

Design Methodology to Explore Hybrid Approximate Adders for Energy-Efficient Image and Video Processing Accelerators

Leonardo Bandeira Soares^{ID}, *Student Member, IEEE*, Morgana Macedo Azevedo da Rosa, Cláudio Machado Diniz^{ID}, *Member, IEEE*, Eduardo Antonio César da Costa, *Member, IEEE*, and Sergio Bampi^{ID}, *Senior Member, IEEE*

Abstract—This paper proposes a new design methodology to explore the state-of-the-art approximate adders for accelerator architectures conceived in the realm of multiplier-less multiple constant multiplication optimization problem. The proposed methodology is composed of: 1) a search heuristic to seek faster and feasible approximate configurations for the architectures under evaluation; 2) low-power techniques regarding hybrid approximate adders design for accelerators based on trees of shift-and-add operations; 3) high-performance evaluation by exploring parallel prefix adders and low power analysis through the use of the adder optimized by a commercial synthesis tool in the precise part of the approximate adders; and 4) energy efficiency analysis by considering both the approximate techniques and voltage over scaling estimation. Furthermore, improvements are proposed for the state-of-the-art approximate adders under evaluation in this paper. Two case studies are considered to assess the proposed methodology: 1) Gaussian image filter and 2) Sobel operator. The precise and approximate image filters were described in very high-speed integrated circuits hardware description language regarding the proposed methodology. Results are shown after synthesis to a 45-nm standard cell-based technology, where energy reductions ranging from 7.7% up to 73.2% were experienced for multiple levels of quality considering the applications under analysis.

Index Terms—Image and video processing, energy efficiency, approximate computing, adders, multiplier-less multiple constant multiplication.

I. INTRODUCTION

THE semiconductor industry faces challenges at each new Complementary Metal Oxide Semiconductor (CMOS) technology node. Power density increase has been experienced due to the observation that Dennard scaling is not feasible

anymore [1]. Furthermore, power and thermal walls bring much more effort to designers, so that digital CMOS design is facing the so-called “Dark Silicon Era” even considering recent Fin Field-Effect Transistor (FinFET) technologies [2]. Therefore, the current and future computing scenario is characterized by the demand for numerous and ubiquitous compute-intensive applications in constrained power budget digital devices. Based on that, energy-efficient techniques (*i.e.* maximize the number of arithmetic operations per energy unit) are paramount to cope with the previously observed challenges. According to [3], two trending energy-efficient techniques are listed as follows: (i) accelerator-rich architectures based on Application Specific Integrated Circuits (ASICs) and (ii) Approximate Computing (AC).

Architectural heterogeneity and the use of ASIC accelerators are energy-efficient techniques to execute the most compute-intensive kernels of an application [3]. On the other hand, the remaining tasks which demand less energy consumption can be scheduled for general-purpose processors. As a result, general-purpose processors’ workload is alleviated due to the use of energy-efficient specific processing cores. Power-management schemes can be implemented to power off accelerators or general-purpose cores when not in use, thus respecting the power and thermal constraints. The works in [4] and [5] show that despite the challenges of architectural integration introduced by this new design paradigm, these accelerator-rich architectures play an essential role in energy efficiency for recent applications. For example, Hameed *et al.* [6] show that an ASIC solution is 500× more energy-efficient than a four core general-purpose processor when considered H.264 video coding application. One essential approach to conceive ASIC implementation of digital filters and transforms is to implement the multiplications by constants in the light of MMCM problem formulated in [7] and [8]. In other words, these architectures are designed by adopting the use of additions, subtractions, and shift in an optimized configuration.

The approximate computing paradigm emerged to increase performance and to reduce power dissipation [9]. The critical approach in approximate hardware is to reduce the computation accuracy in favor of energy-efficiency. In circuit level

Manuscript received July 24, 2018; revised December 5, 2018; accepted January 4, 2019. This work was supported in part by CNPq, in part by CAPES, and in part by FAPERGS Brazilian. This paper was recommended by Associate Editor R. S. Murphy-Arteaga. (*Corresponding author: Cláudio Machado Diniz.*)

L. B. Soares and S. Bampi are with the Graduate Program in Microelectronics, Federal University of Rio Grande do Sul, 91501-970 Porto Alegre, Brazil (e-mail: lbsoares@inf.ufrgs.br; bampi@inf.ufrgs.br).

M. M. A. da Rosa, C. M. Diniz, and E. A. C. da Costa are with the Graduate Program on Electronic Engineering and Computer Science, Catholic University of Pelotas, 96015-560 Pelotas, Brazil (e-mail: claudio.diniz@ucpel.edu.br; eduardo.costa@ucpel.edu.br).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TCSI.2019.2892588

design, this is performed by designing simpler circuits to speed up the critical path timing and/or to consume less power. Approximate computing techniques take advantage of approximation-tolerant applications which do not need high accuracy all the time but only “good enough” or “sufficiently good” results for output perceptual quality. In [10] is stated the following properties to define an approximation-resilient application: (i) there is not a golden or accurate result, but a range of acceptable ones and (ii) robustness to input noisy data. For example, multimedia applications (e.g., video coding, audio filtering, image processing, and so on), highly demanded by current portable devices, are intrinsically related to human senses. The multimedia signals are, in fact, approximation-tolerant applications, since in [11] is stated that human senses process analog information and have difficulty to realize the negative impact of digital approximations. It means that it is possible to adopt approximate computing techniques to improve energy efficiency in multimedia applications by adequately exploring the user experience at different profiles of quality.

The excellent point of approximate computing is that this paradigm can be adopted at any abstraction level from transistor-level up to software application [12]. Furthermore, approximate computing can be an additive design component for accelerator-rich architectures. One can consider that the use of approximate hardware accelerators brings further energy efficiency improvements [12]. In the arithmetic layer of abstraction, works in [11] and [13]–[20] have proposed approximate adders. Adders are basic building blocks for several compute-intensive multimedia applications. Therefore, approximate adders could drive energy efficiency for recent digital compute-intensive and approximation-tolerant applications.

Based on that, this work proposes a design methodology to explore state-of-the-art approximate adders for ASIC implementation of add-and-shift accelerators for image and video processing. Previous works in [22] and [23] examined the use of the state-of-the-art approximate adders for image filters. To explore approximation for the architectures, they adopted simulation-based methodologies in which search heuristics are implemented to seek for energy-efficient approximate configurations. The approximate adders taken by the previously mentioned related works are the Approximate Mirror Adder (AMA) [13] and the Error-Tolerant Adder I (ETAI) [11]. They are divided into precise and approximate parts. In both the works, only the Ripple Carry Adder (RCA) topology is explored in the precise block of those approximate adders. The same observation is valid for the case study explored in [11]: the precise block of ETA-I is only implemented with RCA topology.

This work presents four novel contributions in the scope of approximate computing:

- 1) A faster search heuristic and simulation-based methodology to configure feasible configurations with evaluation of multiple levels of quality;
- 2) A high-performance exploration of approximate hardware accelerators through the use of PPAs and low power evaluation through the optimized adder from the

commercial synthesis tool in the precise part of the approximate adders;

- 3) Combination of different approximate adders to compose hybrid adders for the add-and-shift architectures; Energy-efficiency

- 4) analysis based on the approximate configurations and VOS estimation due to the insertion of PPAs.

Results show that our approach substantially reduces energy consumption ranging from 7.7% up to 73.2% for different levels of quality.

The remainder of this paper is organized as follows: Section II overviews the approximate and precise adders, as well as the background for the Gaussian and Gradient image filters. Section III presents the related works. Section IV presents the proposed design methodology to explore our hybrid approximate adders to conceive low power accelerators design. In Section V the experimental setup and results are shown. In Section VI the conclusions are drawn.

II. BACKGROUND ON APPROXIMATE ADDERS, PARALLEL PREFIX ADDERS, AND IMAGE FILTERS

A. Approximate Adders

The approximate adders can be classified as computational performance- and power-oriented designs. The former is related to adders divided into m independent blocks or sub-adders to speed up the critical path timing. The claim is that, for random and uniformly distributed pairs of operands, more extended carry propagation rarely occurs. Based on that, additional logic is necessary to speculate carry-in for each sub-adder, since this class of approximate adder breaks the carry propagation in many parts. Examples of adders which improve computational performance are the Error-Tolerant Adder II [15], Error Tolerant Adder IV [14], and the Almost Correct Adder (ACA) [16]. This class of approximate adders is also characterized by the presence of infrequent and high magnitude sum errors. Therefore, the works in [16]–[19] proposed accuracy configurable adders to cope with this error characteristics. On the other hand, more logic is added to detect and correct the sum errors.

A different philosophy is to propose power-oriented adders which generally are divided into two parts: (i) the least significant approximate part and (ii) the most significant accurate part. Examples of power-oriented approximate adders can be observed in [11], [13], and [20]. The principal idea in the approximate part is to replace the full adder cells by simpler adder circuits. Therefore, power reduction is the main objective of this class of adders. Besides, these adders also tend to reduce critical path timing, because in the approximate part there is not carry propagation scheme. One can observe that the classical truncation is a type of power-oriented adder which truncates least significant full adder cells. This class of approximate adders is also characterized by the presence of frequent and low magnitude sum errors. Such errors are of low magnitude because the bit-width of the approximate part can be controlled through an approximate parameter k . In this work, the proposed approach is to explore the power-oriented adders to give priority to power-efficiency. It is also ratified




Fig. 1. Approximate adders: (a) Copy adder; (b) ETAI.

in [21] which states that the adders focused on delay reduction cannot be used to explore power-efficiency in structures which demand the massive use of additions like the multipliers. Also, the power-oriented approximate adders enable the exploration of multiple conventional adder topologies in the precise part, which is not true for the approximate adders divided into many blocks.

Therefore, we consider in this study the exploration of the fifth approximate version of AMA [13] and the Error-Tolerant Adder I [11], because they are also explored by related works [22], [23]. The former approximate adder is renamed to “Copy adder” due to its copy function implemented by the buffers. The approximate adders which are explored in this study can be observed in Figure 1.

The “Copy adder” in Figure 1(a) has its k bits long approximate part implemented by buffers to copy the operand a to the sum. This procedure has 50% probability of getting the correct sum for each bit position. Furthermore, the carry-in estimation for the precise part is implemented by the more straightforward assignment of the input operand bit b_{k-1} . This procedure has 75% probability of getting a correct carry-in estimation to the precise block. The use of half adders implements the approximate part of the ETAI in Figure 1(b). The sum is performed in the non-conventional direction, (*i.e.* from the most significant bit $k-1$ to the least significant bit position 0). The control logic block is conceived as follows: if the first carry-generate c is equal to “1,” then all the remaining least significant sum bits are set to “1.” Otherwise, the sum result is the one computed by the propagate signal. The carry-in to the precise part in ETAI is statically set to “0.” This procedure has 50% probability of getting the correct carry-in to the precise part. One can observe that for both the approximate adders, any conventional adder topology can be implemented in the precise part. Related works in [22] and [23] explore the use of the RCA. In this work, we explore the RCA, the high-performance PPAs, and the low power adder fully optimized by the commercial tool used in this study. That is why in the next subsection a brief PPA overview is developed.




Fig. 2. Parallel prefix adder structure.

B. Precise Adders

As previously mentioned, adders are fundamental building blocks in a great variety of computational applications. Based on that, many adder topologies have been proposed to deal with the tradeoff between power and computational performance. The RCA topology is characterized to present low values of power consumption, area, and computational performance. Depending on the high-performance application requirements and given that computational complexity is increasing in nowadays tasks, the critical approach is to accelerate the adder’s critical path delay (*i.e.* carry propagation) at the expense of higher area and power dissipation. Based on that, the Parallel Prefix Adders were proposed to deal with high-performance demands [24].

The carry propagation structure in the PPA adders is implemented by simple logic cells which tend to keep a regular connection. Based on that, the sum computation can be divided into pre-processing, prefix computation and post-processing parts, as can be seen in Figure 2.

In the pre-processing part, the propagate p_i (*i.e.* $a_i \oplus b_i$) and generate g_i (*i.e.* $a_i \wedge b_i$) signals are computed based on the input operands a_i and b_i . In the prefix computation stage, the carry computation is accelerated by the parallel composition of the black cells which implement the group propagate $P_{i:j}$ and generate $G_{i:j}$ signals as well as the gray cells which compute the carry c_i . Finally, the post-processing step is given by the sum $s_i = p_i \oplus c_{i-1}$.

1) *Use of PPA Adders on the Precise Part of the Approximate Adders:* Different configurations between the black and gray cells can be obtained. According to the PPAs taxonomy presented in [25], the different prefix cells configurations allow tradeoffs among (i) the number of logic levels, (ii) the maximum fanout, and (iii) the maximum number of horizontal wire tracks (*i.e.* wire density). All of these aspects affect the adder delay. Based on that, PPA topologies proposed in [26]–[30] are considered in this study. Their main

TABLE I
TAXONOMY OF n -BIT PPA'S

PPA	logic levels	maximum fanout	maximum wiring tracks
Brent-Kung (B-K) [26]	$2\log_2(n) - 1$	2	1
Sklansky (SK) [30]	$\log_2(n)$	$\frac{n}{2} + 1$	1
Kogge-Stone (K-S) [27]	$\log_2(n)$	2	$\frac{n}{2}$
Han-Carlson (H-S) [28]	$\log_2(n) + 1$	2	$\frac{n}{4}$
Ladner-Fischer (L-F) [29]	$\log_2(n) + 1$	$\frac{n}{4} + 1$	1

characteristics concerning logic levels, maximum fanout, and the maximum number of wiring tracks are presented in Table I [25].

As can be seen in Table I, the Brent-Kung adder has the highest number of levels, while presents low values for both fanout and wire density. The Sklansky and Kogge-Stone have the lowest number of logic levels, but the former gives the highest fanout, and the latter has the worst wire density due to the highest number of prefix cells. The Han-Carlson is the hybrid solution between the Brent-Kung and Kogge-Stone. Therefore, this adder balance the tradeoff between the number of logic levels and wiring tracks. The Ladner-Fischer adder is the hybrid approach between Sklansky and Brent-Kung so that the tradeoff is balanced between the number of logic levels and the fanout.

C. Gaussian and Gradient Image and Video Filters

Image and Video filters are of great importance in nowadays compute-intensive applications. With the emerging Internet-of-Things (IoT) scenario, many applications rely on computer vision algorithms to extract features and to provide services in many fields such like agriculture, safety, transportation, and so forth. The challenging point is that image, and video sensors produce a large quantity of data to be processed, and, consequently, increase the computational effort. Furthermore, if considered the ubiquitous environment, then energy efficiency gains much more priority.

The Gaussian filter is a smoothing filter to remove noise. The two-dimension (i.e., x and y directions) Gaussian kernel is obtained as shown in (1).

$$G(x, y) = \frac{1}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{2\sigma^2}} \quad (1)$$

In (1), σ^2 denotes the variance, and this is one of the parameters to obtain different versions of Gaussian kernels. The other parameter is the window size which determines the number of image pixels to be convolved. Stronger smoothing is obtained through the use of larger window sizes. On the other hand, for larger window sizes, the computational cost substantially increases.

The Sobel operator is a prominent example of a gradient filter, whose pair of vertical and horizontal 3×3 convolution masks is presented in (2) and (3), respectively. One of the masks estimates the gradient in the y -direction (rows), while the other estimates the gradient in the x -direction (columns). The square root of the sum of the squares of horizontal and

vertical derivatives calculates the magnitude of the gradient.

$$1/4 \begin{bmatrix} 1 & 2 & 1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{bmatrix} \quad (2)$$

$$1/4 \begin{bmatrix} 1 & 0 & -1 \\ 2 & 0 & -2 \\ 1 & 0 & -1 \end{bmatrix} \quad (3)$$

Both the Gaussian and Sobel operator can be combined to perform edge detection application, which enables feature extraction for computer vision algorithms [31]. Convolution operators tend to demand high computational effort. For instance, given a 512×512 pixels grayscale image, both the 5×5 Gaussian and 3×3 Gradient filters are responsible for more than 11.2 million multiplications and 10.4 million additions. These numbers of arithmetic operations are much higher when considering recent video and image resolutions. Therefore, these filters can be regarded as compute-intensive applications, and hardware accelerator implementation should be investigated.

An energy-efficient approach to conceive hardware accelerators for digital image and video filters is by solving the MMCM problem [7], [8]. Hence, the hardware filters can be implemented by a tree of adder and shifts in a parallel topology increasing computational performance as well as sharing arithmetic operators to reduce area and power. In the next section, a brief review of previously proposed design methodologies, adder topologies exploration in approximate computing scope, and the hybrid approximate adder are presented.

III. RELATED WORK ON DESIGN METHODOLOGIES FOR APPROXIMATE IMAGE FILTERS AND HYBRID APPROXIMATE ADDERS

Given that an effective approach to conceiving image filters is the add-and-shift implementation, the work in [22] explores the use of Copy adder and ETAI inside add-and-shift accelerators for Gaussian and Gradient image filtering process. Therefore, their approach is to adopt a search heuristic, based on the expected output magnitude, to combine different approximate parameters inside the architectures. Each configuration of approximate adders with different approximation parameters is simulated to evaluate the quality response of these filters. After that, quality constraints are imposed to select a set of configurations which are synthesized and analyzed concerning energy efficiency. The work in [22] presents interesting power results for iso-performance analysis. The power reduction for the best case is about 50% considering 45 nm technology. On the other hand, limited quality and real-time evaluation are presented. Their work did not explore maximum frequency analysis, and the precise part of the adders are implemented with only RCA.

In [23] the ETAI is explored, but the authors propose the use of OR gate instead of XOR ones for the approximate part. They adopt a search heuristic based on the number of adder steps in critical path considering the trees of add-and-shift for Finite Impulse Response (FIR) image filters. Their design methodology is also evaluated for different FIR

filters and “Lena” benchmark. Although the authors use a simulation-based analysis, the application quality analysis is not comprehensively explored. Maximum energy reduction of 50.7% is observed for 65 nm ASIC implementation. No observation is performed about different adder topologies exploration in the precise part of the approximate adders. Furthermore, there is not information about real-time analysis of the approximate filters.

From the related works, one can conclude that exploring different conventional adder topologies in precise part is an open research subject. The work in [32] examined different PPA topologies for the computational performance-oriented designs. It occurs because the ACA structure presented in [16] is based on the Kogge-Stone adder. Therefore, in [32] other PPAs are evaluated to carry speculation adders.

In a different point of view, our approach gives priority to control power efficiency and error magnitude by adopting power-based approximate adders and further foresees that performance can be rescued by exploring faster topologies in the precise part of the adders under evaluation.

The work in [33] explores the combination of performance- and power-oriented adders to conceive hybrid approaches. They evaluate these hybrid adders in a 40 nm ASIC implementation of a Single Instruction Multiple Data (SIMD) Central Processing Unit (CPU). Their approach is validated for a Sobel operator where limited real-time analysis is performed. According to the authors, the minimum time of 190 μ s is taken per 512×512 gray scale images, with a maximum energy reduction of 15%. Furthermore, the authors propose a new metric in the arithmetic layer which considers both the probability and magnitude of errors which are directly related to the two classes of approximate adders. On the other hand, quality metrics in edge detection level of abstraction is not observed in their work. The interesting hybrid solution proposed by Najafi *et al.* [33] presents essential contributions to the scope of approximate computing. However, when adopting performance-oriented approximate adders, the adder topology turns out to be static.

Based on that, our proposed approach is not directly in the same scope as the previous work in [33]. Our solution focuses on exploring different precise adder topologies to rescue performance and further evaluate VOS estimation while maintaining the magnitude sum errors at a low and controlled level, which is desired in many image applications. Furthermore, this work proposes the combination of truncation, Copy adder and ETAI (*i.e.* power-oriented adders) to reduce power dissipation in parallel shift-and-add ASIC hardware accelerators. Our first explorations by using PPAs in approximate adders regarding Sobel operation were presented in [34] and [35]. On the other hand, these works did not present a complete design methodology either do not explore a hybrid solution for approximate adders considering power efficient solution for the image filter architectures under evaluation.

IV. PROPOSED DESIGN METHODOLOGY

This section presents the proposed design methodology to explore our hybrid approximate adders. The basic design flow is shown in Figure 3. The first step is to simulate the




Fig. 3. Proposed design methodology.

application under evaluation by adopting MatLab or C models considering both the architectures and approximate adders. The hybrid approximate adder models are implemented by an overloaded function in MatLab and C language. Therefore, add-and-shift filter structures are developed in simulation scripts. The function is called to perform the necessary additions, with the appropriate parameters to select between the type of hybrid approximate adder, the number of bits of the approximate part, and so forth. The quality metrics are generated by considering the use of real test-cases. The next stage is related to the application quality evaluation for all the exercised approximate configurations. After that, different levels of quality are selected to generate the Register Transfer Level (RTL) designs automatically. The logic synthesis is performed for each RTL file, by using the standard cell technology library files (*i.e.*, .lib, .lef, cap tables, and so forth). Then, the mapped gate-level netlist is created to enable the next step of post-synthesis simulation. During the simulation, standard cell technology library files (*i.e.*, the Verilog file with the behavioral model of the standard cells) and real test cases are used to capture switching activity which is saved in Value Change Dump (VCD) or Toggle Count Format (TCF) files. The Verilog gate-level netlist, standard cell technology library files, and VCD or TCF files are then used to estimate power.

In the next subsections the proposed hybrid approximate adders, the accelerator architectures under analysis, and the proposed heuristic adopted to perform simulation are shown.

A. Proposed Hybrid Approximate Adders for Shift-and-Add Architectures

The 5×5 Gaussian and the 3×3 Sobel image filter architectures under evaluation are similar to the ones adopted




Fig. 4. Gaussian image filter architecture.




Fig. 5. Gradient image filter architecture.

by Oliveira *et al.* [22]. The difference is observed in the Gaussian architecture, where the partial terms of the adder tree were reorganized to enable left shift overlapping regions in the operands. It is performed to leverage the power efficiency provided by the proposed hybrid adders and to improve the proposed search heuristic which will be presented in the next subsection.

The Gaussian and Gradient architectures can be observed in the Figure 4 and Figure 5, respectively [22]. One can observe that these architectures are implemented by the shift operations, adders, and subtractors. There are two observable configurations in which the proposed hybrid approximate adders can be adopted: (i) both the operands present overlapped number of left shift operations, (ii) one operand present excessive number of left shift operations than the other operand. These aspects can be observed in Figure 6.

As can be seen in Figure 6, if there is overlapping between the number of left shifts in both the operands, then the proposed approach considers the use of truncation adder in the overlapped and least significant region. The excessive amount of left shift operations in one of the operands, when compared

to the other can be leveraged with the use of the Copy adder. One can notice in the examples in Figure 6(a) and (b) that these procedures do not produce sum errors. The remaining most significant region in the adder can further be explored with the state-of-the-art Copy adder and ETAI (*i.e.*, precise plus approximate adder).

The example in Figure 6(a) shows the proposed hybrid scheme of copy-copy-truncation. The operands are left shifted two and four times, respectively. One can see that there are three approximation parameters: (i) $k_1 = 2$ which controls the truncation in the overlapped shift region, (ii) $k_2 = 2$ which controls the copy adder in the excessive shift region, and (iii) $k_3 = 4$ which controls the approximate part of the copy adder in the non-shifted region. Almost the same observations can be made in Figure 6(b) for the configuration ETAI-copy-truncation. The only difference is that in the non-shifted region the ETAI is adopted instead of the Copy adder. For the ETAI, we adopted the modification proposed by Kang *et al.* [23]: the use of OR logic gates instead of XOR. It occurs because there is not difference regarding produced sum result, while the former gate has less area than the latter. Besides, this work also considers the use of carry-in estimation performed in Copy adder for the ETAI. As previously mentioned, this procedure has more probability of getting correct estimation than statically set the carry-in to "0." If a given adder of the architecture has not left shifted operands, then the approximation is not hybrid (*i.e.*, $k_1 = 0$ and $k_2 = 0$). On the other hand, the copy adder or ETAI can be explored in this adder (*i.e.*, $k_3 \geq 0$).

B. Proposed Search Heuristic

The exhaustive search in simulation-based methodologies tends to be time consuming or prohibitive to find the most energy-efficient configuration. Therefore, the use of search heuristics is essential in this scenario. As previously shown, related works in [22] and [23] proposed search heuristics to seek for energy-efficient accelerators. In this work, the proposed approach is to first establish the k_1 and k_2 parameters




Fig. 6. Proposed hybrid approximate adders. (a) copy-copy-truncation adder. (b) ETAI-copy-truncation adder.

for the overlapped and excessive shift regions. After that, the next step is to seek different k_3 parameters for all the adders in the tree through an iterative process. The algorithm for the heuristic is shown in Algorithm 1. One can observe that the k_1 and k_2 parameters are fixed, while the k_3 is iterated to seek for different approximate configurations. During the initialization step from lines 1 to 9, the k_1 and k_2 parameters are determined solely by the overlapped regions previously shown in Figure 6. The iteration in lines 10 to 12 is used to create and initialize the data structure which stores the k_3 values for each adder in the filter architectures. The search for k_3 parameters is performed in lines 14 to 17. During the search, the quality metric is stored after running the application for each k_3 parameter value. The application quality is evaluated, and objective metric is calculated by considering the use of real test cases.

Based on the heuristic, eight greyscale images (*i.e.*, 7 MatLab built-in images plus the “Lena” standard test image) were used to evaluate the Gaussian and Gradient quality. The quality analysis for different k_3 configurations is shown in Figure 7. One can observe that quality metrics decrease when k_3 increases. The bars represent the average μ , while the error bars represent one standard deviation 1σ . At the top of each error bar, the coefficient of variation (*i.e.*, σ/μ) is denoted in percentage. The gray and white bars represent the copy-copy-truncation and ETAI-copy-truncation hybrid exploration, respectively. The objective metrics are

Algorithm 1 The Simulation-Based Search Heuristic

Input: The array structure T containing the pairs a and b of left shifted bits per adder in the tree (right shifts and no shifts are treated as zero)

Input: The maximum number of configurations N to be tested

Input: The data set Y of real test cases

Output: The array structure S containing the values for k_1 and k_2 per adder node in the tree

Output: The 2-D array Q containing the values for k_3 per adder node in the tree and N tested configurations

Output: The array P with length N containing the application metric per tested configurations

```

1 initialization of the hybrid approximation;
2 for  $i \leftarrow 0$  to  $\text{length}(T)-1$  do
3   if  $T[i].a \geq T[i].b$  then
4      $S[i].k_1 \leftarrow T[i].b$ ;
5      $S[i].k_2 \leftarrow T[i].a - T[i].b$ ;
6   else
7      $S[i].k_1 \leftarrow T[i].a$ ;
8      $S[i].k_2 \leftarrow T[i].b - T[i].a$ ;
9   end
10  for  $j \leftarrow 0$  to  $N-1$  do
11     $Q[j][i] \leftarrow j$ ;
12  end
13 end
14 initialization of the iterative search;
15 for  $j \leftarrow 0$  to  $N-1$  do
16    $P[j] \leftarrow \text{evaluateApplication}(S, Q[j], Y)$ ;
17 end

```

analyzed by considering the precise filter as being the reference.

In the Figure 7(a) is shown the Peak Signal-to-Noise Ratio (PSNR) levels per iterated k_3 parameter. One can observe that high PSNR level is experienced when the $k_3 = 0$. Therefore, in this case, the approximation is only provided by the hybrid approximate adders as can be seen in Figure 4 and Figure 5. Based on that, one can conclude that approximation provided by the k_1 and k_2 parameters insert shallow level of errors. Since the hybrid copy-truncation scheme does not produce addition errors, the low PSNR degradation is due to an error in carry-in estimation for the precise part of the hybrid approximate adders. As expected, the average PSNR level decay for both the copy-copy-truncation and ETAI-copy-truncation when the k_3 increases. Furthermore, the coefficient of variation increases when the approximation is progressively explored. In Figure 7(a), the variability remains lower than 15% for all the cases.

In Figure 7(b) the Performance Conformance is observed for edge detection regarding the Sobel operator in Figure 5. The Performance Conformance is defined as in (4)(5)(6) [36].

$$\text{recall} = \frac{tp}{tp + fn} \quad (4)$$




Fig. 7. Quality Analysis. (a) Gaussian filter. (b) Gradient filter.

$$\text{precision} = \frac{tp}{tp + fp} \quad (5)$$

$$PC = \min(\text{recall}, \text{precision}) \quad (6)$$

In (4), the recall is defined as the number of pixels which are correctly detected as edges, by the approximate solution, over the number of pixels which should be correctly detected as edges. Therefore, tp and fn refer to the true positives and the false negatives. In (5), the precision is defined as the number of pixels correctly detected as edges over the total number of pixels detected as edges by the approximate solution. The term fp refers to the false positives. Therefore, the Performance Conformance PC is defined by the minimum between the recall and precision, where the result can range from 0 to 1.

One can observe in Figure 7(b) that the PC also decays when k_3 increases. The first configuration in which $k_3 = 0$ presents results near the maximum possible level with very low variability. It can be explained because there are only two hybrid adders approximated by the k_2 parameter as can be observed in Figure 5. It occurs because there is not overlapping of left shifted operations in these adders. The degradation in

TABLE II
 k_3 PARAMETERS FOR THE SELECTED QUALITY PROFILES

level	copy-copy-truncation	ETAI-copy-truncation
30 dB	7	5
40 dB	5	3
50 dB	2	1
0.85	5	4
0.95	3	3

terms of PC is more aggressive for $k_3 \geq 6$. At this point, the variability also substantially grows, and these values for k_3 should be avoided in architectural design.

In this study three levels of quality are selected to implement the Gaussian filter architecture: (i) PSNR near and higher than 50dB, (ii) PSNR near and higher than 40dB, and (iii) PSNR near and higher than 30dB. These levels are empirically selected after subjective observation of the output filtered images from the benchmark. For the Gradient filter, the following levels were selected based on the same analysis: (i) PC near and above 0.95 and (ii) PC near and above 0.85. In [36] the maximum achieved PC for the edge detection application is of 0.914. The adopted k_3 parameters which respect the selected quality levels are shown in Table II.

V. RESULTS

In this section, the experimental setup, quality results for the two evaluated case studies, energy efficiency analysis, and area evaluation are presented.

A. Experimental Setup

In order to evaluate the quality results, all the approximate configurations were simulated by considering eight images from the Berkeley Segmentation Dataset benchmark [37]. This new set of images is adopted to compare quality metrics with the previous analysis performed during the approximation search shown in Figure 7.

The precise and approximate designs were described in VHDL based on the approximation parameters shown in Table II. For the Gaussian filter case study, three quality levels are under evaluation. Besides, two versions of approximate hybrid adders are adopted. Therefore, this results in six approximate designs. The proposed approach in this work considers the exploration of seven conventional adder topologies for the precise part (*i.e.*, RCA plus 5 PPA's and the adder optimized by the synthesis tool). Thus, the total number of approximate designs under evaluation for the Gaussian filter is 42. The seven precise architectures are also considered in this work since they are the baselines to compare to the approximate solutions. It results in 49 described designs considering the precise and approximate ones. This same analysis is performed on the Sobel operator case study, where a total of 35 designs are explored. Therefore, the total number of described designs are 84.

All the designs are synthesized by adopting the RTL Compiler tool from Cadence, and they are mapped to the 45 nm Nangate Free PDK. The cells' structures from the designs are preserved to avoid distortions in the adder topologies.

The only exception is for the design where the behavioral adder description implements the precise part (*i.e.*, the “+” operator in Hardware Description Language). For this case, the synthesis tool is allowed to optimize this adder component for low power. In this work, this optimized implementation is identified as the “tool adder.” As presented in the proposed design methodology in Figure 3, the switching activity is extracted by considering 10000 5×5 and 3×3 blocks from images of the tested benchmark. Maximum frequency analysis is performed by the use of the bisection search method. In addition to the maximum frequency, two clock frequencies constraints are also adopted during synthesis: (i) 63 MHz, and (ii) 249 MHz. These frequencies are the minimum clock frequency targets to achieve real-time for Gaussian and Sobel operation considering Full High Definition (1920×1080) and Ultra High Definition 4K (3840×2160) video resolutions at 30 frames per second. The Mean Energy per Operation (MEOp) is obtained based on the period targets and the estimated power dissipation. Energy reduction is also evaluated for all the approximate designs about their correspondent precise versions. The maximum and minimum energy reductions are selected per filter design and frequency target as shown in Figure 9. Therefore, they are obtained considering all the approximate versions under analysis.

B. Gaussian and Gradient Quality Results

As can be seen in Figure 8, the PSNR and Performance Conformance metrics tend to be respected when a different benchmark is adopted to evaluate the approximate configurations. The lowest quality targets of 30dB and 0.85 are the ones which present higher degradation and coefficient of variation. When quality increases for both the applications, the variability is reduced. These results show that the Gaussian and Gradient responses follow the same behavior of the previous analysis to seek approximate parameters shown in Figure 7. The proposed design methodology enabled the comparison between the copy- and ETAI-based hybrid approximations. One can notice that our approach can be used to compare other approximate adders in a simulation-based design flow. One can conclude that the copy-based approximation tends to present better results than the ETAI. The exception is for the 30 dB and 0.85 targets where the ETAI-based presents the better results. It can be explained by the k_3 parameter choice shown in Table II. For example, the $k_3 = 7$ was the selected parameter for the Gaussian filter with a target of 30 dB. On the other hand, higher parameters tend to suffer from higher variability, and this may produce lower or higher results.

C. Energy Efficiency Results

Figure 9(a) and (b) show the Mean Energy per Operation in pJ for the Gaussian and Gradient filters, respectively. The columns represent the precise filter results plus the following approximate versions: (i) ETAI-copy-truncation with PSNR and Performance Conformance targets of 50 dB and 0.95, respectively, and (ii) Copy-copy-truncation with PSNR and Performance Conformance targets of 30 dB, and 0.85, respectively. Those two approximate profiles represent the highest




Fig. 8. Quality Results for the BSD benchmark. (a) Gaussian filter. (b) Gradient filter.

and the lowest Mean Energy per Operation consumption among all the approximate versions for all the cases. The operating frequencies are of 63 MHz and 249 MHz to filter videos at 30 fps regarding the Full HD and Ultra HD 4K video resolutions.

As expected, the copy-copy-truncation presents the higher energy reductions than the ETAI. It occurs due to two significant observations: (i) the ETAI has more logic complexity than the copy adder, and (ii) the copy adder presents less quality distortion than the former, which enables more approximation.

When considering the precise versions as being the baseline architecture, one can observe that a wide range of energy savings can be analyzed. In the Gaussian architecture, the energy reduction is more expressive than in the Gradient filter. It is due to the massive presence of adders in the Gaussian adder tree shown in Figure 4 and Figure 5. The energy reductions in Figure 9(a) for 63 MHz are ranging from 8.3% to 73.2%, while for 249 MHz they are from 7.7% up to 70.9%. The energy reductions in Figure 9(b) for 63 MHz are ranging from 18.7% to 57.2%, while for 249 MHz they are from




Fig. 9. Mean Energy per Operation. (a) Gaussian filter. (b) Gradient filter.

TABLE III

VOLTAGE OVER SCALING ANALYSIS @ 249 MHz FOR GAUSSIAN FILTER

Topology	RCA	B-K	H-C	K-S	L-F	SK
max freq (MHz)	399.9	450.2	425.1	501.1	465.8	453.1
VDD_{scaled} (V)	-	0.98	1.03	0.88	0.94	0.97
P_{Dyn} (mW) @ 1.1V	0.97	1.08	1.12	1.31	1.08	1.12
P_{Dyn} (mW) @ VOS	-	0.85	1	0.83	0.8	0.88

TABLE IV

VOLTAGE OVER SCALING ANALYSIS @ 249 MHz FOR GRADIENT FILTER

Topology	RCA	B-K	H-C	K-S	L-F	SK
max freq (MHz)	507	551.7	571.8	628.3	553	588.3
VDD_{scaled} (V)	-	1.01	0.98	0.89	1.0	0.95
P_{Dyn} (mW) @ 1.1V	0.3	0.33	0.34	0.37	0.33	0.34
P_{Dyn} (mW) @ VOS	-	0.28	0.26	0.24	0.28	0.25

13.4% up to 52.3%. These results clearly show that our design methodology and proposed hybrid approximate adders provide energy efficiency for compute-intensive image/video filters.

The Gaussian and Gradient filters fully implemented by the precise “tool adder” present the lowest energy consumption when compared to the PPA adders and the RCA versions. One can conclude that the synthesis tool makes substantial effort to build low power designs. Also, one can observe that our proposed approximate approach further improves the energy efficiency in this scenario. For instance, energy reductions provided by the hybrid approximate solution range from 6.3% up to 64.4% when considered the Gaussian filter with precise part implemented by the “tool adder” and an operating frequency of 249 MHz. These energy reductions are of up to 66.2% when the clock frequency target to the Gaussian filter is of 63 MHz.

As expected, the PPAs versions consume more energy than the RCA-based version as can be seen in the Figure 9. On the other hand, Table III and Table IV show the maximum

frequency reached by the precise Gaussian and Sobel operator implemented by the 5 PPAs and the RCA topology, respectively. This result shows that the RCA-based filter is the slowest hardware, and, therefore, the PPA-based designs are evaluated under VOS operation considering the frequency of 249 MHz.

In [13] is shown that the VOS technique consists of scaling down the VDD without scaling the clock frequency accordingly. The circuit delay (as in all digital CMOS) is inversely proportional to the voltage supply VDD , as demonstrated in [13]. Therefore, they propose a VOS model shown in (7) to calculate the lower boundary regarding scaled VDD (VDD_{scaled}) which still avoids timing induced errors. We consider this model to estimate additional dynamic power reduction when applying VOS in adder topologies which are faster than the RCA. This reduction is significant, as the dynamic power is directly proportional to VDD^2 . In (7) the term *slack* refers to the difference between the minimum period achieved of a given PPA topology and the baseline RCA. The term T_c

TABLE V
AREA (μm^2) ANALYSIS @ 249 MHz

Gaussian filter analysis						
design	precise		copy-copy-truncation 30 dB		ETAI-copy-truncation 50 dB	
criterion	# of cells	area (μm^2)	# of cells	area (μm^2)	# of cells	area (μm^2)
B-K	2679	3924	782	1166	2359	3533
H-C	2927	4188	803	1188	2563	3750
K-S	3830	5149	952	1347	3292	4525
L-F	2694	3940	782	1166	2365	3539
RCA	2126	3336	697	1075	1948	3095
SK	2914	4174	795	1180	2523	3707
tool	1800	3210	623	1048	1722	2988
Gradient filter analysis						
B-K	702	1208	340	676	538	996
H-C	756	1256	352	687	562	1017
K-S	960	1436	406	735	682	1123
L-F	708	1213	340	676	538	996
RCA	570	1091	322	660	490	953
SK	756	1256	346	681	544	1001
tool	554	1032	304	600	524	947

denotes the period of the operating clock, which in this part of the study is $1/249000000$ seconds. Since at VOS condition the clock frequency is not accordingly scaled when VDD is reduced, there is not performance penalty. Also, most of the faster PPAs presents dynamic power (P_{dyn}) reduction shown in Tables III and IV when compared to the RCA baseline filter. One can conclude that the dynamic power reduction reaches up to 17.4% (Ladner-Fischer version) and 19.3% (Kogge-Stone version) when considering the Gaussian and Gradient operators, respectively. Besides the additional dynamic power savings of the PPA adders, they can accomplish higher frame rates to process higher video resolutions than the RCA. Therefore, the use of PPA adders may be preferable than the RCA, depending on the observed scope.

$$VDD_{scaled} = VDD(1 - \frac{slack}{T_c}) \quad (7)$$

The “tool adder” is not exercised in this context because the commercial synthesis tool tends to push the limits to achieve the highest possible clock frequency. This procedure may conceive a gate-level netlist which is substantially different from the one synthesized for 249 MHz. Based on that, the maximum frequency may not represent a fair analysis considering VOS operation.

D. Area Analysis

The area analysis is shown in Table V for the Gaussian and Gradient filters with a clock frequency of 249 MHz. The number of cells and area (μm^2) are shown for the precise designs plus the same approximate configurations previously shown in Figure 9 and enumerated as follows: (i) the copy-copy-truncation with PSNR and Performance Conformance targets of 30 dB and 0.85, and (ii) the ETAI-copy-truncation with PSNR and Performance Conformance targets of 50 dB and 0.95. These configurations were chosen because they present the maximum and minimum area reductions among all the approximate designs. The results are organized per image filter and approximate configuration. One can observe that the “tool adder” and the Kogge-Stone (KS) topology

implemented in the precise parts of the hybrid approximate adders represent the minimum and maximum area reductions for almost all the cases, respectively. It is expected since the “tool adder” is optimized for low power, while the KS has the highest area. Based on that, these reductions are of 67.4% up to 73.8% for the Gaussian filter, and 39.5% up to 48.8% for the Sobel operator implemented by the copy-copy-truncation. When considered the ETAI-copy-truncation, the reductions are of 6.9% up to 14.1% for the Gaussian filter, and of 8.2% up to 21.8% for the Gradient filter. Following the same conclusions made in the Energy Efficiency analysis, the area reductions ratify the contributions of this study.

E. Energy Efficiency Vs. Application Quality

In this subsection, the main objective is to evaluate the relationship between application quality and energy consumption. One can observe in Figure 10(a) and (b) that energy consumption raises when the application quality is improved. The Figure 10(a) and (b) show all the evaluated approximate configurations for the Kogge-Stone and “tool adder” precise parts. These precise parts were selected because they represent the highest and the lowest energy consumption among all the conventional topologies. The results for the Gaussian filter are shown in Figure 10(a). One can observe that for all the approximate versions, the energy consumption increases when higher PSNR quality is demanded. The same can be observed for the Gradient filter in Figure 10(b). These results are expected since higher quality profiles are associated with designs which are less approximated.

VI. COMPARISON WITH RELATED WORK

As previously mentioned in the related work section, essential contributions were given in [22], [23], and [33]. In [22] the maximum energy reductions of 26.9% and 60% are provided for the 45 nm ASIC implementation of Gaussian and Gradient filters, respectively. On the other hand, the quality analysis is limited, since the Sobel operator is evaluated by adopting PSNR quality metric instead of Performance Conformance which is more appropriate for edge detection scope. Also, only

TABLE VI
COMPARISON WITH RELATED WORK

work	design flow	acceleration method	application	quality metric	maximum energy reduction
[22]	ASIC 45nm	add-and-shift	Gaussian and Sobel filters	PSNR	60%
[23]	ASIC 65nm	add-and-shift	FIR filters	PSNR and Accuracy	50.7%
[33]	ASIC 40nm	SIMD CPU	Sobel filter	Saturated Mean Squared Error	15%
This work	ASIC 45nm	add-and-shift	Gaussian and Sobel filters	PSNR and Performance conformance	73.2%




Fig. 10. Mean energy per operation vs. application quality. (a) Gaussian filter. (b) Gradient filter.

the RCA is explored in the precise part of approximate adders, and no maximum frequency analysis is presented.

In [23] the maximum energy reduction is of 50.7% for a set of 65 nm ASIC FIR filters approximated by their proposed synthesis flow. On the other hand, the quality analysis is given regarding accuracy, and no evaluation at the application layer is observed. Furthermore, the authors did not evaluate real-time scenario for the FIR filters under analysis.

In [33] the maximum energy reduction is of 15% considering the Sobel operator application being processed by their proposed 40nm ASIC SIMD co-processor. On the other hand, no quality evaluation is performed in application quality.

The design methodology proposed in this work uses the same case studies and technology exercised in [22]. Based on that, one can conclude that our proposed approach presents substantial reductions of up to 73.2% and 57.2% for the Gaussian and Gradient filters, respectively.

In the scope of search heuristic and design methodology, this work also introduces new contributions when compared to the related works. The energy efficiency results validate the effort of seeking for hybrid approximate adders inside shift-and-add accelerators. Table VI shows an overall comparison among this work and the related ones.

Among the related works, [33] is the single which proposes a real-time analysis for Sobel operation. According to the authors, the precise approach reaches a maximum performance of 215 μs per 512×512 greyscale frame, while the energy consumption is approximately 9 μJ per frame. When considering the most energy efficient approximate approach, the authors show that the maximum performance is of 195 μs , while the energy consumption is approximately 8.2 μJ per frame. In this work, the most energy-efficient approaches without considering VOS estimation is the RCA topology. One can conclude, by examining the frequency of 249 MHz, that the precise Sobel operator, based on RCA, takes 1052 μs to process a 512×512 gray scale image, while the energy consumption is of 0.34 μJ per frame. The most energy-efficient approximate Sobel architecture is the copy-copy-truncation with Performance Conformance target of 0.85. The approximate approach takes the same time to process 512×512 greyscale images, while the energy consumption is of 0.2 μJ per frame. When comparing precise approaches (*i.e.* same level of quality), our 45 nm accelerator architecture presents energy reduction of 96.2% when compared to the precise 40 nm SIMD co-processor running Sobel operator in [33].

The limitations of this work are related to the filter designs being implemented in application-specific architecture scope. Therefore, changes in filter kernels are not possible after the ASIC is fabricated. When this scenario of general-purpose is required, the work in [33] is of remarkable use and importance.

On the other hand, our approach is of notable contribution in the design of ASIC accelerators which can be integrated into heterogeneous architectures (*i.e.*, integration of general-purpose processors plus ASIC accelerators in a chip). It is because the results of this work indicate that substantial energy reduction is observed every time a predefined and static filter kernel is required during a given application.

VII. CONCLUSION

This work proposed a novel design flow methodology to cope with energy efficiency in CMOS technology. The proposed solution is focused on exploring approximation in add-and-shift architectures to reduce power consumption and increase computational performance. The proposed search heuristic and hybrid approximate adders presented substantial energy reductions of up to 73.2% when compared to the precise and baseline architectures. PPA topologies were explored, to rescue performance, in addition to RCA-based approach, where VOS estimation shows additional dynamic power reduction of up to 19.3%. Area reduction up to 73.8% is also observed in this study. The proposed design methodology

also enabled a more comprehensive observation of application quality by considering average results and variability. Comparison with state-of-the-art related work is provided showing the contributions of this work for low power digital CMOS design and approximate computing scope. Future work and effort are focused on giving configurable capabilities to the filter images under analysis, thus enabling the exploration of different power-performance profiles during the execution time.

REFERENCES

- [1] R. H. Dennard, "Past progress and future challenges in LSI Technology: From DRAM and scaling to ultra-low-power CMOS," *IEEE Solid-State Circuits Mag.*, vol. 7, no. 2, pp. 29–38, 2015.
- [2] J. Henkel, H. Khdr, S. Pagani, and M. Shafique, "New trends in dark silicon," in *Proc. 52nd ACM/EDAC/IEEE Design Automat. Conf. (DAC)*, San Francisco, CA, USA, Jun. 2015, pp. 1–6.
- [3] M. Shafique, S. Garg, J. Henkel, and D. Marculescu, "The EDA challenges in the dark silicon era," in *Proc. 51st ACM/EDAC/IEEE Design Automat. Conf. (DAC)*, San Francisco, CA, USA, Jun. 2014, pp. 1–6.
- [4] R. Iyer, "Accelerator-rich architectures: Implications, opportunities and challenges," in *Proc. 17th Asia South Pacific Design Automat. Conf.*, Sydney, NSW, Australia, Jan./Feb. 2012, pp. 106–107.
- [5] J. Cong, M. A. Ghodrati, M. Gill, B. Grigorian, K. Gururaj, and G. Reinman, "Accelerator-rich architectures: Opportunities and progress," in *Proc. 51st ACM/EDAC/IEEE Design Automat. Conf. (DAC)*, San Francisco, CA, USA, Jun. 2014, pp. 1–6.
- [6] R. Hameed *et al.*, "Understanding sources of inefficiency in general-purpose chips," *ACM SIGARCH Comput. Archit. News*, vol. 38, no. 3, pp. 37–47, Jun. 2010.
- [7] Y. Voronenko and M. Püschel, "Multiplierless multiple constant multiplication," *ACM Trans. Algorithms*, vol. 3, no. 2, p. 11, May 2017.
- [8] L. Aksoy, E. Costa, P. Flores, and J. Monteiro, "Optimization of area and delay at gate-level in multiple constant multiplications," in *Proc. 13th Euromicro Conf. Digit. Syst. Design: Architectures, Methods Tools*, Lille, France, Sep. 2010, pp. 3–10.
- [9] J. Han and M. Orshansky, "Approximate computing: An emerging paradigm for energy-efficient design," in *Proc. 18th IEEE Eur. Test Symp. (ETS)*, Avignon, France, May 2013, pp. 1–6.
- [10] S. Venkataramani, S. T. Chakradhar, K. Roy, and A. Raghunathan, "Approximate computing and the quest for computing efficiency," in *Proc. 52nd ACM/EDAC/IEEE Design Automat. Conf. (DAC)*, San Francisco, CA, USA, Jun. 2015, pp. 1–6.
- [11] N. Zhu, W. L. Goh, W. Zhang, K. S. Yeo, and Z. H. Kong, "Design of low-power high-speed truncation-error-tolerant adder and its application in digital signal processing," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 18, no. 8, pp. 1225–1229, Aug. 2010.
- [12] Q. Xu, T. Mytkowicz, and N. S. Kim, "Approximate computing: A survey," *IEEE Design Test*, vol. 33, no. 1, pp. 8–22, Feb. 2016.
- [13] V. Gupta, D. Mohapatra, A. Raghunathan, and K. Roy, "Low-power digital signal processing using approximate adders," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 32, no. 1, pp. 124–137, Jan. 2013.
- [14] N. Zhu, W. L. Goh, G. Wang, and K. S. Yeo, "Enhanced low-power high-speed adder for error-tolerant application," in *Proc. Int. SoC Design Conf. (ISOCC)*, Seoul, South Korea, Nov. 2010, pp. 323–327.
- [15] N. Zhu, W. L. Goh, and K. S. Yeo, "An enhanced low-power high-speed adder for error-tolerant application," in *Proc. 12th Int. Symp. Integr. Circuits (ISIC)*, Singapore, Dec. 2009, pp. 69–72.
- [16] A. K. Verma, P. Brisk, and P. Jenne, "Variable latency speculative addition: A new paradigm for arithmetic circuit design," in *Proc. Design, Automat. Test Eur. (DATE)*, Munich, Germany, 2008, pp. 1250–1255.
- [17] R. Ye, T. Wang, F. Yuan, R. Kumar, and Q. Xu, "On reconfiguration-oriented approximate adder design and its application," in *Proc. IEEE/ACM Int. Conf. Comput.-Aided Design (ICCAD)*, San Jose, CA, USA, Nov. 2013, pp. 48–54.
- [18] M. Shafique, W. Ahmad, R. Hafiz, and J. Henkel, "A low latency generic accuracy configurable adder," in *Proc. 52nd ACM/EDAC/IEEE Design Automat. Conf. (DAC)*, San Francisco, CA, USA, Jun. 2015, pp. 1–6.
- [19] A. B. Kahng and S. Kang, "Accuracy-configurable adder for approximate arithmetic designs," in *Proc. Design Autom. Conf. (DAC)*, San Francisco, CA, USA, Jun. 2012, pp. 820–825.
- [20] H. R. Mahdiani, A. Ahmadi, S. M. Fakhraie, and C. Lucas, "Bio-inspired imprecise computational blocks for efficient VLSI implementation of soft-computing applications," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 57, no. 4, pp. 850–862, Apr. 2010.
- [21] S. Rehman, W. El-Harouni, M. Shafique, A. Kumar, J. Henkel, and J. Henkel, "Architectural-space exploration of approximate multipliers," in *Proc. IEEE/ACM Int. Conf. Comput.-Aided Design (ICCAD)*, Austin, TX, USA, Nov. 2016, pp. 1–8.
- [22] J. de Oliveira, L. Soares, E. Costa, and S. Bampi, "Exploiting approximate adder circuits for power-efficient Gaussian and Gradient filters for Canny edge detector algorithm," in *Proc. IEEE 7th Latin Amer. Symp. Circuits Systems (LASCAS)*, Florianopolis, Brazil, Feb./Mar. 2016, pp. 379–382.
- [23] Y. Kang, J. Kim, and S. Kang, "Novel approximate synthesis flow for energy-efficient FIR filter," in *Proc. IEEE 34th Int. Conf. Comput. Design (ICCD)*, Scottsdale, AZ, USA, Oct. 2016, pp. 96–102.
- [24] A. Beaumont-Smith and C.-C. Lim, "Parallel prefix adder design," in *Proc. 15th IEEE Symp. Comput. Arithmetic*, Vail, CO, USA, Jun. 2001, pp. 218–225.
- [25] D. L. Harris, "Parallel prefix networks that make tradeoffs between logic levels, fanout and wiring racks," U.S. Patent 7 152 089 B2, Dec. 19, 2006.
- [26] R. P. Brent and H. T. Kung, "A regular layout for parallel adders," *IEEE Trans. Comput.*, vol. C-31, no. 3, pp. 260–264, Mar. 1982.
- [27] P. M. Kogge and H. S. Stone, "A parallel algorithm for the efficient solution of a general class of recurrence equations," *IEEE Trans. Comput.*, vol. C-22, no. 8, pp. 786–793, Aug. 1973.
- [28] T. Han and D. A. Carlson, "Fast area-efficient VLSI adders," in *Proc. IEEE 8th Symp. Comput. Arithmetic*, Como, Italy, May 1987, pp. 49–56.
- [29] R. E. Ladner and M. J. Fischer, "Parallel prefix computation," *J. ACM*, vol. 27, no. 4, pp. 831–838, Oct. 1980.
- [30] J. Sklansky, "Conditional-sum addition logic," *IRE Trans. Electron. Comput.*, vols. EC-9, no. 2, pp. 226–231, Jun. 1960.
- [31] J. Canny, "A computational approach to edge detection," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. PAMI-8, no. 6, pp. 679–698, Nov. 1986.
- [32] D. Esposito, D. De Caro, and A. G. M. Strollo, "Variable latency speculative parallel prefix adders for unsigned and signed operands," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 63, no. 8, pp. 1200–1209, Aug. 2016.
- [33] A. Najafi, M. Weißbrich, G. P. Vayá, and A. Garcia-Ortiz, "Coherent design of hybrid approximate adders: Unified design framework and metrics," *IEEE Trans. Emerg. Sel. Topics Circuits Syst.*, vol. 8, no. 4, pp. 736–745, Dec. 2018.
- [34] M. Macedo, L. Soares, B. Silveira, C. M. Diniz, and E. A. C. da Costa, "Exploring the use of parallel prefix adder topologies into approximate adder circuits," in *Proc. 24th IEEE Int. Conf. Electron., Circuits Syst. (ICECS)*, Batumi, Georgia, Dec. 2017, pp. 298–301.
- [35] L. B. Soares, M. M. A. da Rosa, C. M. Diniz, E. A. C. da Costa, and S. Bampi, "Exploring power-performance-quality tradeoff of approximate adders for energy efficient sobel filtering," in *Proc. IEEE 9th Latin Amer. Symp. Circuits Syst. (LASCAS)*, Puerto Vallarta, Mexico, Feb. 2018, pp. 1–4.
- [36] J. Lee, H. Tang, and J. Park, "Energy efficient canny edge detector for advanced mobile vision applications," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 28, no. 4, pp. 1037–1046, Apr. 2018.
- [37] D. Martin, C. Fowlkes, D. Tal, and J. Malik, "A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics," in *Proc. 8th IEEE Int. Conf. Comput. Vis. (ICCV)*, Vancouver, BC, Canada, Jul. 2001, pp. 416–423.



Leonardo Bandeira Soares (S'12) received the Engineering degree in computer engineering from the Federal University of Rio Grande, Rio Grande, Brazil, in 2010, and the M.Sc. and Ph.D. degrees in microelectronics from the Federal University of Rio Grande do Sul, Porto Alegre, Brazil, in 2013 and 2018, respectively. He is currently a Professor with the Federal Institute of Technology of Rio Grande do Sul. His research interests are very large-scale integration architectures, approximate computing, video coding, digital signal processing, and energy efficiency in complementary metal-oxide-semiconductor design.



Morgana Macedo Azevedo da Rosa received the degree in computer engineering from the Catholic University of Pelotas, Pelotas, Brazil, where she is currently pursuing the master's degree in electronic engineering and computing. Her research interests are arithmetic circuits and very large-scale integration design.



Eduardo Antonio César da Costa (M'13) received the Engineering degree in electrical engineering from the University of Pernambuco, Recife, Brazil, in 1988, the M.Sc. degree in electrical engineering from the Federal University of Paraíba, Campina Grande, Brazil, in 1991, and the Ph.D. degree in computer science from the Federal University of Rio Grande do Sul, Porto Alegre, Brazil, in 2002. Part of his doctoral work was developed at the Instituto de Engenharia de Sistemas Computadores, Lisbon, Portugal. He is currently a Full Professor with the Catholic University of Pelotas (UCPel), Pelotas, Brazil. He is a Co-Founder and a Coordinator of the Graduate Program on Electronic Engineering and Computing, UCPel. His research interests are very large-scale integration architectures and low-power design.



and video processing algorithms, architectures, and very large-scale integration design.

Cláudio Machado Diniz (S'08–M'15) received the degree in computer engineering from the Federal University of Rio Grande, Brazil, in 2007, and the M.Sc. and Ph.D. degrees in computer science from the Federal University Rio Grande do Sul, Brazil, in 2009 and 2015, respectively. He was an Intern Researcher with the Chair for Embedded Systems, Karlsruhe Institute of Technology, Karlsruhe, Germany. He is currently an Assistant Professor with the Catholic University of Pelotas, Pelotas, Brazil. His research interests include image



Sergio Bampi (M'86–SM'17) received the Electronics Engineer degree and the B.Sc. degree in physics from the Federal University of Rio Grande do Sul in 1979 and the M.S.E.E. degree and the Ph.D. degree in electrical engineering from Stanford University in 1982 and 1986, respectively. In 1981, he joined the Informatics Institute, Federal University of Rio Grande do Sul, Brazil, where he is currently a Full Professor. He has published more than 380 research papers in conferences and journals, in the fields of complementary metal–oxide–semiconductor analog, digital, and RF design, video coding algorithms, and dedicated hardware architectures. He was the President of the Brazilian Microelectronics Society and the FAPERGS Brazilian research funding agency, and the CEITEC Technical Director. He was a Distinguished Lecturer of the IEEE CAS Society from 2009 to 2010. He served as the Technical Program Chair for SBCCI in 1997 and 2005, the IEEE LASCAS in 2013, VARI in 2015, and SBMICRO Congress in 1989 and 1995, and served on TPC committees for ICCAD, SBCCI, ICM, LASCAS, VLSI-SoC, ICECS, and many other international conferences.