

Coupling-Driven Bus Design for Low-Power Application-Specific Systems

Youngsoo Shin and Takayasu Sakurai

Center for Collaborative Research and Institute of Industrial Science,
University of Tokyo, Tokyo 106-8558, Japan

Abstract

In modern embedded systems including communication and multimedia applications, large fraction of power is consumed during memory access and data transfer. Thus, buses should be designed and optimized to consume reasonable power while delivering sufficient performance. In this paper, we address bus ordering problems for low-power application-specific systems. A heuristic algorithm is proposed to determine the order in a way that effective lateral component of capacitance is reduced, thereby reducing the power consumed by buses. Experimental results for various examples indicate that the average power saving from 30% to 46.7% depending on capacitance components can be obtained without any circuit overhead.

1 Introduction

As the scale of process technologies steadily shrinks and the size of designs increases, interconnects (especially global interconnects) have increasing impact on the area, delay, and power consumption of circuits [1]. Specifically, reduction in scale causes the lateral component of capacitance to dominate the total capacitance of interconnects. This is because wire-to-wire spacing is shrinking for higher densities and the aspect ratios of interconnects have to be increased to compensate for increasing interconnect resistance, which in turn is due to shrinking wire widths. For example of metal 3 layer in typical $0.35\ \mu\text{m}$ CMOS process, the lateral component of capacitance reaches 5 times the sum of fringing and vertical components when the substrate serves as a bottom plane.

In the domain of embedded systems, an increasing fraction of implementations make use of core processors as basic computational units. In these systems, especially in communication and multimedia applications, large fraction of power is consumed during memory access and data transfer. Thus, *system bus*, which is an essential system component to interconnect subsystems for data transfer, should be designed and optimized to consume reasonable power while providing sufficient performance. Although there has been significant work devoted to reduce power consumption of off-chip buses with coding techniques [2], [3], [4], the overhead of coding logic in terms of delay, area, and power cannot be tolerated if the same techniques are to be used for on-chip buses [5]. Furthermore, the effects from lateral component of capacitance should be taken into account when on-chip buses in deep submicron technologies are of concern.

In this paper, we propose a low-power on-chip bus design technique for embedded application-specific systems, which takes the lateral as well as vertical and fringing components of capacitance into consideration. Specifically, we are given a processor core and the embedded application that runs on it. We assume that capacitance components of buses are available from the layout of the processor core and the address streams from typical runs of the application code. The sequence of the address patterns can be available a priori after the algorithm of an applica-

tion is specified such as in signal and image processing applications. Based on the capacitance components and the address streams, we determine the optimal order of the bus in such a way that power consumption of the bus is minimized. The rationale for ordering is that the effective lateral capacitance is reduced if bus lines, with high probability of switching in the same direction, are located adjacent to each other. The obtained order can be used to slightly modify the processor layout without any circuit overhead.

We present the power model, which incorporates all the capacitance components, and define a bus ordering problem. Then, we propose a heuristic algorithm to determine the bus order. Note that the optimal order can be obtained only if the bus is narrow and the length of address streams is small, which are not the case for most of embedded applications. We also evaluate the proposed heuristic algorithm by comparing it with the simulated annealing algorithm.

The remainder of the paper is organized as follows. In the next section, we present a power model and the definition of a problem followed by a heuristic algorithm. In Section 3, we present results of experiments for several examples, and in Section 4 we draw conclusions.

2 Coupling-Driven Bus Design

2.1 Power Model and Problem Definition

We are given a bus $B=(b_0, b_1, \dots, b_{n-1})$, which transfers a sequence of patterns $B_i = (b_0^i, b_1^i, \dots, b_{n-1}^i)$, where i is the time index, n is the bus width, and b_j^i is the value of a bus line b_j at time i . We shuffle bus lines in B such that effective load capacitances seen from driving ends are minimized. Note that if load capacitances are constant, which is the case when there are only vertical capacitance components, the shuffling has no effects with respect to dynamic power consumption, which is a dominant source of power dissipation in a digital CMOS circuit. However, because of lateral capacitance components, the load capacitances are not constant but depend on signal transitions of neighboring wires.

Figure 1 shows parasitic capacitances involved with adjacent bus lines. C_c denotes a lateral

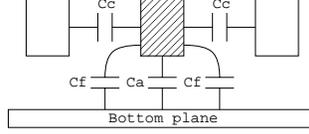


Figure 1. Parasitic capacitances with adjacent lines.

component between electromagnetically coupled lines. C_a and C_f denote the vertical and fringing components, respectively, between a metal line and a bottom plane. We denote the sum of C_a and $2C_f$ as C_l . The ratio between C_c and C_l is denoted by

$$\eta = \frac{C_c}{C_l}. \quad (1)$$

The effect of lateral capacitance (C_c) is that total load capacitance seen by a gate is no longer a constant value, but depends on signal activities of neighboring lines due to the Miller effect [6]. Assume that the *physical* lateral capacitance between two neighboring lines, b_i and b_{i+1} , is C_c . The Miller effect states that if two lines switch in opposite directions, the *effective* lateral capacitance between them is $2C_c$ because the effective voltage swing between them is doubled. On the contrary, the effective lateral capacitance becomes 0 if both lines switch in the same direction.

In a digital CMOS circuit, the dynamic power is proportional to load capacitance and switching activity. Thus, if load capacitance is constant, power consumption is proportional to the number of transitions. In other words, if we have two sets of patterns with the same width and length to be transferred on the same bus, we can compare the power consumption of the bus from each set by comparing the number of transitions. In order to take a similar approach when we incorporate the effect of C_c as well as C_l , we first define the *switching encoding* for j -th bus line as

$$s_j^i = \begin{cases} 1, & \text{if } b_j^{i-1} = 0 \text{ and } b_j^i = 1 \\ -1, & \text{if } b_j^{i-1} = 1 \text{ and } b_j^i = 0 \\ 0, & \text{otherwise.} \end{cases} \quad (2)$$

If $\eta = 0$, we can readily obtain the total number of bus transitions by summing $|s_j^i|$ over j and i .

However, if $\eta \neq 0$, we should take switching polarities of adjacent lines into account to include the effect of C_c . For this purpose, we define the *switching similarity* between adjacent lines (j -th and k -th) at time i , denoted by $\zeta_i(j, k)$, as the amount of effects from C_c seen from j -th line (thus, $\zeta_i(k, j)$ is different from $\zeta_i(j, k)$). For example, if two lines make transitions in opposite directions at time i , $\zeta_i(j, k) = 2$. Then, it can be readily shown that $\zeta_i(j, k)$ is given by

$$\zeta_i(j, k) = |s_j^i| (2 - |s_j^i + s_k^i|). \quad (3)$$

Now, we define the *effective bus transition* of b_j at time i , denoted by α_j^i , as the measure of effective transitions induced both from C_l and C_c , which is normalized to transitions considering only C_l . It can be expressed by

$$\alpha_j^i = \begin{cases} |s_j^i| (1 + \eta \zeta_i(j, j+1)), & \text{if } j = 0 \\ |s_j^i| (1 + \eta \zeta_i(j, j-1)), & \text{if } j = n-1 \\ |s_j^i| (1 + \eta (\zeta_i(j, j-1) + \zeta_i(j, j+1))), & \text{otherwise.} \end{cases} \quad (4)$$

Based on the effective bus transitions, our problem of bus ordering can be defined as follows:

- Given η and a bus $B = (b_0, b_1, \dots, b_{n-1})$, which transfers a sequence of patterns $B_i = (b_0^i, b_1^i, \dots, b_{n-1}^i)$,
- Find the shuffled bus \tilde{B} that minimizes the sum of α_j^i over j and i .

2.2 Coupling-Driven Bus Ordering Algorithm

Because of the exponential number of alternatives for \tilde{B} ($n!$ alternatives for n -b wide bus), the optimal one can be obtained only when the width of bus is narrow and the number of patterns is small, which are not the case for most of embedded applications. In this subsection, we propose a heuristic algorithm based on both switching correlation and transition probability. The *switching correlation coefficient* or simply *switching correlation* for two bus lines (j -th and k -th) is defined by

$$\rho_{jk} = \frac{K_{jk}}{\sigma_j \sigma_k}, \quad (5)$$

Ordering bus lines

Compute transition probability (p_j) of each line;
Compute switching correlation (ρ_{jk}) of each pair of lines;
Find a set of shielding lines $S \leftarrow \{b_j | p_j < \xi\}$;
 $R \leftarrow \{b_0, b_1, \dots, b_{n-1}\} - S$;
 $\Psi \leftarrow \text{build-clusters}(R, p_j, \rho_{jk})$;
 $\{ \} \leftarrow \text{arrange-clusters}(\Psi, S)$;

Figure 2. Heuristic algorithm for bus ordering.

where σ_j is the standard deviation of s_j defined in (2). K_{jk} is the covariance of s_j and s_k and defined by

$$K_{jk} = E\{s_j s_k\} - m_j m_k, \quad (6)$$

where $E\{x\}$ is the expected value of x and m_j is the mean of s_j .

The heuristic algorithm is outlined in Figure 2. Initially, we group bus lines with relatively low transition probability (below some threshold, ξ). These lines serve as *shielding lines* between clusters. The remaining bus lines are subdivided into a group of ordered sets, called *clusters*. The clusters and a set of shielding lines are ordered to result in the final order. Note that orders of lines in each cluster is fixed once each cluster is built, except of the possibility of conditional reverse.

The heuristic to group bus lines into a set of clusters is shown in Figure 3. For each cluster, we first select the line with the highest transition probability among lines not selected and then build a new cluster. At each iteration of inner **while** loop, we select a line (b_k) that maximizes the switching correlation between b_k and the first or the last element of a cluster (recall that cluster is an ordered set) under consideration. This continues until there are no candidate lines having positive switching correlation with the first or the last element of a cluster. In this way, each cluster is formed in such a way that lines with high transition probabilities and high switching correlations are more likely to be grouped together, thereby reducing the effective lateral capac-

```

build-clusters( $R, p_j, \rho_{jk}$ )

  while  $R$  not empty do

    Select  $b_j \in R$  such that  $p_j$  is maximum, and  $R \leftarrow R - \{b_j\}$ ;

    Form a new cluster  $\Psi_i \leftarrow \{b_j\}$ ;

    while true do

      Find  $b_k \in R$  maximizing  $\rho_{kl} > 0$ , where  $b_l$  is the first or the last element of  $\Psi_i$ ;

      If  $b_k$  is not found then exit loop; end if

      If  $b_l$  is the first element of  $\Psi_i$  then  $\Psi_i \leftarrow \{b_k\} \cup \Psi_i$ ;

      else  $\Psi_i \leftarrow \Psi_i \cup \{b_k\}$ ; end if

    end do

     $\Psi \leftarrow \Psi \cup \Psi_i$ ;

  end do

  return  $\Psi$ ;

```

Figure 3. Heuristic algorithm for clustering.

itances. Furthermore, lines located at both ends of each cluster have relatively low transition probabilities that also contributes toward reducing the effective lateral capacitances, which is to be clarified in Figure 4.

From a set of clusters and a set of shielding lines, the final order of bus lines is determined by the heuristic algorithm as shown in Figure 4. First, we select two clusters to be located at both ends of the final order. Because the lines to be located at both ends of the final order will have only one lateral capacitances, they should be lines with high transition probabilities. Then, the remaining clusters are located sequentially with shielding lines between each cluster. Because clusters are built in a way that any combination of two clusters results in negative switching correlation between lines located in the boundaries of clusters (see Figure 3), locating a shielding line in-between clusters decreases the effective lateral capacitances.

```

arrange-clusters( $\Psi, S$ )

  Select  $\Psi_l \in \Psi$  such that  $p(\Psi_l)$  is maximum, and  $\Psi \leftarrow \Psi - \Psi_l$ ;

  Select  $\Psi_r \in \Psi$  such that  $p(\Psi_r)$  is maximum, and  $\Psi \leftarrow \Psi - \Psi_r$ ;

  if the first element of  $\Psi_l$  has  $p(\Psi_l)$  then  $F \leftarrow \Psi_l$ ;

  else  $F \leftarrow \text{reverse}(\Psi_l)$ ; end if

  foreach  $\Psi_i \in \Psi$  do

     $\Psi \leftarrow \Psi - \Psi_i$ , and  $F \leftarrow F \cup \Psi_i$ ;

    Select  $b_i \in S, S \leftarrow S - \{b_i\}$ , and  $F \leftarrow F \cup \{b_i\}$ ;

  end do

  if  $S$  is not empty then  $F \leftarrow F \cup S$ ; end if

  if the last element of  $\Psi_r$  has  $p(\Psi_r)$  then  $F \leftarrow F \cup \Psi_r$ ;

  else  $F \leftarrow F \cup \text{reverse}(\Psi_r)$ ; end if

  return  $F$ ;

```

Figure 4. Heuristic algorithm to find the final order. $p(\Psi_i)$ is the maximum of transition probabilities of the first and the last element of Ψ_i . $\text{reverse}(\Psi_i)$ reverses the order of elements of Ψ_i .

3 Experimental Results

To evaluate the efficiency of the proposed algorithm, we perform experiments for the following set of sample patterns:

- wavelet, linear, laplace, compress, and lowpass: data address patterns in benchmark examples collected from typical image or signal processing algorithms [7]. We assume 16-b wide data address buses for all the programs. Patterns are extracted with the help of Shade [8].
- fft: 7-b wide data address patterns between 128-point complex *fft processor* of an audio decoder [9] and memory. Patterns are extracted through VHDL simulation.

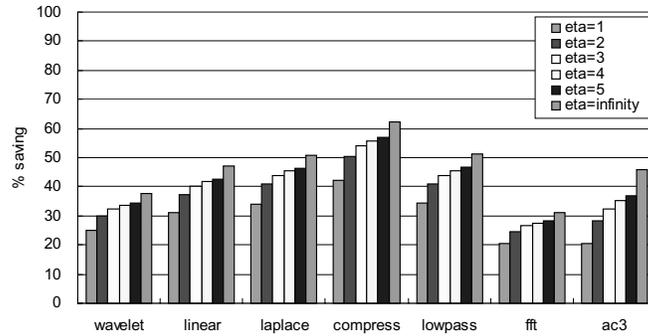


Figure 5. Percentage saving with heuristic algorithm.

- ac3: 16-b wide data address patterns between memory and *parser processor* of the audio decoder, which reads input data stored in a frame memory. Patterns are extracted through VHDL simulation.

We assume that C_c and C_l are constants over all bus lines, and perform experiments with $\eta = 1, 2, 3, 4, 5, \infty$. The resulting percentage saving in power with the proposed heuristic algorithm is shown in Figure 5. Figure 6 corresponds to the result after ordering is done using simulated annealing (SA)¹ [10] instead of the heuristic algorithm, which gives an idea of how good the solutions obtained by the proposed heuristic algorithm are. However, SA itself can be used instead of the proposed heuristic algorithm when the bus width and length is of reasonable size. We also obtain the optimal order for `fft`, which is 7-b wide and consists of 782 patterns. Interestingly, the result is the same as that obtained with SA.

The results of the average percentage saving with the heuristic algorithm are compared to those of SA in Figure 7. The heuristic algorithm gives 30% on the average when $\eta = 1$ up to 46.7% when η is infinity. The difference between heuristic algorithm and SA ranges from 1.9% to 4.4%.

¹Two kinds of moves are used. One is one-to-one exchange between randomly selected two bus lines. Another is group-to-group exchange between randomly selected two groups of bus lines. The move itself is chosen randomly.

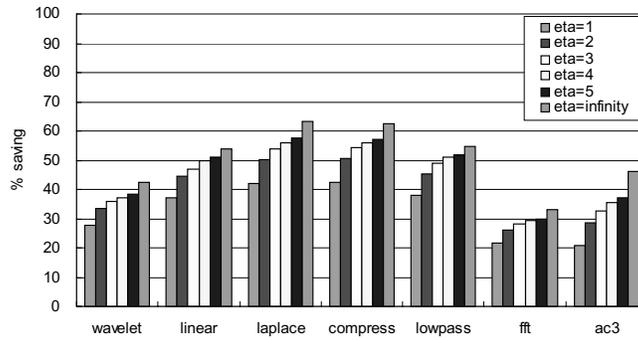


Figure 6. Percentage saving with simulated annealing algorithm.

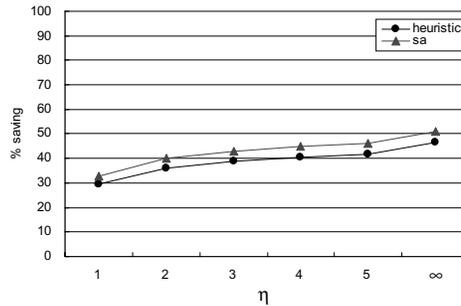


Figure 7. Comparison of heuristic algorithm and simulated annealing.

4 Conclusion

In this paper, we address on-chip bus design technique targeting low-power application-specific systems. In the proposed scheme, we shuffle bus lines in order to minimize the number of effective bus transitions, which includes effects from both lateral and vertical capacitance components, thereby minimizing the power consumed by on-chip buses. We present a heuristic algorithm of shuffling bus lines. The proposed scheme is particularly suitable for address buses in memory-intensive application-specific systems. Experimental results show that savings are substantial for benchmark examples and a large example such as an audio decoder. The performance of the proposed heuristic algorithm is compared to that of simulated annealing.

References

- [1] M. T. Bohr, “Interconnect scaling – the real limiter to high performance ULSI,” in *Proc. IEEE Int’l Electron Devices Meeting*, Dec. 1995, pp. 241–244.
- [2] M. R. Stan and W. P. Burleson, “Bus-invert coding for low-power I/O,” *IEEE Trans. on VLSI Systems*, vol. 3, no. 1, pp. 49–58, Mar. 1995.
- [3] S. Ramprasad, N. R. Shanbhag, and I. N. Hajj, “A coding framework for low-power address and data busses,” *IEEE Trans. on VLSI Systems*, vol. 7, no. 2, pp. 212–221, June 1999.
- [4] L. Benini, G. De Micheli, E. Macii, M. Poncino, and S. Quer, “System-level power optimization of special purpose applications: The Beach Solution,” in *Proc. Int’l Symposium on Low Power Electronics and Design*, Aug. 1997, pp. 24–29.
- [5] P. Sotiriadis and A. Chandrakasan, “Low power bus coding techniques considering inter-wire capacitances,” in *Proc. IEEE Custom Integrated Circuits Conf.*, May 2000, pp. 507–510.
- [6] H. B. Bakoglu, *Circuits, Interconnections and Packaging for VLSI*, Addison-Wesley, 1990.
- [7] P. Panda and N. Dutt, “1995 high level synthesis design repository,” in *Proc. Int’l Symposium on System Synthesis*, 1995.
- [8] R. Cmelik and D. Keppel, “Shade: A fast instruction-set simulator for execution profiling,” Tech. Rep. TR-93-12, Sun Microsystems Laboratories, 1993.
- [9] S. Lee and W. Sung, “A parser processor for MPEG-2 audio and AC-3 decoding,” in *Proc. Int’l Symposium on Circuits and Systems*, June 1997, pp. 2621–2624.
- [10] S. Kirkpatrick, Jr. C. D. Gelatt, and M. P. Vecchi, “Optimization by simulated annealing,” *Science*, vol. 220, no. 4598, pp. 671–680, May 1983.