

Low-Power Design Methodology for an On-chip Bus with Adaptive Bandwidth Capability

Rizwan Bashirullah

Department of Electrical
and Computer Engineering
North Carolina State University
Raleigh, NC, 27606

rbashir@ncsu.edu

Wentai Liu*

Department of Engineering
University of California
1156 High Street
Santa Cruz, CA, 95064-1077

wentai@soe.ucsc.edu

Ralph K. Cavin

Semiconductor Research
Corporation
Research Triangle Park
NC, 27709

Ralph.Cavin@src.org

ABSTRACT

This paper describes a low-power design methodology for a bus architecture based on hybrid current/voltage mode signaling for deep sub-micrometer on-chip interconnects that achieves high data transmission rates while minimizing the number of repeaters by nearly 1/3. The technique uses low-impedance current-mode sensing to increase the data throughput and minimizes the static power dissipation inherent to current-mode signaling by adaptively changing the interconnection bandwidth given a change in input signal activity. Since bandwidth is related to power dissipation, the adaptive bus attains energy efficient data transmission by expending minimum power required to support the bus signal activity.

The design method is based on statistical analysis of address streams extracted for typical benchmark programs using a microprocessor time-based simulator in combination with circuit-level power analysis. Simulation results indicate improvements in power dissipation of up to 65% and 40% over current and voltage mode signaling schemes, respectively.

Categories and Subject Descriptors

B.4.3 [Input/Output and Data Communications]: Interconnections (Subsystems) - *Topology (e.g., bus, point-to-point)*.

General Terms

Performance and Design.

Keywords

Bus, low-power, current-mode, delay, point-to-point, on-chip interconnect.

1. INTRODUCTION

Achieving low propagation delays and high signaling bandwidth in on-chip global interconnects is essential to high-performance microprocessors and embedded systems, an increasingly

*On leave from ECE Dept., North Carolina State University.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

DAC 2003, June 2-6, 2003, Anaheim, California, USA.

Copyright 2003 ACM 1-58113-688-9/03/0006...\$5.00.

challenging task given a 0.7X reverse-interconnect scaling trend, a 14% increase in die size, and doubling of clock operating frequency per technology node [1]. In order to achieve low latency and higher throughput data transfers within computational units on-die, repeaters are systematically inserted in long global busses [2], [3]. Often, however, repeater insertion cannot be achieved due to placement blockages caused by underlying critical processing units. In addition, as the required repeater insertion distance decreases with each technology node due to increased interconnect resistive effects, the overall improvement in delay and bandwidth may be undermined by the exponential increase in the number of repeaters on-die and associated driver/repeater power dissipation [1].

In this paper, we propose an on-chip bus architecture based on hybrid current/voltage mode repeaters to address signal latency and throughput while minimizing the number of repeaters required to achieve these goals. Since reducing the number of repeaters results in fewer placement blockages due to underlying logic, improved design implementation flexibility can be achieved. To compensate for the increase in static power dissipation of current sensing techniques [4], a novel adaptive bus technique is proposed. The adaptive bus is designed to automatically increase or decrease the interconnection bandwidth given a change in bus signal activity. Since bandwidth is related to power, adaptively changing the bandwidth of the interconnects minimizes the overall power dissipation of the bus. Thus, the hybrid current/voltage mode repeater bus operates in current-mode when the signal activity and the required bandwidth is high and shifts to voltage-mode operation as the data activity and the required bandwidth decreases.

To demonstrate the performance gains of the bus architecture, a design methodology based on circuit-level power dissipation characterization and statistical analysis is described. Address streams extracted from typical program benchmarks using an Alpha 21264 time-based simulator are used to obtain probabilities of bit transitions as well as the probability of the number of cycles bit patterns remain unchanged. The rationale for this is that the number of cycles before each transition occurs determines the probability that the bus will operate in current or voltage mode. This information is used to estimate the power dissipation of the bus.

This paper is organized as follows. Section 2 provides a brief overview of current-mode signaling and theoretical models for delay, throughput and power dissipation are presented. In section 3, the proposed adaptive bus concept and architecture is described, focusing on circuit design implementation. Section 4 deals with the design methodology technique used to estimate and

optimize the power dissipation of the adaptive bus. Performance results are discussed in section 5, with concluding remarks presented in section 6.

2. BACKGROUND

2.1. Current-mode Signaling

The key to current-mode signal transporting is the shift in pole position and reduction of the system time constants that result from sensing signals with low impedance nodes [4], [5]. Hence, from hereon after, for the purpose of signaling in on-chip interconnects, *current-mode or current sensing refers to sensing a signal with a low impedance termination at the receive-end which results in a shift in pole position thereby increasing the bandwidth of the line*. To account for the change in system time constants due to the impedance termination of the line, a resistor R_L is added to the receiver, as shown in Fig. 1. If we assume that the driver and interconnect parameters are unchanged, the parallel termination R_L determines the impedance of the receiver and hence the current or voltage mode operation of the line.

2.2. Delay, Throughput and Power Equations

Simple yet accurate closed-form expressions of delay and power dissipation for current-mode (CM) and voltage-mode (VM) signaling have been reported in [5]. In this work, the formulations are extended to take into account the effect of driver source capacitance (C_S),

$$t_v = R_T C_T \frac{\ln \left[\frac{1.2 + 0.5(\eta_L / R_{LT})}{1 - v} \right]}{1 + 1.0058 \ln \left[1.2 + 0.5 \left(\frac{\eta_L}{R_{LT}} \right) \right]} \times \left\{ \frac{\eta_L}{2} \left(1 + \frac{1}{3R_{LT}} \right) + \eta_L R_{ST} \left(1 + \frac{1}{2R_{LT}} \right) + \eta_S C_{ST} (1 + R_{LT}) + \eta_L C_{LT} (1 + R_{ST}) \right\} \quad (1)$$

The delay (t_v) is defined as the time from ($t=0$) to the time when the normalized voltage reaches v at the end of the line. R_T and C_T are the total interconnect resistance and capacitance; $R_S(R_L)$ and $C_S(C_L)$ are the source(load) resistance and capacitance, respectively; $R_{LT}=R_L/R_T$, $R_{ST}=R_S/R_T$, $C_{ST}=C_S/C_T$ and $C_{LT}=C_L/C_T$. $\eta_L=R_L/(R_L+R_S+R_T)$ and $\eta_S=R_S/(R_L+R_S+R_T)$ are defined as voltage loss factors of the load and source, respectively.

The maximum NRZ data rate that can be supported by the line can be expressed as,

$$f_{\max} = \frac{1}{t_{90}} \quad (2)$$

where t_{90} is the 0-90% delay from (1). In similar manner, closed form expressions for dynamic and static power can be written as [5],

$$P_{\text{dyn}} = (\eta_L V_{dd})^2 f \cdot \text{act} \left\{ C_L + C_S \left[1 + \frac{2R_T}{R_L} + \left(\frac{R_T}{R_L} \right)^2 \right] + C_T \left[1 + \frac{R_T}{R_L} + \frac{1}{3} \left(\frac{R_T}{R_L} \right)^2 \right] \right\} \quad (3.a)$$

$$P_{\text{static}} = \frac{\eta_L V_{dd}^2}{R_L} \quad (3.b)$$




Fig. 1. Inverter driven interconnect model with arbitrary receive-end termination for current or voltage mode signaling.

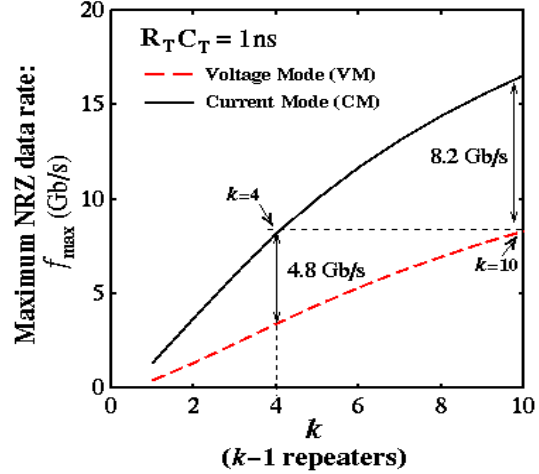


Fig. 2. Data rate comparison for current and voltage mode repeater insertion interconnects with optimally sized drivers.

In (3.a), act is the activity factor. Equations (1)-(3) are useful to determine performance trade-offs between voltage-mode (i.e. $R_L=\infty$) and current-mode (i.e. $R_L\ll\infty$) interconnects. For instance, for given values of R_S , C_S , R_T , C_T and C_L , the maximum NRZ data-rate (f_{\max}) increases significantly as R_L is reduced. As shown in Fig. 2, the improvement in f_{\max} using CM sensing schemes is apparent, achieving target data-rates with nearly 1/3 the number of VM repeaters. For the design example shown in Fig. 2, 3 CM repeaters achieves nearly 4.8Gb/s more NRZ bandwidth than 3 VM repeaters, and exhibits the same data rate performance as 9 VM repeaters.

3. ADAPTIVE BUS ARCHITECTURE

The architecture of the adaptive bus is shown in Fig. 3. It consists of a small FIFO of depth C_p+1 clock cycles, a digital transition detector, a control line and the hybrid voltage/current mode repeaters. The input to the control line (C_{in}) sets the operation of the hybrid repeaters in either voltage or current-mode. In the event of input data transitions ($D_{in}[0]$, $D_{in}[1]$, ..., $D_{in}[N]$), the transition detectors activate the control line to set the bus lines in CM operation mode. Similarly, in the absence of data transitions, the bus lines are set to VM operation mode. Specifically, if the data $D_{in}[0:N]$ does not change for C_p clock cycles, the bus lines automatically shift to VM operation to reduce the static power dissipation. In order to minimize circuit overhead, each control line is shared among $(N+1)$ bus lines.

Fig. 4a shows the hybrid voltage/current mode repeater. The operation is described as follows. When the control voltage (V_{ctrl}) of the input stage is below the threshold voltage of the feedback transistor, the repeater operates as a regular full-swing voltage-mode inverter. As V_{ctrl} increases, the feedback transistor turns on and the repeater operates as a self-biased inverter. The termination (R_L) looking into the repeater decreases as V_{ctrl} increases, thereby

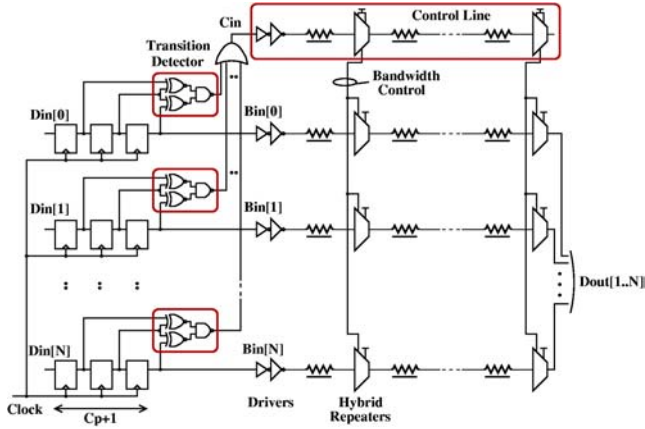


Fig. 3. Architecture of the adaptive bus.




Fig. 4. Hybrid current/voltage mode line interface repeater, (a) circuit schematic; (b) termination resistance (R_L) and interconnection bandwidth vs. V_{ctrl} .




Fig. 5. Timing of the adaptive bus

shifting the pole frequency of the interconnect line which has the effect of increasing the bandwidth, as illustrated in Fig. 4b.

The bus operation in each clock cycle for an arbitrary input data sequence is shown in Fig. 5 (from hereon a clock cycle refers to the system sampling time). In this example, the data is sampled at both positive and negative edges of the clock. For simplicity, we assume that two bus lines $Din[0]$ and $Din[1]$ share the same control line $C0$. As shown in Fig. 5, the input data is delayed by C_p clock cycles to allow for the transition detectors and control line to update the repeater's mode of operation. The minimum required C_p is given by the overall processing delay of the path determined by the transition detectors and control line. Since the control line is identical to the bus lines and continuously operates in CM, only the first repeater of the bus lines needs to be updated before the delayed input data ($Bin[0:N]$) can be launched. As the control signal $C0$ propagates, it updates the subsequent repeater stages of the bus lines, similar to a domino effect. The importance of this is that the latency of the processing delay from CM-to-VM or vice-versa is significantly reduced. In Fig. 5, C_p is assumed to be two cycles long. On the falling edge of the control signal $C0$,

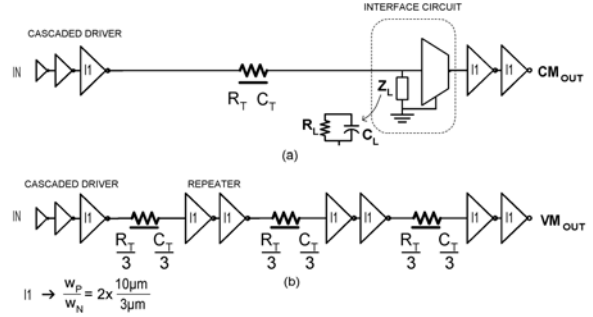


Fig. 6. SPICE simulation benchmark for power analysis, (a) current and (b) voltage mode. The design was based on TSMC 0.35 μ m parameters.




Fig. 7. Power dissipation comparison of current and voltage mode benchmarks depicted in Fig. 6.

the line switches to VM after approximately two cycle delays, indicated by the shaded regions. Notice that the data bus lines switch to CM operation whenever there is an input transition, and remains in VM operation in the absence of transitions for more than C_p cycles.

4. DESIGN METHODOLOGY

4.1. Circuit-level Power Modeling for Current and Voltage Mode Signaling

To evaluate the overall power dissipation performance of current and voltage mode signaling for on-chip interconnects, a circuit-level test benchmark designed in TSMC 0.35 μ m technology with $V_{dd}=3V$ was used, as shown in Fig. 6. The interconnect line is a metal-3 layer wire and metal-2 ground with a length of 1-cm, modeled by a 1000 segment distributed RC line. The resulting total resistance (R_T) and capacitance (C_T), including fringing capacitance, is given by 175 Ω and 2.56pF, respectively. To fairly compare the power dissipation performance of both schemes, we deliberately add inverters "I1" after the current-mode receiver interface circuit. The inverters are sized with $W_p=2 \times 10\mu m$ and $W_n=2 \times 3\mu m$ and minimum drawn length of $L=0.4\mu m$. The target maximum data rate was set at 1Gb/s (i.e. bit time $T_b=1ns$), which requires at least two VM repeaters, whereas no repeaters were required for CM signaling. The circuit topology of the CM receiver and CMOS level swing conversion circuit is shown in Fig. 4a.

Fig 7 shows the overall power dissipation performance of the test benchmark for several i , where i represents the number of cycles

in bit times (T_b) for which the logical level remains unchanged. At relatively large i , the VM line in Fig. 6b exhibits lower overall power dissipation than the CM line in Fig. 6a. This is due to the static power dissipation inherent to parallel resistive termination of CM signaling. However, as i is decreased, the dynamic power dissipation of full-swing VM signaling dominates. For this example, the crossover point occurs at approximately $i=2.5$ or $T_b=2.5\text{ns}$, which is equivalent to a bus frequency of 200MHz (i.e. $1/5\text{ns}$) – relatively small compared to current GHz processors. Notice that the slope at which the power dissipation increases is smaller for the CM signaling case, a result due to the reduced voltage swing in the interconnect line. It should be pointed out, that unlike low-swing VM signaling schemes [6], CM signaling reduces the voltage swing while enhancing the bandwidth of the line. The results depicted in Fig. 7 suggest that CM signaling is beneficial at higher signaling data-rates.

4.2. Bus Statistics

The purpose of the bus statistics analysis is to determine the probability of bit transitions as well as the probability of number of clock cycles that the bit patterns remain unchanged. Given this information, it is possible to infer the power dissipation of the adaptive bus lines.

We simulated an Alpha 21264 machine using SimpleScalar 2.0 [7] and modified the timing simulator “sim-outorder.c” module to extract instruction addresses. Three benchmarks from the SPECINT2000 test suite - MCF (Combinatorial Optimization), PARSER (word-processor) and GZIP (compression) – were used for the simulation results. A total of 100 million 32-bit instruction addresses were collected for each benchmark. The instruction addresses were divided into half-bytes (4-bits) and the number of clock cycles before each 4-bit pattern change was accumulated. The percentage of clock cycles of in-sequence half-bytes is shown in Fig. 8 for each benchmark. In Fig. 8, each bar is divided into 1, 2, 3, 4, 5 and greater than 5 clock cycle bins. For instance, “1” refers to the percentage of total simulated clock cycles in which a 4-bit pattern remains unchanged for 1 cycle; “2” refers to the percentage of total simulated clock cycles in which a 4-bit pattern remains unchanged for 2 cycles; and so forth. The results show a high correlation of switching activity for the lower order bits, whereas the higher order bits remain nearly unchanged for the entire instruction streams.

4.3. Power Estimation Methodology

Let P_{TN_i} denote the RMS power dissipation of N bus lines given that the bits remain unchanged for i clock cycles, and p_{rN_i} denote the probability defined as the percentage of total simulated clock cycles in which the N bus lines remain unchanged for i clock cycles (i.e. as depicted in Fig. 8). Since the adaptive bus operates in CM or VM, the overall power dissipation can be obtained by adding the fraction of power for which the bus operates in current-mode (P_{CM_N}) and the fraction of power for which the bus operates in voltage-mode (P_{VM_N}). Assuming that the adaptive bus requires C_p clock cycles to update the bus lines from CM to VM, the total power dissipation of N bus lines operating in CM when $i \leq C_p$ is,

$$P_{CM_N} = \sum_{i=1}^{C_p} p_{rN_i} P_{TN_i} \quad (4)$$

Similarly, the total power dissipation of N bus lines operating in VM when $i > C_p$ is,

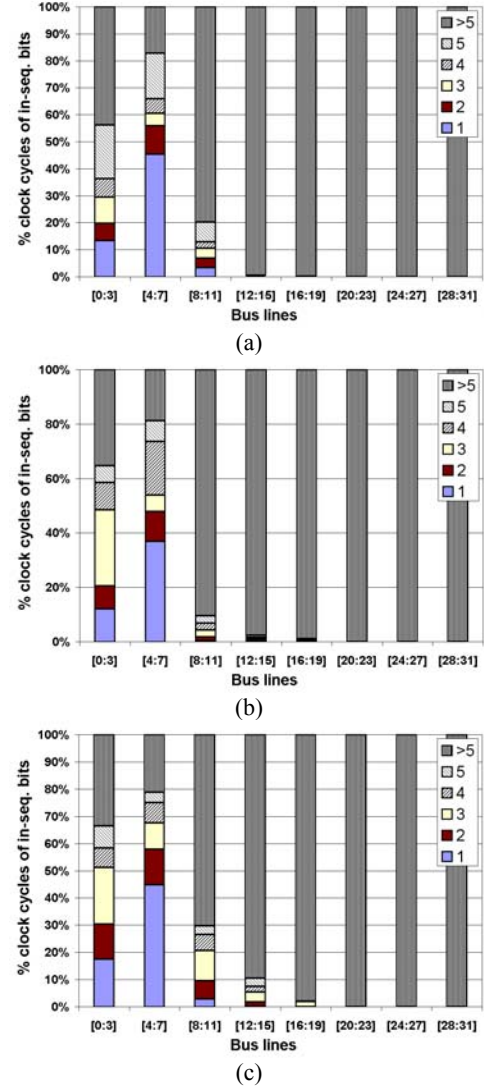


Fig. 8. Bus transition statistics per 4-bit bus lines and percentage of clock cycles each 4-bit pattern remains unchanged. Simulated benchmarks using SPEC2000 test suite (a) PARSER, (b) GZIP and (c) MCF.

$$P_{VM_N} = \sum_{i=C_p+1}^{N_c} \frac{C_p}{i} p_{rN_i} P_{TN_i} \quad (5)$$

where N_c is the total number of simulated clock cycles. Notice that in (5), P_{VM_N} is not assumed to be negligible even though the bus operates in VM. The reason for this is that the bus remains in CM for at least C_p cycles even after the switching to VM, due to the finite update time of C_p cycles. As a result, P_{TN_i} in (5) can be reduced to,

$$P_{TN_i} = P_{TN_{C_p}} \quad i > C_p \quad (6)$$

From (4)-(6), the total power dissipation of N bus lines can be rewritten as,

$$P_{total_N} = \sum_{i=1}^{C_p} p_{rN_i} P_{TN_i} + C_p \cdot P_{TN_{C_p}} \sum_{i=C_p+1}^{N_c} \frac{p_{rN_i}}{i} \quad (7)$$

In (7), p_{rN_i} and N_c are obtained from the simulated bus statistics (i.e. section 4), whereas P_{TN_i} can be extracted from SPICE simulations. Notice that by letting C_p approach N_c , equation (7)

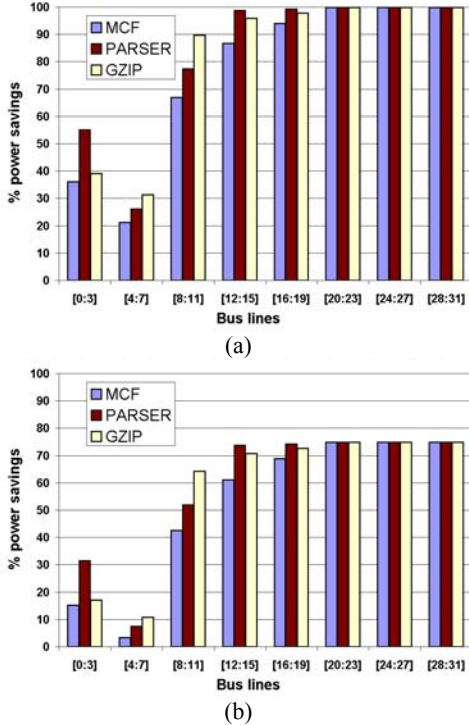


Fig. 9. Percent reduction in power dissipation of the adaptive hybrid current/voltage mode bus technique over current-mode bus. The adaptive bus uses 1 control line per 4 bus lines, (a) performance without power dissipation of control lines, (b) with control lines included.

can also be used to determine the power dissipation of the bus when operating entirely in current or voltage mode. For the purpose of comparison only, we assume that $P_{TNi} = N P_{Ti}$, where P_{Ti} is the power dissipation of a single bus line as depicted in Fig. 4. The definition of P_{TNi} is the worst-case power dissipation since it assumes that all bus lines transition simultaneously.

5. RESULTS

5.1. Power Savings

To verify the savings in power dissipation of the adaptive bus technique over a current-mode bus, results based on (7) for the benchmark tests simulated in section 4 are shown in Fig. 9. In this example, the adaptive bus uses one control line to update the state of four bus lines (i.e. total of 8 control lines for 32 bus lines). The control lines operate in current-mode and are assumed to be identical to the bus lines, with an update time latency of 3 cycles (i.e. $C_p=3$). Fig. 9a and 9b shows the power savings without and with the added power of the control lines, respectively, indicating that higher performance gains could be obtained by minimizing the total number of control lines. The mean power savings of all three simulated benchmarks including control lines is over 50%.

5.2. Bus Switching Activity and Control Line Design

The results depicted in Fig. 9 clearly indicate that the static power dissipation inherent to current sensing techniques - most dominant in bus lines with low switching activity - can be significantly minimized with the proposed adaptive bus. However, address

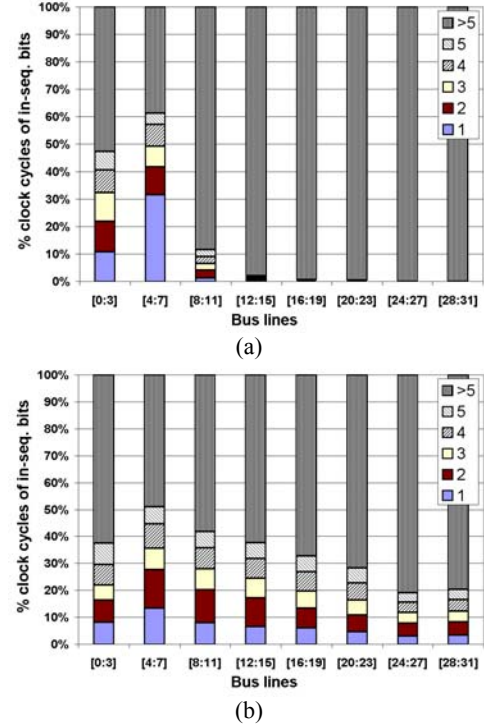


Fig. 10. GCC benchmark bus statistics for (a) instruction and (b) data address streams simulated for 100 million clock cycles.




Fig. 11. Control line design for 32-bit adaptive bus. Type-I uses 8 control lines (1 per 4 bus lines) and Type-II uses 2 control lines (1 per 16 bus lines).

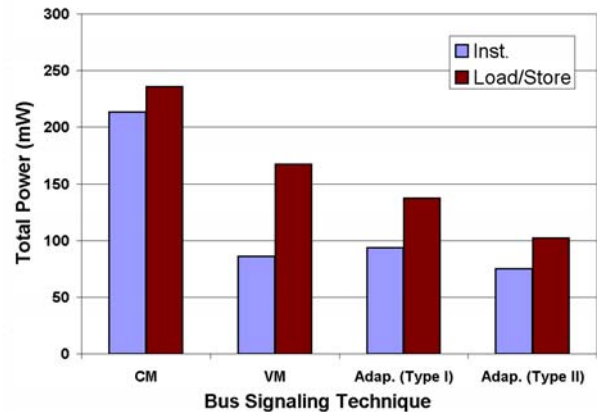


Fig. 12. Total power dissipation comparison for GCC benchmark

busses may also exhibit low probability of in-sequence address streams, as in the case of data addresses (i.e. load/stores). When the probability of sequential addresses is very low, the switching activity of the higher order bits in the bus lines increases. This behavior is illustrated in Fig. 10, where the percentage of clock cycles of in-sequence half-bytes for instruction and data addresses are shown for the GCC benchmark (i.e. C Programming Language Compiler). In Fig. 10a, the instruction addresses exhibit a high

correlation of switching activity for the lower order bits, which indicates a higher spatial locality amongst the address streams since instructions are usually stored in adjacent locations of memory. Conversely, data addresses exhibit a more uniform switching activity distribution within the bus lines – representative of a lower probability of in-sequence address streams.

To examine the effect of varying switching activity distribution within bus lines on power dissipation, the performance of the adaptive bus is compared against both VM and CM signaling schemes. In this example, two designs for control lines are also compared, as shown in Fig. 11. The type-I adaptive bus consists of 8 control lines, each one used to update the signaling state (i.e. CM or VM) of 4 bus lines. Alternatively, the type-II adaptive bus uses 2 control lines, each one updating the state of 16 bus lines. The main difference between the two control line design approaches, apart from the obvious reduction in the number of control lines, is that a type-II bus will shift from CM to VM only when all 16 bus lines remain inactive for more than C_p clock cycles, whereas in a type-I bus only 4 bus lines need to be inactive. Thus, the probability that a type-II adaptive bus will remain in CM operation for a longer fraction of total simulated clock cycles is likely to be higher than the type-I bus.

The overall power dissipation performance of a 32-bit wide bus for simulated statistics of the GCC benchmark is shown in Fig. 12. The following observations can be inferred from these results:

- 1) The CM bus exhibits the highest power dissipation; nearly 2.5 and 1.4 times higher than the VM bus for instruction and data addresses, respectively – this is due to static power dissipation of CM signaling. However, the relative change in power dissipation for instruction and data address streams is only 10% for the CM bus whereas the VM bus changes by 94%. This indicates that CM signaling is more suitable for increasing switching activity, an effect due to the reduction in voltage swings.
- 2) The type-II adaptive bus outperforms the type-I bus for both instruction and data address streams. In fact, the type-II bus remains in CM operation for a longer percentage of total simulated clock cycles – because the probability of all 16 bus lines remaining inactive is likely to be lower than 4 bus lines remaining inactive. However, there is an increase in power dissipation due to the additional control lines of the type-I bus, making the type-II bus more suitable.
- 3) The type-II bus exhibits nearly 13% and 40% improvement over the VM bus for both instruction and data address streams, respectively, and up to 65% power savings over the CM bus.

In addition to the power savings of the adaptive bus technique, an important result that stems from using CM signaling is the reduction in the number of repeaters. As shown in Table I, the 32-bit type-II adaptive bus can achieve the target data rate of 1Gb/s across a 1-cm long wire with 34 instead of 96 repeaters/receivers required for the VM bus.

Number of Repeaters + Receivers	CM	VM	Adaptive (Type-I)	Adaptive (Type-II)
	32	96	32+8=40	32+2=34

Table I. Total number of repeaters and receivers for several bus signaling schemes

6. CONCLUSIONS

A new bus architecture based on hybrid current/voltage mode signaling to achieve high data rates while minimizing the number of required repeaters by nearly 1/3 has been presented. Current-mode signaling uses low-impedance receive-end termination to shift the pole position of the line, thereby achieving high transmission bandwidths. Thus, the attractiveness of current-mode signaling stems from the fact that relatively high data-rates can be attained despite the continuing reverse interconnect scaling trends. To compensate for the increase in static power dissipation inherent to current sensing, the proposed bus technique adaptively changes the mode of operation from current to voltage when the signal activity is low and from voltage to current mode otherwise. Thus, the bus energy expenditure can be minimized to support the required bus signal activity only.

A low-power design methodology based on circuit-level power estimation and statistical analysis of address streams for typical benchmarks extracted using a time-based Alpha 21264 simulator, reveal an improvement in power dissipation of up to 65% and 40% for current and voltage mode signaling, respectively. Overall power dissipation improvement is attained over voltage-mode signaling schemes because at high data rates, the dynamic power dissipation of full-swing signals can become significant. Conversely, the rate at which the power dissipation increases with signaling frequency is much smaller for current sensing, an effect owed to the reduced signal swings.

7. ACKNOWLEDGMENTS

The authors would like to thank Karthik Sundaramoorthy and Dr. Eric Rotenberg for their support and valuable discussions on extracting the bus statistics.

This work is supported in part by the National Science Foundation and Semiconductor Research Corporation under award 983.001.

8. REFERENCES

- [1] R. Krishnamurthy, A. Alvandpour, V. De, S. Borkar “High-performance and Low Power Challenges for Sub-70nm Microprocessor Circuits,” Custom Integrated Circuits Conference, pp. 125-128, 2002.
- [2] H.B. Bakoglu, Circuits, Interconnections, and Packaging for VLSI. Reading, MA: Addison-Wesley, 1990.
- [3] R. McInerney et al., “Methodology for Repeater Insertion Management in the RTL, Layout, Floorplan and Fullchip Timing Databases of the Itanium™ Microprocessor,” ISPD, pp. 99-104, 2000.
- [4] E. Seevinck, P. van Beers, H. Ontrop, “Current-mode techniques for high-speed VLSI circuits with application to current sense amplifier for CMOS SRAM’s,” IEEE J. Solid-State Circuits, vol. 26, no. 4, pp. 525-536, April 1991.
- [5] R. Bashirullah, W. Liu, R. Cavin, “Delay and power model for current-mode signaling in deep submicron global interconnects,” CICC 2002, pp 513-516.
- [6] H. Zhang, V. George, J.M. Rabaey, “Low-Swing On-chip Signaling Techniques: Effectiveness and Robustness,” IEEE Trans. VLSI, vol. 8, no. 3, pp. 264-272, June 2000.
- [7] D. Burger and T. M. Austin, “The SimpleScalar tool set, version 2.0,” University of Wisconsin, Madison, Technical Report CS-TR-97-1342, June 1997.