

Approximate Logic Synthesis via Iterative SMT-based Subcircuit Rewriting

Morteza Rezaalipour, Marco Biasion, Francesco Costa, Cristian Tirelli, Rodrigo Otoni, Lorenzo Ferretti, George A. Constantinides, and Laura Pozzi

Abstract—This paper presents a novel iterative approach to achieve effective and scalable approximate logic synthesis (ALS). The core idea is to perform circuit rewriting in a way that is both local, i.e., is applied piece-wise to selected subcircuits, and extensive, i.e., systematically explores the design space for good solutions. Concretely, we propose SubXPAT, a new Boolean rewriting framework which iteratively employs satisfiability modulo theories (SMT) solving to select and approximate key parts of a circuit. Selection aims at finding subcircuits that at the same time include a significant number of gates and can be efficiently approximated, which is done by searching for large convex subcircuits with a limited number of inputs and outputs. Approximation is guided by the use of a parametric template, structured as a sum of products, which allows for fine-grained control over the subcircuit characteristics. SubXPAT was implemented as an open-source tool and compared against other ALS tools implementing state-of-the-art techniques. Our experimental evaluation used a broad range of arithmetic circuits with different bit-widths and our results indicate that SubXPAT generates approximate circuits that are more area-efficient than those generated by state-of-the-art techniques.

Index Terms—approximate computing, Boolean rewriting, local analysis, SMT solving.

I. INTRODUCTION

APPROXIMATE computing is a new paradigm in computer design postulating that inexact hardware should be used whenever full accuracy is not required by a given application; hence, accuracy should be traded for improvements in metrics such as area and energy performance [1]. Approximate circuits can be employed in any error-tolerant application, including those for which an exact output cannot be distinguished from an inexact one, noisy data inherently prevents full accuracy, or no single golden answer exists.

Approximate logic synthesis (ALS) is the process of deriving an approximate circuit automatically, from a given exact circuit and a tolerated error threshold (ET). There exist several families of ALS methods, categorised in a recent survey [2]

Morteza Rezaalipour, Marco Biasion, Francesco Costa, Cristian Tirelli, Rodrigo Otoni, and Laura Pozzi are all with the Faculty of Informatics, Università della Svizzera italiana, Lugano, Switzerland (e-mail: morteza.rezaalipour@usi.ch; marco.biasion@usi.ch; francesco.costa@usi.ch; cristian.tirelli@usi.ch; otonir@usi.ch; laura.pozzi@usi.ch).

Lorenzo Ferretti is with the Silicon System AI group, Micron Technology, San Jose, USA (e-mail: lferretti@micron.com).

George A. Constantinides is with the Faculty of Engineering, Imperial College London, London, UK (e-mail: g.constantinides@imperial.ac.uk).

This work was partially supported by the Swiss National Science Foundation via grant 200020_188613.

Manuscript received January XX, 20XX; revised January XX, 20XX.

into those based on *netlist modification*, which generate the resulting circuit starting from an existing structure, e.g., by removal of gates or wires, and those based on *Boolean rewriting*, which generate the resulting circuit without reference to an original structure, e.g., by Boolean factorisation. Despite extensive work towards improving ALS techniques [3]–[8], approximation methods can still be further improved in terms of exploiting the flexibility allowed by the given ET.

With the aim of generating approximations that are better than those that can be obtained with current techniques we propose SUBXPAT, a novel Boolean rewriting framework for the synthesis of approximate circuits that iteratively selects and approximates key subcircuits of the exact circuit. The core of SUBXPAT is our XPAT algorithm, which employs a satisfiability modulo theories (SMT) solver to shape a *parametrical template* into an optimised circuit, i.e., a circuit with a lower amount of gates, while remaining sound w.r.t. the given ET [9]. Concretely, the template we use is a sum of products and the role of the solver is to identify, for every circuit output, which products of which input literals must be included in the final synthesis. XPAT by itself, i.e., when it is applied to the whole circuit, yields highly optimised results, but is limited in its scalability, only being practical when applied to small circuits. To address this, the iterative infrastructure of SUBXPAT is used to select subcircuits that are both small enough to be approximated by XPAT in a timely manner and large enough to lead to significant optimisations.

To enable the practical use of SUBXPAT we implemented it as an open-source tool, which also allows for third-parties to make use of our framework. We performed an experimental evaluation using 23 arithmetic circuits and compared SUBXPAT’s results against those of the state-of-the-art ALS methods MUSCAT [7], MECALS [8], and EVOAPPROX [4]. Our results show that SUBXPAT yields better improvements in area in most cases, sometimes of more than 50%.

To summarise, our contributions are the following:

- (1) An SMT-based ALS algorithm, called XPAT, which synthesises approximate circuits via the shaping of a parametrical template into an optimised circuit (originally appeared in [9]).
- (2) An iterative framework for ALS, called SUBXPAT, which enables iterative applications of XPAT to different subcircuits of the original circuit.
- (3) An open-source implementation of the SUBXPAT framework that allows for its usage in practice.

- (4) An extensive evaluation that showcases both the ability of XPAT to yield highly optimised circuits and the scalability gains achieved with SUBXPAT.

The remainder of this article is structured as follows. Section II discusses related work. Sections III, IV, and V respectively present XPAT, SUBXPAT, and iterative SUBXPAT. Section VI details the fine-tuning of our framework. Section VII reports our evaluation results. Finally, Section VIII presents our conclusions and future work.

II. STATE OF THE ART

Techniques for ALS can be categorised over two axes: approximation procedure and approximation target. The procedure can be either monolithic or iterative, depending if the approximation happens in one step or if a number of consecutive approximations is made. The target can be either the whole circuit or a subcircuit, with subcircuits with different characteristics being potentially chosen.

Monolithic approaches commonly target the whole circuit, since they perform approximation only once. Examples of this combination are MUSCAT [7] and CIRCUIT CARVING [5]. Due to having to perform a global analysis over the whole circuit, monolithic approaches have limited scalability, preventing their use in many practical scenarios. One positive aspect, however, is that having a view of the whole circuit at once can in principle allow for more powerful approximations.

Iterative approaches arose as a way to improve scalability, by means of performing local analysis over a subcircuit, allowing for ALS to be applied to significantly larger circuits. Examples of this combination are MECALS [8], GLP [10], and SASIMI [3], all of which select, at each iteration, a subcircuit consisting of a single gate. This restrictive selection aids in the aim of improving scalability, but at the expense of limiting approximation options. These approaches consider local approximation candidates (LACs) which consist, at each iteration, in simple modifications such as removing a gate or altering a wire. Existing iterative approaches are, thus, likely to underperform when compared to monolithic approaches, for the cases in which the monolithic approaches terminate in a reasonable amount of time.

In contrast to considering single modifications to a gate or a wire, the technique adopted by SUBXPAT considers, at each iteration, a powerful logic-rewriting LAC. Our technique goes beyond the state of the art by (1) considering multiple-output subcircuits for approximation, which are significantly more complex than single gates, and by (2) using an SMT-based rewriting guided by a parametrical circuit template.

Two other approaches of note are BLASYS [6], [11], [12] and EVOAPPROX [4], [13]. BLASYS also iteratively applies rewriting to subcircuits, in this case by means of logic factorisation, while EVOAPPROX employs genetic programming to achieve its approximations. Our evaluation (cf. Section VII) quantitatively shows that SUBXPAT outperforms the state of the art, corroborating the qualitative intuition presented here, i.e., that applying our powerful logic-rewriting yields better results than employing previously considered simple LACs, as done by MECALS, and that template-based rewriting can be more efficient than techniques such as genetic programming.

Lastly, note that techniques such as MUSCAT, MECALS, and EVOAPPROX are comparable with ours because they, like the present work, target maximum (worst-case) error. Hence, these techniques produce circuits that never exceed the given ET, for every input configuration. Another category of works was developed [14]–[20] to target, instead, average error metrics, such as error rate and mean absolute error – as such, these techniques are not directly comparable.

III. XPAT

We approach ALS through the lens of Boolean rewriting. This allows us to explore a large design space that is not restricted by the structure of the circuit that we aim to approximate, in principle enabling approximations that would otherwise not be achievable by considering structural methods alone. To guide the search over this large design space, we make use of a *parametrical template*. Concretely, XPAT's template consists of sums of products, with each sum dictating the value of one primary output and the parameters deciding which products of which primary inputs should be included in each sum. We chose this template because of its expressiveness, which guarantees us to be able to represent any Boolean function. The design space exploration is framed in terms of SMT formulas, whose solving leads to the setting of the template's parameters in such a way that the approximate circuit respects the given ET. In the following we detail the template used, in Section III-A, the SMT formulation employed, in Section III-B, and how the design space is explored, in Section III-C.

A. Parametrical Template

Consider a given circuit template with n primary inputs, in_1, \dots, in_n , and m primary outputs, out_1, \dots, out_m . The behaviour of each output out_i is formalised as follows:

$$out_i = \bigvee_{k \in \{1 \dots K\}} Prod_k$$

where K is the number of products included in each sum, $Prod_k$ is a product of selected primary inputs, and $1 \leq i \leq m$, as illustrated in Figure 1a. Each product $Prod_k$ has n inputs, originating from n multiplexers, as illustrated in Figure 1b. Each multiplexer decides, based on parameter p_k^j , if primary input in_j is used as is, is negated, or is replaced by the constant 1, with $1 \leq j \leq n$; the usage of constant 1 indicates that the respective primary input is ignored by the product. Note that the multiplexers are not a part of the generated circuit, since the selected lines are chosen at compile time.

All p_k^j parameters are of 2-bit form, with the first bit indicating if in_j is used or ignored, and the second bit indicating the sign of in_j . Products can thus be described in terms of vectors of form $[p_k^1, \dots, p_k^n]$. For example, having $in_1 = a$, $in_2 = b$, and $in_3 = c$ as primary inputs, the product described by $[11, 10, 01]$ would output ab' , while the product described by $[01, 01, 10]$ would output c' .

With our template of sums of parametric products we can capture all circuit behaviours, except the output of constant 0. To address this, a further 1-bit parameter is added for each

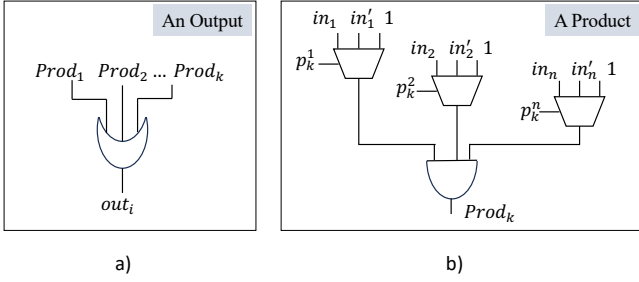


Fig. 1: Our 2-layer sum of products template. (a) A sum of K products. (b) A product of n primary inputs; parameter p_k^j determines how primary input in_j is to be handled.

sum, to select whether the output of the sum or the constant 0 should be used for the respective primary output. Having m outputs, each defined as a sum of K products, with each product having n 2-bit parameters, and considering the extra parameter needed to express the constant 0, the total number of parameters of our template is $m(2nK + 1)$. This represents our design space, which we explore via SMT solving.

B. SMT Formulation

Our formulation of the design space exploration is based on an error miter, which consists of the exact circuit, the parametrical template for approximation, two functions used for error estimation, and a query, as illustrated in Figure 2. Function $map : \mathbb{B}^m \rightarrow S$ does the mapping from the primary output of each circuit to an element of S , with S being a user-defined set containing interpretations of Boolean values. For example, with $m = 3$ and S instantiated to \mathbb{Z} we can have $map(out_3, out_2, out_1) = 4*out_3 + 2*out_2 + 1*out_1$. Function $dist : S \times S \rightarrow Dist$ measures the distance, according to a desired metric and with $Dist$ being a totally ordered set, between the mapped primary outputs of the exact and approximate circuits. For example, with $Dist$ instantiated to \mathbb{Z} we can have $dist(s_1, s_2) = |s_1 - s_2|$.

We aim to generate an approximate circuit that never exceeds the ET. To achieve this the template parameters must be

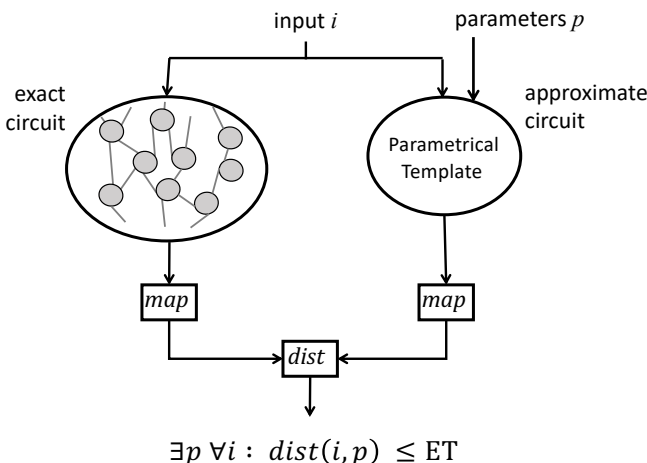


Fig. 2: Miter for XPAT.

instantiated in such a way that, for all possible primary input combinations, the distance between the exact and approximate primary outputs does not exceed the ET. Thus, we ask the SMT solver whether it is possible to assign values to the parameters in such a way that the approximate circuit always respects the ET. Concretely, we ask whether the following formula is satisfiable:

$$\exists p \in P . \forall i \in \mathbb{B}^n . dist \left(\begin{array}{c} map(outExact(i)), \\ map(outAppx(i, p)) \end{array} \right) \leq ET$$

where P is the set containing all possible values for the $m(2nK + 1)$ parameters, and $outExact$ and $outAppx$ are functions respectively capturing the behaviour of the exact and the approximate circuits. A satisfiable result will be accompanied by an assignment to p that ensures that the approximate circuit always respects the ET.

C. Design Space Exploration

Simply asking the SMT solver for a satisfying assignment might lead us to a low quality solution. This is due to the exponential growth of the design space based on the number of template parameters, which enables numerous valid solutions, combined with the fact that the solver cannot differentiate between assignments of different quality, i.e., assignments that will lead to larger or smaller synthesised areas. To address this, we systematically limit the design space by restricting the number of literals per product (LPP) and products per output (PPO) allowed in the approximate circuit. By asking for a satisfying assignment that, in addition to respecting the given ET, also respects an upper bound on LPP and PPO, we can generate approximations that will utilise fewer gates and will thus lead to smaller synthesised areas; intuitively, we aim to consider sets $P_0 \subset P_1 \subset \dots \subset P$ in our satisfiability query.

The additional constraints added might lead to a design space in which no valid solution exists. We must thus consider different pairs of LPP and PPO constraints, which we explore in a grid-like manner as illustrated in Figure 3. The starting point of the exploration has LPP = 0 and PPO = 1, with LPP being 0 symbolising that products are substituted by constants. In case the formula is unsatisfiable LPP is incremented and the solver is called again. If the LPP bound is reached and the formula is still unsatisfiable, we increment PPO, set LPP to 1, and continue as before. The process terminates once either a satisfiable assignment is found or both the LPP and the PPO bounds are reached.

IV. SUBXPAT

Since XPAT requires finding a satisfiable assignment for a formula that quantifies over the whole circuit, it has limited scalability. To address this, SUBXPAT approximates selected *subcircuits*, instead of the whole circuit, as done by XPAT. In the following we first detail the SMT formulation designed for SUBXPAT, in Section IV-A, and then describe how subcircuit selection takes place, in Section IV-B and Section IV-C.

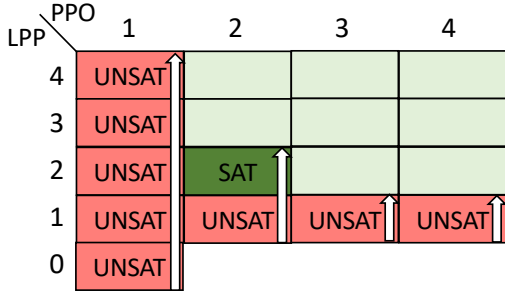


Fig. 3: Design space exploration guided by LPP and PPO; we explore cells 1-3 and 1-4 because they may have alternative (and better quality) solutions to the one found in cell 2-2.

A. SMT Formulation for Local Approximation

Assuming that a subcircuit has been selected for approximation – and how this is done is described in the two following subsections – SUBXPAT can then be used to perform a *local* approximation, i.e., an approximation that rewrites only that part of the exact circuit, and thus avoids the combinatorial explosion arising from applying XPAT to the whole circuit. SUBXPAT’s formulation is a generalisation of the formulation described in Section III-B. Similar to XPAT, a miter is used containing the exact circuit and a parametrical template for approximation. The crucial difference is that the gates not contained in the selected subcircuit remain untouched, as illustrated in Figure 4, which decreases the size of the design space and leads to a formula easier to solve. Concretely, function *outAppx* is updated to only consider parameters for the template replacing the selected subcircuit, with the size of the set P being thus decreased. The size of the selected subcircuit dictates both the approximation’s cost and potential benefits, with a selection containing all gates of the exact circuit making SUBXPAT subsume to XPAT.

B. Subcircuit Selection Criteria

A crucial question in subcircuit selection is: given a subcircuit and a tolerated ET, is an approximation of such subcircuit

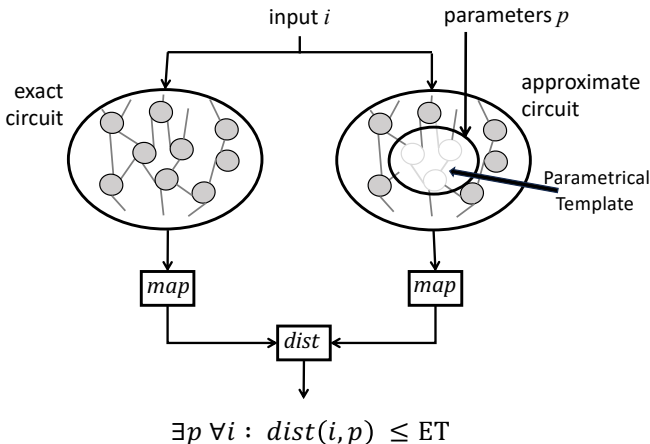


Fig. 4: Miter for SUBXPAT.




Fig. 5: Examples of subcircuits. Subcircuit A drives a single output, and modifying its logic will induce an error of 2^3 . Modifying subcircuit B will induce an error of 2^1 . Quantifying the potential error induced by modifying a multi-output subcircuit, such as C, is non-trivial. Subcircuit D is an example of non-convexity, since its inputs depend on its outputs.

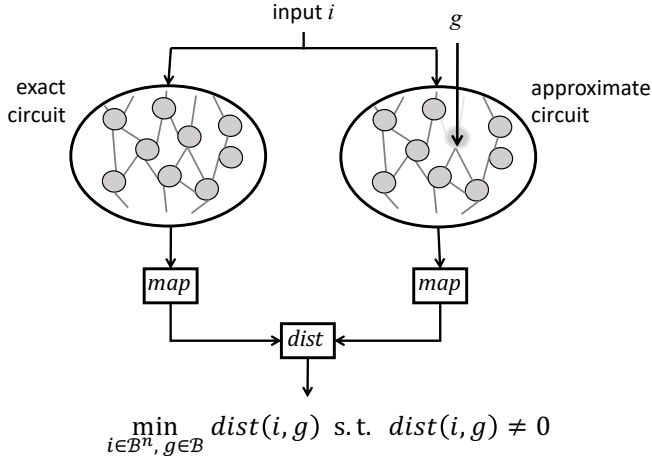
even possible? By this we mean: can the logic of such subcircuit be rewritten in a way that (1) its error is bounded by the ET, and (2) we do not need to resort to semantically-equivalent rewriting? The trivial case of semantically-equivalent rewriting, i.e., a rewriting in which the error is zero for all inputs, is always possible, but is not of interest since the aim of approximation is to *make use* of the tolerated ET to maximise area gains. In practice, we are interested in subcircuits that can be rewritten to make a non-zero error within the ET. Figure 5 illustrates this point. Subcircuit A influences only a single bit of the circuit, with significance 2^3 . Assuming a tolerated error of 4, subcircuit A is not a good choice for approximation, because either it is rewritten in a semantically-equivalent way, and then the error made is 0, or it is rewritten in any other, non-semantically-equivalent, way, and the error made is 16. Hence, no real approximation is possible in this case. We call a subcircuit *feasible*, for a given ET, if a rewriting exists to induce an error larger than 0 and within ET. Subcircuit A is not feasible for an ET of 4, while subcircuit B is.

We are interested in formalising this concept of *feasibility*, in order to only select feasible subcircuits for approximation. To do so, we need to be able to quantify, or somehow bound, the effect that any perturbation of a gate has on the error at the output of a circuit. As we have seen, when a subcircuit only influences a single bit of the output, its effect on error is trivial to define. In the generic case, however, i.e., when a subcircuit has several outputs that have arbitrary downward cones leading to primary outputs, such as subcircuit C, estimating the error caused by a rewriting is non-trivial.

In order to define feasibility, we first define the quantity $mnze_g$ for every gate g in the circuit. This quantity represents the *minimum non-zero error* that can be induced by perturbing in any possible way the logic of *only* gate g . Formally, $mnze_g$ is defined as follows:

$$mnze_g = \min_{i \in \mathbb{B}^n, g \in \mathbb{B}} \text{dist} \left(\text{map}(\text{outExct}(i)), \text{map}(\text{outAppx}(i, g)) \right) \\ \text{s.t. } \text{map}(\text{outExct}(i)) \neq \text{map}(\text{outAppx}(i, g))$$

The $mnze_g$ calculation determines the minimum non-zero distance between the output of the exact circuit, $\text{map}(\text{outExct}(i))$, and the output of an approximate circuit, $\text{map}(\text{outAppx}(i, g))$, where the latter is derived from the

Fig. 6: Miter for $mnze_g$.

exact circuit by perturbing only gate g . This is illustrated in Figure 6; $\text{dist}(i, g) \neq 0$ is ensured by calculating the distance between different points. Note that the minimum is calculated over every input $i \in \mathbb{B}^n$ and every value of $g \in \mathbb{B}$.

We define a *gate* to be feasible w.r.t. a tolerated ET if $mnze_g \leq \text{ET}$. Furthermore, we define a *subcircuit* to be feasible w.r.t. a tolerated ET if *all* its output nodes are feasible. The intuition for $mnze_g$ is given in Figure 7. The calculation of $mnze_g$ for each gate is automated via the use of SMT optimisation [21], on the miter depicted in Figure 6.

Having defined feasibility, we now state our subcircuit selection criteria. We select subcircuits that are (1) feasible, (2) convex, (3) maximal in size, and (4) limited in their number of inputs and outputs. With feasibility already formalised, the rest of the criteria are described and motivated in the following.

Convexity: we must select subcircuits that are convex, in order to prevent the occurrence of cycles in the rewritten circuit. An example of non-convexity is subcircuit D in Figure 5a, where one of the inputs depends on one of the outputs.

Maximality: in order to maximise approximation gain, we select subcircuits that are maximal in size – subject to the constraints given in the following point.

Limited I/O: The number of inputs and outputs of the subcircuit to be rewritten is directly proportional to search space complexity, i.e. to the number of parameters of the resulting template – as discussed in Section III-A. Hence, to achieve scalability, we choose to search for subcircuits within a limited I/O count.

C. Subcircuit Selection Encoding

We now formalise the procedure for subcircuit selection, and detail the encoding that we use. We represent the circuit we want to approximate as a directed graph $G = (V, E, l)$, where V is the set of nodes abstracting the inputs, outputs, and gates in the circuit, $E \subseteq V \times V$ is the set of edges abstracting the circuit's wires, and $l : V \rightarrow \mathbb{N}^+$ is the partial function labelling vertices with their $mnze_g$; function l is partial to improve labelling efficiency (cf. Section VI-B). An SMT formula is used to encode the subgraph selection

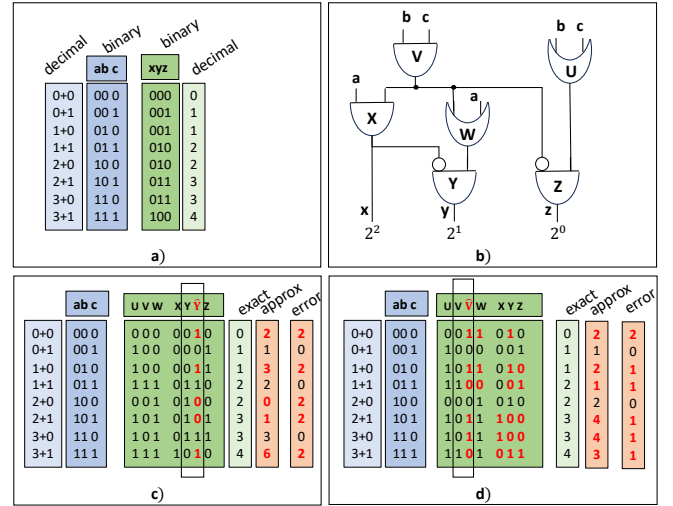


Fig. 7: Motivation for $mnze_g$ calculation. (a) Truth table for an adder of 2 plus 1 bits (3 inputs and 3 outputs), in both binary and decimal. (b) A circuit implementing the truth table. (c) Perturbing the output of gate Y in a specific way (called here \hat{Y}), makes the error at every truth table line be 2 or 0. Considering *every* possible perturbation, which can be done via the miter in Figure 6, makes the value of $mnze_Y$ be derived to 2. (d) The effect of perturbing gate V propagates further through the circuit – a specific perturbation called \hat{V} is shown here, for which the error at every truth table line is 2, 1, or 0; the value of $mnze_V$, considering every possible perturbation, can be derived via the miter in Figure 6 to be 1.

problem. For each node $v \in V$, a Boolean literal x_v is defined, whose value indicates if v is in the subgraph or not. The objective is to maximise the number of literals assigned to *true* while respecting selection constraints. Concretely, we ask the solver for a satisfying assignment to the following formula:

$$\text{maximise } |\{x_v \mid v \in V \wedge x_v \wedge \phi\}|$$

The selection constraints are encapsulated into ϕ as follows:

$$\phi = \text{convex} \wedge \text{limitedIO} \wedge \text{allOutputsFeasible}$$

To aid in the definition of ϕ , we first define constraints that capture if an edge is an input or an output to the subgraph. Concretely, for every edge $e \in E$ we define the following:

$$e_{s,d}^{in} = \neg x_s \wedge x_d \quad e_{s,d}^{out} = x_s \wedge \neg x_d$$

where $e_{s,d}^{in}$ indicates that the source of the edge is not included in the subgraph while its destination is, and $e_{s,d}^{out}$ indicates the opposite scenario. We then define the sets of input and output nodes of the subgraph, I and O , as follows:

$$I = \{d \mid (s, d) \in E \wedge e_{s,d}^{in}\}$$

$$O = \{s \mid (s, d) \in E \wedge e_{s,d}^{out}\}$$

For a subgraph $H = (V_H, E_H, l_H)$ of G , we need to ensure convexity, limitation on the number of inputs and outputs, and feasibility of all outputs. Subgraph H is convex w.r.t. to G if, for every pair of nodes $s, d \in V_H$, all paths $\rho_G(s, d)$ are

contained in E_H . A path in G between nodes s and d , denoted $\rho_G(s, d)$, is a sequence of edges $(s, w_1), (w_1, w_2), \dots, (w_n, d)$. To ensure convexity, we enforce that all paths leading to an input node must not come from V_H . Concretely:

$$\text{convex} = (\forall s \in V. (\exists d \in I. \rho(s, d)) \implies \neg x_s)$$

Given constants I_{max} and O_{max} dictating the maximum number of inputs and outputs desired, the limitation constraint is defined as follows:

$$\text{limitedIO} = (|I| \leq I_{max} \wedge |O| \leq O_{max})$$

Lastly, the feasibility constraint is straightforward, as follows:

$$\text{allOutputsFeasible} = (\forall o \in O. l(o) \leq ET)$$

V. ITERATIVE SUBXPAT

With SUBXPAT we are able to scale XPAT to larger circuits, making it applicable in practice, but limit the potential gains by considering only one subcircuit of the exact circuit. Given that SUBXPAT leaves the parts of the circuit not selected for approximation untouched, we can apply it iteratively, at each step selecting and approximating a different subcircuit, thus achieving scalable and effective approximation.

In pursuing iterative SUBXPAT, which we refer to simply as SUBXPAT, the main concern is which ET to use at each iteration, i.e. how much error to be *consumed* at each iteration. The naive approach is to consider the whole ET at each iteration. This, however, tends to yield poor results, because subcircuit selection can choose in the first iteration a subcircuit feeding primary outputs of high significance and thus severely hinder approximation possibilities in the following iterations. A better approach, which we choose to follow, is to start with a small fraction of the ET at the first iteration and monotonically increase it at each iteration. This effectively guides subcircuit selection to choose subcircuits feeding primary outputs of low, but increasing, significance, allowing for a higher percentage of the whole circuit to be approximated.

Following a monotonic approach enables many ET increase strategies. These can be, for instance, linear or exponential increases at each iteration, or a combination of both. SUBXPAT allows for fine-tuning of this and many other of its elements, as discussed in the next section.

VI. FINE-TUNING SUBXPAT

The SUBXPAT framework is very modular, accepting different strategies to accomplish its many steps. Here we discuss strategies relating to SMT encoding, in Section VI-A, gate labelling, in Section VI-B, subcircuit selection, in Section VI-C, and iterative ET increase, in Section VI-D.

A. SMT Encoding

How a problem is encoded into an SMT formula can have a significant effect on the solver's performance. Here we describe two encoding choices that led to order of magnitude performance improvement to two SUBXPAT steps.

While the structure of the double quantified formula defined in Section III-B is inherent to (SUB)XPAT, the encoding of

some of its functions can be done in different ways. Functions *map* and *dist* are defined over the sets S and $Dist$, which can be instantiated using different SMT theories. The more natural instantiation is to substitute both sets by \mathbb{Z} , yielding formulas over the theory of linear integer arithmetics (LIA). Although intuitive, this encoding requires the solver to engage in the expensive interplay between its SAT and theory engines, required by the conflict-driven clause learning with theory propagation approach followed by modern solvers [22]. Since we are dealing with circuits, an alternative is to substitute S and $Dist$ by \mathbb{B}^m , yielding instead formulas over the theory of fixed-size bit-vectors (BV). Since bit-vectors are very close to plain Boolean values, the solver will flatten, i.e., bit-blast, the formula and forward it to its SAT engine, leading to a computationally cheaper solving that does not need to engage with a theory engine.

The graph-based constraints described in Section IV-C can also be handled in different ways. The naive approach is to encode all nodes and edges as plain Boolean values. This, however, makes the solver need to infer the relations between the constraints during solving. Alternatively, abstract data types (ADTs) can be employed, which directly provide the solver the structural information inherent to the problem and prevent the need for solving inference.

B. Partial Labelling

As stated in Section IV-C, we label each gate with the minimum non-zero error that can be induced by modifying the logic of (only) said gate. An important aspect of SUBXPAT's labelling is that, to reduce computation time, we do not label all the circuit's gates, as this can become a bottleneck when applying SUBXPAT to large circuits. Instead, we only label a relevant subset of gates at each iteration. This is realised by first labelling the gates feeding the primary outputs, and then recursively labelling their predecessors *as long as their label is lower than the given ET*. Stopping when the ET is reached prevents wasting computation on labelling gates that will never be selected for approximation. Lastly, since the labelling of one node is independent of the labelling of other nodes, parallelising the process is a straightforward way of improving performance.

C. Subcircuit Selection Feasibility Criteria

While we found that having all outputs be feasible is a good subcircuit selection strategy in general, SUBXPAT allows for different feasibility criteria. Requiring only a subset of outputs to be feasible – as opposed to all – is a straightforward relaxation of our chosen strategy. Conversely, having many outputs be jointly, rather than individually, feasible, i.e., the sum of their labels respects the ET, is a stricter approach. The flexibility given by this spectrum enables SUBXPAT to be customised for specific use cases if need be. Additionally, soft SMT constraints can potentially be used to establish optional selection objectives, e.g., to allow some outputs to be unfeasible but require their labels to be as low as possible, which enables further customisation based on user needs.

TABLE I: Circuits used in SUBXPAT’s evaluation.

Benchmark	Description	Number of gates
adder i8 o5	4-bit adder	52
adder i12 o7	6-bit adder	82
adder i16 o9	8-bit adder	115
adder i20 o11	10-bit adder	145
adder i24 o13	12-bit adder	182
adder i28 o15	14-bit adder	223
adder i32 o17	16-bit adder	263
absDiff i8 o4	4-bit absolute difference	96
absDiff i12 o6	6-bit absolute difference	167
absDiff i16 o8	8-bit absolute difference	191
absDiff i20 o10	10-bit absolute difference	270
absDiff i24 o12	12-bit absolute difference	427
absDiff i28 o14	14-bit absolute difference	441
absDiff i32 o16	16-bit absolute difference	505
mul i8 o8	4-bit multiplier	163
mul i10 o10	5-bit multiplier	283
mul i12 o12	6-bit multiplier	434
mul i14 o14	7-bit multiplier	618
mul i16 o16	8-bit multiplier	865
mAdd i9 o6	3-bit multiply-and-add	118
mAdd i12 o8	4-bit multiply-and-add	227
mAdd i15 o10	5-bit multiply-and-add	370
mAdd i18 o12	6-bit multiply-and-add	554

D. ET Increase

We discussed in Section V the fact that (1) a different ET value can be chosen at every iteration, and that (2) we choose to select monotonically increasing values. Still, different increase strategies can favour different types of circuits and of selection criteria. We opt for a simple linear increase based on fractions of the whole ET, as we have observed that this leads to the best results overall. In specific cases, however, a combination of large and small increases might lead to better results, since this allows SUBXPAT to utilise approximation potential leftover after a large step. Large and small steps can be either linear or exponential, with different combinations being worth exploring in different contexts.

VII. EVALUATION

We carried out an extensive evaluation of SUBXPAT and compared it against three state-of-the-art methods. We first detail our experimental setup, in Section VII-A, and then present our results related to XPAT, in Section VII-B, and SUBXPAT, in Section VII-C.

A. Experimental Setup

We implemented SUBXPAT in a tool of the same name. The tool¹ is open-source under the MIT licence and contains the whole machinery of SUBXPAT and different strategies for each of its steps. It uses Z3 [23] to solve SMT queries and Yosys [24] as a synthesiser. Besides providing off-the-shelf push-button support for ALS, SUBXPAT’s modularity also allows users to implement custom strategies for each of its steps and to potentially use different solvers and synthesisers.

To evaluate SUBXPAT we used 23 Verilog specifications of arithmetics circuits, listed in Table I. We compared our approach, in terms of circuit area and tool runtime, to




Fig. 8: Areas of approximations of a 2-bit adder.

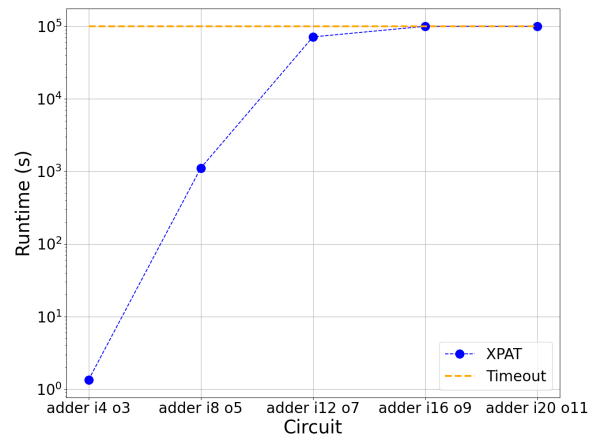


Fig. 9: XPAT runtime results for adders of different size.

that of MUSCAT, MECALS, and EVOAPPROX; we used $dist(s_1, s_2) = |s_1 - s_2|$ for all our miters. All the experiments were carried out on a Linux machine with a 3.30 GHz Intel Core i9 CPU and 256 GB of RAM.

B. XPAT Results

Our prior experiments [9] showed that XPAT yields good results for the benchmarks it can handle, but also showed that it has limited scalability. To highlight this point, we performed a study using adder circuits. Figure 8 shows the solutions found by XPAT, MECALS, and MUSCAT when trying to approximate a 2-bit adder with different ET values; EVOAPPROX solutions are not shown here (and also later for some other benchmarks) because it only has approximations for specific circuits². As can be seen, XPAT yields better solutions than the compared tools. This, however, cannot be maintained for larger benchmarks, since XPAT’s runtime increases exponentially with circuit size and quickly becomes intractable, as shown in Figure 9. The good solutions obtained with XPAT, coupled with its non-scalability, clearly motivate the need for SUBXPAT.

¹See <https://github.com/LP-RG/subxpat>.

²See <https://github.com/ehw-fit/evoapproxlib>.

TABLE II: SUBXPAT results for three different I/O limits for subcircuit selection; runtime is given in minutes.

Benchmark	ADRS per I/O			Runtime per I/O		
	2/1	3/2	6/3	2-1	3-2	6-3
adder i8 o5	0.17	0.15	0.14	1.08	1.06	3.09
adder i12 o7	1.62	0.29	0.00	1.79	2.09	13.85
adder i16 o9	0.34	0.03	0.08	2.89	3.94	233.55
adder i20 o11	0.40	0.05	0.02	4.79	7.33	35.67
adder i24 o13	0.49	0.31	0.01	10.10	8.35	102.46
adder i28 o15	0.38	0.08	0.03	10.81	11.33	160.80
adder i32 o17	1.03	0.12	0.15	18.62	16.38	61.38
absDiff i8 o4	0.13	0.08	0.00	1.10	0.33	5.60
absDiff i12 o6	0.98	0.70	0.04	2.33	1.69	22.48
absDiff i16 o8	0.48	0.69	0.00	4.03	2.70	80.45
absDiff i20 o10	0.90	0.67	0.01	7.67	4.18	106.62
absDiff i24 o12	0.24	0.67	0.02	11.58	8.32	177.18
absDiff i28 o14	0.93	0.24	0.05	22.54	14.99	349.99
absDiff i32 o16	0.49	0.46	0.01	37.42	22.85	519.96
mul i8 o8	0.25	0.17	0.01	3.14	8.35	31.21
mul i10 o10	0.23	0.09	0.02	9.78	21.85	94.66
mul i12 o12	1.46	0.16	0.00	49.01	79.83	308.53
mul i14 o14	8.59	0.06	0.07	107.38	210.10	821.64
mul i16 o16	15.45	0.16	0.02	446.30	581.48	1599.42
mAdd i9 o6	0.44	0.21	0.07	2.16	2.71	12.23
mAdd i12 o8	0.38	0.18	0.01	6.36	10.23	81.40
mAdd i15 o10	0.75	0.08	0.07	23.88	23.88	140.24
mAdd i18 o12	1.48	0.07	0.04	67.56	103.67	382.28
Average	1.63	0.25	0.04	37.06	49.90	232.38

C. SubXPAT Results

To evaluate SUBXPAT we first performed two calibration studies to set its parameters, and then compared it against state-of-the-art methods. Calibration was done in terms of both subcircuit dimensions and the resetting of constants, as will be explained shortly.

We used the Average Distance from Reference Set (ADRS) as a comparison metric in our experiments, which measures the distance between the Pareto curve obtained by one particular approach and the global Pareto curve, i.e., the curve with the best results among all approaches considered. Having area and ET as our axes, the ADRS is formally defined as follows:

$$\text{ADRS}(\bar{P}, P) = \left[\frac{1}{|\bar{P}|} \sum_{\bar{p} \in \bar{P}} \min_{p \in P} \delta(\bar{p}, p) \right] \times 100$$

$$\delta(\bar{p}, p) = \max \left\{ 0, \frac{A(\bar{p}) - A(p)}{A(p)}, \frac{E(\bar{p}) - E(p)}{E(p)} \right\}$$

where \bar{P} and P are sets containing the points of the particular and global Pareto curves, and A and E are functions which extract the area and ET value of a given point. A low ADRS score is desirable, as it indicates that the curve of \bar{P} is close to the curve of P .

1) *Subcircuit Dimensions*: To assess the effect of subcircuit dimensions in solution quality and in runtime, we evaluated our tool with different limits on the number of subcircuit inputs and outputs. We considered input/output (I/O) values of 2/1, 3/2, and 6/3. The results obtained are summarised in Table II. As can be seen, our results indicate that larger I/O limits lead to both better approximations and higher (yet tractable) runtimes. Solutions for a 16-bit adder can be seen in Figure 10. For the following experiments we always use an I/O of 6/3.

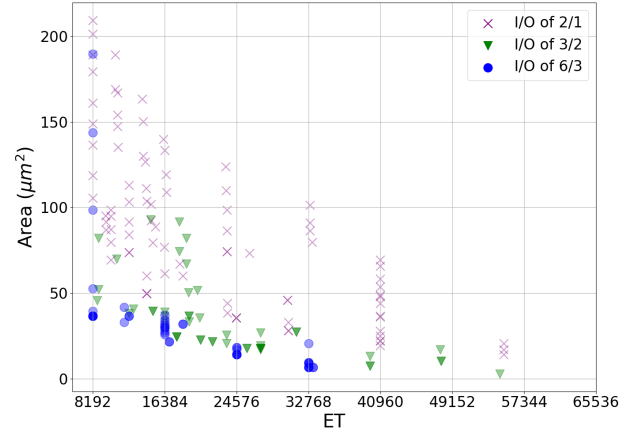


Fig. 10: Areas of SUBXPAT approximations of a 16-bit adder obtained with different I/O limits for subcircuit selection.

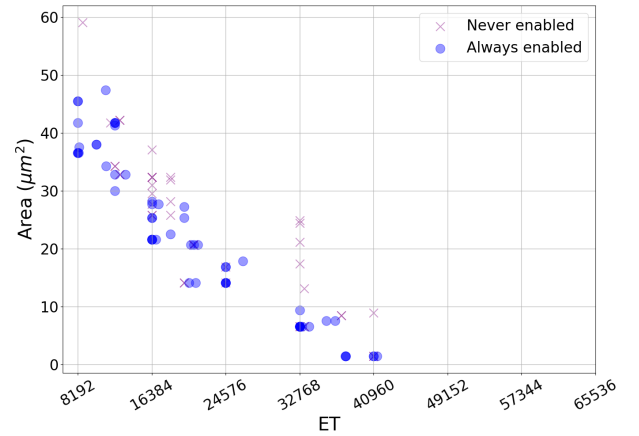


Fig. 11: Areas of SUBXPAT approximations of a 16-bit adder obtained with different constant resetting modes.

2) *Resetting of Constants*: During the execution of its many iterations, SUBXPAT can set some primary outputs to be constant values. By default, this is a final decision and cannot be changed in following iterations. However, to enable greater flexibility, we consider allowing the resetting of such constants at each iteration, i.e., primary outputs previously set to 1 could be changed to 0 and vice versa. To assess the effect of this change, we evaluated our tool in two modes: one in which resetting is never allowed and one in which it is always allowed. The results obtained are summarised in Table III. As can be seen, our results indicate that allowing for constants to be reset improves solution quality at the cost of increased runtime. Solutions for a 16-bit adder can be seen in Figure 11. Note that we achieve the resetting of constants seamlessly and automatically, by having the appropriate primary outputs be free variables in our SMT encoding. For the following experiments we always allow for the resetting of constants.

3) *Comparison Against the State of the Art*: Having observed that an I/O of 6/3 combined with constant resetting leads to the best results with SUBXPAT while maintaining tractable runtimes, we now compare SUBXPAT against the state of the art. The results obtained are summarised in

TABLE III: SUBXPAT results for two constant resetting modes; runtime is given in minutes.

Benchmark	ADRS per mode		Runtime per mode	
	Never	Always	Never	Always
adder i8 o5	0.15	0.08	2.20	0.29
adder i12 o7	0.12	0.00	10.42	8.20
adder i16 o9	0.33	0.02	17.13	36.77
adder i20 o11	0.04	0.01	32.02	99.16
adder i24 o13	0.06	0.03	44.99	488.68
adder i28 o15	0.06	0.02	186.71	653.83
adder i32 o17	0.06	0.00	64.62	413.94
absDiff i8 o4	0.06	0.00	3.48	1.49
absDiff i12 o6	0.03	0.11	21.96	6.30
absDiff i16 o8	0.10	0.01	66.71	81.82
absDiff i20 o10	0.06	0.06	91.10	134.48
absDiff i24 o12	0.08	0.05	148.90	423.64
absDiff i28 o14	0.09	0.04	239.56	437.05
absDiff i32 o16	0.10	0.00	576.92	1062.99
mul i8 o8	0.02	0.09	22.49	24.73
mul i10 o10	0.07	0.04	76.33	156.78
mul i12 o12	0.03	0.03	264.51	371.78
mul i14 o14	0.05	0.02	982.47	976.01
mul i16 o16	0.11	0.02	2157.67	2461.48
mAdd i9 o6	0.17	0.00	15.62	13.28
mAdd i12 o8	0.06	0.03	66.83	127.82
mAdd i15 o10	0.06	0.02	192.24	303.63
mAdd i18 o12	0.04	0.02	369.43	606.42
Average	0.08	0.03	245.84	386.55

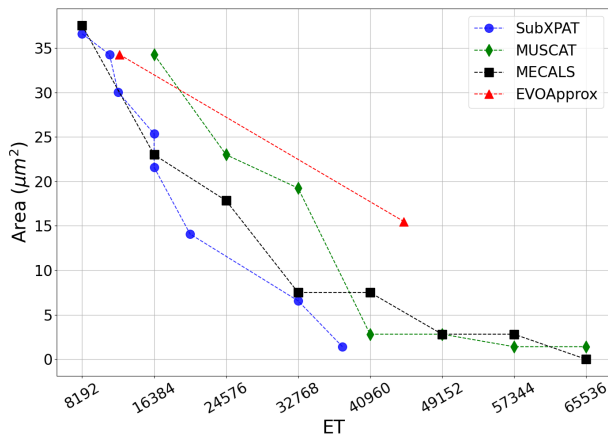


Fig. 12: Areas of approximations of a 16-bit adder.

Table IV. As can be seen, SUBXPAT yields better quality solutions in most cases, sometimes having an ADRS score of less than half than that of the other methods. In terms of runtime, SUBXPAT stands between MECALS and MUSCAT on average; runtimes for EVOAPPROX are not available². The area of the approximated circuits generated by different methods, with different ET values, for a 16-bit adder can be seen in Figure 12.

VIII. CONCLUSIONS

We proposed SUBXPAT, a novel iterative approach for ALS that is both effective and scalable. It employs an SMT solver to search for good approximations of a given exact circuit – operating iteratively on subcircuits of the original circuit, and achieving significant area savings while maintaining scalability. Our evaluation, conducted over a broad range

of arithmetic circuits, shows that SUBXPAT can yield better quality approximations than state-of-the-art methods while having reasonable runtime.

Avenues for future work include investigating (1) the use of alternative parametrical templates, e.g., templates with more than two layers of logic, (2) the formulation of the design space exploration in new ways, e.g., via quantified Boolean formulas (QBF), and (3) improving the runtime of local approximations, e.g., by attempting to avoid the use double quantification in a sound manner.

REFERENCES

- [1] Q. Xu, T. Mytkowicz, and N. Kim, “Approximate Computing: A Survey,” *IEEE Design and Test*, vol. 33, no. 1, pp. 8–22, 2016.
- [2] I. Scarabottolo, G. Ansaloni, G. A. Constantinides, L. Pozzi, and S. Reda, “Approximate Logic Synthesis: A Survey,” *Proceedings of the IEEE*, vol. 108, no. 12, pp. 2195–2213, 2020.
- [3] S. Venkataramani, K. Roy, and A. Raghunathan, “Substitute-and-Simplify: A Unified Design Paradigm for Approximate and Quality Configurable Circuits,” in *Proceedings of the 16th Design, Automation and Test in Europe Conference*, 2013, pp. 1367–1372.
- [4] V. Mrazek, R. Hrbacek, Z. Vasicek, and L. Sekanina, “EvoApprox8b: Library of Approximate Adders and Multipliers for Circuit Design and Benchmarking of Approximation Methods,” in *Proceedings of the 20th Design, Automation and Test in Europe Conference*, 2017, pp. 258–261.
- [5] I. Scarabottolo, G. Ansaloni, and L. Pozzi, “Circuit Carving: A Methodology for the Design of Approximate Hardware,” in *Proceedings of the 21st Design, Automation and Test in Europe Conference*, 2018, pp. 545–550.
- [6] S. Hashemi, H. Tann, and S. Reda, “BLASYS: Approximate Logic Synthesis Using Boolean Matrix Factorization,” in *Proceedings of the 55th Design Automation Conference*, 2018, pp. 1–6.
- [7] L. Witschen, T. Wiersema, M. Artmann, and M. Platzner, “MUSCAT: MUS-based Circuit Approximation Technique,” in *Proceedings of the 25th Design, Automation and Test in Europe Conference*, 2022, pp. 172–177.
- [8] C. Meng, J. Sun, Y. Mai, and W. Qian, “MECALs: A Maximum Error Checking Technique for Approximate Logic Synthesis,” in *Proceedings of the 26th Design, Automation and Test in Europe Conference*, 2023, pp. 1–6.
- [9] M. Rezaalipour, M. Biasion, I. Scarabottolo, G. A. Constantinides, and L. Pozzi, “A Parametrizable Template for Approximate Logic Synthesis,” in *Proceedings of the 53rd Annual IEEE/IFIP International Conference on Dependable Systems and Networks Workshops*, 2023, pp. 175–178.
- [10] J. Schlachter, V. Camus, K. V. Palem, and C. Enz, “Design and Applications of Approximate Circuits by Gate-Level Pruning,” *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 25, no. 5, pp. 1694–1702, 2017.
- [11] S. Hashemi and S. Reda, “Generalized Matrix Factorization Techniques for Approximate Logic Synthesis,” in *Proceedings of the 22nd Design, Automation and Test in Europe Conference*, 2019, pp. 1289–1292.
- [12] J. Ma, S. Hashemi, and S. Reda, “Approximate Logic Synthesis Using Boolean Matrix Factorization,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 41, no. 1, pp. 15–28, 2022.
- [13] V. Mrazek, M. A. Hanif, Z. Vasicek, L. Sekanina, and M. Shafique, “autoAX: An Automatic Design Space Exploration and Circuit Building Methodology utilizing Libraries of Approximate Components,” in *Proceedings of the 56th Design Automation Conference*, 2019, pp. 1–6.
- [14] M. Soeken, D. Große, A. Chandrasekharan, and R. Drechsler, “BDD Minimization for Approximate Computing,” in *Proceedings of the 21st Asia and South Pacific Design Automation Conference*, 2016, pp. 474–479.
- [15] C. Meng, W. Qian, and A. Mishchenko, “ALSRAC: Approximate Logic Synthesis by Resubstitution with Approximate Care Set,” in *Proceedings of the 57th Design Automation Conference*, 2020, pp. 1–6.
- [16] S. Fröhlich, D. Große, and R. Drechsler, “Error Bounded Exact BDD Minimization in Approximate Computing,” in *Proceedings of the 47th International Symposium on Multiple-Valued Logic*, 2017, pp. 254–259.


TABLE IV: Comparison between SUBXPAT and three state of the art methods; runtime is given in minutes.

Benchmark	ADRS per method				Runtime per method			
	SUBXPAT	MUSCAT	MECALs	EVOAPPROX	SUBXPAT	MUSCAT	MECALs	EVOAPPROX
adder i8 o5	0.07	1.39	0.13	-	0.47	2.12	0.02	-
adder i12 o7	0.00	1.41	0.29	-	8.7	55.53	0.11	-
adder i16 o9	0.00	0.41	0.18	0.30	37.48	827.33	0.71	N/A
adder i20 o11	0.00	0.42	0.09	-	99.95	830.69	0.89	-
adder i24 o13	0.02	0.46	0.14	0.31	489.57	138.04	2.13	N/A
adder i28 o15	0.00	0.70	0.13	-	655.2	138.83	4.34	-
adder i32 o17	0.00	0.44	0.15	1.55	415.77	728.32	31.52	N/A
absDiff i8 o4	0.02	0.23	0.16	-	1.86	40.95	0.03	-
absDiff i12 o6	0.06	0.36	0.09	-	7.43	597.67	0.23	-
absDiff i16 o8	0.18	0.78	0.16	-	84.61	831.47	2.38	-
absDiff i20 o10	0.14	0.50	0.09	-	137.51	832.53	6.57	-
absDiff i24 o12	0.04	0.48	0.02	-	438.07	141.38	15.45	-
absDiff i28 o14	0.17	0.73	0.11	-	475.25	142.98	29.64	-
absDiff i32 o16	0.03	0.51	0.03	-	1130.27	143.54	67.51	-
mul i8 o8	0.03	0.68	0.18	-	25.77	831.14	1.03	-
mul i10 o10	0.03	0.41	0.15	-	158.74	832.9	14.07	-
mul i12 o12	0.03	0.84	0.14	-	375.28	835.96	46.88	-
mul i14 o14	0.04	0.76	0.09	1.66	990.48	146.62	289.83	N/A
mul i16 o16	0.01	0.96	0.14	-	2704.7	155.34	621.63	-
mAdd i9 o6	0.01	0.53	0.18	-	14.05	728.23	0.12	-
mAdd i12 o8	0.01	0.71	0.23	-	129.51	832.31	4.65	-
mAdd i15 o10	0.01	1.24	0.18	-	308.83	140.02	19.57	-
mAdd i18 o12	0.02	0.58	0.14	-	615.53	144.22	71.31	-
Average	0.04	0.68	0.14	0.96	404.57	439.05	53.51	N/A


- [17] A. Lingamneni, C. Enz, K. Palem, and C. Piguet, "Synthesizing Parsimonious Inexact Circuits Through Probabilistic Design Techniques," *ACM Transactions on Embedded Computing Systems*, vol. 12, no. 2s, pp. 1–26, 2013.
- [18] G. Liu and Z. Zhang, "Statistically Certified Approximate Logic Synthesis," in *Proceedings of the 36th International Conference on Computer-Aided Design*, 2017, pp. 344–351.
- [19] C. Meng, H. Wang, Y. Mai, W. Qian, and G. D. Micheli, "VACSEM: Verifying Average Errors in Approximate Circuits Using Simulation-Enhanced Model Counting," in *Proceedings of the 27th Design, Automation and Test in Europe Conference*, 2024, pp. 1–6.
- [20] C. Meng, Z. Zhou, Y. Yao, S. Huang, Y. Chen, and W. Qian, "HEDALS: Highly Efficient Delay-Driven Approximate Logic Synthesis," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 42, no. 11, pp. 3491–3504, 2023.
- [21] N. Bjørner, A.-D. Phan, and L. Fleckenstein, "vZ - An Optimizing SMT Solver," in *Proceedings of the 21st International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, 2015, pp. 194–199.
- [22] D. Kroening and O. Strichman, *Decision Procedures - An Algorithmic Point of View*. Springer Berlin, Heidelberg, 2016.
- [23] L. de Moura and N. Bjørner, "Z3: An Efficient SMT Solver," in *Proceedings of the 14th International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, 2008, p. 337–340.
- [24] C. Wolf, "Yosys Open SYnthesis Suite," <https://yosyshq.net/yosys>.




Marco Biasion received his BSc degree in informatics from the Università della Svizzera italiana, Switzerland, in 2023, where he is currently pursuing a MSc degree in informatics with specialisation in computer systems. His research interests include approximate logic synthesis, reconfigurable architectures, and hardware optimisation.




Francesco Costa received his BSc degree in informatics from the Università della Svizzera italiana, Switzerland, in 2023, where he is currently pursuing a MSc degree in computational science. His research interests include computer simulations, computational biology, and bioinformatics.



Morteza Rezaalipour received his MSc degree in computer architecture from the K. N. Toosi University of Technology, Iran, in 2019. He is a team leader at MicroLab in the Department of Computer Engineering at K. N. Toosi University of Technology and is currently pursuing a PhD degree in informatics at the Università della Svizzera italiana, Switzerland. His research interests include approximate logic synthesis, low-power design, and memory systems.




Cristian Tirelli received his MSc degree in computer engineering from the University of Cassino and Southern Lazio, Italy, in 2020. He is currently pursuing a PhD degree in informatics at the Università della Svizzera italiana, Switzerland. He was awarded an SNSF Mobility Grant that allowed him to stay one year (2022/23) as a visiting researcher at the University of California at Los Angeles, USA. His research interests include compiler methods, hardware accelerators, and reconfigurable architectures.




Rodrigo Otoni received his PhD degree in informatics from the Università della Svizzera italiana, Switzerland, in 2023, where he currently holds a position as a Postdoctoral Researcher. His background is in formal methods and his research at present focuses on automated reasoning in the contexts of verification, synthesis, and certification.

Dr. Otoni recurrently serves as a sub-reviewer for conferences such as FMCAD, TACAS, and CAV, and served as a program committee member for the FMCAD'24 Student Forum. He was recently awarded a seed grant from the Hasler Foundation for his independent research.




Lorenzo Ferretti received his PhD degree in informatics from the Università della Svizzera italiana, Switzerland, in 2020, where he then held a position as a Postdoctoral Researcher until 2021. After being awarded an SNSF Postdoc Mobility Fellowship he then joined the University of California at Los Angeles, USA. Currently, he is a Principal AI Architect in the Silicon System AI group at Micron Technology. His research interests include high-level synthesis, graph-learning applied to hardware optimisation, and generative models applied to test generation.

Dr. Ferretti serves as a reviewer for different IEEE conferences and scientific journals and is currently part of the technical program committee of the Design, Automation and Test in Europe Conference.



George A. Constantinides received his PhD degree in electrical and electronic engineering from the Imperial College London, UK, in 2001. Since 2002, he has been with the faculty of Imperial College London, where he is currently a Professor of Digital Computation and the Head of the Circuits and Systems Research Group.

Prof. Constantinides serves on several program committees and has published over 200 research papers in peer-refereed journals and international conferences. He was the General Chair of the ACM/SIGDA International Symposium on Field-Programmable Gate Arrays in 2015 and is a Fellow of the British Computer Society.



Laura Pozzi received her PhD degree in computer engineering from the Politecnico di Milano, Italy, in 2000. She is currently a Professor with the Faculty of Informatics, Università della Svizzera Italiana, Switzerland. She was a Postdoctoral Researcher with the École Polytechnique Fédérale de Lausanne, Switzerland; a Research Engineer with STMicroelectronics, USA; and an Industrial Visitor with the University of California at Berkeley, USA. Her current research interests include automating embedded processor customization, high-performance

compiler techniques, innovative reconfigurable fabrics, high-level synthesis design space exploration, and approximate computing.

Prof. Pozzi has served as an Associate Editor for the IEEE Transactions on Computer-Aided Design and the IEEE Design and Test, and is or has been on the technical program committee of several international conferences in the areas of compilers and architectures for embedded systems.