

# Thermal-Safe Test Scheduling for Core-Based System-on-Chip Integrated Circuits

Paul Rosinger, Bashir M. Al-Hashimi, *Senior Member, IEEE*, and Krishnendu Chakrabarty, *Senior Member, IEEE*

**Abstract**—Overheating has been acknowledged as a major problem during the testing of complex system-on-chip integrated circuits. Several power-constrained test-scheduling solutions have been recently proposed to tackle this problem during system integration. However, we show that these approaches cannot guarantee hot-spot-free test schedules because they do not take into account the nonuniform distribution of heat dissipation across the die and the physical adjacency of simultaneously active cores. This paper proposes a new test-scheduling approach that is able to produce short test schedules and guarantee thermal safety at the same time. Two thermal-safe test-scheduling algorithms are proposed. The first algorithm computes an exact (shortest) test schedule that is guaranteed to satisfy a given maximum temperature constraint. The second algorithm is a heuristic intended for complex systems with a large number of embedded cores, for which the exact thermal-safe test-scheduling algorithm may not be feasible. Based on a low-complexity test-session thermal-cost model, this algorithm produces near-optimal length test schedules with significantly less computational effort compared to the optimal algorithm.

**Index Terms**—Design for testability, manufacturing testing, reliability.

## I. INTRODUCTION

RECENT REPORTS from industry indicate that power consumption during scan testing in some designs can be significantly higher compared to the normal operation mode. A case study involving the Motorola Version 3 ColdFire processor [14] reported a minimum of 3X increase of test power over functional power. In the same experiment, there have also been cases where the test power for at-speed compressed patterns was as high as 8X the functional power. A more recent industrial paper [19] reports test power up to 30X higher compared to the normal operation mode. The elevated levels of power dissipation during test lead inherently to higher die temperatures compared to the normal operation. This creates a number of problems because both soft error rates and aging increase exponentially with temperature. An undesirable consequence of overheating is thermal stress. At high tem-

peratures, transistors fail to switch properly and many failure mechanisms, such as electromigration, are accelerated resulting in an overall decrease in reliability or even permanent damage. These problems are exacerbated for core-based system-on-chip (SOC) designs because, quite often, several embedded cores are tested concurrently in order to reduce the overall test time. A significant amount of research has been devoted to reducing the power consumption during test in order to avoid the overheating of the silicon die during test. Consequently, several low-power solutions targeting core-level design-for-test (DFT), as well as system-level DFT, have been recently proposed. Techniques falling in the first category include low-power scan chain architectures with gated clocks [17], [18], scan cell and test pattern reordering [3], [5], and low-transition test patterns generated by specialized automatic test-pattern generation (ATPG) algorithms [22] and low-transition TPGs [21]. The second category of techniques is mainly based on power-constrained test-scheduling algorithms [1], [2], [6], [7], [9], [11]–[13], [15].<sup>1</sup>

This paper focuses on avoiding overheating during test through appropriate test scheduling. The main contributions of the paper are:

- 1) propose a thermal-aware test as a better alternative to power-constrained test when dealing with overheating during test;
- 2) propose a test-session thermal model that will reduce the thermal simulation effort required to identify a thermal-safe test schedule.

The motivation for this paper is presented in Section II. The basic ideas behind the existing power constrained test-scheduling approaches are examined from the perspective of chip overheating, and it is explained why these approaches cannot guarantee thermal safety. Section III proposes a new test-scheduling approach that overcomes this problem. An exact algorithm that guarantees minimum test times as well as thermal safety is presented in Section III-A. It is shown through experimental results that significantly shorter test schedules can be obtained without increasing the maximum die temperature during test when compared with existing power-constrained test-scheduling approaches. While this algorithm guarantees the optimal solution for a given thermal limit, it may require significant computational effort for complex systems, which is mainly due to the required amount of accurate thermal

Manuscript received May 30, 2005; revised September 9, 2005. This work was supported by the Engineering and Physical Sciences Research Council (EPSRC) under Grant GR/S05557. The work of K. Chakrabarty was supported by the U.S. National Science Foundation under Grant CCR-0204077. This paper was presented in part at the Design, Automation, and Test in Europe (DATE) Conference in March 2005. This paper was recommended by Associate Editor S. M. Reddy.

P. Rosinger and B. M. Al-Hashimi are with the School of Electronics and Computer Science, University of Southampton, Southampton SO17 1BJ, U.K. (e-mail: pmr@ecs.soton.ac.uk; bmah@ecs.soton.ac.uk).

K. Chakrabarty is with the Department of Electrical and Computer Engineering, Duke University, Durham, NC 27708 USA (e-mail: krish@ee.duke.edu).

Digital Object Identifier 10.1109/TCAD.2006.873898

<sup>1</sup>It should be noted that in [13], the term “hot-spot” is used to denote an area of the die where the peak power consumption exceeded a given threshold. In this paper, we use the term “hot-spot” to denote areas where the die temperature exceeds the maximum allowable junction temperature.




Fig. 1. Example chip. (a) Floorplan. (b) Test compatibility graph. (c) Tools.

simulations. Therefore, a fast heuristic algorithm for thermal-aware test scheduling is proposed in Section III-C. This approach uses a low-complexity test-session thermal-cost model in order to speedup the solution space exploration and reduce the thermal-simulation effort required to reach an acceptable solution. The experiments show that this heuristic produces nearly optimal test schedules (and even optimal schedules in some cases) while significantly reducing the thermal simulation effort.

## II. MOTIVATION

In this section, we examine the effectiveness of the power-constrained test scheduling (PCTS) as a means of avoiding die overheating during test. The common idea behind PCTS is to impose a chip-wide maximum allowable limit on the power consumption, which should not be exceeded during the test application. Several recently proposed power-constrained test-scheduling algorithms aim to maximize the number of tests running in parallel without exceeding this limit [1], [2], [6], [7], [9], [11]–[13], [15].

Silicon die hot spots result from localized overheating, which occurs much faster than chip-wide overheating due to the nonuniform spatial on-die power distribution. Recent research supported by industrial observations suggests that spatial temperature gradients exceeding 30 °C are possible even under typical operating conditions [20], which suggests that there are large variations in power density across the die. These gradients, especially between active and inactive blocks, are likely to increase during testing since test power dissipation can be significantly higher compared to functional power [14], [19]. Having large variations in power densities across the die means that constraining the maximum chip-level power consumption is not an effective way for avoiding local overheating. We demonstrate this using the hypothetical system shown in Fig. 2, which serves as an example of nonuniform power distribution. As shown in the test-description table from Fig. 1(c), cores with different sizes, such as  $C1$  and  $C3$  for example, are assumed to consume the same amount of power during test. Let us consider two possible test sessions  $TS1 = \{C2, C4, C5\}$  and  $TS2 = \{C3, C6, C7\}$ . They are both valid

**INPUT:**  $S$ , the core set for the target system  
 TCG: the test compatibility graph  
 $T_{max}$  = maximum tolerable temperature

**OUTPUT:** Thermal-safe schedule as a list of thermal-safe test sessions

```

1  foreach  $C_i \in S$ 
2    thermal_simulation( $C_i$ )
3    if  $\text{MaxTemp}(C_i) \geq T_{max}$ 
4      fix_thermal_violation( $C_i$ )
5    endif
6  endfor

7  TCC = allCliques(TCG)
8  TCS = allSubsets(TCC)
9  violation = TRUE
10 while violation == TRUE do
11   violation = FALSE
12   LP = formulate_min_weight_cover_lp(S, TCS)
13   LPSol = lp_solve(LP)
14   foreach  $TS_i \in LPSol$ 
15     thermal_simulation( $TS_i$ )
16     if  $\text{MaxTemp}(TS_i) \geq T_{max}$ 
17       remove  $TS_i$  from TCS
18       violation = TRUE
19     endif
20   endfor
21 endwhile

22 LPSol holds the optimum thermal-safe test schedule
    
```

Fig. 2. Exact thermal-safe test-scheduling algorithm.

in terms of test compatibility, as can be seen from the associated test-compatibility graph shown in Fig. 1(b). According to the existing power constrained test-scheduling approaches, both test sessions would be acceptable under a power constraint of 15 W or more. We have run thermal simulations using the HotSpot tool presented in [20], on each of these test sessions, and found a large discrepancy in terms of maximum die temperature: 127.19 °C for  $TS1$  while only 66.47 °C for  $TS2$ . This difference is mainly because the power density (power consumed per area unit) varies significantly from the cores such as  $C2$ ,  $C4$ , and  $C5$  to cores such as  $C3$ ,  $C6$ , and  $C7$  (for example, the power density of core  $C2$  is four times higher than that of  $C3$ ). Moreover, the hottest cores in the two test sessions were  $C5$  and

*C7*. This is because of their reduced lateral heat-removal paths: cores *C5* and *C7* have only two “cold” (inactive) neighbors (more specifically, on their “EAST” and “SOUTH” edges), while all other cores have three “cold” neighbors.

This example has shown that imposing a global power constraint during testing cannot guarantee thermal safety because it does not consider power densities across the die nor the clustering of “hot” cores, which can limit the lateral heat removal. In the following section, we present a new test-scheduling approach that overcomes these issues.

### III. THERMAL-SAFE TEST SCHEDULING

The mean time to failure (MTTF)—a commonly used metric in reliability models—is based on the Arrhenius equation, which shows that reliability is decreasing exponentially with the absolute junction temperature:  $MTTF = A \exp[E_a/kT]$ , where  $A$  is an empirical constant,  $E_a$  is the so-called activation energy, and  $k$  is Boltzmann’s constant [20]. The semiconductor industry is currently using commonly accepted limits for the maximum tolerable operating junction temperature based on the device package type. These have been well accepted as numbers relating to reasonable device lifetimes and thus failure rates. For example, for devices fabricated in a molded package, the maximum allowable junction temperature is 150 °C, while for devices assembled in ceramic or cavity dual in-line packages (DIP), the maximum allowable junction temperature is 175 °C [10]. Based on these practices, the thermal-safe test-scheduling approach proposed in this paper aims to produce solutions guaranteeing that the maximum allowable junction temperature will not be exceeded during test. Throughout this paper, the term “hot-spot” will be used to refer to cores that exceed the maximum allowable junction temperature during test. Any tests running below this critical temperature are considered to be “thermally safe.”

In the following, we propose two thermal-safe test-scheduling algorithms. The first one, although computationally expensive, computes the exact solution to the problem, i.e., the shortest test schedule that meets the thermal constraint. The second proposed algorithm takes into consideration the on-chip lateral heat-transfer paths in order to determine a nearly optimal solution with less computational effort. The results obtained using this algorithm are then compared with the solutions obtained using the exact algorithm.

Both proposed test-scheduling algorithms start from the set of cores ( $S$ ) of the target system, the corresponding test compatibility graph (TCG), such as the one shown in Fig. 1(b), and the maximum junction temperature that can be tolerated during test ( $T_{max}$ ). Each core is annotated with the length of its corresponding test. The TCG captures the concurrence compatibility relationships between the system cores: Each node in the TCG corresponds to a core, and an edge between two nodes means that the two corresponding cores can be tested concurrently without causing any resource conflicts. The floorplan of each system is also needed for performing thermal simulations on the generated test sessions. The algorithms return a thermal-safe test schedule as a list of test sessions, where each test session is a group of cores to be tested concurrently.

#### A. Exact Algorithm

In this section, we present an algorithm for determining an exact solution to the thermal-safe test-scheduling problem (see Fig. 2). The algorithm computes the shortest test schedule, which guarantees that the specified  $T_{max}$  will not be exceeded during test. An outline of this algorithm is presented in the following. First, a thermal simulation is performed on each individual core and corresponding test in order to ensure they all comply with the thermal constraint. The shortest test schedule is computed using only the test compatibility relations between the cores. A thermal simulation is then performed to check whether the test schedule complies with the thermal constraint ( $T_{max}$ ). If  $T_{max}$  is violated during any test sessions, these test sessions are discarded and the process is repeated until a thermal-safe test schedule is found.

We will explain the steps of the algorithm using the hypothetical system shown in Fig. 2, assuming a thermal constraint  $T_{max} = 110$  °C. In the first stage (lines 1–6), it is ensured through thermal simulations that each core, when tested individually with all other cores inactive, does not exceed  $T_{max}$ . None of the cores in our example system was found to violate the thermal constraint when tested individually. In case a thermal violation is detected, the designer needs to fix it by appropriate modifications to the core DFT infrastructure and/or test set. If the violation cannot be fixed, it means that  $T_{max}$  is too restrictive, and other means for reducing the core temperature are needed. Possible solutions include redesigning the cooling structures for the chip and reducing the test clock frequency. Once all cores have passed this initial check, the algorithm computes the clique set [4] test compatibility cliques (TCC) for TCG (line 7). For our example, the clique set for the TCG shown in Fig. 1(b) is  $TCC = [[C2, C5, C4], [C5, C6], [C2, C1], [C3, C6, C7]]$ . Since the number of nodes in TCG (number of cores in a design) is reasonably low, we have used a straightforward exhaustive search algorithm for determining TCC (the `all_cliques` function in Fig. 2). Each clique in the TCC represents a maximal group of cores that can be tested concurrently without causing resource sharing conflicts. Consequently, any valid test schedule must consist only of subsets (TCS) of the cliques in TCC (line 8 in Fig. 2). The shortest test schedule can be determined as the minimum weight set cover for TCS, where the weight of each test compatible subset in TCS is the length of its longest test (it is assumed that all tests in a test session start at the same time). For our example, TCS and the corresponding subset weights are  $[C2] = 0.5$ ,  $[C2, C5, C4] = 0.7$ ,  $[C5] = 0.7$ ,  $[C5, C6] = 0.7$ ,  $[C2, C5] = 0.7$ ,  $[C3] = 0.6$ ,  $[C7, C3] = 0.6$ ,  $[C4] = 0.6$ ,  $[C2, C4] = 0.6$ ,  $[C5, C4] = 0.7$ ,  $[C3, C6] = 0.6$ ,  $[C7, C6] = 0.1$ ,  $[C7] = 0.1$ ,  $[C2, C1] = 0.5$ ,  $[C1] = 0.4$ ,  $[C7, C3, C6] = 0.6$ , and  $[C6] = 0.1$ . The minimum weight set cover is determined using an 0-1 integer linear programming (ILP) formulation (the `formulate_min_weight_cover_lp` function in Fig. 2). The ILP formulation for our example is

$$\begin{aligned} \text{Constraint 1: } & X_{TCS_0} + X_{TCS_1} + X_{TCS_4} + X_{TCS_8} + \\ & X_{TCS_{13}} = 1. \\ \text{Constraint 2: } & X_{TCS_1} + X_{TCS_2} + X_{TCS_3} + X_{TCS_4} + \\ & X_{TCS_9} = 1. \end{aligned}$$

TABLE I  
PCTS VERSUS OPTIMAL THERMAL-AWARE TEST SCHEDULING

Design name	Power-constrained test scheduling		Thermal-aware test scheduling (optimal)			
	Test time(s)	Max temp.(°C)	Test time(s)	Sav.(%)	Iterations	Simulation length(s)
asic_z	0.32	70.81	0.28	12.69	1	0.28
kime	3.81	56.51	3.48	8.42	1	3.48
muresan_10	2.4	58.85	2.0	16.66	1	2.0
muresan_20	5.69	181.79	4.2	26.18	6	24.9
system_l	3.05	191.74	2.53	17.04	2	5.06
system_s	12.12	104.48	9.22	23.88	6	54.17

Constraint 3:  $X_{TCS_6} + X_{TCS_{11}} + X_{TCS_{12}} + X_{TCS_{15}} = 1$ .  
 Constraint 4:  $X_{TCS_5} + X_{TCS_6} + X_{TCS_{10}} + X_{TCS_{15}} = 1$ .  
 Constraint 5:  $X_{TCS_{13}} + X_{TCS_{14}} = 1$ .  
 Constraint 6:  $X_{TCS_1} + X_{TCS_7} + X_{TCS_8} + X_{TCS_9} = 1$ .  
 Constraint 7:  $X_{TCS_3} + X_{TCS_{10}} + X_{TCS_{11}} + X_{TCS_{15}} + X_{TCS_{16}} = 1$ .

Minimize:  $0.5X_{TCS_0} + 0.7X_{TCS_1} + 0.7X_{TCS_2} + 0.7X_{TCS_3} + 0.7X_{TCS_4} + 0.6X_{TCS_5} + 0.6X_{TCS_6} + 0.6X_{TCS_7} + 0.6X_{TCS_8} + 0.7X_{TCS_9} + 0.6X_{TCS_{10}} + 0.1X_{TCS_{11}} + 0.1X_{TCS_{12}} + 0.5X_{TCS_{13}} + 0.4X_{TCS_{14}} + 0.6X_{TCS_{15}} + 0.1X_{TCS_{16}}$ .

A value of one for the binary variable  $X_{TCS_i}$  means TCS<sub>*i*</sub> represents a test session in the final test schedule and a value of zero if it does not. A constraint is added to the ILP formulation for each core  $C_i$ , hence seven constraints for our example enforcing that each core should be covered by one and only one TCS in the final solution. The ILP objective is to minimize the sum of the weights (in this case the test lengths) of the test compatible subsets in the minimum weight set cover. The solution for this particular ILP is  $X_{TCS_1} = 1$ ,  $X_{TCS_{14}} = 1$ , and  $X_{TCS_{15}} = 1$ . The corresponding test schedule will be  $[[C2, C5, C4], [C1], \text{ and } [C7, C3, C6]]$ . Once a valid test schedule has been computed, a thermal simulation is performed to check whether  $T_{max}$  is not exceeded during test. The following information is necessary for performing a thermal simulation for a test session: the chip floorplan, the test power values for the cores in the test session, and the corresponding test lengths. The thermal simulator produces temperature traces for each core for the entire duration of the test session based on a user specified time step. The maximum temperature for each core can then be easily computed from these traces. Thermal simulations for the three test sessions in the previously computed test schedule produce the following results:

Test session	[C2,C4,C5]	[C1]	[C3,C6,C7]
Max temperature (°C)	127.19	98.83	66.47

These results show that the maximum die temperature during the first test session ( $[C2, C5, C4]$ ) violates the thermal constraint of 110 °C (line 16). Consequently, this test session is removed from TCS, and a new test schedule is computed based on the updated TCS. The algorithm continues until a test schedule that does not violate the thermal constraint is found. For our example, the shortest thermal-safe test schedule is

TABLE II  
CHIP PHYSICAL CONSTANTS

Chip thickness	3mm
Heatsink side	50mm
Heatsink thickness	6.9mm
Spreader side	30mm
Spreader thickness	0.5mm
Ambient temperature	45°C

found after three such iterations, and it consists of the following test sessions:

Test session	[C2,C4]	[C5]	[C1]	[C3,C6,C7]
Max temperature (°C)	103.2	106.54	98.83	66.47

In order to reach this result, a total of 5.8 s of test-session time had to be thermally simulated.

B. Experimental Results for Exact Algorithm

Table I compares the results obtained using the proposed algorithm with those obtained using the power constrained test-scheduling approach presented in [7]. We have chosen the approach presented in [7] for comparison since it is very recent, has been applied to large designs, and performs well in comparison with other existing power constrained test-scheduling approaches. Details such as floorplan information and realistic test power and time values had to be added or modified in the original design descriptions in order to provide all necessary information for the proposed thermal safe test-scheduling algorithms. The modified design descriptions used in our experiments can be found at [16]. Some of the physical constants used for thermal simulations performed with the HotSpot tool presented in [20] are reported in Table II. The second column shows the test times corresponding to the power-constrained test schedules. Columns four to seven show the results corresponding to the proposed thermal-aware test-scheduling algorithm. For each design, the temperature limit  $T_{max}$  was set to the maximum temperature of the power constrained test schedule in order to see whether shorter test schedules could be obtained within the same thermal limits. Columns four and five show the test times and relative savings obtained using the thermal-safe test-scheduling algorithm when compared to the algorithm presented in [7]. The experimental data show that the proposed algorithm was able to produce up to 26% shorter test schedules without increasing the maximum die temperature during test. It should be noted that the proposed thermal-aware test-scheduling approach does not pose any constraints on the overall power dissipation; hence, it is possible that the resulting test schedules may exhibit higher overall power compared to the

TABLE III  
DUALITY BETWEEN THERMAL AND ELECTRICAL DOMAINS

Thermal domain	Electrical domain
P, heat flow, power (W)	I, current flow (A)
T, temperature difference (K)	V, voltage (V)
$R_{th}$ , thermal resistance (K/W)	R, electrical resistance ( $\Omega$ )

power-constrained test schedules. However, this falls outside the goal of the proposed test-scheduling approach, which is only to keep the die temperature within the safe limits. The last two columns show the number of iterations and the cumulated length of the thermal simulations required in each case to find a thermal-safe test schedule. One iteration consists of the process of computing a test schedule and checking if it meets the thermal constraint. For example, for system\_s, six test schedules need to be computed until a suitable (thermal-safe) solution was found. It should be noted that the most time-consuming part of the algorithm is represented by the thermal simulations, which could take up to a couple of minutes per design depending on the chosen time step. Computing the clique set and solving the ILP for the minimum weight set cover were taking under 1 s of CPU time on a Pentium IV at 1.8-GHz system.

### C. Heuristic Algorithm

Although the algorithm presented in the previous section computes the optimal solution to the thermal-safe test-scheduling problem, it requires significant computational effort, especially because it requires a large amount of thermal simulations. This is mainly because no knowledge of the heat-transfer paths is used while computing the test schedule, and the thermal compliance check is performed only in a postscheduling phase. This implies that for tight thermal constraints (such as in the case of system\_s shown in Table I), several iterations, and, thus, several thermal simulation runs, are required until a valid solution is found. The thermal simulation effort required to identify thermal-safe test schedules can be reduced by exploiting the knowledge of the on-chip heat-transfer paths. There are two predominant paths for heat transfer out of the integrated-circuit package. The first one is from the die to the surrounding package material, then to the package lead frame and on to the printed circuit board, and finally to the ambient air. The second path is from the package to the heat spreader, to the heat sink, and then to the ambient air. Local die temperature is strongly dependent on the proximity with other heat sources because close heat sources means more heat has to flow through the same paths. Therefore, keeping simultaneous heat sources as far apart as possible reduces the probability of hot spots.

In order to capture the thermal interactions between different cores that are tested concurrently, we have derived a thermo-resistive model for the test sessions. The basic idea is to derive some quantitative measure of the lateral heat-removal paths for a core by taking into account the thermal interactions with active neighboring cores.

The duality between the electrical and thermal domains, illustrated in Table III, offers a convenient basis for an architecture-level thermal model. According to this duality relationship, heat flow can be described as a “current” passing




Fig. 3. Lateral thermal resistance between neighboring cores.

through a thermal resistance leading to a temperature difference analogous to a “voltage.” Thermal resistance  $R_{th}$  is directly proportional to the thickness of the material ( $t$ ) and inversely proportional to the cross-sectional area across which the heat is being transferred ( $A$ )

$$R_{th} = \frac{t}{kA} \quad (1)$$

where  $k$  is the thermal conductivity of the material per volume unit (100 W/mK for silicon and 400 W/mK for copper at 85 °C).

In order to clarify how the lateral thermal resistances are computed, consider the two adjacent cores CORE1 and CORE2 shown in Fig. 3. The chip thickness is  $t$  and the core dimensions are  $(L1, W1)$  and  $(L2, W2)$ , respectively. The lateral resistance  $R_{th21}$  is the thermal resistance from the center of Block 2 to the shared edge of cores one and two. In this case, the heat is constricted from CORE1 to CORE2 via the surface areas defined by  $L1*t$  and  $L2*t$ . The constriction thermal resistance can be calculated by assuming the heat source area to be  $L1*t$ , the silicon bulk area that accepts the heat to be  $L2*t$ , and the thickness of the bulk to be  $W2/2$ . With these values found, the spreading/constriction resistance can be computed using the formulas given in [8]. The resistance is of the spreading type if the lateral area of the source is smaller than the bulk lateral area, and it is of the constriction type otherwise. When computing the lateral thermal resistances, each core is assumed to present a thermal resistance toward each neighboring core.

The lateral thermo-resistive representation for the example floorplan system shown in Fig. 1(a), according to the thermo-resistive thermal model presented in [20], is shown in Fig. 4. In the following, we propose a test-session thermal-cost model aiming to capture the thermal effects due to the physical proximity of simultaneously active cores, since these effects can be controlled by the choice of cores that are to be active at the same time. The proposed test-session thermal-cost model is




Fig. 4. Lateral thermo-resistive model.

derived from this lateral thermo-resistive model presented using the following simplifying assumptions.

- 1) Only steady-state temperatures are considered, as they represent upper bounds for the transient thermal profiles of individual cores. Therefore, only the thermal resistances of the generic RC model are used.
- 2) The heat transfer between the two cores tested concurrently is considered to be negligible; hence, the thermal resistance between those cores is ignored for the current test session. This is a valid assumption because the amount of exchanged heat depends on the temperature difference, which is low for cores tested at the same time.
- 3) Inactive cores are assumed to be thermally grounded, i.e., their temperature is assumed to be equal to the ambient temperature and fixed for the entire duration of the test session.

Let us consider a test session consisting of cores  $C2$ ,  $C4$ , and  $C5$  from our example shown in Fig. 2. The lateral thermo-resistive model derived for this test session according to the previous assumptions is shown in Fig. 5(a). The white arrows pointed to the center of the active cores signify the power dissipated by each of the cores, which is pumped away from the core through the lateral thermal resistances. As it can be observed, the thermal resistances between the pairs of nodes corresponding to active cores (such as  $[C2, C5]$  and  $[C4, C5]$ ) are omitted (assumption 2), while all remaining thermal resistances connect the active core nodes to the ambient, i.e., thermal ground (assumption 3). According to this model, the heat-transfer paths from an active core to its cooler surroundings appear as a number of thermal resistances in parallel. For example, core  $C4$  in Fig. 5(a) has three lateral heat-removal paths toward cores  $C1$ ,  $C6$ , and the left chip edge (west edge).  $C5$  does not represent a heat-removal path for  $C4$  since it is itself an active and, thus, “hot” core. A small equivalent lateral spreading resistance associated with an active core represents a good heat exchange between the core and the ambient; consequently, it predicts a lower core temperature during test. On the other hand, a large

lateral thermal spreading resistance means poor heat exchange with the ambient; therefore, it signals a potential hot spot during test for cores with high power consumption. This can be seen by comparing the lateral thermo-resistive models shown in Fig. 3. Each active core in Fig. 5(a) has only three lateral heat-removal paths represented by three thermal resistors. Cores  $C2$  and  $C4$  shown in Fig. 5(b) have both gained an additional lateral heat-removal path through the removal of core  $C5$  from the test session. The equivalent lateral thermal resistances of cores  $C2$  and  $C4$  are lower in this case compared to the scenario for test session  $[C2, C4, C5]$  shown in Fig. 5(a). The thermal simulations performed on the two test sessions shown in Fig. 3 yielded a  $103.20\text{ }^\circ\text{C}$  maximum temperature for  $[C2, C4]$  and a  $127.19\text{ }^\circ\text{C}$  maximum temperature for  $[C2, C4, C5]$ , which supports our earlier observations.

The thermal-cost model we are proposing for a core is basically the value of the equivalent thermal resistance toward cooler surroundings weighted by the power dissipated by that core, as shown in (2). This is necessary in order to account for the actual power density of the core as well as for the lateral heat-removal paths

$$\text{ThCost}(C_i, \text{TS}) = R_{\text{th}}(C_i, \text{TS}) \times P(C_i), \quad C_i \in \text{TS}. \quad (2)$$

In order to assess the impact of the lateral heat exchange on the core temperature, we have performed the following experiment. We randomly generated a number of test sessions. For each core in each test session, we computed its thermal cost according to (2), and the worst case (i.e., maximum) values were correlated with the maximum temperature reached during the execution of each test session. The high-correlation coefficients obtained for several designs, shown in Table IV, suggest that lateral heat spreading has a significant influence on the maximum core temperature. The maximum core temperature during test was determined through thermal simulations using the HotSpot tool [20].

Based on the results of the previous experiment, we are extending our thermal-cost model to test the sessions as follows:

$$\text{ThCost}(\text{TS}) = \max \text{ThCost}(C_i, \text{TS}) \quad C_i \in \text{TS}. \quad (3)$$

In the following, we are presenting a fast heuristic algorithm for computing the thermal-safe test schedules, which uses the proposed test-session thermal-cost model in order to reduce the required amount of thermal simulations (see Fig. 6). As in the exact algorithm presented in Section III-A, the heuristic algorithm starts by checking whether the individual cores comply with the maximum-allowable-temperature limit  $T_{\text{max}}$  (lines 1–6). Once all cores have passed this test, they are marked as available (line 7) and are arranged in the descending order of their test lengths (line 8). While there still are unscheduled cores, the algorithm tries to assign them to the current test session (TS). The TS is initially empty (line 11), and cores are added to it until no core can be added due to resource conflicts (lines 12–22). A thermal simulation is performed on TS to verify if it complies with  $T_{\text{max}}$ . In our experiments, we have used the Hotspot tool presented in [20]; however, any other thermal simulator could be used for this purpose. Consequently, the accuracy of the results is dependent on




Fig. 5. Lateral thermo-resistive models for two test sessions. (a)  $(C2, C4, C5)$ . (b)  $(C2, C4)$ .

TABLE IV  
CORRELATION BETWEEN TEST-SESSION THERMAL COST  
AND MAXIMUM CORE TEMPERATURE

Design name	Correlation coefficient
asic_z	0.98
kime	0.82
muresan_10	0.94
muresan_20	0.74
system_l	0.98
system_s	0.98
<b>Average</b>	<b>0.91</b>

the chosen thermal simulator. According to the data reported in [20], the simulation-accuracy error of Hotspot is at most 5.8% with respect to FloTherm, the commercial thermal simulator from FloWorks (<http://www.floworks.com>). If the maximum temperature for TS complies with  $T_{\max}$ , the test session is added to the test schedule and the process is repeated for the unscheduled cores. If the maximum temperature during TS exceeds  $T_{\max}$ , the flag `GotCostLimit` is set to true, and a thermal-cost limit is computed based on the thermal cost of TS and the fraction by which  $T_{\max}$  was exceeded (lines 34–37). The thermal cost of TS is computed according to (3). The test-session cost limit is computed as the thermal cost of TS, scaled down linearly by the fraction by which  $T_{\max}$  was exceeded. The thermal-cost adjustment factor is computed as

$$\text{ThCostAdjust}(\text{MaxTemp}(\text{TS}), T_{\max}) = \frac{(\text{MaxTemp}(\text{TS}) - T_{\max})}{T_{\max}} \times K + 1 \quad (4)$$

where  $K \in (0, 1]$  is a user specified constant used to relax the thermal-cost limit (`ThCostLimit`). In our experiments, we have used  $K = 0.5$ . Once a thermal-cost limit had been computed, a core is added to the current test session only if it does not increase the test-session thermal cost over `ThCostLimit` (lines 23–38). This way, it is ensured that once a thermal violation has been detected, test sessions with similar or worse

lateral-heat-exchange capabilities are avoided without requiring lengthy thermal simulations.

We are illustrating the steps of this algorithm using the example system shown in Fig. 2. The same thermal constraint  $T_{\max} = 100^\circ\text{C}$  used for the exact algorithm in Section III-A will be used here as well. After the initial core check, the available array is initialized with all cores in the system arranged in the descending order of their test lengths (Line 7)

$$\text{Available} = [C5, C3, C4, C2, C1, C6, C7]. \quad (5)$$

`GotCostLimit` is set to false (Line 8) and an empty test-session TS is created (Line 11). The first core added to the TS is  $C5$ . The next available core,  $C3$ , cannot be added to TS because it is not test compatible with  $C5$  [see Fig. 1(b)] (line 15). This process continues until no more cores can be added to TS. At this moment,  $\text{TS} = [C5, C4, C2]$ . A thermal simulation is performed on TS (Line 31) to determine the maximum temperature reached during TS, in this case  $127.19^\circ\text{C}$ . This violates the thermal constraint of  $110^\circ\text{C}$ ; therefore, the algorithm proceeds by setting `GotCostLimit` to true and computing the thermal-cost limit based on the maximum temperature reached during TS and the thermal-cost value of TS:  $\text{ThCostLimit} = 35.43$  (line 36). The algorithm continues by discarding TS and building a new test session, this time checking also that the thermal cost of the test session does not exceed the previously computed thermal-cost limit. The first core added to TS is  $C5$ . At this point, the thermal cost for TS is 26.88. This is below the imposed limit; therefore, the algorithm continues to add test compatible cores to TS (line 26). The next core to be added is  $C4$ , which arises the thermal cost of TS to 32.01. This is still below the imposed limit, so a new core,  $C2$ , is added to TS. The thermal cost of TS becomes 39.57, which is over the imposed limit. Consequently,  $C2$  is removed from TS. Since no more cores can be added to TS, a thermal simulation is performed on TS (Line 31). The maximum temperature is found to be  $117.04^\circ\text{C}$ , which violates the thermal constraint. The

```

INPUT:  $S$ , the core set for the target system
         TCG: the test compatibility graph
          $T_{max}$  = maximum tolerable temperature

OUTPUT: Thermal-safe schedule as a list of
           thermal-safe test sessions

1  foreach  $C_i \in S$ 
2    thermal_simulation( $C_i$ )
3    if MaxTemp( $C_i$ )  $\geq T_{max}$ 
4      fix_thermal_violation( $C_i$ )
5    endif
6  endfor

7  Available = {  $C_i | C_i \in S$  }
8  sort_descending( Available )
9  Hsol =  $\emptyset$ 
10 GotCostLimit = FALSE
10 while Available  $\neq \emptyset$  do
11   TS =  $\emptyset$ 
12   foreach CandidateCore  $\in$  Available do
13     IsCompatible = TRUE
14     foreach  $C \in TS$  do
15       if  $C$  not compatible with CandidateCore
16         IsCompatible = FALSE
17         gotoline 20
18       endif
19     endfor
20     if IsCompatible = TRUE
21       if GotCostLimit = FALSE
22         addCandidateCore to TS
23       else
24         ThCost = computeThCost( TS  $\cup$  {CandidateCore} )
25         if ThCost  $\leq$  ThCostLimit
26           add CandidateCore to TS
27         endif
28       endif
29     endif
30   endfor

31   thermal_simulation( TS )
32   if MaxTemp(TS)  $\leq T_{max}$ 
33     add TS to Hsol
34   else
35     GotCostLimit = TRUE
36     ThCostLimit = computeThCost( TS )  $\times$  ThCostAdjust( MaxTemp(TS),  $T_{max}$  )
37   endif
38 endwhile

39 Hsol holds the thermal-safe test schedule

```

Fig. 6. Heuristic thermal-safe test-scheduling algorithm.

thermal-cost limit is readjusted to  $\text{ThCostLimit} = 30.45$  (Line 36), and TS is discarded.  $C_5$  is added to a new empty test session, arising its thermal cost to 26.88.  $C_4$ ,  $C_5$ , and  $C_6$  cannot be added to TS since the thermal cost for TS would exceed the new thermal-cost limit. A thermal simulation is run for  $TS = [C_5]$ , and TS is added to the test schedule and removed from the available cores since its maximum temperature of 106.54 °C meets the thermal constraint. The GotCostLimit is set to false, and the algorithm continues to schedule the remaining cores. In the next iterations, the algorithm adds  $[C_3, C_6, C_7]$ ,  $[C_2, C_4]$ , and  $[C_1]$  to the test schedule. This matches the test schedule produced by the exact algorithm; however, the thermal simulation effort was reduced from 5.8 s of test-session time to 3.7 s. The complexity of computing a test session using this approach is  $O(N^3)$ , where  $N$  is the

number of cores. However, the thermal simulation (line 31 in Fig. 6), which is the most computationally expensive part of the algorithm, is performed in the outermost loop, which has only a complexity of  $O(N)$ . This is a considerable improvement in terms of the required computational effort over the exact algorithm described in Section III-A, which has an exponential complexity due to the NP-hard nature of the optimization problem.

#### D. Experimental Results for Heuristic Algorithm

A number of experiments has been performed in order to assess the performance of the proposed test-scheduling heuristic. The first set of experiments was used to compare the proposed heuristic with the power-constrained test-scheduling approach

TABLE V  
POWER CONSTRAINED TEST SCHEDULING VERSUS HEURISTIC THERMAL-AWARE TEST SCHEDULING

Design name	Power-constrained test scheduling		Thermal-aware test scheduling (heuristic)			
	Test time(s)	Max temp.(°C)	Test time(s)	Sav.(%)	Violations	Simulation length(s)
asic_z	0.32	70.81	0.28	12.69	0	0.28
kime	3.81	56.51	3.48	8.42	0	3.48
muresan_10	2.4	58.85	2.0	16.66	1	2.4
muresan_20	5.69	181.79	4.89	14.05	1	6.0
system_l	3.05	191.74	2.87	5.09	0	2.87
system_s	12.12	104.48	9.22	23.88	1	17.67

TABLE VI  
TEST TIMES FOR DIFFERENT TEMPERATURE CONSTRAINTS

Design name	Max temp.(°C)	Test time(s)	Simulation length(s)
asic_z	71.15	0.28	0.28
	76.15	0.28	0.28
	81.15	0.28	0.28
	86.15	0.28	0.28
	91.15	0.28	0.28
kime	56.84	3.48	3.48
	61.84	3.48	3.48
	66.84	3.48	3.48
	71.84	3.48	3.48
	76.84	3.48	3.48
muresan_10	59.18	2.0	2.4
	64.18	2.0	2.0
	69.18	2.0	2.0
	74.18	2.0	2.0
	79.18	2.0	2.0
muresan_20	182.25	4.89	6.0
	187.25	4.89	6.0
	192.25	4.49	4.49
	197.25	4.49	4.49
	202.25	4.49	4.49
system_l	194.90	2.87	2.87
	199.90	2.87	2.87
	204.90	2.87	2.87
	209.90	2.87	2.87
	214.90	2.87	2.87
system_s	104.85	9.22	17.67
	109.85	9.22	17.67
	114.85	8.44	8.44
	119.85	8.44	8.44
	124.85	8.44	8.44

presented in [7]. The results of these experiments are reported in Table V. The maximum temperature during test corresponding to the power constrained test schedules was used as a thermal constraint ( $T_{max}$ ) for the proposed test-scheduling algorithm. This way, it is guaranteed that the resulting test schedules will not lead to higher temperatures than that using the power-constrained approach. From the fourth and fifth columns, it can be observed that the proposed heuristic algorithm outperformed the power-constrained test-scheduling approach for all designs, which produced up to 24% shorter test schedules. Moreover, for three out of the six designs considered, the heuristic algorithm produced the same test schedules as the exact algorithm presented in Section III-A. The last column in Table VI shows significant reductions in terms of thermal simulation effort when compared to the exact algorithm. For example, the simulation length was reduced from 54 to 18 s for system\_s.

Another set of experiments was performed to analyze the effect of different maximum temperature limits on the test time

TABLE VII  
COMPARISON BETWEEN EXACT AND HEURISTIC THERMAL-AWARE TEST-SCHEDULING ALGORITHMS

Design name	Test time increase (%)	Simulation effort reduction (%)
asic_z	0	0
kime	0	0
muresan_10	0	-20
muresan_20	16.4	75.9
system_l	13.4	43.28
system_s	0	67.38

and simulation effort. From Table VI, it can be observed that, as expected, both test time and simulation effort decrease as the thermal constraint becomes more relaxed. For example, the test time is reduced by 25% when the thermal constraint is increased by only 5 °C from 187.25 °C for muresan\_20, and it is reduced from 9.22 to 8.44 s when the thermal constraint is increased from 109.85 to 114.85 °C for system\_s. Even more reductions are obtained in terms of the simulation effort. For example, for system\_s, increasing the thermal constraint from 109.85 to 114.85 °C reduced the simulation effort by half, from nearly 18 s to less than 8.5 s.

Table VII compares the proposed heuristic and the exact thermal-safe test-scheduling algorithms. As mentioned earlier, the heuristic determines the optimum solution for the four out of the six designs considered. In only one case, the required thermal-simulation effort exceeded that required by the exact algorithm, while in all other cases up to 75% reductions have been obtained.

#### IV. CONCLUSION

Overheating has been acknowledged as a major problem during the testing of complex SOC integrated circuits. In this paper, we have outlined the need for thermal-safe testing and explained that existing power-constrained test-scheduling approaches cannot guarantee a thermal safety during test. Next, we have proposed a new test-scheduling approach that produces short test schedules and guarantees thermal safety during testing at the same time. Two possible algorithms have been developed for the proposed thermal-safe test-scheduling approach. The first proposed algorithm, although computationally expensive, provides an optimal solution to the thermal-safe test-scheduling problem. The second algorithm uses a fast heuristic based on a low-complexity test-session thermal model in order to reduce the required computational effort while producing optimal or near-optimal test schedules. Experimental results show that up to 24% shorter test schedules can be obtained using the proposed approach without increasing the

maximum temperature during test application, when compared to power constrained test-scheduling approaches. The proposed approach provides an effective solution to the problems arising from chip overheating during test.

ACKNOWLEDGMENT

The authors would like to thank E. Larsson, Linköping University, Linköping, Sweden, for providing the code and designs used for the work presented in [7].

REFERENCES

[1] K. Chakrabarty, "Design of system-on-a-chip test access architectures under place-and-route and power constraints," in *Proc. IEEE/ACM DAC*, 2000, pp. 432–437.

[2] R. Chou, K. Saluja, and V. Agrawal, "Scheduling tests for VLSI systems under power constraints," *IEEE Trans. VLSI Syst.*, vol. 5, no. 2, pp. 175–184, Jun. 1997.

[3] P. Flores, J. Costa, H. Neto, J. Monteiro, and J. Marques-Silva, "Assignment and reordering of incompletely specified pattern sequences targeting minimum power dissipation," in *Proc. 12th Int. Conf. VLSI Des.*, 1999, pp. 37–41.

[4] A. Gibbons, *Algorithmic Graph Theory*. Cambridge, U.K.: Cambridge Univ. Press, 1985.

[5] P. Girard, C. Landrault, S. Pravossoudovitch, and D. Severac, "Reducing power consumption during test application by test vector ordering," in *Proc. ISCAS*, 1998, pp. 296–299.

[6] V. Iyengar and K. Chakrabarty, "System-on-a-chip test with precedence relationships, preemption and power constraints," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 21, no. 9, pp. 1088–1094, Sep. 2002.

[7] E. Larsson, K. Arvidsson, H. Fujiwara, and Z. Peng, "Efficient test solutions for core-based designs," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 23, no. 5, pp. 758–775, May 2004.

[8] S. Lee, S. Song, V. Au, and K. Moran, "Constricting/spreading resistance model for electronics packaging," in *Proc. ASME/JSME Therm. Eng. Conf.*, 1995, pp. 199–206.

[9] V. Muresan, X. Wang, V. Muresan, and M. Vladutiu, "A comparison of classical scheduling approaches in power-constrained block-test scheduling," in *Proc. IEEE ITC*, 2000, pp. 882–891.

[10] National Semiconductor *Understanding Integrated Circuit Package Power Capabilities*. (2000, Apr.). [Online]. Available: [http://www.national.com/ms/UNUNDERSTANDING\\_INTERGRATED\\_CIRCUIT\\_PACKAGE\\_POWER\\_CA.pdf](http://www.national.com/ms/UNUNDERSTANDING_INTERGRATED_CIRCUIT_PACKAGE_POWER_CA.pdf)

[11] N. Nicolici and B. Al-Hashimi, "Power conscious test synthesis and scheduling for BIST RTL data paths," in *Proc. IEEE ITC*, Oct. 2000, pp. 662–671.

[12] M. Nourani and J. Chin, "Power-time trade off in test scheduling for SoCs," in *Proc. IEEE ICCD*, Oct. 2003, pp. 548–553.

[13] —, "Test scheduling with power-time tradeoff and hot-spot avoidance using MILP," *Proc. Inst. Elect. Eng.—Comput. Digit. Techn.*, vol. 151, no. 5, pp. 341–355, Sep. 2004.

[14] B. Pouya and A. Crouch, "Optimization trade-offs for vector volume and test power," in *Proc. ITC*, 2000, pp. 873–881.

[15] C. P. Ravikumar, G. Chandra, and A. Verma, "Simultaneous module selection and scheduling for power-constrained testing of core based systems," in *Proc. 13th Int. Conf. VLSI Des.*, 2000, pp. 462–467.

[16] P. Rosinger. (2005). *Sample designs for validating thermal-aware test solutions*. [Online]. Available: [http://www.ecs.soton.ac.uk/pmr/thermal\\_test\\_sample\\_designs.zip](http://www.ecs.soton.ac.uk/pmr/thermal_test_sample_designs.zip)

[17] P. Rosinger, B. Al-Hashimi, and N. Nicolici, "Scan architecture with mutually exclusive scan segment activation for shift and capture power reduction," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 23, no. 7, pp. 1142–1154, Jul. 2004.

[18] J. Saxena, K. M. Butler, and L. Whetsel, "An analysis of power reduction techniques in scan testing," in *Proc. IEEE ITC*, 2001, pp. 670–677.

[19] C. Shi and R. Kapur. (2004, Sep. 15). How power aware test improves reliability and yield, *EETimes*. [Online]. Available: <http://www.eetimes.com/news/design/features/showArticle.jhtml?articleId=47208594&kc=4235>

[20] K. Skadron, M. Stan, W. Huang, S. Velusamy, K. Sankaranarayanan, and D. Tarjan, "Temperature-aware microarchitecture," in *Proc. ISCA*, 2003, pp. 2–13.


[21] S. Wang and S. K. Gupta, "DS-LFSR: A new BIST TPG for low heat dissipation," in *Proc. IEEE Int. Test Conf.*, 1997, pp. 848–857.

[22] —, "ATPG for heat dissipation minimization during test application," *IEEE Trans. Comput.*, vol. 47, no. 2, pp. 256–262, Feb. 1998.



**Paul Rosinger** received the B.Sc. degree in computer science from the Technical University of Timisoara, Timisoara, Romania, in 1999, and the Ph.D. degree in electronics and computer science from the University of Southampton, Southampton, U.K., in 2003.

He is currently a Postdoctoral Research Fellow at the University of Southampton. His current research interests include testing of digital systems, low-power embedded systems, and reconfigurable architectures.



**Bashir M. Al-Hashimi** (M'99–SM'01) received the B.Sc. degree with first-class classification in electrical and electronics engineering from the University of Bath, Bath, U.K., in 1984, and the Ph.D. degree from York University, U.K., in 1989.

After receiving the Ph.D. degree, he worked in the semiconductor industry and designed integrated-circuit signal-processing applications and developed computer-aided design (CAD) tools for simulation and synthesis of analog and digital circuits. In 1999, he joined the School of Electronics and Computer

Science, University of Southampton, Southampton, U.K., where he is currently a Professor of computer engineering. He has authored one book on SPICE simulation (CRC, 1995) and coauthored two books: *Power Constrained Testing of VLSI Circuits* (Kluwer, 2002) and *System-Level Design Techniques for Energy-Efficient Embedded Systems* (Kluwer, 2004). Recently, he edited the book *System-on-Chip: Next Generation Electronics* (IEE Press, 2006). He has authored and coauthored over 150 technical papers. His research and teaching interests include low-power system-level design, system-on-chip test, and very large-scale integration CAD.

Dr. Al-Hashimi is a Fellow of the IEE. He is the Editor-in-Chief of the *IEEE Proceedings: Computers and Digital Techniques*, an editor of the *Journal Electronic Testing: Theory and Applications (JETTA)*, and is a member of the editorial board of the journal *Low Power Electronics*, and the journal of *Embedded Computing*. He is the General Chair of the 11th IEEE European Test Symposium (Southampton 2006) and the General Chair of DATE Friday Workshops (2005 and 2006). He was the coauthor of the James Beausang Best Paper Award at the 2000 IEEE International Test Conference relating to low power BIST for RTL data paths.



**Krishnendu Chakrabarty** (S'91–M'92–SM'00) received the B.Tech. degree from the Indian Institute of Technology (IIT), Kharagpur, India, in 1990 and the M.S.E. and Ph.D. degrees from the University of Michigan, Ann Arbor, in 1992, and 1995, respectively, all in computer science and engineering.

He is currently an Associate Professor of electrical and computer engineering, Duke University, Durham, NC. His current research interests include design and testing of SOC integrated circuits, design automation of microfluidic-based biochips, microfluidic-based chip cooling, distributed sensor networks, and embedded real-time systems. He is a coauthor of two books: *Microelectrofluidic Systems: Modeling and Simulation* (CRC, 2002) and *Test Resource Partitioning for System-on-a-Chip* (Kluwer, 2002), and the Editor of *SOC (System-on-a-Chip) Testing for Plug and Play Test Automation* (Kluwer, 2002). He is also a coauthor of the forthcoming book *Scalable Infrastructure for Distributed Sensor Networks* (Springer). He has published 200 papers in journals and refereed conference proceedings, and he holds a U.S. patent in BIST.

Dr. Chakrabarty is a recipient of the National Science Foundation Early Faculty (CAREER) award and the Office of Naval Research Young Investigator Award. He is a recipient of a Best Paper Award at the 2005 IEEE International Conference on Computer Design and 2001 IEEE Design, Automation and Test in Europe (DATE) Conference. He is also a recipient of the Humboldt Research Fellowship awarded by the Alexander von Humboldt Foundation, Germany. He is an Associate Editor of IEEE TRANSACTIONS ON COMPUTER-AIDED DESIGN OF INTEGRATED CIRCUITS AND SYSTEMS, an Associate Editor of IEEE TRANSACTIONS ON VERY LARGE INTEGRATION SYSTEM (VLSI) SYSTEMS, an Associate Editor of the Association of Computing Machinery (ACM), *Journal on Emerging Technologies in Computing Systems*, an Editor of *Journal of Electronic Testing: Theory and Applications (JETTA)*, and a member of the Editorial Board for *Sensor Letters and Journal of Embedded Computing*. He serves as a Subject Area Editor for the *International Journal of Distributed Sensor Networks*. He has also served as an Associate Editor of IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS II: ANALOG AND DIGITAL SIGNAL PROCESSING. He is a member of ACM and ACM Special Interest Group on Design Automation (SIGDA), and a member of Sigma Xi. He serves as Vice Chair of Technical Activities in IEEE's Test Technology Technical Council and is a member of the program committees of several IEEE/ACM conferences and workshops. He chaired the emerging technologies subcommittee for the IEEE International Conference CAD 2005, and he is designated Subcommittee Chair for new, emerging, and specialized technologies for the 2006 Design Automation Conference. He served as the Tutorial Co-Chair for the 2005 IEEE International Conference on VLSI Design and is the designated Program Co-Chair for the 2005 IEEE Asian Test Symposium.