comparison, we present in Table 1 the dielectric constant values calculated using the Wheeler-Hammerstad formula from the same measured S_{21} spectrum. Labels 1–6 correspond to the resonant mode numbers. The dielectric constant value determined by the Wheeler-Hammerstad formula varies from 9.05 to 10.09 for different modes; however, such a strong frequency dependence is unjustified for alumina. This example has demonstrated the systematic error introduced in the dielectric constant determination using the Wheeler-Hammerstad method.

In conclusion, we have successfully demonstrated a new combined method for determining the substrate permittivity using the FDTD simulation and measurement of the S_{21} spectra of microstrip ring resonators. The error can be as small as 1% and it is much more self-consistent than the Wheeler-Hammerstad formula.

Acknowledgment: This research was sponsored by the Center for Dielectric Studies of the Pennsylvania State University. The authors also like to thank A. Baker for sample fabrication and S. Perini for device measurement.

13 March 2000
Electronics Letters Online No: 20000717
DOI: 10.1049/el:20000717

E. Semouchkina, W. Cao and M. Lanagan (*Materials Research Laboratory, The Pennsylvania State University, University Park, PA 16802, USA*)

E-mail: eas203@psu.edu

References

1 AMEY, D., and HOROWITZ, S.J.: 'Microwave material characterization'. Proc. ISHM Int. Symp. Microelectronics, 1996, pp. 494–499
2 CHANG, K.: 'Microwave ring circuits and antennas' (John Wiley & Sons, New York, NY, 1996)
3 HAMMERSTAD, E.O.: 'Equations for microstrip circuit design'. Proc. European Microwave Conf., Hamburg, Germany, September 1975, pp. 268–272
4 WHEELER, H.A.: 'Transmission line properties of parallel strips separated by a dielectric shield', *IEEE Trans.*, 1965, **MTT-13**, pp. 172–185
5 SCHNEIDER, M.V.: 'Microstrip lines for microwave integrated circuits', *Bell Syst. Tech. J.*, 1969, **48**, pp. 1421–1444
6 THROUGHTON, P.: 'High Q-factor resonator in microstrip', *Electron. Lett.*, 1968, **4**, pp. 520–522
7 TAFLOVE, A.: 'Computational electrodynamics' (Artech House, Boston, London, 1995)
8 KENT, G.: 'Nondestructive permittivity measurement of substrates', *IEEE Trans.*, 1996, **IM-45**, pp. 102–106

Low-power data representation

Sunghwan Kim and Jihong Kim

A method of changing a data representation to reduce the switching activity from datapath elements is described. The proposed method finds a low-power data representation so that more frequently consecutive data values have a smaller Hamming distance between their binary codes. The experimental results show that a switching activity reduction of up to 22.6% over a standard data representation scheme can be achieved.

Introduction: Power consumption has become an important parameter in the design of mobile embedded systems such as cellular phones. In a CMOS VLSI circuit (that uses well-designed logic gates), switching activity accounts for over 90% of the total power consumption of the circuit [1]. To reduce the switching activity, various encoding methods have been proposed in the architecture level. For example, Gray code addressing [2] and bus-invert coding [3] minimise the switching activity in the address bus and data bus, respectively.

Previously proposed encoding methods, however, were limited in reducing the system-wide switching activity because they were usually targeted for a single datapath element. For example, bus-invert coding reduces the switching activity from a system bus only. In this Letter, we describe a method of changing a data representation to reduce the switching activity simultaneously from several datapath elements, including a register file, a data cache bus, a data cache, and a main memory bus. Our experiment shows that a significant amount of the system-wide switching activity can be reduced by using a low-power data representation.

Data representation for low-power systems: Reading and writing data values cause the switching activity in datapath elements. The switching activity from a datapath element is directly proportional to the number of bits switched between the successive data accesses in the datapath element. Our method is based on the simple observation that if we can measure the data value transition frequencies in advance and its distribution is highly skewed, the switching activity can be reduced by changing a data representation so that more frequently consecutive data value pairs have a smaller Hamming distance between their binary codes. The proposed method is most useful in designing standalone mobile devices such as a portable MPEG player, where the information on data value transitions of a target application can be computed in advance.

To find a data representation that minimises the switching activity from the datapath elements, we should measure data value transition distributions between all the pairs of data values. For an n-bit data representation, a data value transition distribution from a datapath element D can be represented by a data value transition graph (DVTG) $G_D = (V, E, w_D)$ where V is a set of all the n-bit numbers, E is a set of the undirected edges between all the elements in V, and w_D is a probability density function that maps each edge $e = (n_1, n_2)$ in E to a real number between 0 and 1. $w_D(e)$ indicates the relative frequency of the transitions between n_1 and n_2. Since each datapath element has different load capacitances and usage frequencies, the global data value transition graph G_g (which represents a data value transition distribution for the total datapath) is constructed by merging each datapath element's DVTG using a weight equal to the load capacitance multiplied by the usage frequency of an individual datapath element.

Given a global data value transition graph $G_g = (V, E, w_g)$, a set C of binary codes of length n, and a data encoding function f: $V \rightarrow C$, our power metric $asa(G_g, f)$, the average switching activity per data value transition in G_g under f, is defined as follows:

$$asa(G_g, f) = \sum_{e=(v_1,v_2)\in E} w_g(e) \times h(f(v_1), f(v_2))$$

where h is a function that returns the Hamming distance between two binary codes. For low-power data representation, our goal is to find an optimal data encoding function f_{opt} that minimises $asa(G_g, f)$.

Architectural support for low-power data representation: Changing a data representation introduces a compatibility problem to conventional architectures. For example, a standard ALU (that assumes the 2's complement representation) will not work as expected under low-power data representation. To solve the incompatibility problem, the existing architecture needs to be augmented with several converters. The converters are used to change a standard data representation (e.g. the 2's complement notation) to a low-power data representation and vice versa.

Fig. 1 shows the overall block diagram of a generic architecture augmented with converters: **L2N** changes a low-power data representation to a standard data representation and **N2L** converts a standard data representation to a low-power data representation. As shown in Fig. 1, input data for the ALU is converted to a standard representation, and the result of the ALU is converted back to a low-power data representation. The shaded blocks in Fig. 1 store data in a low-power format while non-shaded blocks store data in a standard format. Since the converters may introduce the speed overhead and area overhead to the architecture, the conversion unit size, n, should be carefully selected. In this Letter, we consider only two cases where $n = 2$ and $n = 4$ to avoid performance degradation.

Experimental results: To validate the switching activity reduction over the 2's complement data representation, we performed experiments using the SimpleScalar tool set. We assumed that the over-

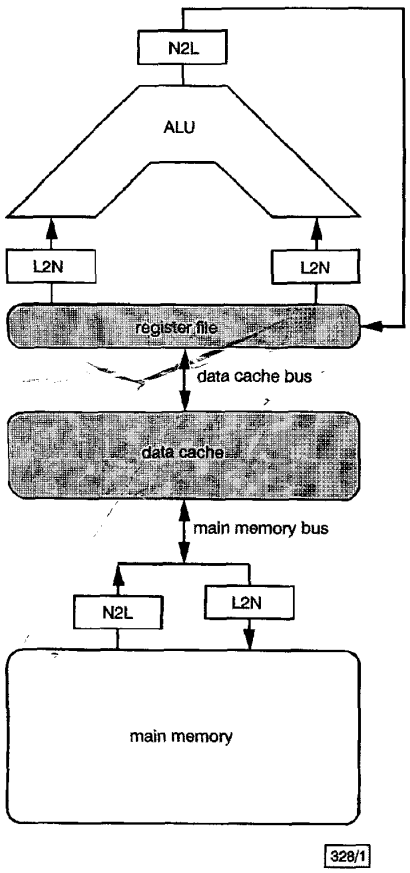

Fig. 1 *Generic architecture with data format converters*

all architecture is similar to the one shown in Fig. 1, having a 32 × 32 register file, a 32-bit data cache bus, a 256kbytes direct-mapped data cache with 32-byte lines, and a 32-bit main memory bus. An MPEG-1 decoder was used as a target application for the experiments. To find a low-power data representation, we used one set of three MPEG-1 bitstreams as a reference set. Based on a global data value transition graph constructed by the reference set, we computed two low-power data representations, one for $n = 2$ and the other for $n = 4$. For the 4-bit conversion unit size, eight 4-bit global data value transition graphs were constructed and eight data encoding functions computed (because the word size in the SimpleScalar is 32 bits). Since it is an NP-hard problem to find an optimal data encoding function, we used a local search heuristic, 2-opt [4], for an approximate solution.

Table 1: Switching activity reduction results by low-power data representations

MPEG bitstreams	4-bit unit size		2-bit unit size	
	INPUT1	INPUT2	INPUT1	INPUT2
	%	%	%	%
Blade Runner	25.5	26.1	6.7	6.7
Highlander	21.0	20.5	8.4	8.4
Psycho	22.2	22.3	7.0	7.0
The Wall	23.0	22.6	8.2	8.2
Wallace & Gromit	21.1	20.4	7.8	7.8
Average	22.6	22.4	7.6	7.6

As expected, the global data value transition distributions were skewed. For example, when $n = 4$, ~9% of the total data value pairs accounted for ~95% of the total data value transitions. Table 1 summarises the switching activity reduction results by using low-power data representations obtained from the MPEG-1 decoder. The reduction results shown in Table 1 include the overhead switching activity from the data format converters. We used two different reference sets, INPUT1 and INPUT2, to determine how the selection of a reference set affects the switching activity reduction. For both the 4-bit and 2-bit conversion unit sizes, our results show that the switching activity reduction is not significantly affected by the selection of a reference set. On average, for the 4-bit conversion unit size, the switching activity was reduced by 22.6% and 22.4% from the INPUT1 and INPUT2 reference sets, respectively. For the 2-bit conversion unit size, the Gray code representation was an optimal low-power data representation for both the INPUT1 and INPUT2 reference sets. On average, the switching activity was reduced by 7.6% for the 2-bit conversion unit size.

Conclusion: A method of changing a data representation to reduce the switching activity from the total datapath is presented. Unlike previous encoding methods, the proposed method decides a low-power data representation based on the weighted data value transition frequencies from several datapath elements. Our method is based on the observation that the distribution of the data value transitions is highly skewed, and exploits this observation in choosing a data representation so that more frequently consecutive data value pairs are encoded to have a smaller Hamming distance between their codes. The experimental results show that, on average, a switching activity reduction of up to 22.6% over a standard data representation can be achieved.

Acknowledgment: This research was supported in part by S.N.U. Posco Research Fund (98-09-2093).

17 February 2000
Electronics Letters Online No: 20000722
DOI: 10.1049/el:20000722

Sunghwan Kim and Jihong Kim (*Department of Computer Science, Seoul National University, Seoul 151-742, Korea*)

References

1 CHANDRAKASAN, A., SHUNG, T., and BRODERSEN, R.W.: 'Low power CMOS digital design', *J. Solid-State Circuits*, 1992, **27**, (4), pp. 473–484
2 SU, L., TSUI, C.Y., and DESPAIN, A.M.: 'Low power architecture design and compilation techniques for high-performance processors'. Proc. COMPCON'94, 1994, pp. 489–498
3 STAN, M.R., and BURLESON, W.P.: 'Bus-invert coding for low-power I/O', *IEEE Trans. VLSI Syst.*, 1995, **3**, (1), pp. 49–58
4 AART, E., and LENSTRA, J.K.: 'Local search in combinatorial optimization' (John Wiley & Sons Ltd., 1997)

Experimental evaluation of TeraLight™ resistance to cross-nonlinear effects for channel spacings down to 50GHz

S. Gurib, A. Bertaina, S. Gauchard, S. Bigo, J.-P. Hamaide, J.-L. Beylat, L.-A. de Montmorillon, R. Sauvageon, P. Nouchi, J.-C. Rousseau and J.-F. Chariot

An investigation is presented into the impact of cross-nonlinear effects over TeraLight™ fibre through an analysis of a 400km-long 32 × 10Gbit/s WDM transmission with channel spacing of 100 and 50GHz, span loss of 28dB and optimised dispersion management.

Introduction: The choice of fibre infrastructure is an important issue in the design of high capacity optical networks. The growing demand for capacity leads to ever lower channel spacings, while input powers are kept relatively high in order to bridge long distances. However, this enhances the cross-nonlinear effects that limit transmission quality. Nonzero dispersion shifted fibres (NZDSFs) have been extensively studied at 10Gbit/s transmission rate [1], but their applications are limited when high input powers and low channel spacings are involved. Standard singlemode fibre