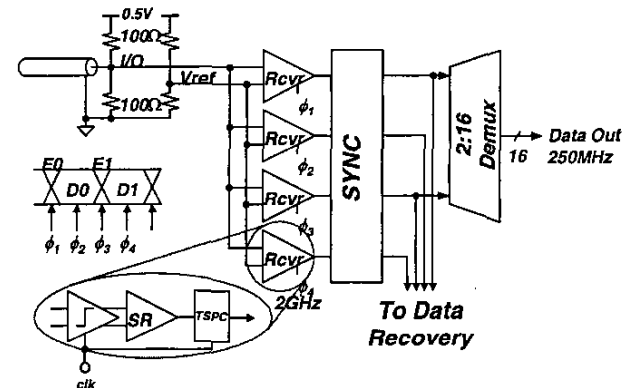


Fig. 2: Block diagram of receiver

In the right side of Fig 1, the two operation modes, transmit mode and receive mode, of the transceiver are shown. In transmit mode, Rx is disabled. To reduce the signal reflection at the end of the channel, the driver is 50Ω source terminated in both pull up and pull down direction to ensure impedance matching with the channel. In receive mode, the Tx is disabled. Both pull up and pull down paths are turned on to form a voltage divider with 50Ω input impedance. Consequently, the inputs of the receiver are biased at 0.25VDC. Because the termination of the receiver is fully differential, noise that is common to input and V_{ref} will not affect the performance. The transmitter design can be found in [4]. The driver's resistive switches are implemented using only NMOS devices with their gate voltage adaptively controlled by a voltage regulator. The regulator biases the resistance to be 50Ω across process corners. Since the gate voltage tracks the process variation, it is conveniently used as the supply voltage of the remaining digital logic. The logic is then regulated to operate with no process, voltage, or temperature (PVT) dependence.