

Adaptive Bayesian Diagnosis of Intermittent Faults

Rodríguez Gómez, Laura; Cook, Alejandro; Indlekofer, Thomas; Hellebrand, Sybille; Wunderlich, Hans-Joachim

Journal of Electronic Testing: Theory and Applications (JETTA) Vol. 30(5) 30
September 2014

url: <http://link.springer.com/article/10.1007/s10836-014-5477-1>

doi: <http://dx.doi.org/10.1007/s10836-014-5477-1>

Abstract: With increasing transient error rates, distinguishing intermittent and transient faults is especially challenging. In addition to particle strikes relatively high transient error rates are observed in architectures for opportunistic computing and in technologies under high variations. This paper presents a method to classify faults into permanent, intermittent and transient faults based on some intermediate signatures during embedded test or built-in self-test. Permanent faults are easily determined by repeating test sessions. Intermittent and transient faults can be identified by the amount of failing test sessions in many cases. For the remaining faults, a Bayesian classification technique has been developed which is applicable to large digital circuits. The combination of these methods is able to identify intermittent faults with a probability of more than 98%.

Preprint

General Copyright Notice

This article may be used for research, teaching and private study purposes. Any substantial or systematic reproduction, re-distribution, re-selling, loan or sub-licensing, systematic supply or distribution in any form to anyone is expressly forbidden.

This is the author's "personal copy" of the final, accepted version of the paper published by *Springer Science+Business Media New York*.

©2014 Springer Science+Business Media New York

Adaptive Bayesian Diagnosis of Intermittent Faults

Laura Rodríguez Gómez¹, Alejandro Cook¹, Thomas Indlekofer²,
Sybille Hellebrand², Hans-Joachim Wunderlich¹

¹Institute of Computer Architecture and Computer Engineering, University of Stuttgart, Germany

²Institute of Electrical Engineering and Information Technology, University of Paderborn, Germany

Abstract

With increasing transient error rates, distinguishing intermittent and transient faults is especially challenging. In addition to particle strikes relatively high transient error rates are observed in architectures for opportunistic computing and in technologies under high variations. This paper presents a method to classify faults into permanent, intermittent and transient faults based on some intermediate signatures during embedded test or built-in self-test.

Permanent faults are easily determined by repeating test sessions. Intermittent and transient faults can be identified by the amount of failing test sessions in many cases. For the remaining faults, a Bayesian classification technique has been developed which is applicable to large digital circuits. The combination of these methods is able to identify intermittent faults with a probability of more than 98 %.

1 Introduction

In recent years, volume diagnosis has gained increasing attention as a key contributor to fast yield ramp-up. A number of efficient techniques have been developed mainly assuming permanent faults as the root cause of test failures [Rajski99, Ghosh00, Liu02, Wohl02, Cheng06, Holst09, Tang07, Wang09, Elm10, Cook11a]. Continuous technology scaling, however, has led to a more complex scenario, and diagnosis must now support a refined quality assessment as described in the sequel. Firstly, increasing parameter variations in nano-scale CMOS have expedited new strategies for “underdesigned and opportunistic computing” [Ernst04, Gupta13]. To fully exploit the potential of technology scaling, these approaches avoid an overly pessimistic design with large guard bands. Instead, they include mechanisms to detect and compensate a certain amount of transient faults caused by parameter variations or external noise [Nicolaidis99, Ernst04, Nicolaidis07]. During structural test, however, transient faults may indicate failures, even if they could be compensated during system operation. Secondly, marginalities in the design may lead to intermittent faults depending on certain activation conditions, such as for example power droop [Tirumurti04, Polian07]. Intermittent faults may repeatedly occur at the same location or in its neighborhood. Even though they are not permanent, they may severely affect the system functionality, indicate potential early life failures, or reduce robustness against transient faults.

Dealing with transient and intermittent faults is particularly challenging. On the one hand, transient faults are uncritical for appropriately designed “robust” circuits, and test failures due to transient faults cause unnecessary yield loss in this case. On the other hand, intermittent faults impact quality, but they may lead to similar observations during test as transient faults. Nevertheless, depending on the quality requirements, a limited amount of intermittent faults may still be acceptable. To control the trade-off between yield and quality, test and diagnosis procedures must be able to distinguish between transient, intermittent, and permanent faults. As a first step in this direction, the integrated test and diagnosis scheme in [Cook11b] partitions the test into several shorter sessions and immediately repeats failing sessions. This way, permanent faults can be identified quickly, as they will lead again to the same faulty behavior. If a repeated session

shows a fault free behavior or a different faulty behavior, this indicates a non-permanent fault, but a classification into transient or intermittent fault is not obvious, and thus the main focus of this paper.

So far, only little work has been published on the diagnosis of intermittent faults. In the context of online monitoring, the intermittent faults can be identified based on the observed failure rates in the system [Fechner09], but during offline test the observation time is much shorter making a pure statistical analysis impractical. A recent approach for system diagnosis applies Bayesian reasoning to deal with intermittent faults [DeKler09]. As a proof of concept the author uses it, in particular, for localizing intermittent faults in logic circuits. However, this approach assumes full observability of internal nodes and is not directly applicable to realistic circuits.

This paper presents an adaptive scheme for test and diagnosis, which effectively combines the window-based diagnosis in [Cook11b] with Bayesian reasoning. The rest of the paper is organized as follows: Section 2 briefly summarizes the necessary background on fault modeling, window-based and Bayesian diagnosis. Subsequently, in Section 3 the new approach for adaptive Bayesian diagnosis is explained in detail. The experimental results presented in Section 4 show that the combination of these methods can classify intermittent faults with a confidence of more than 98 %.

2 Background

2.1 Fault Modeling

In the following, the typical properties of transient and intermittent faults are summarized, and the fault model used in this work is explained. Transient faults randomly appear at victim nodes v and last at most for one clock cycle. They can be caused by external noise [Baumann05], but even more likely, delay problems occurring at clock boundaries or dynamic parameter variations such as power supply and interconnect noises, electromagnetic interferences and electrostatic discharges can lead to violations of timing safety margins which manifest themselves as transient faults [Constantinescu03, Borkar05].

Intermittent faults can be traced back to unstable or marginal hardware and are activated by specific environmental conditions, like increasing or decreasing temperature or voltage. Their observable behavior is similar to that of transient faults. But as they may also evolve into permanent faults, they can be as critical as permanent faults. One example is high frequency power droop, which results from power starvation when multiple cells connected to the same power grid segment suddenly increase their current demand [Tirumurti04, Polian07]. Thus, a high frequency power droop occurs when several nodes in the neighborhood of a victim node v switch in the same direction as v . Other root causes for intermittent faults include low frequency power droop or cross talk.

Standard fault models like stuck-at, delay or bridging faults are not sufficient to appropriately deal with transient or intermittent faults, because activation and/or timing conditions cannot be taken into account. For this reason, the conditional stuck-at fault model is applied [Holst09].

Definition 1 (conditional stuck-at fault): Let v denote a circuit line and $cond$ a Boolean or timing condition. The conditional stuck-at zero fault $cond_0_v$ forces v to zero, if $cond$ is met, and the conditional stuck-at one fault $cond_1_v$ forces v to 1, if $cond$ is satisfied.

For instance, $(v=1)_0_v$ is a permanent stuck-at-0 fault, and $(v_{-1}=0 \wedge v=1)_0_v$ describes a slow-to-rise fault. To model transient faults which occur at a certain time step only, a particular pattern in a sequence $P = (p_1, \dots, p_n)$ is specified in $cond$. The expression $(p_i | P)_0_v$ means that the line v is set to '0' when the i -th pattern of the sequence P is applied. To characterize intermittent faults, their activation conditions are encoded in $cond$. For example, a high frequency power droop is activated, if many nodes in the neighborhood switch in the same direction as the victim node v . This can be described by $(v_{-1} = 0) \wedge (v = 1) \wedge (|\{w \in \mathcal{N}(v) : (w_{-1} = 0) \wedge (w = 1)\}| \geq \tau)$ or $(v_{-1} = 1) \wedge (v = 0) \wedge (|\{w \in \mathcal{N}(v) : (w_{-1} = 1) \wedge (w = 0)\}| \geq \tau)$, where $\mathcal{N}(v)$ denotes the neighborhood of v and τ a user-defined threshold value. For an exact definition of

neighborhoods layout data are required. As an approximation the circuit topology at gate level is analyzed in this work. The circuit netlist is mapped to an undirected graph $G = (V, E)$, the nodes of which correspond to circuit nodes and the edges of which correspond to direct connections between circuit nodes.

Definition 2 (neighborhood): Let $G = (V, E)$ be an undirected graph, $v \in V$, and $r > 0$ a natural number. Then the set $\mathcal{N}^r(v) := \{w \in V : \text{there is a path from } v \text{ to } w \text{ of length at most } r\}$ is called the neighborhood of radius r of v .

Figure 1 shows an example. The shaded area depicts the neighborhood of radius 2 of v .




Figure 1: Topological neighborhood of a circuit node v .

2.2 Window-based Diagnosis

Volume diagnosis is an extremely challenging task, as it should interfere with the test flow as little as possible. Thus, the diagnostic procedures must be compatible with standard test architectures and response compaction schemes, as for example the widely accepted STUMPS architecture shown in Figure 2 [Bardell82].




Figure 2: STUMPS architecture.

For built-in self-test (BIST), a test pattern generator (TPG) and a multiple-input signature register (MISR) are added to the circuit under test (CUT). In each clock cycle a “slice” of a test pattern is loaded into the scan chains. When a complete test pattern is shifted in, the test response is captured in the flip-flops. While the response is shifted out and compressed by the MISR, the next pattern can be shifted in. For embedded test, the TPG is replaced by an on-chip decompressor receiving encoded test data from the automatic test equipment (ATE), and the compacted test responses are sent back to the ATE. If a transient fault occurs in the CUT during the capture cycle, a faulty test response may be obtained and stored in the flip-flops. The fault effect is preserved in the flip-flops during shifting and may lead to a faulty signature. However, even if the transient fault lasts longer than one clock cycle, only the fault effect during the capture cycle has impact on the signature. Therefore, in the sequel transient faults are assumed to have a duration of a single clock cycle only. In contrast to that, intermittent faults can be activated by several patterns.

Because of the limited bandwidth in embedded test and the limited storage capacities in built-in self-test, the amount of response data to be analyzed must be minimized for either test strategy. It is beyond the scope of this paper to give a complete review of the state of the art in this extensively studied field. Instead, this section focuses on direct diagnosis only, which supports test response compaction by signature analysis and

derives the fault locations directly from the faulty signatures [Cheng06]. In the following it will be briefly summarized how direct diagnosis can be combined with enhanced test response compaction as well as with mechanisms to distinguish between permanent and non-permanent faults.

While first approaches for direct diagnosis with extreme response compaction are not compatible with the STUMPS architecture [Elm10], the window-based diagnosis in [Cook11a] overcomes this problem by partitioning the test into N contiguous subsequences (windows). As illustrated in Figure 3, each window is characterized by a single cumulative signature.




Figure 3: Test sequence with repeated windows and cumulative signatures s_1, \dots, s_N .

Unlike the diagnosis scheme in [Wohl02], which repeats windows with failing signatures in a special diagnostic mode without response compaction, the approach in [Cook11a] can determine the fault location directly from the cumulative signature of a window. Compared to the standard scheme with one signature per pattern this provides an additional reduction of the response data directly proportional to the window size. The diagnostic algorithm is based on the conditional stuck-at fault model.

Definition 3 (Deviation vectors): Let P denote the set of all patterns in a window, and let v be a candidate fault location. For a pattern $p \in P$ the deviation vectors $d((p | P)_0_v)$ and $d((p | P)_1_v)$ represent the deviations from the fault free signature in the presence of the conditional stuck-at faults $(p | P)_0_v$ and $(p | P)_1_v$.

The deviation vectors are precomputed for each pattern p_1, \dots, p_n in the window and each candidate fault location. From this information about single activations of faults, the linear equations

$$c_1 d((p_1 | P)_0_v) \oplus \dots \oplus c_n d((p_n | P)_0_v) = S_{obs}(P) \oplus S_{ref}(P) \quad (1)$$

and

$$c_1 d((p_1 | P)_1_v) \oplus \dots \oplus c_n d((p_n | P)_1_v) = S_{obs}(P) \oplus S_{ref}(P) \quad (2)$$

are built, where c_1 to c_n are variables over $GF(2)$, and $S_{obs}(P)$ and $S_{ref}(P)$ denote the observed and the reference signature, respectively. As the MISR signatures are m -bit vectors, each of the equations (1) and (2) corresponds to a system of linear equations over $GF(2)$ with n variables and m equations. Such a system of equations has a solution, if the fault location v can explain the faulty behavior. Hereby it is assumed that for a given fault location v , it is sufficient to consider only conditional stuck-at faults of the same polarity (no line flips).

To guarantee a unique solution, the number of variables n , and thus the number of patterns in a window must not exceed m . However, the analysis of spurious failure effects requires additional considerations to avoid ambiguous diagnostic results. Due to aliasing the solution of the equations may point to fault location v , although it does not uniquely match the observed behavior. For example, consider a stuck-at fault at location v , which is detected by every pattern in the window. If $n = m$ and the deviation vectors $d((p_1 | P)_0_v), \dots, d((p_n | P)_0_v)$ are linearly independent, then $d((p_1 | P)_0_v), \dots, d((p_n | P)_0_v)$ span the complete m -dimensional space $GF(2)^m$ and the corresponding system of equations is solvable for any

observed faulty behavior $S_{obs}(P)$. If $n < m$, then n linearly independent deviation vectors $d((p_1 | P)_0_v), \dots, d((p_n | P)_0_v)$ can explain $2^n - 1$ of $2^m - 1$ possible faulty behaviors. Assuming that the faults are statistically independent and equally probable, this provides an aliasing probability of

$$\frac{2^n - 1}{2^m - 1} \approx 2^{n-m}. \quad (3)$$

To reduce the aliasing probability, the number of bits in the MISR can be made larger than the number of patterns in the test window.

The response data can be further reduced by skillfully exploiting the MISR properties [Indlekofer10, Cook11b]. In fact, it is sufficient to observe only a few bits of the signature and guarantee the propagation of fault effects with the help of a shadow MISR. To distinguish between permanent and non-permanent faults, a window is repeated immediately in case of a mismatch between the observed and expected response data. If the repeated test again provides the same faulty signature, then a permanent fault has been detected, and the faulty signature is stored in the fail memory. Otherwise a non-permanent fault must have been the root cause of failure.

2.3 Bayesian Diagnosis

Bayesian networks provide an engineering framework for the analysis of statistical data [Pearl88, Agosta04, Ben-Gal07, Barber12]. They are widely used in classification problems, such as medical diagnosis or troubleshooting applications. The problem is modeled by a collection of random variables and the dependencies between them, and the Bayesian network provides a graph representation of this model. More precisely, a Bayesian network is a directed acyclic graph, the vertices of which correspond to the random variables. The edges represent the dependencies between random variables, which are typically cause-effect relationships. They are labeled with conditional probabilities as illustrated in Figure 4 for a diagnosis problem.




Figure 4: Bayesian network for fault diagnosis.

The vertices f_1 to f_4 correspond to faults in a system. They are characterized by their “a priori” probabilities of occurrence $p(f_1)$ to $p(f_4)$. The vertices s_1 , s_2 , and s_3 represent the “symptoms” observed during a test. The edges are labeled with conditional probabilities, where $p(s | f)$ denotes the probability that symptom s is observed under the condition that fault f has occurred. In contrast to that, $p(f | s)$ is the probability that the fault f is a correct diagnosis for the symptom s . The goal of Bayesian inference is to deduce these “a posteriori” probabilities of faults and use them to guide the diagnosis. The deduction rules are based on the mathematical laws for conditional probabilities, in particular on Bayes Theorem stating

$$p(s | f) \cdot p(f) = p(f | s) \cdot p(s) \Leftrightarrow p(s | f) = \frac{p(f | s) \cdot p(s)}{p(f)}. \quad (4)$$

In general, exact inference in a Bayesian network is an NP-hard problem, but there are some efficient algorithms for restricted classes of networks such as, for example, message passing.

The network in Figure 4 is typical for a diagnosis problem where multiple faults can occur simultaneously. If only single faults are considered, then a simpler network can be used with only one multi-valued random variable representing the possible faults (cf. Fig. 4) [Przytula00].

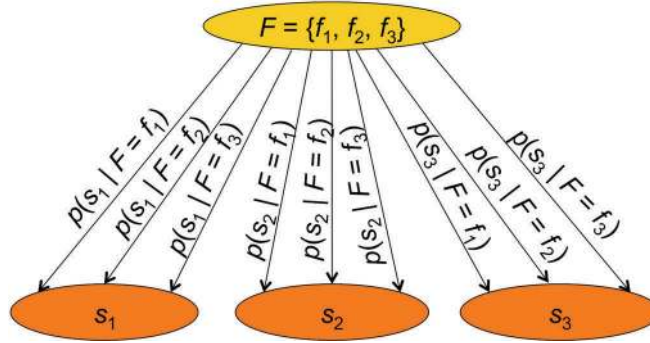


Figure 5: Bayesian network for the diagnosis of single faults.

In the context of electronic testing, Bayesian networks have already been successfully used for the diagnosis of analog and power circuits, e.g. [Aminian01, Liu06, Ye08, Krishnan10]. Also at the system and board level some pioneering approaches are available for coarse-grained diagnosis [Barford04, O’Farrill05, Zhang10]. So far, Bayesian reasoning has not yet been exploited for digital diagnosis at the gate level, although some research on improving volume diagnosis by other machine learning techniques has been done [Tang07, Wang09]. However, all these approaches address permanent faults and do not consider intermittent faults.

In [DeKler09] a general approach for system diagnosis in the presence of intermittent faults is described. In this work, intermittent faults are represented by pairs of a priori probabilities: the probability of occurrence and the probability of faulty behavior in the presence of the fault. The diagnosis procedure observes the system for several time steps. At each time step the Bayesian network is updated according to the observed behavior. The diagnosis strategy selects internal observation points, such that the a posteriori probabilities for wrong diagnoses decrease quickly. As a proof of concept the presented technique is applied to some of the smaller ISCAS85 benchmark circuits [Brglez85], but it is not directly applicable to larger industrial circuits.

3 Adaptive Bayesian Diagnosis

3.1 Diagnostic Flow

To deal with intermittent faults, the developed diagnostic procedure combines Bayesian reasoning with the window-based diagnosis of Section 2.2. As described above, the test is partitioned into N windows with n patterns each. In case of a mismatch between the observed and the expected response data, a window is immediately repeated one or several times until the correct signature is obtained or a user-defined limit R_{max} of repetitions is reached. As shown in [Amgalan08] the quality of the test increases with increasing R_{max} . However, for the sake of simplicity, in the following $R_{max} = 2$ is assumed without loss of generality.

The “adaptive behavior” is stored together with the failing signatures in the fail memory. Table 1 summarizes the different outcomes and their interpretations. The expected reference signature is denoted by s_{REF} , and the observed signatures in the first and second run of the window are denoted by s and s' .

TABLE I. ADAPTIVE BEHAVIOR FOR $R_{MAX} = 2$

Behavior	Code	Interpretation
$s \neq s_{REF}$ $s' = s_{REF}$	10	Transient or intermittent fault
$s \neq s_{REF}$ $s' \neq s_{REF}, s = s'$	00	Permanent or intermittent fault
$s \neq s_{REF}$ $s' \neq s_{REF}, s \neq s'$	01	a) Intermittent problem activated by different patterns or affecting several nodes b) Another transient fault

In this case the entries in the fail memory (symptoms) consist of three components: the observed faulty signature s , the window id and the code for the adaptive behavior. If $s \neq s'$, an additional entry with s' is added. This is illustrated in Figure 6 for embedded test.

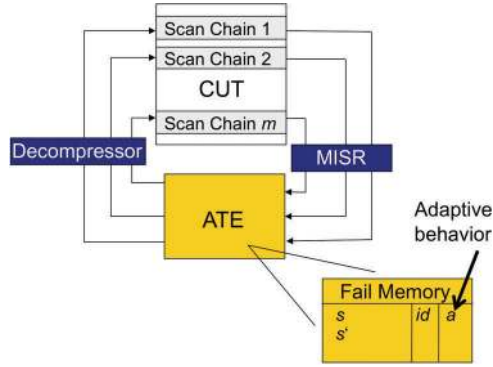


Figure 6: Hardware architecture for adaptive diagnosis.

After the test, the fail memory is read out and analyzed. In a first step, several rules for immediate classification are checked to reduce the computational effort:

(1) If the same faulty signature has been observed in all repetitions for one or more windows (code 00), then a permanent or intermittent fault is assumed.

(2) Similarly, if different faulty signatures have been observed in all repetitions for one or more windows (code 01), then an intermittent fault is assumed, too. Although the same behavior can also be caused by several transient faults, the pessimistic classification is preferred for the following reason: As the faulty behavior does not vanish after repeating the windows, the faults targeted in these windows cannot be excluded. Increasing the number of repetitions R_{max} alleviates this problem. Then it is more likely to finally observe a fault free behavior in the case of transient faults (code 10).

(3) The overall number of faulty sessions is used as an indicator for intermittent faults as follows. Let μ denote the expected transient error rate, then the probability that a single test session with n patterns fails due to transient errors is $1 - (1 - \mu)^n$. Depending on the number of repetitions, the overall number \tilde{N} of sessions is between N and $2N$. The probability that T out of \tilde{N} sessions fail due to transient faults is given by

$$\binom{\tilde{N}}{T} \left(1 - (1 - \mu)^n\right)^T \left((1 - \mu)^n\right)^{\tilde{N}-T} \quad (5)$$

Based on this, a bound T_{max} is selected, such that the probability of T_{max} faulty sessions caused by transient faults is below a user-defined threshold. If the number of fault sessions exceeds T_{max} , then an intermittent fault is assumed.

In all three cases described above, the chip is discarded as faulty, and the window-based diagnosis of section 2.2 is used for localizing the fault. For the remaining cases Bayesian reasoning is applied to

distinguish between transient and intermittent faults. The network is built automatically as described in section 3.2.

3.2 Building the Bayesian Network

In this work it is assumed that failures are caused by single intermittent or permanent faults, therefore the analysis relies on a single multi-valued random variable F representing the faults as illustrated in Figure 5. However, transient faults are considered as “background noise”, i.e. a transient fault may randomly appear at any time and also modify or mask the effect of an intermittent fault. To determine the range of F , the window-based diagnosis of section 2.2 is run as a pre-processor. A ranked list of k candidate fault locations is extracted, and for each fault location j an intermittent fault f_j is added to the range of F . To represent the possible occurrence of a transient fault in a circuit without intermittent or permanent faults, an extra fault f_{trans} is added. So the overall range of F is given by $\{f_1, \dots, f_k, f_{trans}\}$. The a priori probabilities for the faults in F can be based on expert knowledge.

The symptoms S correspond to the signatures observed for the N test sessions. For each test session T_i , $i = 1, \dots, N$, the observed signature s_i is added to S , and in case of a repeated window also the signature s_i' is added. The resulting network is illustrated in Figure 7 (a shaded signature s_i' is only present, if s_i is faulty).

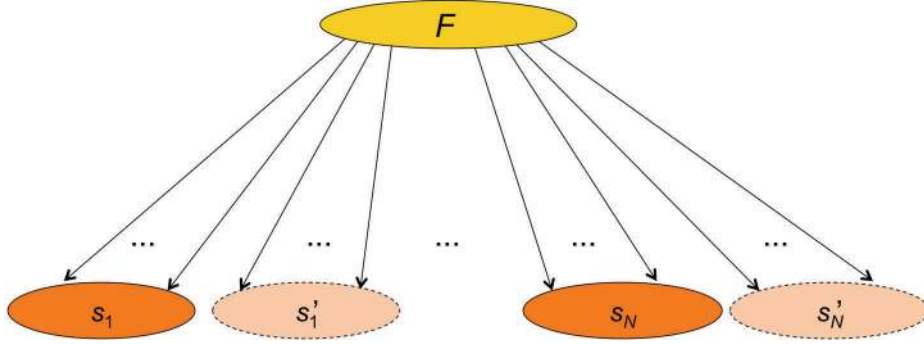


Figure 7: Bayesian network for adaptive diagnosis.

The computation of the conditional probabilities $p(s | f)$ is based on both a probabilistic and deterministic characterization of faults. Transient faults are described by their probability of occurrence $\mu = \mu_a \cdot \mu_p$, where μ_a is the probability that the fault is activated, and μ_p is the probability that the fault is propagated to the outputs. Intermittent faults are characterized by an activation probability λ , the patterns which can propagate them to the outputs, and their fault activation profile during the test.

Definition 4 (detecting patterns): Let P denote the set of all patterns in a window, and let v denote a candidate fault location. Furthermore let $b \in \{0, 1\}$ be the polarity of conditional stuck-at faults, then the set $P_d(v) := \{p \in P | d((p | P)_b_v) \neq 0\}$ is called the set of detecting patterns for v and $n_d(v) := |P_d(v)|$ is the number of detecting patterns for v .

Definition 5 (activation profile of a signature): Let P denote the set of all patterns in a window, let s be the observed signature, and v a candidate fault location. Furthermore let $b \in \{0, 1\}$ be the polarity of conditional stuck-at faults, then $P_d(v, s) := \{p \in P_d(v) | (p | P)_b_v \text{ must have been active to explain } s\}$ and $P_i(v, s) := \{p \in P_d(v) | (p | P)_b_v \text{ must have been inactive to explain } s\}$ describe the activation profile for v to explain s . The respective cardinalities are denoted by $n_d(v, s) := |P_d(v, s)|$ and $n_i(v, s) := |P_i(v, s)|$.

To illustrate these definitions the small circuit depicted in Figure 8 is used as an example. Assume that a complete test is performed in two sessions with four patterns each and after preprocessing there is only one candidate fault location at node v . Table II shows the circuit behavior for the first window in the fault free case and the in the presence of the faults d stuck-at zero (0_v) and d stuck-at one (1_v).




Figure 8: Example for activation profile.

TABLE II: CIRCUIT BEHAVIOR FOR SESSION 1

Test Pattern	Input	Fault Free Outputs	0_v	1_v
	a b c	x y	x y	x y
p ₁	000	00	00	10
p ₂	001	10	10	11
p ₃	010	00	00	10
p ₄	011	10	10	11

For the sake of simplicity assume that no further response compaction is applied and a symptom of a test session is described by the sequence of output vectors $(x(p_1) y(p_1) x(p_2) y(p_2) x(p_3) y(p_3) x(p_4) y(p_4))$, i.e. the fault free case is described by the symptom (00 10 00 10). The fault 0_v provides exactly the same symptom and cannot be detected in this window. The fault 1_v leads to the symptom (10 11 10 11), i.e. 1_v can be detected by each pattern and $P_d(v) = \{p_1, p_2, p_3, p_4\}$ and $n_d(v) = 4$. However, as we are dealing with intermittent faults with activation probability $\lambda < 1$, 1_v may not be active for all the patterns. Assume that the symptom $s = (00 10 10 10)$ is observed after the session. To explain this symptom, 1_v must have been active for p_3 and inactive for the remaining patterns, i.e. $P_a(v, s) = \{p_3\}$ and $P_i(v, s) = \{p_1, p_2, p_4\}$. The respective cardinalities are $n_a(v, s) = 1$ and $n_i(v, s) = 3$. In contrast, the symptom $s = (01 10 00 10)$ cannot be explained at all. One or several transient faults must have caused the deviation of the fault free signature or modified the fault effects of an intermittent fault.

Conditional Probabilities for Faulty Signatures

For a test session with patterns P , a faulty signature $s \neq s_{REF}$ and an intermittent fault f at location v , the preprocessing step provides information about the detecting patterns and the fault activation profile. If f can explain the faulty signature, the fault effect has not been modified by transient noise, and the conditional probability $p(s \neq s_{REF} | f)$ is given by

$$p(s \neq s_{REF} | f) = \lambda^{n_a(v,s)} (1 - \lambda)^{n_i(v,s)} (1 - \mu)^n, \quad (6)$$

which is the product of the probabilities that f is activated by the $n_a(v, s)$ patterns in $P_d(v, s)$, f is not activated by the remaining $n_i(v, s)$ detecting patterns, and no transient fault has occurred during the complete session with n patterns. For the small example discussed above, the conditional probability that $s = (00 10 10 10)$ is observed when 1_v is present is $p(s | v) = \lambda (1 - \lambda)^3 (1 - \mu)^4$.

If the intermittent fault f at location v cannot explain the signature s , as $s = (01 10 00 10)$ in the example, then the fault effect must have been modified by transient noise, or the fault has not been activated at all and the faulty signature results from a transient fault only. The conditional probability $p(s \neq s_{REF} | f)$ is then given by

$$p(s \neq s_{REF} | f) = \left(1 - (1 - \lambda)^{n_d(v)}\right) \left(1 - (1 - \mu)^n\right) + (1 - \lambda)^{n_d(v)} \left(1 - (1 - \mu)^n\right) = 1 - (1 - \mu)^n, \quad (7)$$

where $(1 - \lambda)^{n_d^{(v)}}$ is the probability that the fault has not been activated at all, and $1 - (1 - \lambda)^{n_d^{(v)}}$ is the probability that the fault has been activated by one or more patterns. Similarly, $1 - (1 - \mu)^n$ is the probability that at least one transient fault has occurred in the session.

For f_{trans} the term $p(s \neq s_{REF} | f_{trans})$ describes the probability that a transient fault, which has been activated in the circuit, is propagated to the outputs by at least one of the n patterns. This can be computed as $1 - (1 - \mu_p)^n$.

Conditional Probabilities for Fault Free Signatures

In the presence of an intermittent fault f , a correct signature $s = s_{REF}$ can only be obtained, if the fault effect has been masked by transient noise, or the fault has not been activated at all and there is no transient noise. The probability $p(s = s_{REF} | f)$ is therefore computed as

$$p(s = s_{REF} | f) = \left(1 - (1 - \lambda)^{n_d^{(v)}}\right) \left(1 - (1 - \mu)^n\right) + (1 - \lambda)^{n_d^{(v)}} (1 - \mu)^n = 1 - (1 - \mu)^n. \quad (8)$$

If a transient fault only has occurred in the circuit, the correct signature can only be obtained, if it does not propagate to the outputs. The conditional probability $p(s = s_{REF} | f_{trans})$ is therefore computed as $(1 - \mu_p)^n$.

Extension to Multiple Faults

The approach can be extended to multiple faults, but then a network as depicted in Figure 4 with individual nodes for each fault location must be built. Inference becomes more complex, because joint probabilities for the faults must be considered. In addition the basic probabilities $p(s | f)$ also joint conditional probabilities $p(s | f_1, \dots, f_k)$ must be determined and added to the network. For preprocessing the diagnosis technique for multiple faults in [Cook14] can be used.

3.3 Calibrating the Network

When the Bayesian diagnosis is applied in practice the model parameters are not known and the network must be properly calibrated. As mentioned above, the a priori probabilities can be based on expert knowledge. If this is not available the faults are assumed being uniformly distributed. As shown below the diagnostic procedure is robust against this approximation.

Similarly, the transient error rate $\mu = \mu_a \cdot \mu_p$ can be set according to previous observations, and the propagation probability μ_p of a transient fault at a random location can be approximated by fault simulation.

Estimating the activation rate λ for intermittent faults is more challenging. As the defect mechanism causing the intermittent fault is not known beforehand, λ cannot be based on previous observations. Instead, the classification is performed for several different λ values, and for each fault the solution with the highest probability $p(f | s)$ is selected. This allows the method to work for different intermittent fault mechanisms, even without detailed statistics for every possible defect mechanism.

4 Simulation Results

To validate the presented approach, a simulation study was performed with several industrial circuits kindly provided by NXP. The fault-simulation environment used in earlier work was extended to handle intermittent faults as described in Section 2.1 [Cook11a, Cook11b]. In three different experiments, intermittent faults, transient faults, and intermittent faults with transient background noise were injected randomly. Then a mixed-mode test was run with 10,000 pseudo-random patterns and deterministic patterns for the remaining hard to test faults. The test was partitioned into windows of 32 patterns, and for response compaction a 48-bit MISR was used. For fault detection during test, only 8 bits of the MISR were observed. As described in Section 3.3 the classification was performed with several different values of λ , here 0.4, 0.1, and 0.001 were used. The specific details and results of the experiments are described in the following subsections.

4.1 Intermittent Faults

To inject an intermittent fault at a randomly chosen location v , high frequency power droop was simulated with varying parameters for the radius r of the neighborhood $\mathcal{N}^r(v)$ and the threshold value τ for the switching activity. The radius r ranged from 1 to 3, and the threshold value τ was set to 15 %, 30 % or 50 % of nodes in the neighborhood. To model the unpredictable behavior of intermittent faults, the fault was only activated with probability 0.5, if the conditions were met. Overall, 120 randomly selected candidate locations were analyzed for all parameter combinations. As the available netlists represented the circuits before technology mapping, faults at gates with extremely large fanouts and unrealistically large neighborhoods were not considered. To correlate the experimental data with the network model described in Section 3.2, the actual fault activation rate λ_{exp} achieved in the experiments was estimated by logic simulation. The observed activation rates ranged from zero over low rates in the order of 10^{-5} to high rates larger than 0.5.

For the third immediate classification rule in section 3.1, the limit T_{max} for the number of faulty sessions that can be caused by transient errors was set to $T_{max} = 10$. Assuming a transient error rate $\mu = 10^{-5}$ this keeps the probability of T_{max} faulty sessions due to transient errors below 10^{-13} . Table III shows the results after immediate classification.

TABLE III: IMMEDIATE CLASSIFICATION ($T_{MAX} = 10$)

Circuit	Experiments	Fault Free	Permanent Failures	Changing Failures	Above T_{max}	Bayesian Classification
p45k	1044	171	692	140	1	40
p100k	1059	112	807	110	1	29
p141k	1014	116	761	107	2	28
p239k	1050	84	827	116	2	21
p259k	966	54	835	55	0	22
p267k	1053	56	733	234	3	27
p269k	1041	151	741	119	0	30
p279k	1044	227	671	108	0	38
p286k	891	101	693	69	0	28
p295k	1047	271	663	60	0	53

The first column shows the circuit names, and the overall number of evaluated experiments is listed in the second column. The following columns three to five show the number of experiments without any fault effect, the number of experiments with at least one session showing the same faulty signature in all repetitions, the number of remaining experiments with at least one session having different faulty signatures, the number of remaining experiments with the number of faulty sessions exceeding T_{max} . Finally, the sixth column reports the number of experiments where Bayesian classification was needed. The results show that the immediate classification helps to reduce the computational effort for Bayesian classification drastically. The results of the overall classification procedure are summarized in Table IV.

To evaluate the quality of the overall classification procedure, experiments without any fault effect were discarded. The second column therefore shows the number of experiments with at least one faulty session. Column three summarizes the results of immediate classification, and column four lists the number of experiments where the Bayesian network identified an intermittent fault. As both permanent and intermittent faults are considered as critical failures, the overall number of experiments with classification permanent or intermittent is shown in column five. Accordingly, if an intermittent fault is classified as critical failure this is considered as correct classification. Therefore the values in column five are divided by the number of experiments with failures to obtain the percentage of correctly classified failures in column six. The ratio of

correctly classified faults is above 95% for all circuits, but obviously the Bayesian network could not classify all injected faults correctly. The results of Bayesian classification are analyzed in more detail in Table V.

The analysis distinguishes between experiments with a single faulty session only and experiments with 2 to 9 faulty sessions. For each case the number of experiments is shown as well as the number of experiments with classification “transient” (T) and “intermittent” (I). The last column presents the ratio of correctly classified experiments considering only those cases with 2 to 9 faulty sessions. If only a single faulty session is observed during test, then the Bayesian network cannot collect enough information to distinguish the observed failure from a transient failure. However, if more faulty sessions occur, then the Bayesian network provides the correct classification in almost all cases.

TABLE IV: OVERALL CLASSIFICATION ($T_{MAX} = 10$)

Circuit	Experiments with Failures (F)	Critical Failures after Immediate Classification (IC)	Classified as Intermittent by Bayesian Classification (BC)	Overall Critical Failures (CF = IC + BC)	Correctly Classified (CF/F)
p45k	873	833	11	844	0.967
p100k	947	918	13	931	0.983
p141k	898	870	7	877	0.977
p239k	966	945	15	960	0.994
p259k	912	890	13	903	0.990
p267k	997	970	10	980	0.983
p269k	890	860	12	872	0.980
p279k	817	779	17	796	0.974
p286k	790	762	10	772	0.977
p295k	776	723	17	740	0.954

TABLE V: BAYESIAN CLASSIFICATION ($T_{MAX} = 10$)

Circuit	Experiments	1 Faulty Session			2 to 9 Faulty Sessions			
		Experiments	T	I	Experiments	T	I	Percentage Classified as Intermittent
p45k	40	29	29	0	11	0	11	1.000
p100k	29	15	15	0	14	1	13	0.929
p141k	28	21	21	0	7	0	7	1.000
p239k	21	6	6	0	15	0	15	1.000
p259k	22	9	9	0	13	0	13	1.000
p267k	27	17	17	0	10	0	10	1.000
p269k	30	18	18	0	12	0	12	1.000
p279k	38	21	21	0	17	0	17	1.000
p286k	28	18	18	0	10	0	10	1.000
p295k	53	34	34	0	19	2	17	0.895

4.2 Transient Faults

The injection of a transient fault at location v is randomly controlled by the activation probability μ_a . In the reported experiments, probabilities $2 \cdot 10^{-3}$, $2 \cdot 10^{-4}$, and $2 \cdot 10^{-5}$ were chosen. It should be noted that these probabilities are higher than the probabilities typically reported for transient errors caused by radiation. But here also transient errors caused by parameter variations are taken into account, which may occur with higher probabilities. This makes it more difficult to distinguish transient faults from intermittent faults.

Transient faults were assumed to behave like stuck-at faults being active for exactly one clock cycle. Overall 140 experiments were performed for each circuit. Table VI shows the results of immediate classification using again $T_{max} = 10$ and the same format as in Section 4.1.

TABLE VI: IMMEDIATE CLASSIFICATION WITH $T_{MAX} = 10$

Circuit	Experiments	Fault Free	Permanent Failures	Changing Failures	Above T_{max}	Bayesian Classification
p45k	140	11	0	43	14	72
p100k	140	15	0	15	15	95
p141k	140	16	0	22	29	73
p239k	140	13	0	36	32	59
p259k	140	14	0	15	18	93
p267k	140	14	0	22	50	54
p269k	140	13	0	32	45	50
p279k	140	16	0	31	23	70
p286k	140	16	0	14	14	96
p295k	140	13	0	14	17	96

While $T_{max} = 10$ worked well for the experiments in section 4.1, where a transient error rate of $\mu = 10^{-5}$ was assumed, for the higher rates used in this experiment, it results in a wrong classification for a considerable number of experiments (column 6). In fact, in all the experiments with at least T_{max} faulty sessions the injection rate was $2 \cdot 10^{-3}$. If expert knowledge about the expected transient error rate is available, then T_{max} can be better tuned to the considered scenario. Assuming for example a transient error rate of $\mu = 10^{-4}$, the threshold T_{max} must be set to 21 to keep the probability of T_{max} faulty sessions caused by transient errors below 10^{-13} . Table VII updates the results for $T_{max} = 21$. It can be observed, that in this case the number of faulty session is always below T_{max} , and consequently the Bayesian network must be applied more often.

TABLE VII: IMMEDIATE CLASSIFICATION WITH $T_{MAX} = 21$

Circuit	Experiments	Fault Free	Permanent Failures	Changing Failures	Above T_{max}	Bayesian Classification
p45k	140	11	0	43	0	86
p100k	140	15	0	15	0	110
p141k	140	16	0	22	0	102
p239k	140	13	0	36	0	91
p259k	140	14	0	15	0	111
p267k	140	14	0	22	0	104
p269k	140	13	0	32	0	95
p279k	140	16	0	31	0	93
p286k	140	16	0	14	0	110
p295k	140	13	0	14	0	113

The results of the overall classification procedure for $T_{max} = 21$ are summarized in Table VIII using the same format as in Section 4.1. As transient faults were injected in this experiment, the classification “critical failure” is wrong, and the number of correctly classified experiments is obtained by subtracting column 5 (CF) from column 2 (F). The overall ratios of correctly classified experiments are lower than in section 4.1. This is partly due to the immediate classification rule 2, which considers a session with changing faulty signatures as an indicator of a critical failure. However, as pointed out in section 3.1 this rule is necessary to

ensure a high product quality, and its pessimistic effect can be reduced by allowing more than just one repetition of a test session ($R_{max} = 2$) like here.

TABLE VIII: OVERALL CLASSIFICATION ($T_{max} = 21$)

Circuit	Experiments with Failures (F)	Critical Failures after Immediate Classification (IC)	Classified as Intermittent by Bayesian Classification (BC)	Overall Critical Failures (CF = IC + BC)	Correctly Classified ((F - CF)/F)
p45k	129	43	1	44	0.659
p100k	125	15	4	19	0.848
p141k	124	22	1	23	0.815
p239k	127	36	1	37	0.709
p259k	126	15	2	17	0.865
p267k	126	22	1	23	0.817
p269k	127	32	2	34	0.732
p279k	124	31	0	31	0.750
p286k	124	14	0	14	0.887
p295k	127	14	0	14	0.890

The quality of the Bayesian classification is studied in more detail in Table IX. In the same way as in Section 4.1 the ratio of correctly classified experiments is only determined for experiments with 2 to 20 faulty sessions. Again, this rate is very high for all circuits ranging from 96.2 % to 100 %.

TABLE IX: BAYESIAN CLASSIFICATION ($T_{max} = 21$)

Circuit	Experiments	1 Faulty Session			2 to 20 Faulty Sessions			
		Experiments	T	I	Experiments	T	I	Percentage Classified as Transient
p45k	86	6	6	0	80	79	1	0.988
p100k	110	6	6	0	104	100	4	0.962
p141k	102	6	5	1	96	96	0	1.000
p239k	91	3	3	0	88	87	1	0.989
p259k	111	4	4	0	107	105	2	0.981
p267k	104	2	1	1	102	102	0	1.000
p269k	95	5	4	1	90	89	1	0.989
p279k	93	1	1	0	92	92	0	1.000
p286k	110	6	6	0	104	104	0	1.000
p295k	113	6	6	0	107	107	0	1.000

4.3 Intermittent Faults with Transient Background Noise

In this experiment intermittent faults were injected as described in Section 4.1, and in addition to that transients faults were randomly injected with probabilities $2 \cdot 10^{-3}$, $2 \cdot 10^{-4}$, and $2 \cdot 10^{-5}$ as described in Section 4.2. The threshold T_{max} was set to 10 to allow a comparison with the experiments in Section 4.1, where intermittent faults without background noise were considered. Tables X and XI show the results of immediate classification as well as the results of the overall classification procedure.

It can be observed that the ratios of correctly classified experiments are still very high, yet they are a little bit lower than in the first experiment without background noise. This can be explained by the relatively high activation rates for the background noise. In the same experiment both intermittent and transient faults with

comparable activation rate may be present, which makes it extremely difficult for the Bayesian network to distinguish between the two types of faults.

TABLE X: IMMEDIATE CLASSIFICATION ($T_{MAX} = 10$)

Circuit	Experiments	Fault Free	Permanent Failures	Changing Failures	Above T_{max}	Bayesian Classification
p45k	1044	51	696	157	25	115
p100k	1788	44	1369	255	19	101
p141k	1014	36	761	133	30	54
p239k	1050	31	817	143	32	27
p259k	966	32	832	62	7	33
p267k	1053	65	734	159	52	43
p269k	1041	54	737	178	36	36
p279k	1044	73	670	159	36	106
p286k	891	57	693	76	10	55
p295k	1047	80	665	85	41	176

TABLE XI: OVERALL CLASSIFICATION ($T_{MAX} = 10$)

Circuit	Experiments with Failures (F)	Critical Failures after Immediate Classification (IC)	Classified as Intermittent by Bayesian Classification (BC)	Overall Critical Failures (CF = IC + BC)	Correctly Classified (CF/F)
p45k	993	878	20	898	0.904
p100k	1744	1643	47	1690	0.969
p141k	978	924	30	954	0.975
p239k	1019	992	12	1004	0.985
p259k	934	901	13	914	0.979
p267k	988	945	27	972	0.984
p269k	987	951	19	970	0.983
p279k	971	865	55	920	0.947
p286k	834	779	19	798	0.957
p295k	967	791	28	819	0.847

Table XII shows that, despite these problems, the Bayesian network still considerably helps to adjust the results of immediate classification. As before, the experiments with 2 to 9 faulty sessions give a deeper insight into the quality of the Bayesian classification. As expected the percentage of correctly classified intermittent faults is lower than before, because fault effects of intermittent faults can now be modified by transient faults and the explanation of symptoms is not as clear as in the previous experiments. However, compared to a diagnosis approach without Bayesian classification this can still help in two ways. Without Bayesian classification, the immediate classification must also provide a rule for dealing with changing faulty signatures. On the one hand, a pessimistic rule would sort out all devices with this result, and the Bayesian classification reduces unnecessary yield loss. On the other hand, an optimistic rule would assume transient faults and accept the respective devices. In this case, the Bayesian classification considerably improves the product quality. Hence, in average around 50 % of the intermittent faults are finally filtered out by the Bayesian classification in table XII, and only for two circuits the additional gain is below 20 %.

TABLE XII: BAYESIAN CLASSIFICATION ($T_{MAX} = 10$)

Circuit	Experiments	1 Faulty Session			2 to 9 Faulty Sessions			
		Experiments	T	I	Experiments	T	I	Percentage Classified as Intermittent
p45k	115	8	8	0	107	87	20	0.187
p100k	101	7	7	0	94	47	47	0.500
p141k	54	2	2	0	52	22	30	0.577
p239k	27	4	4	0	23	11	12	0.522
p259k	33	4	4	0	29	16	13	0.448
p267k	43	2	2	0	41	14	27	0.659
p269k	36	4	4	0	32	13	19	0.594
p279k	106	3	3	0	103	48	55	0.534
p286k	55	8	8	0	47	28	19	0.404
p295k	176	10	10	0	166	138	28	0.169

5 Conclusions

For innovative technologies, for opportunistic computing schemes and for some yield enhancement strategies, the amount of both transient and intermittent faults will increase significantly. This paper presented a unified method for classifying faults into permanent, intermittent and transient ones during volume testing even after test response compaction. The method can be seamlessly integrated into the standard embedded test and built-in self-test schemes. If logic diagnosis and Bayesian classification are combined, it is possible to identify intermittent faults with a confidence of more than 98%. By adjusting a priori probabilities and decision values it is possible to find appropriate trade-offs between yield by reducing false-positives for intermittent faults and product quality by reducing false-positives for transient faults.

6 References

- [Agosta04] J. M. Agosta, T. Gardos, "Bayes Network "Smart" Diagnostics," Intel Technology Journal, Vol. 8, No. 4, November 2004.
- [Amgalan08] U. Amgalan, C. Hachmann, S. Hellebrand, Hans-Joachim Wunderlich, "Signature Rollback - A Technique for Testing Robust Circuits," Proceedings IEEE VLSI Test Symposium (VTS'08), San Diego, CA, USA, May 2008, pp. 125-130.
- [Aminian01] F. Aminian and M. Aminian, "Fault Diagnosis of Analog Circuits Using Bayesian Neural Networks with Wavelet Transform as Preprocessor," Journal of Electronic Testing (JETTA), Vol. 17, No. 1, February 2001, pp. 29-36.
- [Barber12] D. Barber, "Bayesian Reasoning and Machine Learning," New York: Cambridge University Press, 2012.
- [Bardell82] P. H. Bardell and W. H. McAnney, "Self-Testing of Multichip Logic Modules," Proc. IEEE Int. Test Conference (ITC'82), Philadelphia, PA, USA, Nov. 1982, pp. 200-204.
- [Barford04] L. Barford, V. Kanevsky, and L. Kamas, "Bayesian fault diagnosis in large-scale measurement systems," Proceedings IEEE Instrumentation and Measurement Technology Conference (IMTC'04), Como, Italy, 2004, Vol. 2, pp. 1234-1239.
- [Baumann05] R. Baumann, "Soft errors in advanced computer systems," IEEE Design & Test of Computers, Vol. 22, No. 3, 2005, pp. 258-266.
- [Ben-Gal07] I. Ben-Gal, "Bayesian Networks," in F. Ruggeri, F. Faltin, and R. Kenett, "Encyclopedia of Statistics in Quality & Reliability," Wiley & Sons, 2007.
- [Borkar05] S. Borkar, "Designing Reliable Systems from Unreliable Components: The Challenges of Transistor Variability and Degradation," IEEE Micro, Nov. 2005, pp. 10-16.
- [Brglez85] F. Brglez, et al., "Accelerated ATPG and fault grading via testability analysis," Proc. IEEE International Symposium on Circuits and Systems (ISCAS'85), Kyoto, 1985, pp. 695-698.

- [Cheng06] W.-T. Cheng, M. Sharma, T. Rinderknecht, L. Lai, and C. Hill, "Signature based diagnosis for logic BIST," in Proc. IEEE Int. Test Conference (ITC'06), Santa Clara, CA, USA, 2006, pp. 1-9.
- [Constantinescu03] C. Constantinescu, "Trends and Challenges in VLSI Circuit Reliability," IEEE Micro, Vol. 23, No. 4, July 2003, pp. 14-19.
- [Cook11a] A. Cook, M. Elm, H.-J. Wunderlich, U. Abelein, "Structural In-Field Diagnosis for Random Logic Circuits," Proc. European Test Symposium (ETS'11), Trondheim, May 2011, pp. 111-116.
- [Cook11b] A. Cook, S. Hellebrand, T. Indlekofer, H.-J. Wunderlich, "Diagnostic Test of Robust Circuits," Proceedings Asian Test Symposium (ATS'11), New Delhi, India, November 2011, pp. 285-290.
- [Cook14] A. Cook and H.-J. Wunderlich, "Diagnosis of Multiple Faults with Highly Compacted Test Responses," Proc. 19th IEEE European Test Symposium (ETS'14), Paderborn, Germany, May 2014, pp. 1-6
- [DeKler09] Johan De Kleer, "Diagnosing multiple persistent and intermittent faults," Proceedings 21st International Joint Conference on Artificial Intelligence (IJCAI'09), Pasadena, CA, USA, 2009, pp. 733-738.
- [Elm10] M. Elm and H.-J. Wunderlich, "BISD: Scan-Based Built-In Self-Diagnosis," Proc. Design Automation and Test in Europe (DATE'10), Dresden, Germany, March 8-12, 2010.
- [Ernst04] D. Ernst, et al., "Razor: Circuit-Level Correction of Timing Errors for Low Power Operation," IEEE Micro, Vol. 24, No. 6, Nov.-Dec. 2004, pp. 10-20.
- [Fechner09] Bernd Fechner, "A Dynamic Fault Classification Scheme", Proceedings European Safety and Reliability Conference (ESREL'08), Valencia, Spain, September 2008, pp. 147-153.
- [Ghosh00] J. Ghosh-Dastidar and N. A. Toubia, "A rapid and scalable diagnosis scheme for BIST environments with a large number of scan chains," Proc. 18th IEEE VLSI Test Symposium (VTS'00), Montreal, Canada, 2000, pp. 79-85.
- [Gupta13] P. Gupta, Y. Agarwal, L. Dolecek, N. Dutt, R. K. Gupta, R. Kumar, S. Mitra, A. Nicolau, T. S. Rosing, M. B. Srivastava, S. Swanson, D. Sylvester, "Underdesigned and Opportunistic Computing in Presence of Hardware Variability," IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, Vol.32, No.1, Jan. 2013, pp. 8-23.
- [Holst09] S. Holst and H.-J. Wunderlich, "Adaptive Debug and Diagnosis Without Fault Dictionaries," Journal of Electronic Testing: Theory and Applications (JETTA), Vol. 25, No. 4-5, pp. 259-268.
- [Indlekofer10] T. Indlekofer, M. Schnittger, S. Hellebrand, "Efficient Test Response Compaction for Robust BIST Using Parity Sequences," Proc. 28th IEEE Int. Conference on Computer Design (ICCD'10), Amsterdam, The Netherlands, October 2010, pp. 480-485.
- [Krishnan10] S. Krishnan, K. D. Doornbos, R. Brand, H. G. Kerkhoff, "Block-Level Bayesian Diagnosis of Analogue Electronic Circuits," Proceeding Design, Automation and Test in Europe (DATE'10), Dresden, Germany, March 2010, pp. 1-6.
- [Liu02] C. Liu, K. Chakrabarty, and M. Goessel, "An interval-based diagnosis scheme for identifying failing vectors in a scan-BIST environment," Proc. Design, Automation and Test in Europe (DATE'02), Paris, France, 2002, pp. 382-386.
- [Liu06] F. Liu, P. K. Nikolov, and S. Ozev. "Parametric Fault Diagnosis for Analog Circuits Using a Bayesian Framework," Proceedings 24th IEEE VLSI Test Symposium (VTS'06), Berkeley, CA, USA, April 30 - May 4, 2006, pp. 272-277.
- [Nicolaidis99] M. Nicolaidis, "Time Redundancy Based Soft-Error Tolerant Circuits to Rescue Very Deep Submicron," Proc. 17th IEEE VLSI Test Symposium, San Diego, CA, USA, April 1999.
- [Nicolaidis07] M. Nicolaidis, "GRAAL: A New Fault Tolerant Design Paradigm for Mitigating the Flaws of Deep Nanometric Designs," Proceedings IEEE International Test Conference (ITC'07), San Jose, CA, USA, October 2007, pp. 1-10.
- [O'Farrill05] C. O'Farrill, M. Moakil-Chbany, and B. Eklow, "Optimized reasoning-based diagnosis for non-random, board-level, production defects," Proceedings IEEE International Test Conference (ITC'05), Austin, TX, USA, November 2005, pp. 173-179.
- [Pearl88] J. Pearl, "Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference," Revised 2nd Printing, San Francisco: Morgan Kaufmann Publishers, 1988.
- [Polian07] I. Polian, A. Czutro, S. Kundu, B. Becker, "Power Droop Testing," IEEE Design & Test of Computers, Vol. 24, No. 3, May-June 2007, pp. 276-284.
- [Przytula00] K. W. Przytula, D. Thompson, "Construction of Bayesian networks for diagnostics," Proceedings 2000 IEEE Aerospace Conference, Big Sky, Montana, USA, March 2000, Vol. 5, pp.193-200.
- [Rajski99] J. Rajski and J. Tyszer, "Diagnosis of scan cells in BIST environment," IEEE Trans. on Computers, Vol. 48, No. 7, July 1999, pp. 724-731.
- [Tang07] H. Tang, S. Manish, J. Rajski, M. Keim, B. Benware, "Analyzing Volume Diagnosis Results with Statistical Learning for Yield Improvement," Proceeding 12th IEEE European Test Symposium (ETS'07), Freiburg, Germany, May 2007, pp. 145-150.

- [Tirumurti04] C. Tirumurti, S. Kundu, S. Sur-Kolay, Y.-S. Chang, "A Modeling Approach for Addressing Power Supply Switching Noise Related Failures of Integrated Circuits," Proceedings Design, Automation and Test in Europe (DATE'04), Paris, France, February 2004, pp. 1078-1083.
- [Wang09] S. Wang, W. Wei, "Machine Learning-Based Volume Diagnosis," Proceedings Design, Automation and Test in Europe (DATE'09), Nice, France, April 2009, pp. 902-905.
- [Wohl02] P. Wohl, J. A. Waicukauski, S. Patel, and G. Maston, "Effective diagnostics through interval unloads in a BIST environment," Proc. 39th Design Automation Conference (DAC'02), New Orleans, LA, USA, 2002, pp. 249-254.
- [Ye08] B. Ye, Z. Luo, W. Zhang, and C. Piao, "Fault diagnosis for power circuits based on SVM within the Bayesian framework," Proceedings World Congress on Intelligent Control and Automation, Chongqing, China (WCICA'08), 2008, pp. 5125-5129.
- [Zhang10] Z. Zhang, Z. Wang, X. Gu, and K. Chakrabarty, "Board-Level Fault Diagnosis using Bayesian Inference," Proceedings 28th IEEE VLSI Test Symposium (VTS'10), Santa Cruz, CA, USA, April 2010, pp. 244-249.