Value-Based Bit Ordering for Energy Optimization of On-Chip Global Signal Buses*

Krishnan Sundaresan and Nihar R. Mahapatra

Dept. of Elect. & Comp. Eng., Michigan State University, East Lansing, MI 48824-1226, U.S.A.
E-mail: {sundare2, nrm}@egr.msu.edu

Abstract

In this paper, we present a technique that exploits the statistical behavior of data values transmitted on global signal buses to determine an energy-efficient ordering of bits that minimizes the inter-wire coupling energy and also reduces total bus energy. Statistics are collected for instruction and data bus traffic from eight SPEC CPU2K benchmarks and an optimization problem is formulated and solved optimally using a publicly-available tool. Results obtained using the optimal bit order on large non-overlapping test samples from the same set of benchmarks show that, on average, adjacent inter-wire coupling energies reduce by about 35.4% for instruction buses and by about 21.6% for data buses using the proposed technique.

1. Introduction

As fabrication technology scales down, on-chip wiring complexity has continued to grow rapidly, leading to ever higher interconnect energy dissipation. Energy dissipation in global signal bus lines/wires depends on the number of self and coupling transitions occurring in individual wires and between pairs of adjacent wires constituting the bus, respectively. Low power bus encoding schemes attempt to reduce the number of self and/or coupling transitions but with some area and latency tradeoffs. Such schemes incur overheads in the form of extra area and latency for encoding/decoding circuitry and extra wires for control signals. In prior work, we have shown that most bus encoding techniques do not perform as expected on bit-correlated traffic, like those found in SPEC CPU 2000 benchmarks for wide microprocessor buses [6].

A highly efficient option to bus encoding that involves almost no cost and/or performance overheads is *value-based bit ordering*, proposed in this work, in which we rearrange bits of the data to minimize the inter-wire coupling energy. Data statistics obtained by profiling bus traffic from several SPEC CPU2000 benchmarks are used to construct the training data and an optimization problem is formulated to find an optimal minimum energy bit ordering (MEBO) solution for the bus. Results show that, on average, coupling energy reduces by about 35.4% for instruction buses and by about 21.6% for data buses on large non-overlapping test samples drawn from the same set of benchmarks using this technique. The total bus energy also reduces by 29.3% and 16.5% for instruction and data buses, respectively.

Related work that can be collectively called *low-power bus layout optimization* includes the permutation-based (PB) code [2], non-uniform wire placement [3], and a combination of the two [4]. Recently, a technique using a genetic algorithm and a fast optimization algorithm to minimize bus coupling energy was proposed [5]. Our work differs significantly from prior work as outlined next. We propose a simpler formulation that can be solved to optimality using widely-used solvers compared to previous work that used heuristic-based solutions [2] and/or specialized optimization techniques [5]. Unlike prior work, we incorporate the sidewall fringing energy explicitly in our formulation to help optimize energy more fully.

The organization of the rest of the paper is as follows. In Sec. 2, we discuss our methodology, proposed technique, and results. Then, we conclude in Sec. 3.

2. Minimum Energy Bit Ordering (MEBO)

2.1. Methodology

We used the SimpleScalar/Alpha microarchitecture-level simulator, which models a 4-issue Alpha 21264-like processor with a 64-bit data (load/store) bus and a 128-bit instruction bus, to design and evaluate our proposed technique. Wire geometry parameters were obtained from ITRS-2001 and we used FastCap, a three-dimensional capacitance extraction program, to estimate parasitic wire capacitances. We used eight benchmark programs from the SPEC CPU2000 suite to collect data for our optimization. We first ran each program for one billion committed instructions to skip through the initialization phase, and then recorded the bit-wise energy dissipation characteristics for each bit (of the instruction or data bus) over the next 500 million committed instructions. This sample is the *training set*, i.e., the data used for obtaining the minimum energy bit ordering using our optimization technique described in Sec. 2.2. In particular, for each benchmark, we recorded the total coupling energy for all possible bit-pairs, i.e., when any bit i of the data is placed next to any other bit j ($i \neq j$) on the bus. This gave us a *coupling energy/cost matrix*. We also recorded the fringing component of the self energy dissipated by each bit (due to side wall flux) when placed on one of the edge wires of the bus.

2.2. Problem Formulation

We formulate the *minimum energy bit ordering* (MEBO) problem as an instance of the well-known traveling salesman problem (TSP) for which optimal solvers exist. The formal definition follows. Given a set of n nodes (in our case, n bits of the data) and distances between every pair of nodes (in our case, the coupling energy/cost matrix described earlier), find a minimum-distance tour (an arrangement of bits that dissipates minimum total energy) while visiting each node (bit) exactly once. Solving this optimization problem will yield a bit ordering that minimizes the

*This research was supported by US National Science Foundation grant # 0102830.

cost function, i.e., the total energy dissipated due to coupling transitions between adjacent bits. We introduce a variation in this formulation to account for the extra energy dissipated in the fringing capacitance in bits that are placed on wires at the edges of the bus. For this, we add a dummy node to our node set and an additional row and column containing the fringe energies dissipated by each bit, when it is placed at one of the edges, to our coupling energy/cost matrix. The problem now is to find the optimal tour over $n + 1$ nodes. Once the optimal tour (Hamiltonian cycle) is obtained, it can be broken at the dummy node and the resulting linear sequence gives the minimum energy bit order for the bus.

2.3. Results

After collecting the training data from all the benchmarks and representing them in the standard TSPLIB format, we used the publicly-available Concorde tool for TSP to obtain the optimal solution [1]. Then, we used the optimization results to simulate a re-ordered bus and obtained the bus energies for a 200 million sample from each program in our benchmark set. This is our *test set*. Note that, although both our training and test sets are derived from the same set of eight benchmarks, there was no overlap between the sets as we chose the test set beyond the first two billion instructions from the start of each program's execution. The results we obtained are discussed next.

Figs. 1 and 2 show the energy reductions we obtained using the MEBO optimization technique on the instruction and data buses, respectively, for each of the benchmarks. The percentage energy reductions that we report here were calculated with respect to the energy dissipated in an unencoded bus with minimum spacing between the wires.

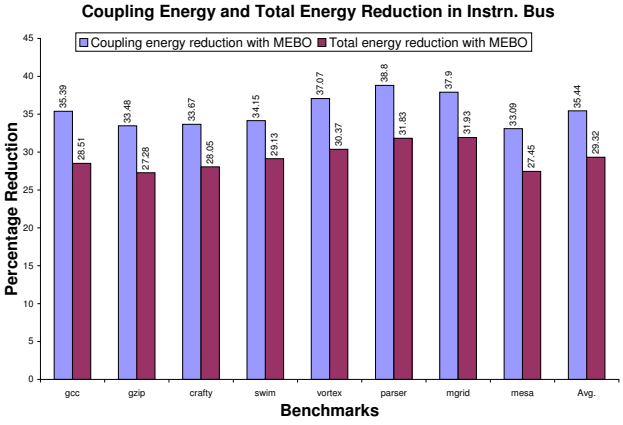

Figure 1: Percentage energy reduction obtained using our proposed MEBO optimization technique on the test set for the 128-bit instruction bus.

We note that substantial energy reductions are possible for both instruction and data buses using our bit ordering technique. For data buses (Fig. 1), we observed an average coupling energy reduction of about 21.57% for MEBO across the benchmarks we analyzed, and for instruction buses (Fig. 2), the average reduction obtained was about 35.44% for MEBO. In our results, we observe that coupling energy reductions for instruction buses are consistently above 30% for all benchmarks and this means that bit ordering schemes seem to be more effective for instruction

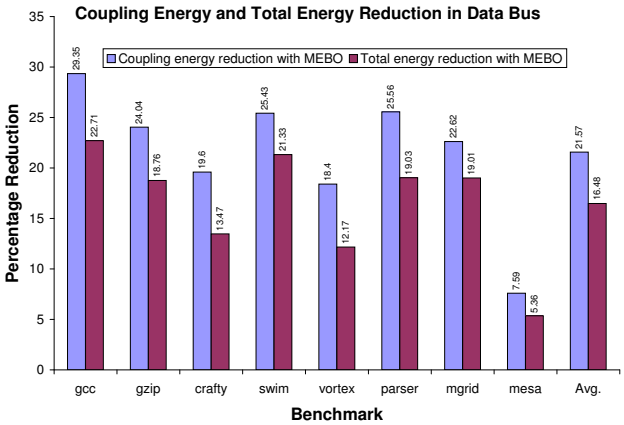

Figure 2: Percentage energy reduction obtained using our proposed MEBO optimization technique on the test set for the 64-bit data bus.

buses. Since MEBO inherently finds the minimum coupling energy dissipation between bits, it may seem intuitive to combine it with an encoding scheme that reduces self transitions, like bus-invert (BI) encoding, to yield further energy benefits. For the data bus traffic that we evaluated, we found that this approach does not result in significantly better energy reduction than using MEBO alone. We attribute this to the low frequency of occurrence of the invert mode and the energy overheads imposed by the extra bit line that is required by BI. However, for instruction bus traffic where the bits are likely to be less correlated, MEBO followed by BI may result in energy benefits.

3. Conclusion

In this paper, we presented a value based technique called minimum energy bit ordering (MEBO) that exploited the statistical behavior of data to find an ordering of bits to reduce coupling as well as total bus energy. We obtained total coupling energy reductions of about 35.4% for instruction buses and about 21.6% for data buses on average across eight SPEC CPU2000 benchmarks for non-overlapping training and test data samples. The total bus energy also reduced by 29.3% and 16.5% for instruction and data buses, respectively. Our bit ordering solution is an efficient low power bus design alternative as it incurs no extra delay, energy, or area overheads compared to low-power bus encoding schemes that require extra hardware area and latency.

References

[1] Concorde TSP Solver. `http://www.tsp.gatech.edu/concorde.html`.

[2] L. Macchiarulo, E. Macii, and M. Poncino. Low-Energy Encoding for Deep-Submicron Address Buses. In *Proceedings of International Symposium on Low Power Electronics and Design*, pages 176–181, October 2001.

[3] L. Macchiarulo, E. Macii, and M. Poncino. Wire Placement for Crosstalk Energy Minimization in Address Buses. In *Proceedings of the 2002 Design, Automation and Test in Europe Conference and Exposition (DATE 2002), Paris, France*, pages 158–162, March 2002.

[4] E. Macii, M. Poncino, and S. Salerno. Combining Wire Swapping and Spacing for Low-Power Deep-Submicron Buses. In *Proceedings of the Great Lakes Symposium on VLSI (GLSVLSI'03)*, pages 198–202, April 2003.

[5] E. Naroska, S.-J. Ruan, and U. Schwiegelshohn. An Efficient Algorithm for Simultaneous Wire Permutation, Inversion, and Spacing. In *Proceedings of the International Symposium on Circuits and Systems*, pages 109–112, May 2005.

[6] K. Sundaresan and N.R. Mahapatra. Accurate Energy Dissipation and Thermal Modeling for Nanometer-Scale Buses. In *Proceedings of the Eleventh International Symposium on High-Performance Computer Architecture (HPCA-11), San Francisco, CA, USA*, pages 51–60, February 2005.