ELSEVIER

Journal of Systems Architecture 48 (2002) 125–135

JOURNAL OF SYSTEMS ARCHITECTURE

www.elsevier.com/locate/sysarc

On the design of low power BIST for multipliers with Booth encoding and Wallace tree summation ☆

Dimitris Bakalis *, Emmanouil Kalligeros, Dimitris Nikolos, Haridimos T. Vergos, George Alexiou

Department of Computer Engineering and Informatics, University of Patras, 26 500, Patras, Greece, Computer Technology Institute, 61 Riga Feraiou Street, 26221 Patras, Greece

Abstract

Low power dissipation (PD) during testing is emerging as one of the major objectives of a built-in self-test (BIST) designer. In this paper we examine the testability of multipliers based on Booth encoding and Wallace tree summation of the partial products and we present a methodology for deriving a low power BIST scheme for them. We propose several design rules for designing the Wallace tree in order to be fully testable under the cell fault model. The proposed low power BIST scheme for the derived multipliers is achieved by: (a) introducing suitable test pattern generators (TPGs), (b) properly assigning the TPG outputs to the multiplier inputs and (c) significantly reducing the test set length. Results indicate that the total power dissipated, the average power per test vector and the peak PD during testing can be reduced up to 73%, 27% and 36% respectively with respect to earlier schemes, depending on the implementation of the basic cells and the size of the multiplier. The test application time is also significantly reduced, while the introduced BIST scheme implementation area is small.

Keywords: Low power testing; Built-in self-test; Booth multipliers; Wallace trees

1. Introduction

In complex integrated circuits, low controllability and observability of embedded blocks impose severe testability problems. In order for the final chip to become a viable product such embedded blocks must be well tested. Built-in self-test (BIST) structures [17] are well suited for testing such embedded blocks, since they cut down the cost of testing by eliminating the need of external testing as well as they can apply the test vectors at speed.

High fault coverage (FC), small area overhead and small application time have been the traditional objectives of BIST designers. Even though these objectives still remain important, a new BIST design objective, namely low power dissipation

☆ Based on "Low power BIST for Wallace tree-based fast multipliers" by D. Bakalis, E. Kalligeros, D. Nikolos, H.T. Vergos, G. Alexiou, which appeared in Proceedings of 1st IEEE International Symposium on Quality Electronic Design, San Jose, CA, USA, pp. 433–438 © 2000 IEEE.

* Corresponding author. Tel.: +30-610-960356; fax: +30-610-991909.

E-mail address: bakalis@cti.gr (D. Bakalis).

1383-7621/02/$ - see front matter
PII: S1383-7621(02)00121-2

(PD) during testing, is emerging [2,3,5–8,10,18, 19,21,22], and is expected to become one of the major objectives in the near future [15].

There are several issues that make the power dissipated during test application an important factor:

Reliability issues: Although there is high correlation between consecutive vectors applied to a circuit during its normal operation, the correlation between consecutive test vectors is significantly lower. Therefore the switching activity of a circuit can be substantially higher during testing than that during its normal operation [18,19]. Since heat and power dissipation in CMOS circuits are proportional to switching activity, the increased switching activity may either cause a circuit under test to be permanently damaged due to excessive heat dissipation or give rise to metal migration (electromigration) that causes the erosion of conductors and leads to subsequent failure of circuits [20]. This is even more severe in circuits equipped with BIST since such circuits are tested frequently in the field.

Technology issues: The multi-chip module technology, that is becoming highly popular, requires sophisticated probing to bare dies for fully testing them [12]. Absence of packaging of these bare dies precludes the traditional heat removal techniques. In such cases, the power dissipated during testing can adversely affect the overall yield, increasing the production cost.

Cost issues: Consumer electronic products typically require a plastic package which imposes a strong limit on the energy that can be dissipated. Excessive dissipation during testing may also prevent periodic testing of battery operated systems that use an on-line testing strategy.

The reader can refer to [7,19] for a more detailed presentation of the motivations for low PD during test application.

Several works have been presented in the past on the reduction of the power dissipated during testing. In [19] a modified PODEM algorithm was presented which derives a test set with reduced switching activity between consecutive test vectors. In [5] a technique for test vector ordering is presented to reduce the switching activity of the circuit during testing. In [18], a test pattern generator (TPG) for reducing heat dissipation has been presented based on the use of two LFSRs operating at different speeds. [10] describes a method for synthesizing a counter in order to reproduce on chip a set of pre-computed test patterns, derived for hard to detect faults, so that the total heat dissipation is minimized. However, a test set targeting the hard to detect faults of a circuit, has some characteristics not available to a test set targeting all faults. In [21] a programmable low power TPG is proposed. It can be implemented using linear cellular automata, with external weighting logic, which realizes the optimal signal probabilities and activities. The drawback of this method is the high implementation cost. In a BIST scheme some vectors generated by the TPG are not useful for testing purposes. A technique that inhibits such consecutive test vectors, by the use of a three-state buffer and the associated control logic has been proposed in [6]. The drawbacks of this method are that it fails to reduce test application time and suffers from high implementation cost. Finally, the partitioning of a circuit into two subcircuits is proposed in [8] so that each subcircuit can be successively tested through two different BIST sessions.

The above mentioned techniques try to solve the general case problem. However there are cases where exploiting the inherent properties of a class of circuits may lead to a more efficient low power BIST scheme. Multipliers are commonly used as embedded blocks in both general purpose datapath structures and specialized digital signal processors. Effective low power BIST schemes for both carry save array multipliers and modified Booth multipliers with carry save addition of the partial products have recently been proposed in [2,3].

Wallace tree summation along with Booth encoding are the most common techniques for designing fast multiplier blocks. Booth encoding aims to reduce the number of partial products whereas Wallace tree summation and carry lookahead (CLA) addition in the final stage of the multiplier aim at the faster addition of the partial products. A BIST scheme for such multipliers has recently been proposed in [13]. This BIST scheme does not take into account the low PD objective.

We will use this BIST scheme as the basis for our comparisons.

In this paper we will first introduce several rules for designing a Wallace tree that is fully testable under the cell fault model [9] when it receives test vectors produced by an 8-bit binary counter. The cell fault model is also used for all other modules except the CLA adder where single stuck-at faults are considered. Next, starting from the basis BIST [13], we will describe a methodology that leads to a new BIST targeting low PD during test. Our methodology is based on (a) suitably modifying the original TPG, (b) properly assigning the TPG outputs to the multiplier inputs and (c) significantly reducing the test set length.

2. Easily testable fast multipliers

We consider $n \times n$ multipliers with inputs $A = (A_{n-1}, \ldots, A_0)$ and $B = (B_{n-1}, \ldots, B_0)$. An 8×8 multiplier is given in Fig. 1. The multipliers that we consider in this paper consist of three units:

(a) The Booth encoding unit for the multiplier encoding and the partial products formation. In this work we use 2-bit encoding on the B operand and therefore reduce the number of partial products in half.
(b) The Wallace tree unit which sums the partial products and produces the sum and carry vector, and
(c) The CLA unit that produces the final result.

Since these multipliers are used for 2's complement multiplication, a sign extension method must be used. We consider the sign generate method which adds a constant value to the partial products in order to produce the final result.

The TPG proposed in [13] for such multipliers (see Fig. 2) consists of an 8-bit counter whose three outputs are repeatedly used to feed the A input of the multiplier while the remaining 5 bits are repeatedly used to form the input B of the multiplier. Multiplexers are used to select between normal inputs and BIST inputs.

The design of the Wallace tree unit can be done in several ways (using 3-2 compressors, 4-2

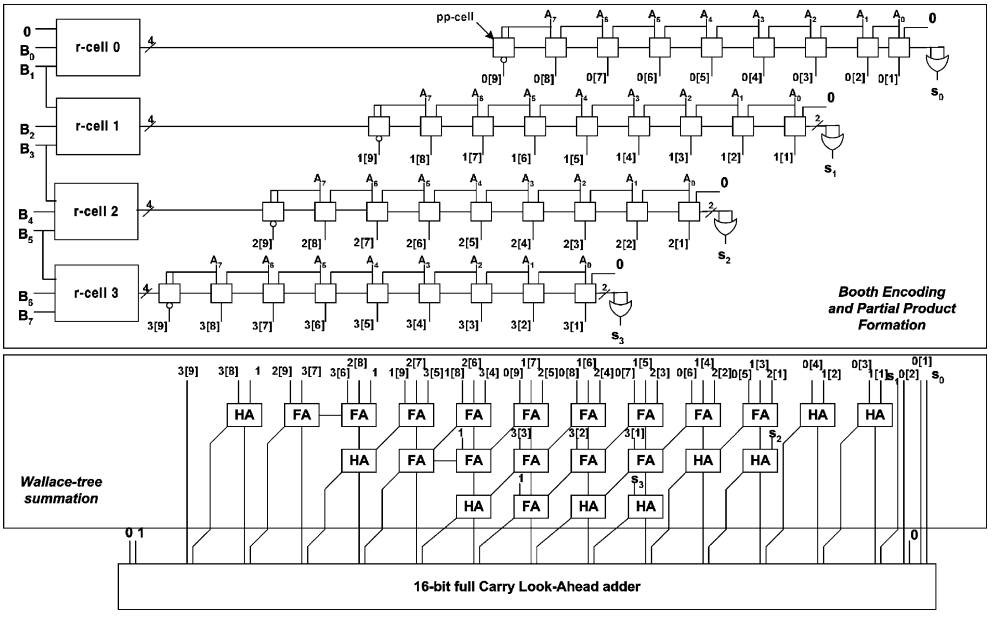

Fig. 1. An 8×8 multiplier with Booth encoding and Wallace tree summation.

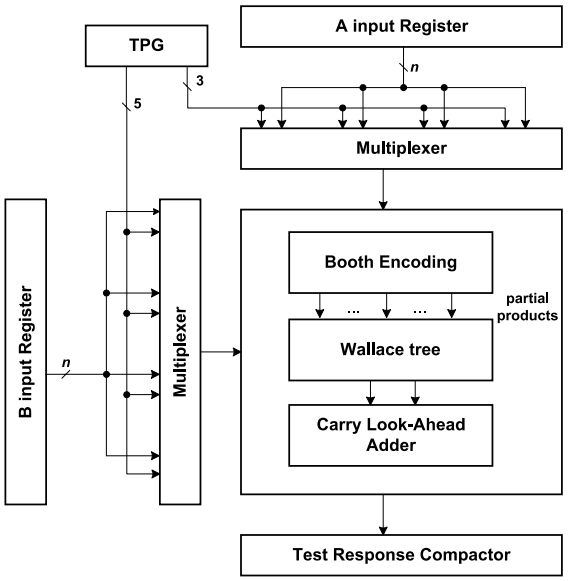

Fig. 2. The BIST scheme.

compressors, etc.). The authors of [13] claim that the 256 vectors of their proposed 8-bit counter TPG are capable of providing all possible input combinations to the inputs of every full or half adder cell. Since they do not propose a specific method for designing the Wallace tree unit we could conclude that the above is valid regardless of the structure used for the Wallace tree. However this conclusion is incorrect. In [4] we have indicated several Wallace tree structures, based on 3-2 compressors, for which the scheme of [13] cannot achieve the described goal. In other words, the Wallace tree structure must follow certain design rules in order for its cells to receive all possible input combinations under the application of the vectors of the 8-bit counter. These design rules are:

Rule I: The partial product bits (PP bits) are grouped in triplets and summed at the first level of each Wallace tree. If the number of PP bits modulo 3 is not equal to zero then the remaining PP bits are summed at next levels of the Wallace tree along with carry bits.

Rule II: If a carry occurs at the $i-1$ level of a certain Wallace tree, then this carry should be inserted at a level k adder of the succeeding most significant tree, such that $k \geqslant i$.

Rule III: Every Wallace tree has at most one half adder which either resides at the last or the previous to the last level of the tree.

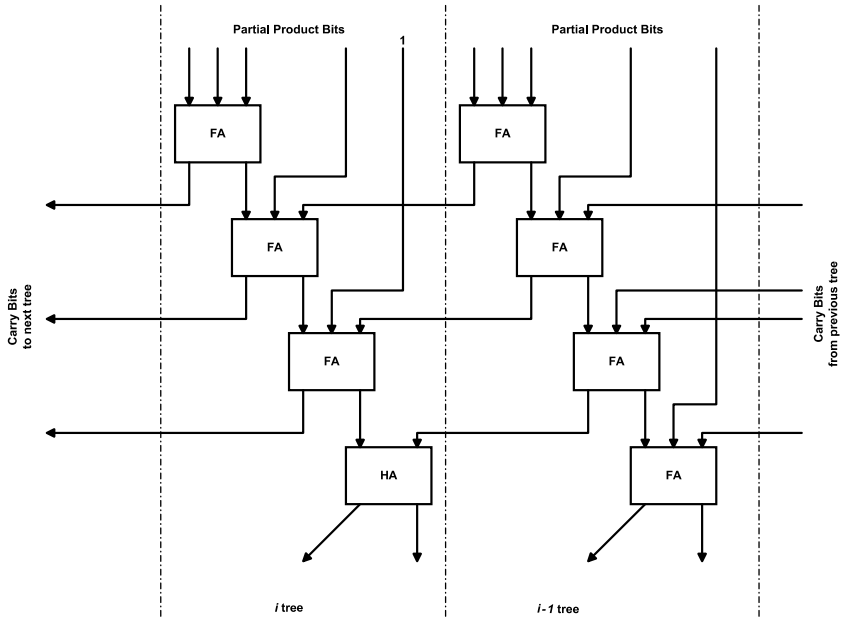

Fig. 3. Part of a Wallace tree unit designed according to Rules I–IV.

Rule IV: The sign extension bits are summed either at the last or the previous to the last level of each tree. In the latter case a half adder should not be used.

Rule V: Carry bits that are the outputs of trees which sum a lot of carry bits of less significance should be propagated at the highest possible level of the succeeding tree and if possible added with the outputs of subtrees that receive only a very small number of carry bits.

An example demonstrating Rules I–IV is given in Fig. 3. We have verified the validity of these design rules by constructing various multipliers (with operand sizes of $n = 8, 12, 16, 24, 32$). Their HDL descriptions can be found in [4].

3. Low power dissipation during testing

3.1. Preliminaries

The primary source of PD in CMOS circuits is charging and discharging of the capacitances. It is expected that by reducing the number of transitions at the primary inputs of a circuit, the total number of transitions at its internal lines will also be reduced, leading to decreased PD. However, depending on the circuit structure, the transitions at some primary inputs may cause more transitions at internal lines than those at other primary inputs.

In [10,18] a procedure has been presented for identifying those primary inputs that cause more transitions at internal lines. Let $f(l)$ denote the function of line l, and $\partial f(l)/\partial \mathrm{in}_i$ the Boolean difference of $f(l)$ with respect to primary input in_i. The latter function indicates whether $f(l)$ is sensitive to changes of input in_i. Let $P(\partial f(l)/\partial \mathrm{in}_i)$ denote the probability that function $\partial f(l)/\partial \mathrm{in}_i$ evaluates to 1. The PD is then estimated as:

$$PD = (1/2)V_{dd}^2 \sum_l C_l \sum_i P\left(\frac{\partial f(l)}{\partial \mathrm{in}_i}\right) T(\mathrm{in}_i) \qquad (1)$$

with C_l denoting the capacitance of line l, V_{dd} the power supply voltage and $T(\mathrm{in}_i)$ the number of transitions of the primary input in_i. Therefore, if we assign a weight w to every primary input in_i such that $w(\mathrm{in}_i) = \sum_l C_l P(\partial f(l)/\partial \mathrm{in}_i)$, then these weights are a good metric of how many lines of the circuit, weighted by their associated capacitance, are affected by primary input in_i.

Relation (1) implies that the PD can be reduced by cutting down the number of transitions at the inputs of the circuit. The reduction is larger when the number of transitions at the inputs with greater weights is reduced. Therefore, the assignment of the TPG outputs to the circuit inputs is significant. The reduction of the cardinality of the test set will also reduce the total number of transitions and thus the PD.

3.2. Assignment of the TPG outputs to the multiplier inputs

In this subsection, we address the problem of properly assigning the TPG outputs to the multiplier inputs for achieving low PD. The error aliasing calculation of the test response compactor (TRC) circuit and the estimation of the PD during testing enforce us to consider specific cell implementations. Furthermore, since we assume the cell fault model, more than one cell implementations must be taken into account. Specifically, three distinct implementations of the half and full adder cells, presented respectively in [11,20,1] were considered. We will refer to these implementations as Cell 1, Cell 2 and Cell 3 respectively. The considered implementations for the Booth encoding and the partial product formation logic are given in Figs. 4 and 5 respectively.

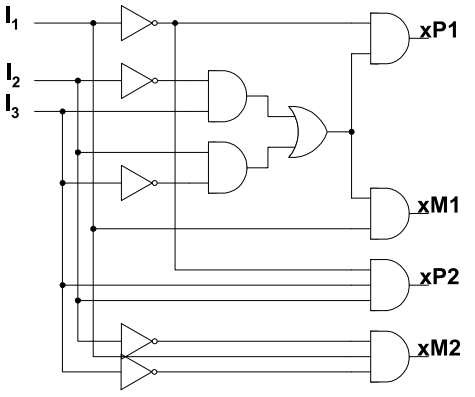

Fig. 4. The implementation of the Booth encoding cell (r-cell).

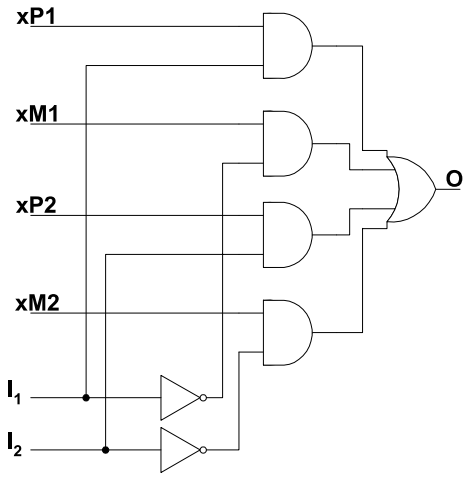

Fig. 5. The implementation of the partial product formation cell (pp-cell).

For each one of the specific cell implementations we computed the primary inputs weights for multipliers of various sizes and we have verified that their distribution is independent of the specific implementation. Comparing any possible pair of inputs, the input with the larger weight contributes more than the other to the PD. Since the sum of weights of B inputs is greater than the sum of weights of A inputs, the five most significant outputs of the TPG should drive the B inputs while its three least significant outputs should drive the A inputs.

The outputs of the TPG are repeatedly assigned to both A and B multiplier inputs. Therefore in order to assign them to specific multiplier inputs we sum the weights of the inputs that receive the same TPG output bit. The results for the sum of weights for $n \times n$ multipliers with $n = 8$, 16 or 32 are listed in Table 1.

For maximum reduction of the PD, the TPG outputs with the smallest number of transitions should be assigned to the inputs with the largest sum of weights. This assignment is denoted as "best assignment". In Table 2 we present the savings in PD of the "best assignment" against a "random assignment". Depending on the specific cell implementation and the size of the multiplier, "best assignment" can lead to PD savings from 2.0% up to 13.1%. For obtaining the results of Table 2, we have used the zero-delay gate level power simulator developed in [2]. This power simulator estimates the PD of the whole circuit consisting of both the multiplier and the BIST circuitry.

We can further reduce the number of transitions at the primary inputs of the multiplier by using as TPG a Gray instead of a binary counter. Therefore we decided to encode in Gray both the three and five output bits of the 8-bit binary counter that drive A and B multiplier inputs respectively. Table 2 lists the savings in PD that can be accomplished by using both "best assignment" as well as a Gray

Table 1
Sum of weights of the multiplier inputs

	Sum of weights for input B					Sum of weights for input A		
	$\sum_{i=0}^{n/5} w(b_{(5i+4)})$	$\sum_{i=0}^{n/5} w(b_{(5i+3)})$	$\sum_{i=0}^{n/5} w(b_{(5i+2)})$	$\sum_{i=0}^{n/5} w(b_{(5i+1)})$	$\sum_{i=0}^{n/5} w(b_{(5i+0)})$	$\sum_{i=0}^{n/3} w(a_{(3i+2)})$	$\sum_{i=0}^{n/3} w(a_{(3i+1)})$	$\sum_{i=0}^{n/3} w(a_{(3i+0)})$
8 × 8 multiplier:								
Cell 1	88	115	167	183	192	172	244	241
Cell 2	129	159	234	251	268	243	337	332
Cell 3	116	151	219	235	247	227	315	310
16 × 16 multiplier:								
Cell 1	757	821	781	844	1002	1211	1196	1395
Cell 2	1222	1307	1253	1354	1621	1947	1910	2210
Cell 3	1096	1190	1137	1226	1455	1755	1731	2002
32 × 32 multiplier:								
Cell 1	4372	4463	4499	5011	5019	6851	7375	7389
Cell 2	7614	7772	7857	8722	8755	11 901	12 752	12 822
Cell 3	6800	6960	7019	7785	7791	10 629	11 400	11 415

Table 2
Percentages of total PD reduction using best assignment and Gray encoding

	8×8			16×16			32×32		
	Cell 1	Cell 2	Cell 3	Cell 1	Cell 2	Cell 3	Cell 1	Cell 2	Cell 3
Best assignment	13.1	12.7	12.3	8.1	7.8	7.5	2.4	2.0	2.1
Best assignment + Gray encoding	30.5	29.7	29.0	28.7	27.0	26.5	24.7	22.0	22.0

counter TPG. As can be observed by Table 2 the use of Gray code can increase the PD savings by approximately 20% in all cases.

3.3. Test length reduction

It is well known that, in BIST schemes, some vectors generated by the TPG circuits are not useful for testing purposes. Therefore another way for reducing the power dissipated during the test application is to reduce the number of vectors applied to the circuit under test.

A straightforward approach to this problem is to use an ATPG tool along with a test set compression program. The result that would be obtained by this method would be an optimal test set in terms of cardinality but also totally inappropriate for implementation as a TPG circuit. This is because the vectors that would be selected would have no straight correlation between them. Therefore, a specially designed circuit would be required for their generation (in most cases this would be a finite state machine). Such circuits apart from requiring large implementation area may also destroy any possible PD gains attained by the elimination of the original redundant test vectors. Another approach, that is not so straightforward, is to modify the original TPG so as to only go through some of its states. This solution does not guarantee to generate an optimal number of test vectors but, as long as the original state sequence is not disrupted a lot, it can be implemented with low area overhead over the original BIST TPG.

In this paper we follow the second approach. Our aim is to reduce the number of test vectors required for exhaustively testing the basic cells of the multiplier (Booth encoding cells and Wallace tree full and half adders) except for the CLA adder at the last stage. By using a logic simulator, we construct a table with 256 columns and X rows, with X indicating the number of cells used in the multiplier. The content of each cell of the table indicates the input combination that a test vector applies to the corresponding cell. The "best assignment" and the Gray encoding of the TPG define the order of the test vectors (columns of the table). Distinct orderings of the test vectors obviously lead to distinct results.

The next goal is to find a subset of columns in this table, able to apply all the input combinations at every cell. Since the cardinality of the initial test set is small (256 test vectors) the computational complexity of such a procedure is also small. Furthermore, since in our case the original TPG is a counter, this means that we must select large subsets of consecutive columns. To do so, we divide the test vector sequence into n groups $(G_1, G_2, \ldots, G_n)$ and mark each one of them as redundant or non-redundant (see Fig. 6). We start by dividing the 256 columns into 8 groups of 32 consecutive vectors each. The idea behind this selection is that if a whole group of 32 vectors can be omitted, then this corresponds to the skipping of four consecutive states of the 5-bit counter driving input B of the multiplier. Therefore it can be easily implemented without increasing substantially the area of the TPG circuit. Then, the same procedure is executed for groups of 16 or 8 vectors each. We stopped our iterations at groups consisting each of 8 vectors. Smaller groups would require prohibitively large implementation area for the TPG. The above leads to a TPG area vs. PD reduction tradeoff. For example, for the 16×16 multiplier one can make a selection between the removal of 1 group of 32 vectors or the removal of 5 groups of 16 vectors (see Table 3). The first solution will result in a smaller area overhead but will also lead to a smaller reduction in the PD.

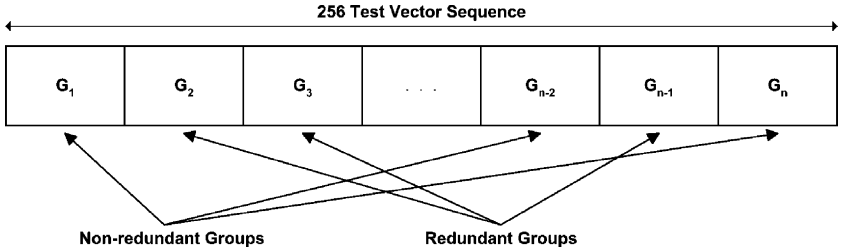

Fig. 6. The n groups of the 256 vector sequence are marked as redundant or non-redundant.

Table 3
Number of redundant groups vs. group size for 8×8, 16×16 and 32×32 multipliers

Group size	Number of groups	Number of redundant groups		
		8×8 multiplier	16×16 multiplier	32×32 multiplier
32	8	0	1	0
16	16	6	5	6
8	32	20	20	18

Examining the various alternatives produced by the above procedure, we found that the removal of groups of eight consecutive vectors leads to good results with a reasonable area increase. For the 16×16 and 32×32 multipliers there are respectively 20 and 18 redundant groups of 8 vectors while the additional area overhead of the TPG circuit is negligible (1.0% in 16×16 and 0.3% in 32×32 multiplier). Therefore, the test set cardinality can be reduced from the original 256 to 96 and 112 vectors respectively.

Although the reduced test sets guarantee exhaustive testing of the basic cells, the FC may drop due to increased error aliasing and undetected faults of the CLA unit for which we consider the single stuck-at fault model. Therefore, we need to verify that the FC obtained, with respect to single stuck-at faults, by the reduced test sets remains at high levels. Table 4 lists the FC achieved, assuming either a rotate carry adder or cascaded compaction [14] as the TRC, for 8×8, 16×16 or 32×32 multipliers and the three cell implementations.

From Table 4 we can see that there is a slight increase in the number of undetected faults with the proposed test set compared to the application of the 256 initial vectors. These are located at the CLA unit. Moreover, we can see that due to in-

Table 4
FC percentages

	8×8			16×16			32×32		
	Cell 1	Cell 2	Cell 3	Cell 1	Cell 2	Cell 3	Cell 1	Cell 2	Cell 3
Test set with 256 vectors									
FC without compaction	99.82	99.84	99.84	99.45	99.55	99.55	99.53	99.64	99.64
FC with rotate carry adder	98.72	98.89	98.89	98.99	99.18	99.18	99.29	99.46	99.46
FC with cascaded compaction	99.36	99.44	99.44	99.26	99.40	99.40	99.49	99.61	99.61
Test set with 96 vectors for 8×8 and 16×16 and 112 vectors for 32×32									
FC without compaction	99.54	99.60	99.60	99.20	99.35	99.35	99.38	99.52	99.52
FC with rotate carry adder	98.67	98.89	98.89	98.71	98.97	98.95	99.08	99.33	99.31
FC with cascaded compaction	99.31	99.40	99.40	99.16	99.30	99.33	99.37	99.52	99.52
Test set with 168 vectors for 8×8, 16×16 and 32×32									
FC without compaction	99.73	99.76	99.76	99.30	99.43	99.43	99.44	99.57	99.57
FC with rotate carry adder	98.81	98.97	98.97	98.99	99.18	99.18	99.16	99.36	99.36
FC with cascaded compaction	99.45	99.52	99.52	99.24	99.38	99.38	99.43	99.56	99.56

creased error aliasing the FC may drop below the acceptable level of 99% when a rotate carry adder is used as TRC. The use of a more effective accumulator-based TRC such as cascaded compaction solves the aliasing problem in all cases.

One may also want to have a uniform test set for all examined $n \times n$ multipliers ($n = 8$, 16 or 32) generated by a uniform TPG. We derived such a test set consisting of 168 test vectors. The FC attained by this test set is presented also in Table 4 and is in all cases above the acceptable level of 99% when cascaded compaction is used. We have to note that in 32×32 multipliers, the rotate carry adder test response compaction scheme provides always FC greater than 99%.

Table 5 presents the savings in total, average per vector and peak PD when the reduced test sets along with "best assignment" and Gray encoding are used. The total power reduction achieved varies from 64.8% to 72.8%. The average PD reduction per vector varies from 19.6% to 27.4%, while the reduction of the peak PD varies from 16.8% to 36.0%. The test application time is also reduced from 56.3% to 62.5%.

To obtain the above comparison results, our gate-level simulator assumes a zero-delay gate model. The authors of [16] report that there is a correlation between the PD of a circuit assuming a zero-delay and the PD assuming a general delay model. Hence, using a zero-delay approximation is reasonable. The reductions in the total power dissipated are expected to be even greater if glitches are also taken into account, since the switching activity of the internal nodes of the multiplier is reduced during the application of the reduced test sets.

Table 6 presents the savings in PD that can be obtained by the uniform test set when combined

Table 5
PD reduction percentages using best assignment, Gray encoding and test set reduction

Multiplier (number of vectors)		Total power reduction (%)	Average power reduction per vector (%)	Peak power reduction (%)
8×8 (96)	Cell 1	72.8	27.4	17.5
	Cell 2	72.4	26.3	24.8
	Cell 3	72.1	25.5	16.8
16×16 (96)	Cell 1	72.6	26.8	24.1
	Cell 2	72.0	25.3	27.1
	Cell 3	71.7	24.4	22.4
32×32 (112)	Cell 1	66.1	22.5	31.0
	Cell 2	64.9	19.8	36.0
	Cell 3	64.8	19.6	30.2

Table 6
PD reduction percentages using best assignment, Gray encoding and uniform test set of 168 vectors

Multiplier		Total power reduction (%)	Average power reduction per vector (%)	Peak power reduction (%)
8×8	Cell 1	52.5	27.6	24.0
	Cell 2	51.9	26.6	25.4
	Cell 3	51.2	25.7	21.1
16×16	Cell 1	51.7	26.4	24.1
	Cell 2	50.6	24.7	27.1
	Cell 3	50.1	23.9	22.4
32×32	Cell 1	48.9	22.2	28.4
	Cell 2	47.1	19.4	34.3
	Cell 3	47.0	19.2	27.5

with "best assignment" and Gray encoding. Comparing Tables 5 and 6 we can see that average power per vector and peak power savings are similar in both cases. We can also observe that the total PD savings in this case are smaller but the uniform test set is still able to reduce the total PD by half with respect to the original BIST scheme [13].

4. Conclusions

Fast multipliers based on Booth encoding and Wallace tree summation of the partial products are used as embedded blocks in both general purpose datapath structures and specialized digital signal processors. In this paper we analyzed their testability characteristics and proposed several design rules in order to make their Wallace tree summation unit fully testable under the cell fault model when the multiplier inputs are driven by an 8-bit counter TPG. As quality and cost related issues make the low power feature a necessity for BIST schemes, we presented a methodology for deriving a novel low power BIST scheme. We showed how the low power objective can be achieved by: (a) proper assignment of the TPG outputs to the multiplier inputs, (b) the use of Gray counter TPG and (c) the reduction of the test set. Our methodology is able to provide the BIST designer with alternate solutions, each having distinct power reduction, implementation area and FC characteristics.

References

[1] M. Abramovici, M. Breuer, A. Friedman, Digital Systems Testing and Testable Design, Computer Science Press, 1990.

[2] D. Bakalis, D. Nikolos, On low power BIST for carry save array multipliers, in: Proceedings of the 5th International On-Line Testing Workshop, 1999, pp. 86–90.

[3] D. Bakalis, H.T. Vergos, D. Nikolos, X. Kavousianos, G. Alexiou, Low power dissipation in BIST schemes for modified Booth multipliers, in: Proceedings of the International Symposium on Defect and Fault Tolerance in VLSI Systems, 1999, pp. 121–129.

[4] D. Bakalis, E. Kalligeros, D. Nikolos, H.T. Vergos, G. Alexiou, Low power BIST for Wallace tree-based fast multipliers, Computer Technology Institute Technical Report (TR99/09/07), Patras, Greece, 1999.

[5] V. Dabholkar, S. Chakravarty, I. Pomeranz, S. Reddy, Techniques for minimizing power dissipation in scan and combinational circuits during test application, IEEE Transactions on CAD of Integrated Circuits and Systems 17 (12) (1998) 1325–1333.

[6] P. Girard, L. Guiller, C. Landrault, S. Pravossoudovitch, A test vector inhibiting technique for low energy BIST design, in: Proceedings of the VLSI Test Symposium, 1999, pp. 407–412.

[7] P. Girard, Low power testing of VLSI circuits: problems and solutions, in: Proceedings of the 1st International Symposium on Quality Electronic Design, 2000, pp. 173–179.

[8] P. Girard, L. Guiller, C. Landrault, S. Pravossoudovitch, Low power BIST design by hypergraph partitioning: methodology and architectures, in: Proceedings of the International Test Conference, 2000, pp. 652–661.

[9] W. Kautz, Testing for faults in cellular logic arrays, in: Proceedings of the 8th Annual Symposium on Switching and Automata Theory, 1967, pp. 161–174.

[10] X. Kavousianos, D. Nikolos, S. Tragoudas, On-chip deterministic counter-based TPG with low heat dissipation, in: Proceedings of the Southwest Symposium on Mixed-Signal Design, 1999, pp. 87–92.

[11] M. Morris Mano, Digital Design, second ed., Prentice-Hall, 1991.

[12] R. Parkar, Bare die test, in: Proceedings of the Multi-Chip Module Conference, 1992, pp. 24–27.

[13] A. Paschalis, D. Gizopoulos, N. Kranitis, M. Psarakis, Y. Zorian, An effective BIST architecture for fast multiplier cores, in: Proceedings of the Design, Automation and Test in Europe Conference, 1999, pp. 117–121.

[14] J. Rajski, J. Tyszer, Arithmetic Built-In Self-Test for Embedded Systems, Prentice-Hall, 1998.

[15] R. Roy, Scaling Towards Nanometer Technologies: Design for Test Challenges, Panel, Design Automation and Test in Europe Conference, 1999.

[16] A. Shen, A. Ghosh, S. Devadas, K. Keutzer, On average power dissipation and random pattern testability of CMOS combinational logic networks, in: Proceedings of the International Conference on CAD, 1992, pp. 402–407.

[17] A. Steininger, Testing and built-in self-test—A survey, Journal of Systems Architecture 46 (9) (2000) 721–747.

[18] S. Wang, S.K. Gupta, DS-LFSR: a new BIST TPG for low heat dissipation, in: Proceedings of the International Test Conference, 1997, pp. 848–857.

[19] S. Wang, S.K. Gupta, ATPG for heat dissipation minimization during test application, IEEE Transactions on Computers 47 (2) (1998) 256–262.

[20] N.H.E. Weste, K. Eshragian, Principles of CMOS VLSI Design: A Systems Perspective, second ed., Addison-Wesley, 1992.

[21] X. Zhang, K. Roy, Design and synthesis of programmable low power weighted random pattern generator, in: Proceedings of the 5th International On-Line Testing Workshop, 1999, pp. 82–85.

[22] Y. Zorian, A distributed BIST control scheme for complex VLSI devices, in: Proceedings of the VLSI Test Symposium, 1993, pp. 4–9.

Dimitris N. Bakalis received the Diploma degree in 1995, the M.Sc. degree in 2000 and the Ph.D. degree in 2001 in Computer Engineering, all from the Department of Computer Engineering and Informatics at the University of Patras in Greece. His Ph.D. Thesis was in the area of built-in self-test. His research interests include low power VLSI design and test. He is a member of the IEEE.

Emmanouil I. Kalligeros received the Diploma degree in 1999 and the M.Sc. degree in 2001 in Computer Engineering, both from the Department of Computer Engineering and Informatics at the University of Patras in Greece. He is now pursuing his Ph.D. in the area of built-in self-test at the same department. His research interests also include delay fault testing and VLSI design.

Dimitris Nikolos received the B.Sc. degree in Physics, the M.Sc. degree in Electronics and the Ph.D. degree in Computer Science, all from the University of Athens. He is currently a Professor in the Computer Engineering and Informatics Department of the University of Patras, Director of the Hardware and Computer Architecture division and head of the Technology and Computer Architecture Laboratory. He has served as Program Co-Chairman of five (1997–2001) IEEE International On-Line Testing Workshops. He also served on the program committees for the IEEE International Symposiums Defect and Fault Tolerance in VLSI systems (1997–1999), for the Third and Fourth European Dependable Computing Conference and for the DATE (2000–2002) Conferences. He was guest Co-Editor of the June 2002 special issue of Journal of Electronic Testing, Theory and Applications (JETTA) devoted to the 2001 IEEE International On-Line Testing Workshop. His main research interests are fault-tolerant computing, computer architecture, VLSI design, test and design for testability. Prof. Nikolos has authored or co-authored more than 110 scientific papers and was co-recipient of the Best Paper Award for his work "Extending the Viability of IDDQ Testing in the Deep Submicron Era" presented at the 3rd IEEE International Symposium on Quality Electronic Design (ISQED 2002). He is a member of the IEEE.

Haridimos T. Vergos received his Diploma in Computer Engineering and his Ph.D. degree in 1991 and 1996 respectively. In 1998, he worked on the development of the first IEEE 802.11 wireless network standard compliant MAC processor for the Multimedia and Networking Group of Atmel Inc. He currently, holds a Lecturer position in the Computer Engineering and Informatics Department, of the University of Patras, Greece. Dr. Vergos has authored more than 30 scientific papers on Computer Architecture, Computer Arithmetic and on the Design and Test of Fault Tolerant Digital Systems. He is a member of the Greek Computer Engineering Society and the Technical Chamber of Greece.

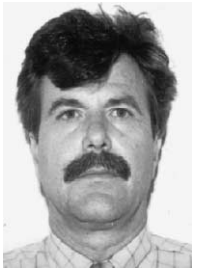

George Ph. Alexiou received a B.Sc. in Physics in 1976 and a Ph.D. in Electronics in 1980, both from the University of Patras, Greece, where he is now a professor in the Department of Computer Engineering and Informatics and Director of the Microelectronics (VLSI) Lab. He has served on all program committees of IEEE International Symposium of Quality Electronics Design (ISQED) starting 2000. He also served in the program committee of IEEE Rapid System Prototyping Workshop starting 2000. He is publishing papers in a number of international journals and conference proceedings. His research interests include VLSI design, VLSI CAD tools, signal processing, digital systems and RF data communications. He is a member of the IEEE and ACM.