

Spatial Encoding Circuit Techniques for Peak Power Reduction of On-Chip High-Performance Buses

Himanshu Kaul, Dennis Sylvester, Mark Anders[†], Ram Krishnamurthy[†]

Electrical Engineering and Computer Science,
University of Michigan, Ann Arbor, MI-48109, USA
{hkaul, dmcs}@umich.edu

[†]Circuit Research Labs, Intel Corporation,
Hillsboro, OR-97124, USA
{mark.a.anders, ram.krishnamurthy}@intel.com

ABSTRACT

We propose various low-latency spatial encoder circuits based on bus-invert coding for reducing peak energy and current in on-chip buses with minimum penalty on total latency. The encoders are implemented in dual-rail domino logic with interfaces for static inputs and static buses. A spatial and temporally encoded dynamic bus technique is also proposed for higher performance targets. Comparisons to standard on-chip buses of various lengths with optimal repeater configurations at the 130nm node show the energy-delay and peak current-delay design space in which the different encoder circuits are beneficial. A 9mm spatially encoded static bus exhibits peak energy gains beyond that achievable through repeater optimization for a single cycle operation at 1GHz, with delay and energy overhead of the encoding included. For throughput constrained buses, the spatially encoded static bus can provide up to 31% reduction in peak energy, while the spatially and temporally encoded dynamic bus yields peak current reductions of more than 50% for all bus lengths.

Categories and Subject Descriptors

B.4.3 [Input/Output and Data Communications]: Interconnections (Subsystems)

General Terms: Performance, Design.

1. INTRODUCTION

Increasing clock frequencies coupled with degrading interconnect delays have resulted in a larger portion of the chips power consumption being used in on-chip communication. This is especially prominent in high-performance designs, where interconnects are heavily buffered to deal with degraded intrinsic wire delays and signal slew rates. For a particular wiring pitch the total wire capacitance and resistance are fixed and it is the size and number of the repeaters on the buses that primarily determine the optimal energy-delay curve of the bus. Another variable for controlling the energy-delay curves of these buses is the maximum number and direction of transitions on the bus in any given clock cycle. Reducing the maximum number of simultaneous transitions on a bus is also important for power supply integrity. As opposed to average power consumption, which determines battery life, the

peak power consumption dictates the packaging and thermal regulation mechanisms that determine the reliability of high-performance chips. Repeaters draw enormous currents due to their large size and necessitate decoupling capacitors near repeater blocks to reduce IR drops. With increasing processor complexity, the number of on-chip buses has increased (accompanied by increased datapath widths). As a result, the peak energy and current demand that these buses place on the power grid has increased drastically. A reduction in peak current demand results in decreased decoupling capacitance area in repeater blocks, which in turn yields smaller gate leakage.

Techniques to reduce the maximum number of simultaneous transitions typically employ encoder and decoder circuits on the bus. A number of encoding methods have been proposed to reduce bus power consumption through reduction in transitions for uncorrelated data (with [1]-[6] being some of the representative work). In this work we do not consider address bus encoding since those encoders rely on the correlation of data in consecutive cycles. We note that none of the previous work deals with the efficiency of the proposed encoding circuits for high-performance buses in comparison to standard repeater tuning strategies.

We are interested in reducing worst-case switching behavior for buses with random data, while introducing minimum encode/decode latency. The simplest spatial encoding (from a circuits perspective) for random input data is the bus-invert code [1]. Bus-invert coding begins by detecting the number of transitions that would occur on the bus if the input data was passed without encoding. If the number of transitions is more than half of the total bits, the data is inverted before being transmitted on the bus and an extra control signal is set high to indicate that the data on the bus is inverted. At the receiver end, the decode circuit inverts the data if the control signal is high. This limits the maximum number of transitions to half the size of the bus width and can result in potential gains of 50% in energy and peak current in the worst case.

Bus-invert code was targeted at off-chip buses, where the bus latencies and capacitive loads are much greater than on-chip buses. Though the bus-invert code can ideally result in a 50% reduction in energy, the encoding and decoding circuits result in increased bus latency. The bus latency can also be increased by changing the number and size of repeaters to reduce peak switching energy and current. The extra control bit for every set of encoded bits also results in increased routing area. Given the same routing area as the encoded bus, the standard bus wiring pitch can be re-optimized (primarily through spacing) to reduce delay as well as energy for the same repeater configuration. These are some of the key trade-

This work was supported by an Intel PhD fellowship and a MARCO/DARPA Gigascale Systems Research Center grant.

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.

ISLPED'04, August 9–11, 2004, Newport Beach, California, USA.

Copyright 2004 ACM 1-58113-929-2/04/0008...\$5.00.

offs that need to be analyzed when comparing encoding techniques to standard approaches for improving the energy-delay curves of on-chip buses. To date, this type of analysis has been ignored.

Controlling the direction of switching can also result in more optimal energy-delay curves. If the encoding is such that the neighboring wires never switch in opposite directions, the wire delay reduces along with switching energy since the cross-coupling capacitance is limited to nominal (rather than twice) in the worst-case. In [7], the authors propose an encoder circuit that addresses this issue. Though they claim improvements in the interconnect delay, they do not include the delay of their encoder and decoder circuits into the total delay. They also do not make any comparison to repeater tuning strategies. Such monotonic switching behavior can be achieved in a dynamic bus at the cost of the switching activity being input state-dependent rather than switching activity dependent. In [8] it has been shown that dynamic buses can easily be encoded to achieve switching activity dependent switching behavior while retaining most of the performance benefits of dynamic buses by switching the bus in the evaluate phase only when input data switches from the previous cycle. In this paper we integrate spatial encoding into the earlier transition encoded dynamic bus to achieve further performance gains on the encoded bus. In this case the gains from avoiding the energy and delay penalties of the doubled coupling capacitance must be weighed against the fact that a transition encoded dynamic bus switches twice in a cycle while static input data switches only once.

The paper is organized as follows. Efficient encoder circuits are proposed for minimal encode /decode latency for static buses and transition encoded dynamic buses in Section 2. The design space where spatial encoding can result in gains in peak energy and current beyond basic repeater optimization is shown in the context of global bus optimizations for an 0.13 μ m technology in Section 3. The energy, current, and performance trade-offs of the different encoder circuits and techniques are also discussed (in Section 3) and analyzed in the context of the various bus optimizations. Section 4 concludes the paper.

2. ENCODER CIRCUITS

The schematics of the standard and encoded buses are shown in Fig. 1. The *in_prev* signal into the first XOR of the encoder represents the data that was sent onto the bus in the previous cycle. The data is assumed to travel from the output of a flop and arrive at the input of a flop at the end of the interconnect. We are interested in the optimal energy-delay curves for transmitting data from the node *in* to the node *flop_in*. For bus-invert coding the majority voter circuit is the most critical from a performance and energy perspective. Various decisions in designing the majority voter give rise to different trade-offs. The most important decision is the size of the encoded set of bits. The logic depth of the majority voter increases as the number of encoded bits is increased. As the number of encoded set of bits is made smaller the area penalty for encoding a fixed bus width increases since a larger number of control signals are required. The criteria used for detecting a majority can also affect the energy-delay optimality for bus-invert coding. The original bus-invert code allowed for a maximum of 50% of the bits to switch on the bus and is energy optimal. The logic depth of the majority voter could be decreased by relaxing the 50% constraint to achieve lower latency and, possibly, energy-delay optimality for tighter delay constraints. We now more closely examine these degrees of freedom.

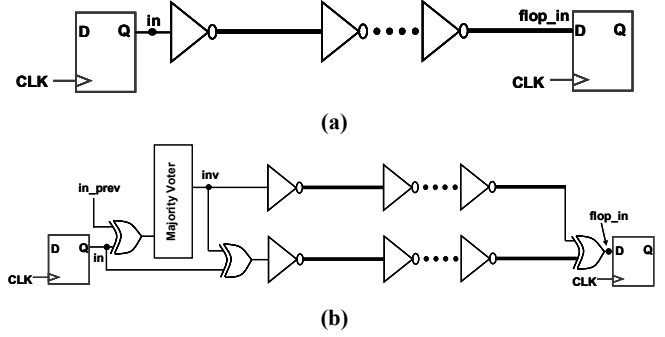


Figure 1. Schematic of (a) original bus with repeaters and (b) encoded static bus with repeaters.

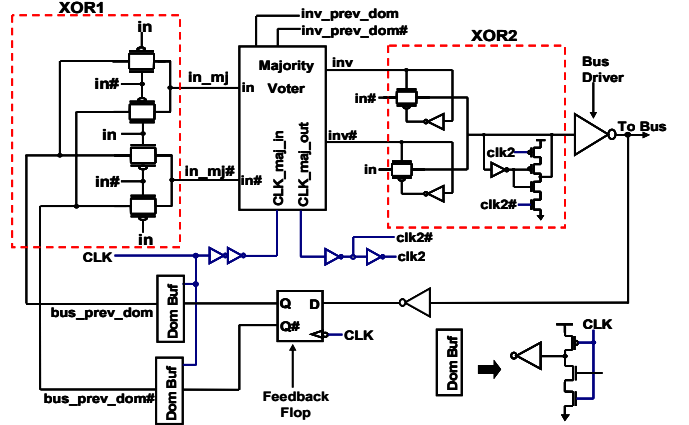


Figure 2. Schematic of encoder circuit for static bus. The domino-static and static-domino interfaces are embedded in the XORs of the encoder.

The bus-invert encoder in [1] has a flop inserted between the encoder and first bus driver to prevent glitches at the bus driver input that can travel down the bus and result in considerable energy consumption. To remove the flop latency while retaining glitch-free signals at the driver input, the majority voters in the encoders have been implemented in dual-rail domino. Domino logic also reduces the encoder latency. The input of the encoder is assumed to be a static signal along with the input to the bus driver. The encoding circuit schematic is shown in Fig. 2. The first XOR gate (XOR1) also serves as a static-to-domino interface and the second XOR gate (XOR2) serves as a domino-to-static interface. XOR2 also has a cross-coupled inverter pair that maintains the state of the input to the driver when the majority voter is in precharge. A similar domino-static interface is used for the *inv* signal for transmission on the static bus. The flop in the feedback path captures the last data sent on the bus and is minimum size since it has a half cycle for its output to settle. However, this flop does place a restriction that the output of the first driver must settle in half a cycle. Domino buffers convert the static output of the feedback flop to a domino compatible signal (*bus_prev_dom*) for input to the XOR1 gate. The previous state of the *inv* signal is also held in a feedback flop and supplied directly to the majority voter as a domino signal (*inv_prev_dom*). The XOR1 gate also places a restriction that the domino output for the previous cycle's data cannot evaluate before the static signal (*in*) has settled to ensure glitch-free inputs to the majority voter during the evaluate phase. This results in increased encoder latency and can be reduced by making the output of the driver flop in Fig. 1 domino compatible.

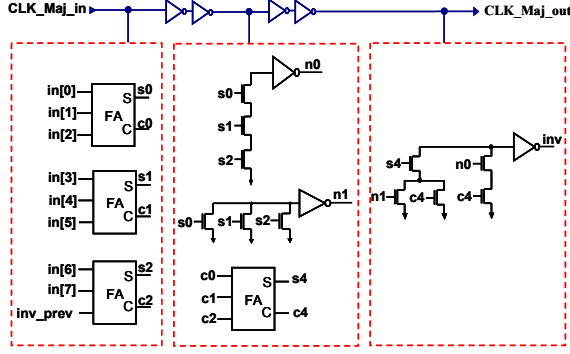


Figure 3. Majority voter for 8/4 encoder.

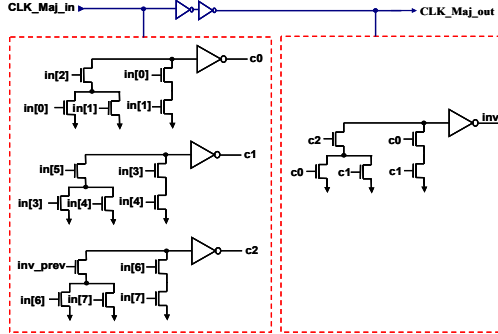


Figure 4. Majority voter for 8/5 encoder.

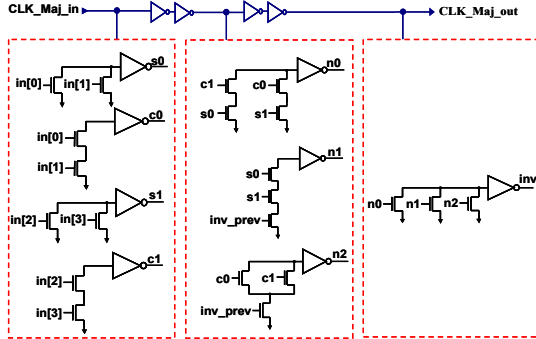


Figure 5. Majority voter for the 4/2 encoder.

The nomenclature for the encoders in our work is such that if the encoder takes n bits of the original data and allows a maximum of m bits on the bus to switch, it is called an n/m encoder. The primary difference between the encoders is the majority voter. The majority voter for an 8/4 encoder is shown in Fig. 3, along with the clock distribution. Only the circuit that produces the true signal is shown due to space constraints. The pull-up PMOS, cross-coupled keepers and keeper latches have not been shown. The full-adder (FA) circuit is implemented as a mirror-type adder. This circuit has a critical path delay of four fanout-of-four delays (FO4s). The same majority voter in [1] has a delay of five FO4s.

The majority voter for the 8/5 encoder (Fig. 4) is implemented with four 2-out-of-3 majority voters and has a delay of two FO4s. From an energy perspective, it ideally allows a maximum reduction of 37.5% but it could allow energy gains at lower overall delays due to its reduced latency. The majority voter for the 4/2 encoder (Fig. 5) has a critical path delay of three FO4s. Though it can ideally achieve the same energy reduction as the 8/4 encoder, it has lower latency at the expense of increased routing area overhead.

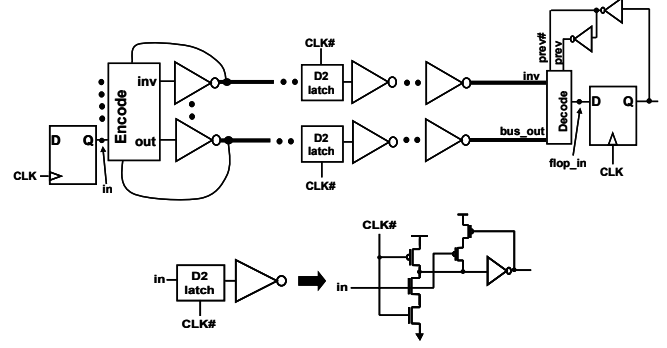


Figure 6. Setup for spatial and transition encoded dynamic bus.

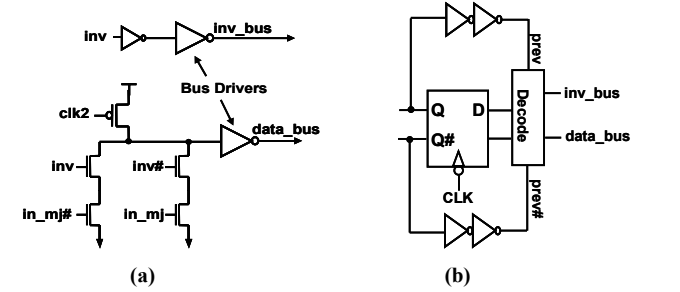


Figure 7. (a) Domino XOR gate used to drive the bus driver for dynamic bus. (b) Feedback flop circuit used for holding previous data sent on the bus.

To enable spatial encoding for peak current reduction at tighter delay targets, a spatial and temporal encoded dynamic bus technique is implemented by embedding the spatial bus-invert coding into the transition encoded dynamic bus technique of [8]. The overall bus structure is shown in Fig. 6. The D2 latch used here has been shown in [8] to provide good noise immunity, making it practical for on-chip dynamic buses. The decode circuit (Fig. 6) implements the logic function $prev \otimes inv \otimes bus_out$. When the inv signal does not switch in the evaluate phase, the input to the flop on a line that does switch toggles. A switching inv signal in the evaluate phase indicates that the inputs to the flops whose lines do not switch need to be toggled. The receiver end of the domino bus can switch either low or high during the evaluate phase, depending on the number of repeaters after the D2 latch. The decoder shown in Fig. 6 provides minimum latency for the critical path in the evaluate phase when the end of the line switches high during evaluate. For decoders placed at the end of the lines that switch low during the evaluate phase, a modified decode circuit has been used for minimizing critical path latency. The $flop_in$ node has static behavior since it does not change when the bus goes into pre-charge.

There are two main changes to the encoder circuit from the static bus encoder to make it compatible for the domino bus. The encoder circuits XOR2 is changed to the one shown in Fig. 7a, where the pull-down stack directly drives the bus driver during the evaluate phase. The data input to XOR2 does not come from the static input signal (in) but from the output of XOR1 (in_mj). XOR1 of Fig.2 does the transition encoding for the domino bus while the majority voter and modified XOR2 perform the spatial

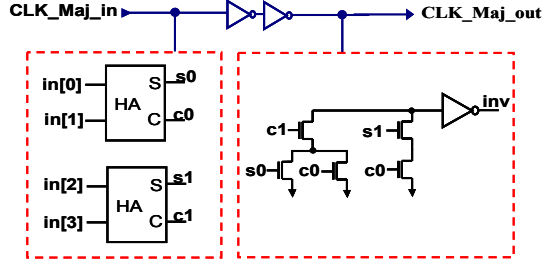


Figure 8. Majority voter for 4/2 encoder for the dynamic bus approach.

encoding on top of that. The *inv* signal does not have a domino-static interface (as in the encoder in Fig.2). For the dynamic bus encoder the previous state of the *inv* signal is not stored but the data on the bus is decoded and stored in the feedback flop (Fig. 7b). The feedback flop needs to hold the same data as the flop at the end of the bus. Though the input data for the feedback flop could be provided directly by the static input signal, taking the signal from the bus allows this encoder circuit to work with buses with multiple-source drivers. The majority voter circuits for the encoder with the domino bus do not use the *inv_prev* signal as an input and the number of inputs to the majority voters reduces by one. For the 8/4 and 8/5 encoders, the majority voter circuits are slightly modified from Figs. 3 and 4, with critical path delays remaining the same. The critical path delay for the majority voter of the 4/2 encoder (Fig. 8) reduces by 1 FO4. To prevent energy consumption (in the encoder) in unused bus cycles, an actual design should gate the global clock with a valid signal, which usually accompanies all data in processors, and use this gated clock signal in the clock distribution within the encoder.

3. ENERGY/CURRENT-DELAY COMPARISONS

An industrial 0.13 μ m CMOS technology is used for the device models with a VDD of 1.2V. All simulations are performed at a temperature of 100C. A typical global metal layer for a 0.13 μ m technology node [9] is used for routing the bus, with a minimum pitch of 1.2 μ m (Fig. 9). Capacitance extraction is performed with a 2D field-solver. The number and size of the repeaters on the buses are optimized to obtain a delay target at minimum worst-case energy. Since repeater insertion for on-chip buses is also used to constrain slew rates for signal integrity and reliable repeater delay purposes, we used a constraint that the 10-90% transition time should not exceed 200ps ($\sim 3X$ the transition time at the output of an inverter driving a FO4 load) at any point on the bus. The bus is used to transmit 16 bits of data from the driver flops to receiver flops (as shown in Fig. 1). The original 16-bit bus is assumed to be routed at minimum pitch. When converted to an encoded bus (also routed at minimum pitch), there is an area penalty (depending on the type of encoder used). The 8/4 and 8/5 encoding result in a control signal being inserted every 8 bits (i.e., a 12.5% increase in routing area). For a fair comparison to the standard bus, the wire width and spacing of the standard bus are re-optimized for minimum energy within the increased routing area due to the spatial encoding. This re-optimization always resulted in increased spacing since the coupling capacitance falls off rapidly, reducing both energy and delay of the bus. The repeater configurations are re-optimized whenever RC parasitics of the wire change. Similarly, the 4/2 encoder results in 25% increased routing area and the bus pitch for the standard bus and the 8/4 and 8/5 encoded bus are re-optimized for comparison to the 4/2 encoder at the same routing

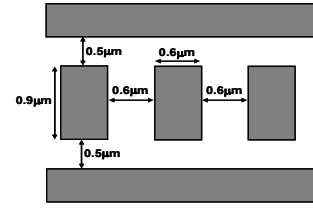


Figure 9. Interconnect structure for global metal layer.

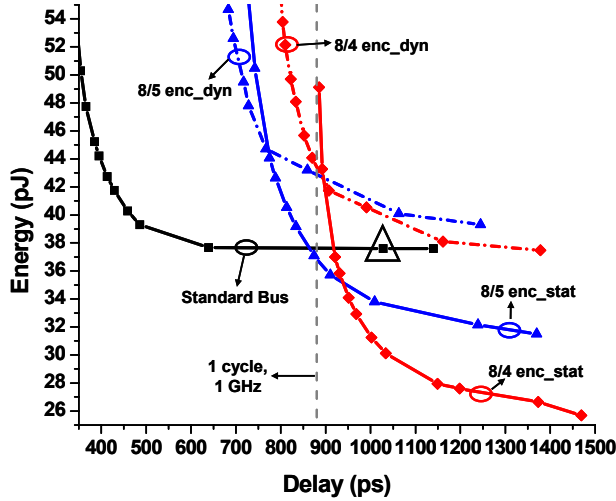
area. For all delay targets, the encoders are sized so as to operate at the “knee” of their respective energy-delay curves for the load capacitance (bus driver) that they drive. The knee points typically resulted in

energies that were 10-20% higher than the minimum energy of the encoder circuit

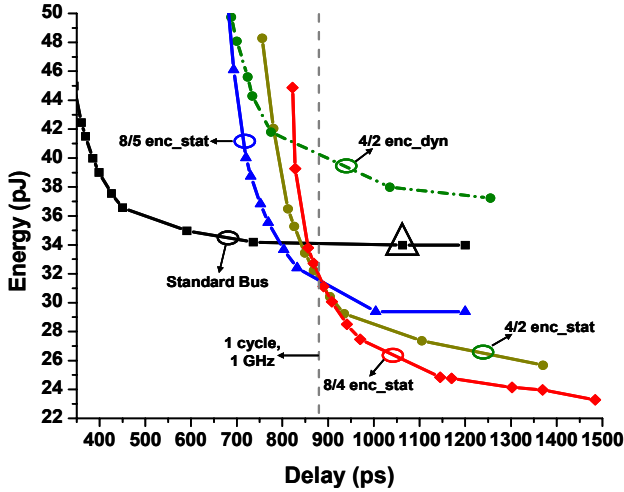
and yielded nearly constant latencies for a particular encoder circuit over a range of load capacitances. The decoder sizes are fixed since they always drive the same load, namely the D input of the receiver flop.

As line lengths increase, the energy and delay overhead of the encoding decreases. Buses of lengths 3mm, 6mm, and 9mm are optimized for the standard and encoded buses to determine the line lengths at which the encoders are effective when their energy and delay overheads are considered. At the various line lengths we also determine the delay targets at which the encoding techniques provide actual gains in energy. The worst case energies are plotted against worst-case total delays for the bus length of 9mm in Fig. 10. The energy numbers include the energy consumption of the encoder (along with the requisite clock network) and the delay numbers include the delay through the encoder and decoder. The energy of the standard bus cannot be reduced beyond a certain point due to slew rate constraints (typically when inter-repeater distance is 1mm in our setups) and this data point for the standard bus has been marked in Figs. 10 and 11 with a Δ symbol. The plot for standard buses has been extended beyond this point in Fig. 10 for the sake of minimum energy comparisons. The spatial encoding with the dynamic bus results in lower energy at tighter delay targets when compared to the bus-invert coding with a static bus. However, the encoded dynamic bus provides these gains in the energy-delay region where the standard bus is preferred. As interconnections become more resistive with scaling, the lower energy benefits at tighter delay targets with the spatial and temporally encoded dynamic bus can be expected to occur in the region where it is more optimal than the standard bus. The spatial and transition encoded dynamic buses consistently result in reduced peak current over all delay targets compared to their static bus counterparts (Fig. 11). Note that peak energy is an integral of the current drawn from the power supply in a clock cycle, while peak current is defined as the maximum instantaneous current drawn from the power supply.

Similar peak energy and current vs. delay curves were obtained for the 3mm and 6mm buses. They have not been shown due to space constraints. The relevant data points for all the bus lengths have been tabulated in Tables 1 and 2. For latency-constrained buses, Table 1 shows the minimum delays at which the different encoders for the static bus start to show gains in worst-case energy over the standard bus. For throughput-constrained buses, where reducing worst-case energy and peak current is the primary goal, the 8/4 encoded static bus and 8/4 encoded dynamic bus always result in the maximum reduction in peak energy and current, respectively. The maximum reductions in peak energy with the 8/4 encoded static bus and peak current with the 8/4 encoded dynamic bus for the different cases are shown in Table 2. Note that peak current reduction with the encoded dynamic bus can result in savings greater than the ideal 50% point since the minimum energy point with the dynamic buses results in smaller repeater sizes



(a)

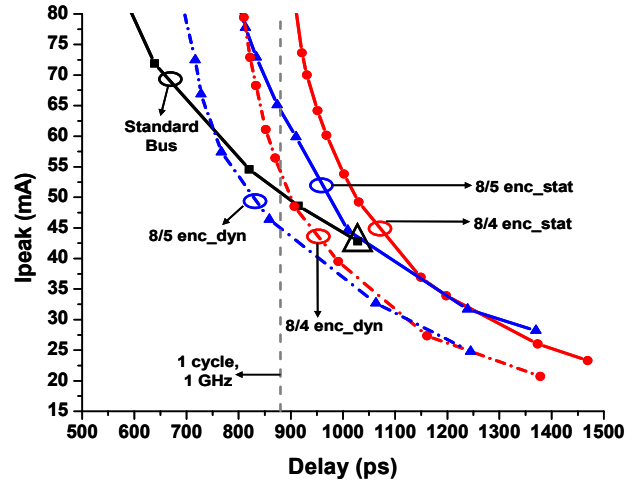


(b)

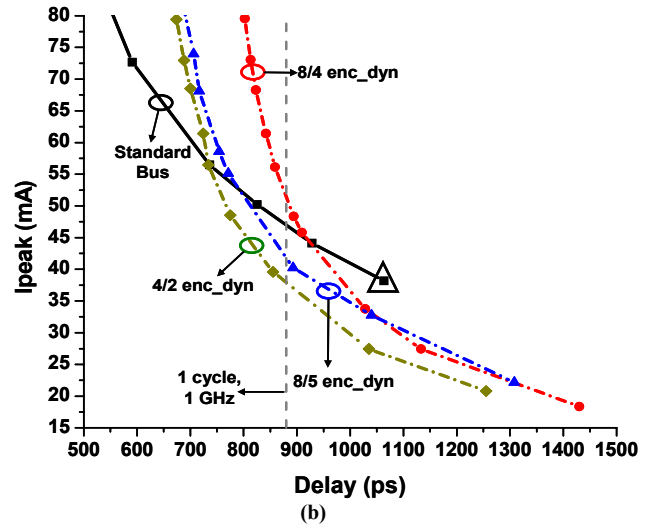
Figure 10. Peak Energy vs Delay curves for a 9mm 16-bit bus with a routing area constraint of (a) 12.5% and (b) 25% over that of minimum pitch. The labels for the plots for dynamic and static encoded buses have 'dyn' and 'stat' post-fixed to them.

(which are a first order determinant of peak current) than the minimum energy point of the standard static bus. As bus lengths decrease, the energy overhead of the encoder reduces the energy gains that result from a reduction in maximum switching bits. As a result, for the 3mm line length only the 8/4 encoder for the static bus yields improvements, with just a 3.6% maximum energy reduction. Since peak current is a transient phenomena, the 8/4 encoder for the dynamic bus results in peak current reductions over all line lengths.

Due to the lower encoder latency of the 8/5 encoder, this encoding typically results in energy gains at lower delay targets than the other encoder implementations. As the delay constraints are loosened, the 8/4 encoder provides higher energy gains since it allows fewer switching bits in the worst case. For the same reason the 4/2 encoder shows energy gains at more aggressive delay targets than the 8/4 encoder. The 8/4 encoder results in larger energy gains than the 4/2 encoder at higher delay targets (though both should ideally result in the same 50% energy gains) since at the same area constraint the bus pitch for the 8/4 encoder is more relaxed than the 4/2 encoder due to fewer control signals. In summary, the potential for the various encoding techniques to show



(a)



(b)

Figure 11. Peak current vs. delay curves for a 9mm 16-bit bus with a routing area constraint of (a) 12.5% and (b) 25% over that of minimum pitch.

peak energy and current reductions at tighter delay targets are in the following order (highest performance to lowest performance) – 8/5 enc_dyn, 4/2 enc_dyn, 8/4 enc_dyn, 8/5 enc_stat, 4/2 enc_stat, 8/4 enc_stat. For throughput-constrained buses, maximum peak energy and peak current reductions can be achieved with the 8/4 encoded static buses and the 8/4 encoded dynamic buses, respectively.

The peak energy-delay curves of the encoded buses can be viewed as the energy-delay curve of the standard bus that has been shifted vertically (energy reduction + encoder energy over-head) and to the right (encoding latency). The region of feasibility for the encoded buses is to the right of the cross-over point of the standard and encoded buses. The total capacitance for a fixed length global wire does not change appreciably with scaling [10]. For the same delay, the increased resistance (from scaling) necessitates a larger number of repeaters, resulting in a net rise in line capacitance. The switching capacitance of the encoder circuit, which reduces with technology scaling, therefore decreases as a fraction of the line capacitance. The encoding latency also reduces with technology scaling (from improvements in FO4 delay), shifting the energy-delay curves of the encoded buses to the left. The reduced latency and energy overhead of the encoders (relative to the wire) will

Table 1. Delays (in ps) at which the encoder types are more beneficial than a standard bus. Delays are grouped into bins for area constraints (defined as a multiple of the area of a 16-bit bus routed at minimum pitch). A “-” entry indicates that the particular encoder is never superior to the standard bus for the given configuration.

Line Length	8/5 Encoder		8/4 Encoder		4/2 Encoder	
	Area Constraint (X Min. Pitch)		Area Constraint (X Min. Pitch)		Area Constraint (X Min. Pitch)	
	1.125	1.25	1.125	1.25	1.125	1.25
3mm	-	-	773	771	-	-
6mm	785	724	821	803	-	736
9mm	865	796	919	855	-	842

Table 2. Maximum reduction in energy with 8/4 encoder (static bus) and peak current (Ipeak) with the 8/4 encoder (dynamic bus) over the standard bus. The total latency at which these reductions are achieved is shown along with the increased latency over the minimum energy point for the standard bus in parentheses.

Line Length	Energy Gains	Latency	Ipeak Gains	Latency
3mm	3.6%	884 ps (+613 ps)	65.5%	799 ps (+528ps)
6mm	21.5%	1134 ps (+476 ps)	60.3%	1164 ps (+506ps)
9mm	31.7%	1484 ps (+441 ps)	51.6%	1378 ps (+350ps)

allow encoders to be feasible at tighter delay targets (for latency-constrained buses) and increase the potential energy reductions for a fixed delay target as processes are scaled below 130nm.

Though the efficiency of the various spatial encoder circuits have been compared with respect to their ability to reduce peak power consumption for on-chip buses, it is important to analyze their impact on average power as this influences the context in which they can be used. For a given input vector the spatially encoded bus can result in energy savings only if the number of switching bits on the encoded bus is less than the number of switching bits on the standard bus. Otherwise, it actually consumes more energy due to the switching energy within the encoder. This increased energy consumption worsens for latency-constrained buses since the bus consumes more energy due to the fact that it has been redesigned to achieve lower latency (to make up for the encoder delay). The ability of the encoder to reduce average power depends on the probability of energy-saving input vectors occurring. In general, at higher switching activities it is much more probable that a majority of the bits will switch simultaneously.

We compare the average power consumption of the throughput-constrained 8/4 and 4/2 encoded static bus of length 9mm at their respective minimum energy points. Random input data is generated for 1000 cycles such that the average switching activity on each input bit is constant but the cycles in which it switches are independent of the other input bits. This is carried out for various switching activities with the results shown in Fig. 12. The average energy is normalized to the minimum energy configuration of the standard bus under the same routing area constraint as the encoded buses. The 4/2 encoded bus should ideally result in energy savings at lower switching activities than the 8/4 encoded bus since the probability of the energy saving cycles (more than 2 bits switching) is higher than the probability of the energy saving cycles for the 8/4 encoded bus (more than 4 bits switching). However, the energy gains are marginal due to the fact that the standard bus energy can be reduced further under the increased area penalty (25%) of the 4/2 encoder. Both encoded buses show average energy gains only

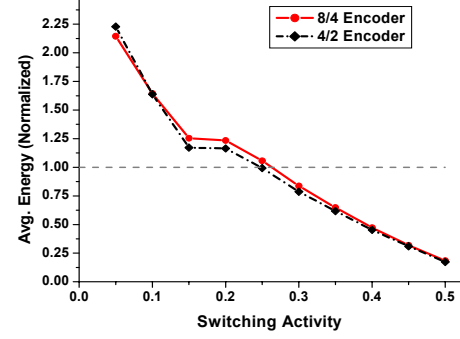


Figure 12. Average energy for an encoded bus vs. switching activity. The average energy is normalized to the standard bus to which the encoded bus is compared.

if the switching activity is higher than 0.25, while switching activities as low as 0.05 result in more than 2X higher average energy than the standard bus. If reducing average power is the only criteria, it needs to be ensured that the average switching activity for the target application is sufficiently high for the encoded bus to result in gains.

4. CONCLUSIONS

Low latency spatial encoder circuits have been proposed for both static buses and transition encoded dynamic buses to achieve peak energy and current reduction in high-performance on-chip buses. The design space in which the specific encoders are optimal has been discussed and shown in the context of bus optimizations for various bus lengths at the 0.13μm technology node. As processes scale, the encoder latencies will account for a smaller fraction of the bus latency. This will increase the feasible design space (in terms of line lengths and effective gains) for using the spatial encoders. For throughput-constrained buses, the 8/4 encoder for static buses resulted in maximum gains in peak energy of 31% for the 9mm bus. The spatial and temporally encoded dynamic bus consistently resulted in the lowest peak currents (51% - 65% gains) for all line lengths. The spatial encoders require high switching activities to achieve average energy gains and this will have to be considered in designs that are targeted solely for minimizing average power.

REFERENCES

- [1] M. Stan & W. Bursleson, “Bus-invert coding for low-power I/O”, *IEEE Tran. VLSI Systems*, Vol. 3, pp. 49-58, Mar 1995.
- [2] M. Stan & W. Bursleson, “Low-Power Encodings for Global Communication in CMOS VLSI”, *IEEE Tran. VLSI Systems*, Vol. 5, pp. 444-455, Dec 1997.
- [3] Y. Shin, *et. al.*, “Partial Bus-Invert Coding for Power Optimization of System Level Bus”, *Proc. ISLPED*, pp. 127-129, 1998.
- [4] A. Boglioli, “Encodings for High-Performance Energy-Efficient Signaling”, *Proc. ISLPED*, pp. 170-175, 2001.
- [5] P. P. Sotiriadis, *et. al.*, “Analysis and Implementation of Charge Recycling for Deep Sub-micron Buses”, *Proc. ISLPED*, pp. 364-369, 2001.
- [6] Y. Zhang, *et. al.*, “Odd/Even Bus Invert with Two-Phase Transfer for Buses with Coupling”, *Proc. ISLPED*, pp. 80-83, 2002.
- [7] K.-H. Baek, *et. al.*, “A Low Energy Encoding Technique for Reduction of Coupling Effects in SoC Interconnects”, *Proc. IEEE Midwest Symp. on Circuits and Systems*, pp. 80-83, 2000.
- [8] M. Anders, *et. al.*, “A transition-encoded dynamic bus technique for high-performance interconnects”, *IEEE Journal of Solid-State Circuits*, Vol. 38, pp. 709-714, May 2003.
- [9] S. Tyagi, *et. al.*, “A 130 nm generation logic technology featuring 70 nm transistors, dual Vt transistors and 6 layers of Cu interconnects”, *IEDM Technical Digest*, pp. 567-570, 2000.
- [10] R. Ho, *et. al.*, “The Future of Wires”, *Proc. IEEE*, pp. 490-504, Apr 2001.