



HAL
open science

Scan Cell Ordering for Low Power Scan Testing

Yves Bonhomme, Patrick Girard, Christian Landrault, Serge Pravossoudovitch

► **To cite this version:**

Yves Bonhomme, Patrick Girard, Christian Landrault, Serge Pravossoudovitch. Scan Cell Ordering for Low Power Scan Testing. ETW: European Test Workshop, May 2002, Corfu, Greece. lirmm-00269337

HAL Id: lirmm-00269337

<https://hal-lirmm.ccsd.cnrs.fr/lirmm-00269337>

Submitted on 21 Jan 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Scan Cell Ordering for Low Power Scan Testing

Y. Bonhomme P. Girard C. Landrault S. Pravossoudovitch

*Laboratoire d'Informatique de Robotique et de Microélectronique de Montpellier,
Université Montpellier II / CNRS*

161 rue Ada, 34392 Montpellier Cedex 5 France

Email: <name>@lirmm.fr

URL :http://www.lirmm.fr/~w3mic

Abstract

Power consumption during scan testing is becoming a primary concern. In this paper, we present a novel approach for scan cell ordering which significantly reduces the power consumed during scan testing. The proposed approach is based on the use of a two steps heuristic procedure that can be exploited by any chip layout program before flip-flops placement and routing. The proposed approach works for any conventional scan design and offers numerous advantages compared with existing low power scan techniques. Reductions of average and peak power consumption during scan testing are up to 34% and 18% respectively for experimented ISCAS benchmark circuits.

1. Introduction

The System-on-Chip (SOC) revolution has brought some new challenges to both design and test engineers. Among these challenges, the power dissipation is one of the most important issues [1]. Generally, a circuit may consume more power in the test mode than in the normal mode due to the following reasons [2]. First, the design-for-testability (DfT) circuitry embedded in a circuit to reduce the test complexity is often idle during normal operations but may be intensively used in the test mode. Second, the test efficiency has been shown to have a high correlation with the toggle rate; hence in the test mode, the switching activity of all nodes is often several times higher than the activity during normal operations. Third, in a SOC, parallel testing is frequently employed to reduce test application time, which may result in excessive energy and power dissipation. This elevated test power may be responsible for several kinds of problems: instant circuit damage, increased product costs, decreased system reliability, performance degradation, reduced autonomy of portable systems and decrease of overall yield. A survey of these problems is given in [3].

The focus of this paper is on the problem of minimizing power dissipation during external scan testing, *i.e.* from an ATE (Automatic Test Equipment). Scan architectures are very popular and are commonly used to test digital circuitry in integrated circuits (ICs) or cores. However, scan-based architectures are expensive in power consumption as each test pattern requires a large number of shift operations with a high circuit activity [4]. Of

course, it is always possible to reduce average power during scan testing by simply scanning at a lower frequency. However, this increases test application time. Another solution is to add logic to hold the output of the scan cells at a constant value during scan shifting [5]. The drawbacks of this approach are the area overhead and the performance degradation that it incurs. Some other solutions have been proposed recently to cope with the power problem during scan testing: low power ATPGs [6,7], a scan path segmentation technique [8,9], a static compaction technique [10], two clock scheme modification techniques [11,12], an interleaving scan architecture for multiple-scan circuits [13], a test data compression technique for SOC [14], test scheduling techniques [15,16], ...

Actually, a simple alternative solution for minimizing power consumption during scan testing is to use test vector ordering or scan cell ordering techniques. Test vector ordering has been investigated in [17,18,19] with the objective to define the order in which test vectors of a deterministic test set have to be applied to the circuit or core under test (CUT) to minimize the overall switching activity. Scan cell ordering has been investigated only in [17], where two heuristics are proposed to determine the order in which the scan flip-flops of a given scan chain have to be connected: a random ordering heuristic and a simulated annealing algorithm. Experimental results on small benchmark circuits show that scan cell ordering can reduce test power by 10-25% with no change in terms of fault coverage and test length. Larger benchmark circuits have not been experimented because they are *intractable* with the proposed heuristics, but an important comment reported in [17] is that for large circuits with a huge number of scan elements, scan cell ordering combined with test vector ordering is much more efficient than test vector ordering alone. Hence, scan cell ordering must be selected as the first test power minimization strategy.

In this paper, we propose a novel approach for scan cell ordering which significantly reduces the power consumed during scan testing. The proposed approach is based on the use of a simple procedure that can be exploited by any chip layout program before flip-flops placement and routing. The inputs to this procedure are i) a given set of scan flip-flops and ii) a sequence of deterministic test vectors with the corresponding output responses. The output is an ordered scan chain with minimum average and

peak power consumption. To tackle this NP-hard problem efficiently, the heuristic procedure operates in two steps: the first one consists in determining the chaining of the scan cells so as to minimize the occurrence of transitions in the scan chain during shifting operations, the second one consists in identifying the input and output scan cells of the scan chain to limit the propagation of transitions during scan operations. An important feature of the proposed approach is that both scan-in and scan-out transitions (and not only scan-in transitions due to test vectors) are considered for solving the scan cell ordering problem. A comparison of the heuristic procedure with an exact method (on small benchmark circuits) demonstrates the efficiency of the proposed solution.

Compared with existing low power scan techniques, our solution offers numerous advantages. The proposed approach works for any conventional scan design - no extra *DfT logic* is required. The fault coverage and the IC test time are left unchanged by the proposed approach. There is almost no penalty on the circuit performance. The area overhead, which is due to routing area, may be kept acceptably low when managed by an efficient chip layout program. The proposed approach is very easy to use in a classical DfT flow and has therefore a very low impact on the system design time. Reductions of average and peak power consumption during scan testing are up to 34% and 18% respectively for experimented ISCAS circuits. Note that the approach described in this paper can be used for full scan cores or ICs having one or multiple scan chains. It can also be combined with a scan vector ordering technique to provide additional reductions in test power.

The remainder of the paper is organized as follows. Section 2 describes how power is estimated in the proposed approach. Section 3 is devoted to the description of the proposed scan cell ordering procedure. Results obtained on the benchmark circuits are reported in Section 4. A discussion on layout considerations during scan cell ordering are finally given in Section 5.

2. Power estimation for scan vectors

In a CMOS circuit, the predominant fraction of the power is dissipated when the circuit elements switch from logic 1 to 0 or vice versa. Static power consumption, which is caused by leakage current, is usually negligible and therefore ignored. During test, the elements in the circuit switch when the primary inputs change value or when the scan cells change value. In this paper, we assume that the primary inputs are controllable directly from the tester and hence are held constant during shifting operations. All the switching activity in the circuit is due to transitions in the scan cells.

Consider a CUT with five flip-flops and the test vector 10001 being scanned in. Let the scan flip-flops be initially set to 00000. After the first clock cycle when the first bit has been scanned in, the scan flip-flops will be in state 10000. The state of the first flip-flop has changed from 0

to 1. This change will cause gates in the CUT to switch. The number of circuit elements that switch depends on the actual circuit. Next, the second bit is scanned in and hence both the first and second flip-flops switch. This process continues until the complete test vector has been scanned in. The test vector is then applied to the CUT and the output response is captured back in the scan chain. As the next test vector is scanned in, the output response from the previous vector is scanned out to the tester. Transitions in the output response being scanned out will also cause switching activity. Hence, we can divide the power dissipation during scan testing into two parts. The scan-in power, which is due to transitions in the test vectors during scan-in, and the scan-out power, which is due to transitions in the output responses during scan-out.

To estimate the scan-in power dissipated by a test vector (or the scan-out power dissipated by the corresponding output response), a weighted transitions metric has been introduced in [10]. This weighted transitions metric models the fact that the scan-in power for a given test vector (or the scan-out power for the output response) depends not only on the number of transitions in it but also on their relative positions. For example, consider the test vector $b_1b_2b_3b_4b_5 = 10001$ where the bit b_5 is first loaded into the scan chain. This vector has two transitions: transition t_1 between b_5 and b_4 , and transition t_2 between b_2 and b_1 . When this vector is scanned in, transition t_1 passes through the entire scan chain. This transition dissipates power at every flip-flop in the scan chain. On the other hand, transition t_2 dissipates power only at the first scan flip-flop during scan-in and hence causes less switching activity in the scan chain than transition t_1 .

This weighted transitions metric has been shown to be strongly correlated to the switching activity at the internal nodes of the CUT during scan-in and scan-out operations. It was shown experimentally in [10] that scan vectors (test vectors or output responses) with higher weighted transitions dissipate more power in the CUT. This metric is therefore a good way to accurately estimate the power consumed during scan testing and hence avoid time-consuming and size limited simulations.

In the proposed ordering procedure, when we need to estimate the scan-in power and the scan-out power dissipated by a complete test sequence, we use the number of weighted transitions as a measure. More formally, the number of weighted transitions in a test vector or in an output response is given by:

$$\text{Weighted_Transitions} = \sum (\text{Size_of_Scan_Chain} - \text{Position_of_Transition})$$

Note that according to one considers a test vector or an output response, the intrinsic value of *Position_of_Transition* changes. For example, consider the scan vector $b_1b_2b_3b_4b_5 = 00001$. In the case where this vector is a test vector, the transition between b_5 and b_4 is in position 1 (b_5 is the first bit being scanned in). In the case where

this vector is an output response, the position of the transition is 4 (b5 is the first bit being scanned out).

The scan-in power can be estimated by summing the weighted transitions of all test vectors in the test sequence. The scan-out power can be estimated by summing the weighted transitions of all output responses. The test power dissipated by the test sequence can be therefore estimated by adding the values of the scan-in power and the scan-out power.

For the estimation of the test power to be correct, an additional term has to be considered. This term corresponds to the power due to the transition that occurs when the first bit of a test vector differs from the last bit of the previous output response. In this case, the transition propagates through the entire scan chain, and the weight assigned to it is equal to the size of the scan chain.




Figure 1: An example set of test vectors and output responses

For example, consider the test sequence shown in Figure 1, which is composed of three test vectors and the corresponding three output responses. The scan chain has four flip-flops, hence scan vectors are four-bit long. The number of weighted transitions for test vector V1 is $(4-3) + (4-1) = 4$, and that of output response R1 is $(4-1) + (4-2) = 5$. Note that the intrinsic value of the position is different for test vectors and output responses. For test vector V2 and output response R2, the numbers of weighted transitions are 6 and 3 respectively. For test vector V3 and output response R3, these numbers are 1 and 5 respectively. In addition, a transition will propagate through the entire scan chain when the first bit of test vector V2 will be scanned in. This comes from the difference between the value of the first bit of V2 and the value of the last bit of output response R1. The weight associated to this transition is 4, which corresponds to the size of the scan chain. Assuming an initial state 0000 for the scan flip-flops, the total number of weighted transitions produced by the test sequence is $(4+5+6+3+1+5) + (4) = 28$. For comparison purpose in the ordering procedure, this number gives an estimation of the test power dissipated by a test sequence.

3. Scan cell ordering procedure

In this section, we show how test power can be minimized by appropriately ordering the scan cells of a given scan chain. The inputs to the proposed procedure are i) a given

set of scan flip-flops and ii) a sequence of deterministic test vectors with the corresponding output responses. The output is an ordered scan chain with minimum test power. To tackle this NP-hard problem efficiently, the heuristic procedure operates in two steps: the first one consists in determining the chaining of scan cells, the second one consists in identifying the input and output scan cells of the scan chain. These two steps are now described in details.

3.1. Determining the scan cell chaining

The first step of the scan cell ordering procedure consists in determining the order in which the scan cells have to be connected to minimize the occurrence of transitions in the scan chain during scan-in and scan-out operations. To this end, we first consider the set of scan vectors (test vectors and output responses) used during scan testing, and we assume an initial order for the scan flip-flops. More formally, we assume that flip-flop 1 corresponds to bit 1 of each scan vector, flip-flop 2 corresponds to bit 2 of each scan vector, ..., flip-flop n corresponds to bit n of each scan vector. For example, consider the test sequence shown in Figure 2, which is composed of four test vectors (V1 to V4) and four output responses (R1 to R4). The scan chain has four flip-flops, hence scan vectors are four-bit long. The initial order of the scan cells in the scan chain is depicted on the figure. According to the above description, flip-flop 1, denoted as ff1, corresponds to bit 1 in each scan vector, flip-flop 2, denoted as ff2, corresponds to bit 2, and so on (Figure 2.a).




Figure 2: The example test sequence and the weighted graph

Next, we calculate the total number of bit differences between each pair of scan flip-flops, which represents the number of transitions that may be generated in the corresponding portion of the scan chain by connecting these two flip-flops together. For the example in Figure 2, calculating the total number of bit differences between each pair of flip-flops (on the complete sequence of test vectors and output responses) provides the following results: $d(ff1, ff2) = 6$, $d(ff1, ff3) = 4$, $d(ff1, ff4) = 2$, $d(ff2, ff3) = 5$, $d(ff2, ff4) = 4$, $d(ff3, ff4) = 5$.

From these values of the bit differences between flip-flops, it is then possible to construct a complete undirected graph in which each vertex represents a flip-flop and each edge represents a possible connection between two flip-flops (Figure 2.b). The weight on each edge of the graph

represents the total number of bit differences between two flip-flops for the complete test sequence, and reflects the number of transitions that may be generated in the corresponding portion of the scan chain by connecting these two flip-flops together.

From this weighted graph, the problem then amounts to finding an Hamiltonian cycle of minimum cost in the graph. The cost of a cycle is obtained by summing the weights on edges belonging to this cycle. This problem is equivalent to the well known traveling salesman problem, which is well known to be NP-hard (the number of possible solutions is $(n-1)!/2 - n$ being the scan chain length) and for which different polynomial-time approximation algorithms can be used [20]. Among these solutions, greedy algorithms represent a good tradeoff between computation time and efficiency of the computed solution. We therefore implemented an heuristic solution based on a greedy algorithm (**with a complexity equal to $O(n^2)$**) to find the scan cell chaining that minimizes the occurrence of transitions in the scan chain during scan-in and scan-out operations. The greedy algorithm starts from an initial state which is always scan cell ff1 in our case (it is reported in [21] that the choice of the initial state is not so crucial in a greedy algorithm considering the number of vertices in the graph to be sufficiently high). Next, the algorithm operates in such a way that, at each stage of decision, a subset of the scan cells is dealt with and considered as definitively assigned.

For the example test sequence of Figure 2.a, the solution found by the greedy algorithm is the following chaining of scan cells: ff1-ff4-ff2-ff3-ff1. This solution represents the best solution to minimize the occurrence of transitions in the scan chain during shifting operations.

3.2. Identifying input and output scan cells

From a given chaining of the scan flip-flops, the second step of the ordering procedure consists in defining both the input scan cell and the output scan cell of the scan chain. Appropriately defining the input and output scan cells allows to minimize the propagation of transitions in the scan chain during shifting operations.




Figure 3: The oriented cyclic graph

Let us consider again a scan chain composed of n scan cells. The chaining obtained from step 1 of the ordering procedure can now be represented by a simple oriented cyclic graph in which each vertex represents a scan cell and each edge represents the connection between two scan cells. For illustrative purpose, the oriented cyclic graph for the example test sequence considered in Figure 2.a is depicted in Figure 3. At this stage of the procedure, defining the input and output scan cells of the scan chain

can be done by cutting the graph on a selected edge. The number of possible solutions to cut the graph is obviously equal to n , which also corresponds to the number of different scan chains that can be formed at this stage of the ordering procedure. Each cutting solution differs in the number of transitions that can be generated and propagated in the corresponding scan chain.

In order to find the scan chain that leads to the lowest test power, and hence identify both an input scan cell and an output scan cell, we evaluate the n possible solutions in terms of weighted transitions generated during scan-in and scan-out operations, and then select the solution having the lowest number of weighted transitions. This solution represents the optimal solution to the scan cell ordering problem, hence allowing to definitively identify the less power consuming scan chain.

To this end, we use the weighted transitions metric discussed in Section 3, and evaluate the n possible sets of scan vectors that each corresponds to a given cutting solution. For example, considering the test sequence given in Figure 2.a, the four possible sets of scan vectors which are evaluated to determine the less power consuming scan chain are shown in Figure 4. Note that in each set, the first, third, fifth and seventh vectors are the test vectors. The remaining scan vectors are the output responses.

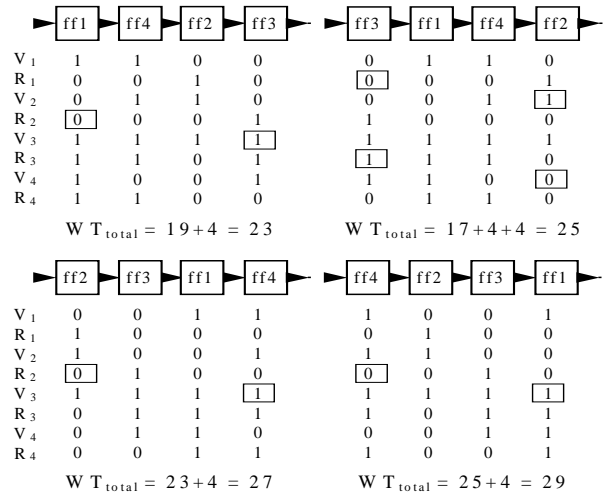


Figure 4: Identification of input and output scan cells

For each set of scan vectors in Figure 4, *i.e.* for each possible scan chain, the total number of weighted transitions has been evaluated and is reported at the bottom of the figure. This number represents the number of weighted transitions in test vectors and output responses plus the number of weighted transitions due to opposite values between the first bit of a test vector and the last bit of the previous output response (*cf.* Section 3). In this example, the lowest value of WT_{total} is 23 and is obtained from the first possible scan chain. Hence, this scan chain (ff1-ff4-ff2-ff3) is the less power consuming scan chain, and ff1 is identified as the input scan cell and ff3 as the output scan cell.

Compared with existing low power scan techniques, our solution offers numerous advantages. The proposed approach works for any conventional scan design - no extra DfT logic is required. The fault coverage and the IC test time are left unchanged by the proposed approach. There is almost no penalty on the circuit performance. The area overhead, which is due to routing area, may be kept acceptably low when managed by an efficient layout synthesis program. The proposed approach is very easy to use in a classical DfT flow and has therefore a very low impact on the system design time.

4. Experimental results

The benchmarking process described here was performed on circuits of the ISCAS'89 [22] benchmark suites. Power consumption in each circuit was estimated by using PowerMill, a dynamic simulator provided by Synopsys [23], assuming a clock frequency equal to 200 MHz and a power supply voltage of 2.5 V. Experiments performed on each circuit have been done with technology parameters extracted from a 0.25 μ m digital CMOS standard cell library. The goal of the experiments we performed has been to measure the savings in test power obtained from the proposed scan cell ordering procedure.

Circuit	# scan cells	# gates	Test length	FC
s298	14	119	33	100 %
s344	15	160	27	100 %
s420	16	196	59	100 %
s510	6	211	68	100 %
s641	19	321	40	100 %
s713	19	351	48	95.48 %
s1423	74	749	44	98.99 %
s5378	179	2225	145	99.05 %
s9234	228	4678.5	249	93.99 %
s13207	669	6395.5	354	98.99 %
s15850	597	7987	279	97.84 %
s35932	1728	16726.5	112	100 %
s38417	1636	20446.5	325	99.70 %

Table 1: Main characteristics of the experimented circuits

First, structural characteristics and test parameters of the experimented circuits are reported in Table 1. All experiments are based on deterministic testing from the ATPG tool "TestGen" of Synopsys [24]. The missing faults in the fault coverage (FC) column are the redundant or aborted faults. The first part of Table 1 shows the number of scan cells and the number of gates for each benchmark circuit. The primary inputs and primary outputs were not included in the scan chain, but were assumed to be held constant during scan-in and scan-out operations. In the second part, we report the test length of each test sequence and the corresponding fault coverage. Note that these results are not influenced by the scan cell ordering procedure described in this paper. The results in terms of power savings achieved by the proposed ordering technique are summarized in Table 2.

For each circuit, we have reported the peak power and the average power obtained first from a random ordering (first part of Table 2) and next with the proposed ordering technique. In all cases, the scan flip-flops have been initialized with random logic values. Results concerning random ordering are an average over ten runs. Peak power and average power are expressed in milliWatts. The last part in Table 2 shows the reduction in peak power and average power dissipation expressed in percentages. These results on benchmark circuits show that peak power reduction of up to 18% and average power reduction of up to 34% can be achieved with the proposed ordering technique. Concerning computing CPU time, ordering solutions are always obtained in less than one hour.

Circuit	Ave. Random Ordering		Proposed Ordering			
	peak [mW]	average [mW]	peak [mW]	average [mW]	peak reduction	average reduction
s641	30.31	1.01	26.55	0.71	12.4 %	29.6 %
s713	33.90	1.09	27.53	0.72	18.8 %	34.1 %
s1423	79.91	3.65	677.25	2.76	15.3 %	24.2 %
s5378	250.5	12.69	250.4	8.97	0.5 %	29.3 %
s9234	395.15	29.45	353.67	22.65	10.5 %	23.1 %
s13207	651.39	48.97	609.9	39.55	6.4 %	19.2 %
s15850	595.63	57.50	555.8	48.76	6.7 %	15.2 %
s35932	1662.02	108.32	1610.5	96.03	3.1 %	11.3 %
s38417	1899.56	135.01	1746.5	126.03	7.9 %	16.5 %

Table 2: Power savings in the CUT

An important comment on these results is that the peak power reduction appears to be irregular and sometimes poor. This is due to the fact that in some cases, the highest current in the CUT appears during application of the test pattern and capture of the corresponding response. As this operation is unchanged with scan cell re-ordering, the reduction in peak power appears to be low. In fact, it is higher than it appears. The reason is that an elevated current consumed in a design can cause damage to the circuit if it occurs during more than one clock cycle. As the switching activity is reduced in the clock cycles preceding and following the capture clock cycle, the peak power is in fact much lower than the peak power without scan cell re-ordering.

Circuit	Branch and Bound		Proposed Ordering		ratio	
	peak [mW]	average [mW]	peak [mW]	average [mW]	peak	average
s298	28.94	0.477	29.27	0.477	-1.13 %	0.0 %
s344	21.25	0.755	21.03	0.755	1.05 %	-0.04 %
s420	29.01	0.169	28.96	0.170	-0.16 %	-0.71 %
s510	30.22	1.097	30.22	1.097	0 %	0 %

Table 4: Comparison between a Branch and Bound technique and the proposed ordering

For the reader's interest, we report in Table 4 comparative results between an exact method for scan cell ordering, *e.g.* the Branch and Bound technique [25], and the

proposed ordering procedure. Results are given only for small circuits because applying the Branch and Bound technique to larger circuits is intractable from a CPU time point of view. Table 4 reports the average and peak power consumed in the CUT during scan testing first from a scan cell ordering produced by a Branch and Bound technique, and next from a scan cell ordering produced by the proposed technique. The two last columns in Table 4 show the ratio between the two techniques. As can be seen, the proposed heuristic procedure based on a greedy algorithm provides almost the same results than those achieved by using an exact method such as the Branch and Bound technique. Considering real-size circuits composed of thousands of gates and flip-flops, for which the scan cell ordering problem is intractable with an exact method, these results highlight the efficiency of the proposed ordering technique.

5. Discussion

In an era of sub-micron technology, routing is becoming a dominant factor in area, timing, and power consumption. Therefore, considering routing impacts when we are looking for a solution to solve the scan cell ordering problem is needed. To incur minimum area overhead and power consumption due to routing, flip-flops should be selected so that chaining of them can lead to a minimum between distance. This is usually the criterion which is employed in any classical scan insertion technique (in addition to the delay criterion). In order to include the routing criterion in our proposed ordering procedure, we need to modify the two metrics used in each step of the procedure: the number of bit differences between each pair of scan flip-flops, and the number of weighted transitions of scan vectors. Modifying these metrics may consist for example in adding a weighting parameter that represents both (or either) area and power consumption overhead. Of course, this concerns only situations in which flip-flops placement is known a priori. In the case where flip-flops placement is not known a priori, the solution is to consider a chip layout program that can accept a fixed scan cells order (produced by our procedure) and in which the closest neighbor criterion is relaxed (having in mind in most situations, the main criterion for scan flip-flops placement relies on timing). Considering routing in our problem will be performed in a future work.

Acknowledgment

The authors would like to thank Professor B. Rouzeyre from the LIRMM for his help on the way to implement the Branch and Bound technique used in our comparisons.

References

[1] A. Crouch, "Design-for-Test for Digital IC's and Embedded Core Systems", Prentice Hall ISBN 0-13-084827-1, 1999.
 [2] Y. Zorian, "A Distributed BIST Control Scheme for Complex VLSI Devices", IEEE VLSI Test Symp., pp. 4-9, 1993.

[3] P. Girard, "Survey of Low Power Testing of VLSI Circuits", IEEE Design & Test of Computers, pp. 2-12, May-June 2002.
 [4] M.L. Bushnell and V.D. Agrawal, "Essentials of Electronic Testing", Kluwer Academic Publishers, ISBN 0-7923-7991-8, 2000.
 [5] A. Hertwig and H.J. Wunderlich, "Low Power Serial Built-In Self-Test", IEEE European Test Workshop, pp. 49-53, 1998.
 [6] S. Wang and S.K. Gupta, "ATPG for Heat Dissipation Minimization for Scan Testing", ACM/IEEE Design Auto. Conf., pp. 614-619, 1997.
 [7] F. Corno, P. Prinetto, M. Rebaudengo and M. Sonza Reorda, "A Test Pattern Generation Methodology for Low Power Consumption", IEEE VLSI Test Symp., pp. 453-459, 1998.
 [8] L. Whetsel, "Adapting Scan Architectures for Low Power Operation", IEEE Int. Test Conf., pp. 863-872, 2000.
 [9] J. Saxena, K.M. Butler and L. Whetsel, "A Scheme to Reduce Power Consumption During Scan Testing", IEEE Int. Test Conf., pp. 670-677, 2001.
 [10] R. Sankaralingam, R. Oruganti and N. Touba, "Static Compaction Techniques to Control Scan Vector Power Dissipation", IEEE VLSI Test Symp., pp. 35-42, 2000.
 [11] R. Sankaralingam, R. Oruganti and N. Touba, "Reducing Power Dissipation During Test Using Scan Chain Disable", IEEE VLSI Test Symp., pp. 319-324, 2001.
 [12] P. Girard, L. Guiller, C. Landrault, S. Pravossoudovitch and H.J. Wunderlich, "A modified Clock Scheme for a Low Power BIST Test Pattern Generator", IEEE VLSI Test Symp., pp. 306-311, 2001.
 [13] K-J. Lee, T-C. Huang and J-J. Chen, "Peak-Power Reduction for Multiple-Scan Circuits during Test Application", IEEE Asian Test Symp., pp. 453-458, 2000.
 [14] A. Chandra and K. Chakrabarty, "Combining Low-Power Scan Testing and Test Data Compression for System-on-a-Chip", ACM/IEEE Design Auto. Conf., pp. 166-169, 2001.
 [15] R.M. Chou, K.K. Saluja and V.D. Agrawal, "Power Constraint Scheduling of Tests", IEEE Int. Conf. on VLSI Design, pp. 271-274, 1994.
 [16] V. Iyengar and K. Chakrabarty, "Precedence-Based, Preemptive, and Power-constrained Test Scheduling for System-on-a-Chip", IEEE VLSI Test Symp., pp. 368-374, 2001.
 [17] V. Dabholkar, S. Chakravarty, I. Pomeranz and S.M. Reddy, "Techniques for Reducing Power Dissipation During Test Application in Full Scan Circuits", IEEE Transactions on CAD, Vol. 17, N° 12, pp. 1325-1333, December 1998.
 [18] P. Girard, C. Landrault, S. Pravossoudovitch and D. Severac, "Reducing Power Consumption during Test Application by Test Vector Ordering", IEEE Int. Symp. on Circuits and Systems, CD-Rom proceedings, 1998.
 [19] P. Girard, L. Guiller, C. Landrault and S. Pravossoudovitch, "A Test Vector Ordering Technique for Switching Activity Reduction during Test Operation", IEEE Great Lakes Symp. on VLSI, pp. 24-27, 1999.
 [20] D.S. Johnson and L.A. McGeoch, "The Traveling Salesman Problem: A Case Study in Local Optimization", in Local Search algorithms in Combinatorial Optimization, E.H.L. Aarts and J.K. Lenstra, eds. John Wiley and Sons, 1996.
 [21] M. Gondran and M. Minoux, "Graphes et Algorithmes", Editions Eyrolles, ISSN 0399-4198, 1979.
 [22] F. Brglez, D. Bryant and K. Kozminski, "Combinational Profiles of Sequential Benchmark Circuits", IEEE Int. Symp. on Circuits and Systems, pp. 1929-1934, 1989.
 [23] PowerMill, 5.1 User Guide, Epic Technology Group, Synopsys Inc., 1998.
 [24] TestGen, Tg 3.0.2 User Guide, Synopsys Inc., 1999.
 [25] M. Held and R.M. Karp, "The Traveling Salesman Problem and Minimum Spanning Trees", Operations Research, 18, pp. 1138-1162, 1970.