



**HAL**  
open science

## Maximizing Yield for Approximate Integrated Circuits

Marcello Traiola, Arnaud Virazel, Patrick Girard, Mario Barbareschi, Alberto Bosio

► **To cite this version:**

Marcello Traiola, Arnaud Virazel, Patrick Girard, Mario Barbareschi, Alberto Bosio. Maximizing Yield for Approximate Integrated Circuits. DATE 2020 - 23rd Design, Automation and Test in Europe Conference and Exhibition, Mar 2020, Grenoble, France. pp.810-815, 10.23919/DATE48585.2020.9116341 . lirmm-03036002

**HAL Id: lirmm-03036002**

**<https://hal-lirmm.ccsd.cnrs.fr/lirmm-03036002v1>**

Submitted on 2 Dec 2020

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

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

# Maximizing Yield for Approximate Integrated Circuits

Author(s) omitted for blind review

**Abstract**—Approximate Integrated Circuits (AxICs) have emerged in the last decade as an outcome of Approximate Computing (AxC) paradigm. AxC focuses on efficiency of computing systems by sacrificing some computation quality. As the AxICs spread, consequent challenges to test them arose. On the other hand, the opportunity to increase the production yield emerged in the AxICs context. Indeed, some particular defects in the manufactured AxIC might not catastrophically impact the final circuit quality. Therefore, some defective AxICs might still be acceptable. Efforts to detect favorable conditions to consider defective AxICs as acceptable – with the goal to increase the production yield – have been done in last years. Unfortunately, the final achieved yield gain is often not as high as expected. In this work, we propose a methodology to actually achieve a yield gain as close as possible to expectations, by proposing a technique to suitably apply tests to AxICs. Experiments carried out on state-of-the-art AxICs show yield gain results very close to the expected ones (i.e., between 98% and 100% of the expectations).

**Index Terms**—Approximate Computing, Approximate Circuits, Testing, Signature analysis

## I. INTRODUCTION

In last decades, the demand of computing efficiency has been growing constantly. On one side, the relevance of new-generation power-consuming applications increases. On the other side, low-power portable devices are more and more deployed in the consumer market. Therefore, new computing paradigms are necessary to cope with the competing requirements introduced by modern technologies [1]. In recent years, several studies on *Recognition, Mining and Synthesis (RMS)* applications have been conducted [1]–[4]. A very interesting peculiarity has been identified, i.e. the *intrinsic resiliency* of those applications. Such a property allows RMS applications to be highly tolerant to errors. This is due to different factors, such as noisy data processed by these applications, non-deterministic algorithms used, and possible non-unique answers [1]. These properties have been exploited by a new and increasingly established computing paradigm, namely the Approximate Computing (AxC) [1], [2].

AxC cleverly profits from the RMS intrinsic resiliency to achieve gains in terms of power consumption, run time, and/or chip area. Indeed, by introducing selective relaxations of non-critical specifications, some parts of the target computing system can be simplified, at the cost of a slight accuracy reduction. Additionally, AxC is able to target different layers of computing systems, from hardware to software [2]. In this work, we focus on Approximate Integrated Circuits (AxICs), which are the outcome of AxC application at hardware level, specifically on Integrated Circuits (ICs). In particular, we focus on IC *functional approximation*. Functional approximation has

been employed in the last few years to design efficient AxICs (in terms of area, timing, and power consumption) by systematically modifying IC functional behavior, thus introducing controlled errors [5]–[14]. To measure the error produced by an AxIC, several error metrics have been proposed in the literature [15]. As a consequence of the AxICs increasing relevance, it becomes important to address the new challenges to test such circuits. In this respect, some previous works [16]–[23] drew attention to the challenges that functional approximation entails for testing procedures. At the same time, opportunities come with IC functional approximation. More in details, the concept of *acceptable circuit* changes: while conventionally a circuit is *good* if its responses are never different from the expected ones, in the AxIC context some unexpected responses might be still acceptable, according to the error metric maximum threshold. Therefore, some *acceptable defects* may be left undetected, ultimately leading to a production yield gain (i.e., the percentage of acceptable circuits, among all fabricated circuits, increases).

In recent years, several works have been presented to classify AxIC faults into *non-redundant* and *ax-redundant* (i.e., catastrophic and acceptable, respectively) [19]–[22] according to an error threshold (i.e., maximum tolerable amount of error). As a result of this classification, two lists of faults are obtained (i.e., non-redundant and ax-redundant). Consequently, the Automatic Test Pattern Generation (ATPG) targets only the non-redundant faults. Obtained tests prevent catastrophic failures from occurring, by detecting non-redundant defects. However, to actually achieve the expected yield gain, test patterns must avoid detecting the ax-redundant faults. Otherwise, *defective yet acceptable* AxICs are rejected, resulting in some yield loss.

Some works focused on generating test patterns respecting some properties. For instance, works in [20] and [21] focused on generating test patterns leading to different output errors depending on the class of faults affecting the AxIC. Unfortunately, these propositions are limited to specific error metrics. Afterwards, the work in [24] focused on generating test patterns to detect all the non-redundant faults and as less ax-redundant faults as possible, regardless of the error metric and of the fault classification technique. While the results in [24] are promising, in some cases the achieved yield gain does not correspond to the expected one. Indeed, usually the structure of the AxIC is such that it is impossible to detect all the non-redundant faults without detecting also some ax-redundant ones.

In this work, we provide the following contributions:

- we show conventional testing problems that prevent the achievement of the expected yield gain;

Vector $i$	Input		$*O^{precise}$	Fault-free $\dagger O^{approx}$	Faulty $\dagger O^{approx}$											
	A	B			Sa1@a	Sa1@b	Sa0@c	Sa1@c	Sa1@d	Sa1@e	Sa0@f	Sa1@f	Sa1@g	Sa1@h	Sa0@i	Sa1@i
0	0	0	0	0	0	0	0	4	0	0	0	2	0	0	0	1
1	0	1	0	0	0	0	4	0	0	0	2	0	1	0	1	
2	0	2	0	0	0	4	0	4	0	2	0	2	0	0	1	
3	0	3	0	0	0	4	0	4	0	2	0	2	0	1	1	
4	1	0	0	0	0	0	0	4	2	0	0	2	1	0	1	
5	1	1	1	1	1	1	1	5	3	1	1	2	1	1	0	1
6	1	2	2	2	2	6	2	6	2	2	0	2	3	2	2	3
7	1	3	3	3	3	7	3	7	3	3	1	3	3	3	2	3
8	2	0	0	0	4	0	0	4	0	0	0	2	0	0	0	1
9	2	1	2	0	4	0	0	4	0	0	0	2	0	1	0	1
10	2	2	4	4	4	4	0	4	4	6	4	6	4	4	4	5
11	2	3	6	4	4	4	0	4	4	6	4	6	4	5	4	5
12	3	0	0	0	4	0	0	4	2	0	0	2	1	0	0	1
13	3	1	3	1	5	1	1	5	3	1	1	3	1	1	0	1
14	3	2	6	6	6	6	2	6	6	6	4	6	7	6	6	7
15	3	3	9	7	7	7	3	7	7	7	5	7	7	7	6	7

Fault classification (according to WCE<sup>‡</sup>): non-red. non-red. non-red. non-red. ax-red. ax-red. non-red. ax-red. ax-red. ax-red. non-red. ax-red.

\*Precise output; †Approximate output;

‡Worst Case Error (WCE) = 2

ax-red. = ax-redundant fault ( $|O_{i(faulty)}^{approx} - O_i^{precise}| \leq 2, \forall i$ )

non-red. = non-redundant fault ( $\exists i : |O_{i(faulty)}^{approx} - O_i^{precise}| > 2$ )

$\boxed{Value}$ : vector  $i$  detects the fault  $\implies Value$  is different from fault-free  $O_i^{approx}$   
 $\boxed{Value}$ : approximate circuit output.  $\implies Value$  is different from  $O_i^{precise}$ .

TABLE I: Truth table (in integer format) of the example circuit [25] for different cases: precise (see Fig 1a), fault-free approximate (see Fig 1b), and faulty approximate with different Stuck-at faults.

- we propose a new *test application technique* to actually achieve the expected yield gain. The technique is applicable regardless of the specific error metric and of the specific test pattern generation technique used.

The remainder of the paper is organized as follows. Section II introduces a motivating example and shows the problem. Section III details the proposed approach. Experimental results are discussed in Section IV and Section V draws conclusions.

## II. PROBLEM STATEMENT AND RELATED WORK

As mentioned in the introduction, the proper structure of an AxIC usually makes impossible for a test set to avoid the detection of some ax-redundant faults [24]. Let us refer to the example below to provide further details. We resort to the same example shown in [24], in order to highlight the mentioned problem and show how to solve it.

### A. Motivating example




Fig. 1: 2-bit Multiplier presented in [25]. Precise circuit (a) and Approximate circuit (b)

Let us consider the 2-bit approximate multiplier (ax-multiplier) proposed in [25] and shown in Figure 1b, obtained from the 2-bit precise multiplier in Figure 1a. In the approximate version, the longest path is reduced (from 3 to 1 logic gates, i.e.  $\sim 66\%$ ), as well as the area (from 8 to 3 logic gates, i.e.  $\sim 63\%$ ). On the other hand, some errors are introduced at output. In the left part of Table I, we report the outputs of the precise IC ( $O^{precise}$ ) and of the AxIC ( $O^{approx}$ ), for each input vector  $i \in [0, 15]$ . Values are reported as integer (e.g., 0000 = “0”, 0001 = “1”, etc.). To measure the error, we used the Worst Case Error (WCE) metric. For the convenience of the reader, we report below the WCE equation:

$$WCE = \max_{\forall i \in \mathcal{I}} |O_i^{approx} - O_i^{precise}|, \quad (1)$$

where  $\mathcal{I}$  is the set of all input combinations, and  $O_i^{precise}$  and  $O_i^{approx}$  are respectively the precise and the approximate circuit’s output values corresponding to the  $i$ -th input application. The WCE in the example is 2. This is the threshold value, which must not be altered by the presence of defects.

In order to illustrate the problem, we report in the right part of Table I the impact of each stuck-at fault on the AxIC output. The fault list was generated with a commercial tool [26] with the fault collapsing option active. We use the notation  $SaX@N$  to indicate a “stuck-at-X affecting the net N”, where X can be either the value 1 or 0 and N is the label of the net. Please, refer to Figure 1b for the net labels.

Firstly, based on the difference between the obtained faulty outputs (faulty  $O_i^{approx}$ ) and the precise output ( $O_i^{precise}$ ), faults are classified. If any absolute difference is greater than the threshold ( $\exists i : |O_{i(faulty)}^{approx} - O_i^{precise}| > 2$ , in the example), the fault is non-redundant. Otherwise, if for all the input vectors the absolute difference is lower than or equal to the threshold ( $|O_{i(faulty)}^{approx} - O_i^{precise}| \leq 2 \forall i$ , in the example),

the fault is ax-redundant. We report the class of each fault in the last row of the table.

Secondly, in the table we report in red solid-bordered boxes the faulty  $O_i^{approx}$  values that differ from the fault-free  $O_i^{approx}$  ones. Thanks to this output difference, in test application phase we can detect whether a fault affects the AxIC or not. In conventional IC test, each difference between actual and expected outputs leads to reject the circuit. When it comes to AxICs, we have to reconsider this mechanism. Indeed, a test vector intended to detect a non-redundant fault can also detect an ax-redundant one, ultimately rejecting a still-acceptable circuit. For example, in Table I, we can remark that the vector 8 detects the Sa1@a and Sa1@c non-redundant faults, but also the Sa1@f and Sa1@i ax-redundant ones.

In [24], authors proposed a technique to generate test patterns which detect all the non-redundant faults but also minimize the number of detected ax-redundant faults. Unfortunately, it is often impossible to avoid the detection of some ax-redundant faults. For instance, we can easily note that, among all the possible test sets, one of the best is the tuple  $\{7, 8, 15\}$ . The three vectors detect 100% of the non-redundant faults (i.e., six faults). Nevertheless, the same patterns detect also 33% of ax-redundant faults (two out of six). Specifically, as mentioned above, the Sa1@f and Sa1@i ax-redundant faults are detected by vector 8. Therefore, while the *expected yield gain* is of six faults out of twelve (i.e., the six ax-redundant faults), by using the classic test application, we still detect two ax-redundant faults. In other words, from 50% expected yield gain (six ax-redundant faults avoided, out of twelve total faults) we drop to 33% (four ax-redundant faults avoided, out of twelve total faults). The phenomenon due to which a good product is considered as faulty by the test process is commonly referred to as *over-testing*.

### B. Over-testing issue

To avoid the over-testing, we need to reconsider the test application phase. In details, after the application of the test patterns to the AxIC under test, we need to verify that the actual output meets some specific conditions and not only whether it differs from the expected output. To show the idea, let us resort again to Table I. As already mentioned, vector 8 detects four faults, two non-redundant and two ax-redundant. The *faulty*  $O_i^{approx}$  differs from the expected *fault-free*  $O_i^{approx}$  output in different ways for different faults. The faulty  $O_i^{approx}$  output is 4 when non-redundant faults occur, 2 or 1 when ax-redundant faults occur. Therefore, by observing the actual output value it is very simple to conclude (i) if a fault occurred and (ii) whether it was non-redundant or ax-redundant.

A first attempt to address over-testing problem was made in [27]. Authors introduced the *threshold testing* principle and applied to conventional ICs in order to increase the production yield. Briefly, the technique identifies *catastrophic faults* according to the WCE metric (see Equation 1) by generating input vectors causing output errors higher than the threshold. A test vector detecting a catastrophic fault could still detect an acceptable one. Therefore, authors compare

test responses with the ones from the precise circuit. If the difference is lower than the threshold, the circuit is considered still acceptable, otherwise it is rejected.

Threshold testing can be considered as a special case of AxIC testing [16]. In fact, threshold testing can be applied to AxICs only if some conditions are met:

- 1) precise circuit test responses are available;
- 2) the considered metric is the WCE;
- 3) test patterns are systematically generated to produce an error greater than the threshold when a non-redundant fault occurs.

Unfortunately, not always the three mentioned conditions are met. For example, as reported in Table I, we can notice that vector 7 detects two non-redundant faults respecting the third condition (i.e., Sa1@b and Sa1@c), but also other two without respecting the condition (i.e., Sa0@f and Sa0@i). Only vector 15 can detect Sa0@f and Sa0@i while respecting the third condition. Therefore, constraints on conventional ATPG need to be applied to produce test patterns respecting the third condition. Some example have been presented in the literature [20], [21].

In this work we present a test application technique for AxICs which does not need the three mentioned conditions. Therefore, we propose a technique to efficiently apply tests to AxICs (i) without knowing the precise circuit test responses, (ii) without requiring a specific error metric, and (iii) without necessitating special test patterns. In the next section we present the technique.

### III. A NEW APPROXIMATION-AWARE TEST APPLICATION TECHNIQUE

The limitations of threshold testing technique [27] motivate us to propose a new *approximation-aware test application* technique to mitigate the over-testing effect. We drew our inspiration from a concept introduced in late seventies, the *signature analysis* [28], which is mostly used in self-testing hardware techniques. In particular, *Built-In Self-Test (BIST)* approach compacts test responses together into a signature, which is used to verify whether the Unit Under Test (UUT) is faulty or not. In detail, when the test mode is activated, test patterns are applied to UUT and a signature is generated. Then, the latter is compared with the golden signature, which was generated by the fault-free circuit and stored within the BIST architecture. If the two signatures are identical, the circuit is considered fault-free. Otherwise, a malfunction is detected. Different compaction methods can be used to produce the signature. An extensive review of those methods can be found in [29].

Basically, we propose to generate multiple signatures, one for each ax-redundant fault, and compare them with the test response signature. If there is at least one match, then the AxIC is considered acceptable. Otherwise, the circuit is rejected. The underlying assumption is modeling defects with the single fault model. This is a very widely adopted assumption in practice [29].

The proposed technique is based only on the analysis of the AxIC's test responses. Therefore – unlike the threshold




Fig. 2: Proposed test application technique

testing technique [27] – the proposed technique can be applied without knowing the precise circuit test responses, regardless of the error metric used to classify faults, and regardless of the technique used to generate test patterns.

The proposed technique is intended to be used for external test (i.e., test are applied by using an Automatic Test Equipment (ATE)). Of course, it can be also used in a BIST context.

#### A. Proposed technique

To apply the proposed technique we assume, as preconditions, to have (i) ax-redundant and non-redundant fault lists (obtained by the *fault classification* with any metrics [19]–[22]) and (ii) test patterns (generated with any technique).

As depicted in Figure 2, the proposed test application technique is composed of two phases:

**At design time** we simulate test patterns with the AxIC netlist and compact the responses together to form a *golden signature* (1.1). Then, we perform the same procedure while injecting, one by one, all the ax-redundant (axR) faults into the AxIC netlist. This results in *ax-redundant signatures* (1.2). Hence, we apply the same process to non-redundant (nR) faults, in order to obtain *non-redundant signatures* (1.3). Finally, we perform the union between golden and ax-redundant signatures, hence we remove signatures in common with non-redundant ones (if any) (1.4). The output of this phase is what we call *ax-aware signature set*.

**At test time** after applying test patterns to the manufactured AxIC, we compact test responses and compare the *actual signature* with all the signatures in the ax-aware signature set. If at least one of the comparisons matches, than the test passes, otherwise the circuit is rejected.

As mentioned at the beginning of the section, different response compaction methods can be used. Moreover, the proposed technique can be used for both external testing and self-testing. Concerning external testing, the *Automatic Test Equipment (ATE)* software can be modified to implement any compaction (e.g., hashing algorithm such as MD5, SHA, etc.). On the other hand, concerning self-testing hardware approaches as the BIST, other techniques exist, such as *one-count*, *transition count*, *Linear Feedback Shift Register (LFSR)*, etc [29].




Fig. 3: Example to show aliasing effect.

Vector $i$	a	b	c	$O_1$	$O_0$	int	Sa1@a	Sa1@b	...
0	0	0	0	0	0	0	2	2	...
1	0	0	1	0	0	0	3	2	...
2	0	1	0	1	0	2	2	2	...
3	0	1	1	1	0	2	3	2	...
4	1	0	0	1	0	2	2	2	...
5	1	0	1	1	1	3	3	3	...
6	1	1	0	1	0	2	2	2	...
7	1	1	1	1	1	3	3	3	...

Fault classification: non-red. ax-red ...

**Value**: vector  $i$  detects the fault. *Value* is different from ‘int’

TABLE II: Truth table of the circuit in Figure 3

#### B. Signature aliasing problem

In conventional test, the overlapping phenomenon of two signatures is referred to as *aliasing*. In details, as reported in [29], during the test response compaction, a signature of a faulty circuit can match the fault-free circuit one. This is due to the loss of information caused by the compaction itself.

In this work we extend the meaning of the aliasing in the context of AxIC testing. Let us resort to a tiny example to show the issue. In Figure 3, we depict a hypothetical circuit where some logic produces three signals (a,b,c) which drive the circuit outputs ( $O_1O_0$ ) through two logic gates. Figure II reports the truth table of the two output signals as function of a, b, and c. The column ‘int’ reports the integer representation of the fault-free circuit output. Let us assume that the faults Sa1@a and Sa1@b are classified as non-redundant and ax-redundant respectively. To test these two faults we can use different vectors (e.g., *vector 0* or *vector 1*). If the test pattern generator selects the *vector 0* to test the two faults, the signature will be identical for both Sa1@a and Sa1@b. This will lead our technique to reject the circuit even when Sa1@b (ax-redundant) occurs. Therefore, we extend the definition of *aliasing* as follows:

**Aliasing** During the test response compaction, a non-redundant signature can match an ax-redundant one.

Circuit	*Relative Yield Gain (%) with conventional test	*Relative Yield Gain (%) with proposed technique		Execution Time (s) <sup>†</sup>	*Relative Yield Gain (%) with the technique in [24]
		Single detection	Double detection		
add8_051	0.00%	100.00%	-	0.648	100.00%
add8_036	20.00%	100.00%	-	0.636	93.65%
add8_012	0.00%	100.00%	-	0.532	97.98%
add8_045	0.00%	100.00%	-	0.496	100.00%
GeAr_N8_R2_P2	26.67%	100.00%	-	0.636	74.03%
ACA_I_N8_Q5	31.86%	99.46%	100.00%	0.764	73.61%
GDA_St_N8_M8_P3	18.70%	100.00%	-	0.704	66.34%
GeAr_N16_R6_P4	12.75%	100.00%	-	0.724	56.60%
ACA_II_N16_Q8	30.83%	100.00%	-	0.772	64.36%
ETAII_N16_Q8	28.92%	98.58%	100.00%	0.924	65.82%
GDA_St_N16_M4_P4	10.10%	100.00%	-	1.324	51.64%
GDA_St_N16_M4_P8	23.36%	100.00%	-	1.276	51.73%
GeAr_N16_R4_P4	30.83%	100.00%	-	0.78	64.36%
GeAr_N16_R4_P8	30.00%	99.40%	100.00%	0.852	56.78%
GeAr_N16_R2_P4	34.48%	99.62%	100.00%	1.016	57.82%
ACA_II_N16_Q4	24.15%	100.00%	-	1.152	48.31%
ETAII_N16_Q4	24.15%	100.00%	-	1.22	48.31%
ACA_I_N16_Q4	24.88%	100.00%	-	1.476	46.70%
<b>Average</b>	<b>20.65%</b>	<b>99.84%</b>	<b>100.00%</b>	<b>0.89 s</b>	<b>67.67%</b>

\*Higher is better

<sup>†</sup>Execution time for the proposed technique to generate the ax-aware signature set

TABLE III: Relative Yield Gain (RYG) obtained with the proposed technique, compared to conventional test

A simple solution to reduce the aliasing probability is to generate test patterns to detect faults multiple times. Nevertheless, this also increments the final test length, thus the cost. Another solution is to impose some constraints to the test pattern generator to systematically select patterns to avoid aliasing. In the example shown, selecting *vector 1* instead of *vector 0* would solve the problem. Indeed, the faulty output when applying *vector 1* is different for the two faults, thus the signatures will differ, as well.

#### IV. EXPERIMENTAL RESULTS

To evaluate the technique effectiveness, we applied it to a set of AxICs taken from the literature. Specifically, we used Accuracy-Configurable Approximate (ACA) adders from [10], Gracefully-Degrading Adders (GDA) from [14], Generic Accuracy configurable (GeAr) adders from [13], Error Tolerant Adders (ETAII) from [30], and some EvoApprox8b library AxICs [31] (add8\_051, add8\_036, add8\_012, add8\_045). Without loss of generality, we used the technique in [21] to perform the fault classification, by resorting to the WCE (Equation 1) as error metric. In this way, for each AxIC, we obtained ax-redundant and non-redundant fault lists. Then, we generated test patterns we resorted to the test flow used in state-of-the-art works [19]–[22]: we used a commercial ATPG tool [26], instrumented with the classic options (static and dynamic compaction), and we targeted only the non-redundant fault list. We refer to this as *conventional test*.

Then, we applied the proposed technique and evaluated the gains, compared to conventional test. In details, we simulated the obtained test patterns with the AxIC netlist while injecting the different faults and compacted the responses to obtain the ax-aware signature set, as shown in Figure 2a. To compact test responses into signatures, we used a software approach. Specifically, once collected test responses into regular computer files, we used the *md5sum* computer program to calculate MD5 hashes out of them. This constituted the *ax-aware signature set*. In the actual test phase, after the AxIC manufacture,

the ax-aware signature set has to be employed, as shown in Figure 2b.

To measure the technique effectiveness, we introduce a metric, namely Relative Yield Gain (RYG), expressed as follows:

$$RYG = 1 - \frac{\text{detected ax-redundant faults}}{\text{total ax-redundant faults}} \quad (2)$$

The RYG measures the part of expected yield gain that is actually achieved as a result of the approximation-aware test process. RYG values range from 0 to 1.  $RYG = 0$  means that all the ax-redundant faults are detected by test procedure; thus all the faulty, yet acceptable, AxICs are rejected.  $RYG = 1$  means that the detection of all ax-redundant faults is avoided, thus the yield gain is as high as expected. As mentioned in Section II, the expected yield gain is determined in the fault classification phase as the percentage of faults classified as ax-redundant. To count the number of ax-redundant faults still detected, for conventional test we performed a fault simulation and for the proposed technique we enumerated the ax-redundant signatures overlapping the non-redundant ones.

#### A. Discussion of the results

In Table III, we show experimental results. In the first column we report the name of the analyzed circuits. In the second column we report the percentage of ax-redundant faults detected with the conventional test (i.e., without our technique). Then, third column reports results obtained with the proposed technique. As it can be seen, the relative yield gain was drastically improved. On average, we achieved 99.84% RYG. For fourteen circuits out of eighteen (~ 77%) the obtained relative yield gain was 100%. For the remaining four circuits, the RYG was always greater than 98%. Such RYG reduction was due to the phenomenon described in Subsection III-B, i.e. aliasing. To mitigate the aliasing effect, we generated test patterns to detect faults twice. In details, we instrumented the ATPG with the option *-ndetects 2*. As reported in the fourth

column of Table III, the aliasing phenomenon was correctly overcome for all the four circuits. The cost of detecting the faults twice was to double the number of test patterns. Clearly, ad hoc methods can be implemented to overcome aliasing. As an example, some techniques mentioned in the first part of the paper generate test patterns that intrinsically avoid the aliasing phenomenon [20], [21], [23], [27]. Indeed, those techniques generate test patterns that always produce error values greater than the threshold when detecting non-redundant faults. On the contrary, error values lower than the threshold are produced when detecting ax-redundant faults. Therefore, non-redundant signatures cannot overlap ax-redundant ones. As already discussed, the mentioned techniques are limited to specific error metrics and depend on the availability of the precise version of the circuit. Finally, table's fifth column shows that the runtime to generate ax-aware signatures was always smaller than 1.5 seconds (0.89 seconds, on average).

Furthermore, to extend the comparison, in the table's last column we report results obtained by the study in [24]. Although relevant results were achieved in [24], the proposed technique shows significant improvements (from 67% to 99% RYG, on average).

## V. CONCLUSION

Approximate Computing (AxC) applied to integrated circuits allowed for a wide range of new design strategies for the scientific community. AxC introduced also the opportunity to achieve gains in production yield. By suitably adapting test procedures, defective circuits – yet still able to provide satisfactory results – can be accepted, thus increasing the yield. In this work, we analyzed the issues preventing approximation-aware test techniques from reaching a satisfactory yield gain. Moreover, we proposed an error-metric-independent *signature-analysis-based* technique to efficiently overcome the discussed problems. Experimental results on state-of-the-art approximate circuits showed very good yield gain results. Indeed, we achieved actual yield gains between 98% and 100% of the expected ones.

## REFERENCES

- [1] Q. Xu, T. Mytkowicz, and N. S. Kim, "Approximate computing: A survey," *IEEE Design & Test*, vol. 33, no. 1, pp. 8–22, Feb 2016.
- [2] S. Mittal, "A survey of techniques for approximate computing," *ACM Comput. Surv.*, vol. 48, no. 4, pp. 62:1–62:33, Mar. 2016. [Online]. Available: <http://doi.acm.org/10.1145/2893356>
- [3] J. Han and M. Orshansky, "Approximate computing: An emerging paradigm for energy-efficient design," in *2013 18th IEEE European Test Symposium (ETS)*, May 2013, pp. 1–6.
- [4] V. K. Chippa, S. T. Chakradhar, K. Roy, and A. Raghunathan, "Analysis and characterization of inherent application resilience for approximate computing," in *Proceedings of the 50th Annual Design Automation Conference*. ACM, 2013, p. 113.
- [5] P. Kulkarni, P. Gupta, and M. Ercegovac, "Trading accuracy for power with an underdesigned multiplier architecture," in *2011 24th International Conference on VLSI Design*, Jan 2011, pp. 346–351.
- [6] S. Rehman, B. S. Prabakaran, W. El-Harouni, M. Shafique, and J. Henkel, *Heterogeneous Approximate Multipliers: Architectures and Design Methodologies*. Springer International Publishing, 2019, pp. 45–66.
- [7] H. Jiang, J. Han, and F. Lombardi, "A comparative review and evaluation of approximate adders," in *Proceedings of the 25th Edition on Great Lakes Symposium on VLSI*, ser. GLSVLSI '15. New York, NY, USA: ACM, 2015, pp. 343–348. [Online]. Available: <http://doi.acm.org/10.1145/2742060.2743760>
- [8] V. Gupta, D. Mohapatra, S. P. Park, A. Raghunathan, and K. Roy, "Impact: Imprecise adders for low-power approximate computing," in *IEEE/ACM International Symposium on Low Power Electronics and Design*, Aug 2011, pp. 409–414.
- [9] J. Miao, K. He, A. Gerstlauer, and M. Orshansky, "Modeling and synthesis of quality-energy optimal approximate adders," in *2012 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, Nov 2012, pp. 728–735.
- [10] A. B. Kahng and S. Kang, "Accuracy-configurable adder for approximate arithmetic designs," in *DAC Design Automation Conference 2012*, June 2012, pp. 820–825.
- [11] V. K. Chippa, S. Venkataramani, S. T. Chakradhar, K. Roy, and A. Raghunathan, "Approximate computing: An integrated hardware approach," in *Asilomar Conference on Signals, Systems and Computers*. IEEE, 2013, pp. 111–117.
- [12] A. Momeni, J. Han, P. Montuschi, and F. Lombardi, "Design and analysis of approximate compressors for multiplication," *IEEE Transactions on Computers*, vol. 64, no. 4, pp. 984–994, April 2015.
- [13] M. Shafique, W. Ahmad, R. Hafiz, and J. Henkel, "A low latency generic accuracy configurable adder," in *2015 52nd ACM/EDAC/IEEE Design Automation Conference (DAC)*, June 2015, pp. 1–6.
- [14] R. Ye, T. Wang, F. Yuan, R. Kumar, and Q. Xu, "On reconfiguration-oriented approximate adder design and its application," in *2013 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, Nov 2013, pp. 48–54.
- [15] J. Liang, J. Han, and F. Lombardi, "New metrics for the reliability of approximate and probabilistic adders," *IEEE Transactions on Computers*, vol. 62, no. 9, pp. 1760–1771, Sept 2013.
- [16] I. Polian, "Test and reliability challenges for approximate circuitry," *IEEE Embedded Systems Letters*, vol. 10, no. 1, pp. 26–29, March 2018.
- [17] L. Anghel, M. Benabdenbi, A. Bosio, M. Traiola, and E. I. Vatajelu, "Test and reliability in approximate computing," *Journal of Electronic Testing*, vol. 34, no. 4, pp. 375–387, Aug 2018. [Online]. Available: <https://doi.org/10.1007/s10836-018-5734-9>
- [18] A. Chandrasekharan, D. Große, and R. Drechsler, *Design Automation Techniques for Approximation Circuits: Verification, Synthesis and Test*. Springer, 2019.
- [19] A. Chandrasekharan, S. Eggersglüß, D. Große, and R. Drechsler, "Approximation-aware testing for approximate circuits," in *2018 23rd Asia and South Pacific Design Automation Conference (ASP-DAC)*, Jan 2018, pp. 239–244.
- [20] A. Gebregiorgis and M. B. Tahoori, "Test pattern generation for approximate circuits based on boolean satisfiability," in *Design, Automation Test in Europe Conference Exhibition (DATE)*, March 2019.
- [21] M. Traiola, A. Virazel, P. Girard, M. Barbareschi, and A. Bosio, "Testing approximate digital circuits: Challenges and opportunities," in *2018 IEEE 19th Latin American Test Symposium (LATS)*, March 2018, pp. 1–6.
- [22] —, "Investigation of mean-error metrics for testing approximate integrated circuits," in *2018 IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems (DFT)*, Oct 2018, pp. 1–6.
- [23] —, "On the comparison of different atpg approaches for approximate integrated circuits," in *2018 IEEE 21st International Symposium on Design and Diagnostics of Electronic Circuits Systems (DDECS)*, April 2018, pp. 85–90.
- [24] —, "A test pattern generation technique for approximate circuits based on an ilp-formulated pattern selection procedure," *IEEE Transactions on Nanotechnology*, pp. 1–1, 2019.
- [25] B. Garg and G. K. Sharma, "Acm: An energy-efficient accuracy configurable multiplier for error-resilient applications," *J. Electron. Test.*, vol. 33, no. 4, pp. 479–489, Aug. 2017.
- [26] Tetramax. [Online]. Available: <https://www.synopsys.com/>
- [27] Z. Jiang and S. K. Gupta, "An atpg for threshold testing: obtaining acceptable yield in future processes," in *Proceedings. International Test Conference*, 2002, pp. 824–833.
- [28] R. A. Frohwerk, "Signature analysis: a new digital field service method," 1977.
- [29] M. L. Bushnell and V. D. Agarwal, *Essentials of Electronic Testing for Digital, Memory, and Mixed-Signal VLSI Circuits*, 01 2000.
- [30] N. Zhu, W. L. Goh, and K. S. Yeo, "An enhanced low-power high-speed adder for error-tolerant application," in *Proceedings of the 2009 12th International Symposium on Integrated Circuits*, Dec 2009, pp. 69–72.
- [31] V. Mrazek, R. Hrbacek, Z. Vasicek, and L. Sekanina, "Evoapprox8b: Library of approx adders and multipliers for circuit design and benchmarking of approximation methods," in *Design, Automation Test in Europe Conference Exhibition (DATE)*, March 2017, pp. 258–261.