

PATH: Evaluation of Boolean Logic Using Path-Based In-Memory Computing Systems

Sven Thijssen^{1b}, Muhammad Rashedul Haq Rashed^{1b}, Sumit Kumar Jha, *Senior Member, IEEE*,
and Rickard Ewetz^{1b}, *Member, IEEE*

Abstract—In-memory computing using nonvolatile memory is a promising pathway to accelerate data-intensive applications. While substantial research efforts have been dedicated to executing Boolean logic using digital in-memory computing, the limitation of state-of-the-art paradigms is that they heavily rely on repeatedly switching the state of the nonvolatile resistive devices using expensive WRITE operations. In this article, we propose a new in-memory computing paradigm called path-based computing (PATH) for evaluating Boolean logic. Computation within the paradigm is performed using a one-time expensive compilation phase and a fast and efficient evaluation phase. The key property of the paradigm is that the execution phase only involves cheap READ operations. First, we define an analogy between binary decision diagrams (BDDs) and one-transistor one-memristor (1T1M) crossbars that allows Boolean functions to be mapped into crossbar designs. When such crossbar design becomes too large to be physically realizable, we propose to synthesize the Boolean function into a PATH system. A PATH system consists of a topology of staircase structures. A staircase structure is a cascade of hardwired crossbars, which minimizes intercrossbar communication. We evaluate the proposed paradigm using ten circuits from the Revlib benchmark suite, eight control circuits of the EPFL benchmark suite, and eight ISCAS85 benchmarks. Compared with state-of-the-art digital in-memory computing paradigms, PATH improves energy and latency with 1006× and 10× on average, respectively.

Index Terms—Design automation, logic design, reconfigurable logic.

I. INTRODUCTION

THE GROWTH of digital data accelerates at a high pace. In 2025, the total amount of digital data is expected to be 175ZB [1]. This growth is driven by a variety of factors, one being the collection of sensor data using IoT devices [2]. The development of 5G and 6G networks will only accelerate the amassment of this data further [3]. Another contributing factor is the emergence of data-driven technologies, such as

Manuscript received 17 March 2023; revised 28 June 2023 and 5 October 2023; accepted 25 November 2023. Date of publication 19 December 2023; date of current version 23 April 2024. This work was supported in part by NSF under Grant 2319399, Grant 2404036, and Grant 2408925. This article was recommended by Associate Editor A. Gamatie. (*Corresponding author: Sven Thijssen.*)

Sven Thijssen is with the Department of Computer Science, University of Central Florida, Orlando, FL 32816 USA (e-mail: sven.thijssen@ucf.edu).

Muhammad Rashedul Haq Rashed and Rickard Ewetz are with the Department of Electrical and Computer Engineering, University of Central Florida, Orlando, FL 32816 USA (e-mail: muhammad.rashed@ucf.edu; rickard.ewetz@ucf.edu).

Sumit Kumar Jha is with the Department of Computer Science, Florida International University, Miami, FL 33199 USA (e-mail: jha@cs.fiu.edu).

Digital Object Identifier 10.1109/TCAD.2023.3344523

TABLE I
COMPARISON OF IN-MEMORY LOGIC STYLES IN TERMS OF UNDERLYING OPERATION AND EVALUATED LOGIC COMPLEXITY

Digital logic style	Representative Studies	Operations in each phase	
		Compile	Execute
IMPLY	[9], [10]	WRITE	WRITE+READ
MAGIC	[11], [12]	WRITE	WRITE+READ
MAJORITY	[13], [14]	WRITE	WRITE+READ
FLOW	[15], [16]	WRITE	WRITE+READ
Path-based	(this paper)	WRITE	READ

deep neural networks [4] and foundational AI models, which require Internet-scale amounts of digital data for unsupervised pretraining [5]. Unfortunately, these data-intensive techniques suffer from the Von Neumann bottleneck [6]. The bottleneck denotes the energy-inefficiency of a bus to transfer data between a computer's memory and computing units. Several other factors, such as the End of Moore's law [7] and the End of Dennard Scaling [8], are challenging the performance of these data-intensive applications.

Processing in-memory using nonvolatile memory has recently attracted significant attention to mitigate the aforementioned limitations [17]. Nonvolatile memory technology includes memristors, resistive random access memory (ReRAM) [18], phase change memory (PCM) [19], and spin-transfer torque magnetic random access memory (STT-MRAM) [20]. Analog in-memory computing is well-known for performing matrix-vector multiplication at high speed and with low-energy consumption. These computations are carried out in dense crossbar arrays. Unfortunately, analog in-memory computing is limited to matrix-vector multiplication, and related arithmetic operations [21]. Some efforts have been made to improve accuracy while maintaining these energy and latency advantages [22]. Unfortunately, despite these efforts, analog in-memory computing cannot deliver the deterministic precision required for high-assurance applications. However, digital computing is more robust due to the clear distinct states for a logical zero and one [23]. For comprehensive reviews on in-memory computing, we refer to [24], [25], and [26].

Several noteworthy digital in-memory computing paradigms are IMPLY [9], MAGIC [11], MAJORITY [13], and FLOW [15]. These in-memory computing paradigms more or less have the following in common: the paradigms consist of two broad phases. First, there is a one-time compilation phase and, second, an execution phase that is performed for each function input. In Table I, we show the READ and WRITE operations performed in each phase for the different logic styles. It can be observed that all previous paradigms use WRITE operations in the execution phase.

WRITE operations are orders of magnitude more expensive than READ operations [27]. Further, WRITE operations are detrimental to the endurance of the memristor's lifetime [28]. In contrast, the proposed path-based computing (PATH) paradigm evaluates Boolean logic using READ operations in the execution phase, mitigating the high-energy consumption for the WRITE operations and thus extending the system's lifetime.

Further, design automation tools are essential to map computation into hardware designs. Hardware-software co-design is a trending approach in a variety of novel computing schemes, including photonic computing [29], [30], quantum computing [31], [32], and in-memory computing [33], [34], to optimize the hardware resources. In this work, we explore a tight hardware-software co-design for 1T1M crossbars and Boolean functions. To achieve this strong relation between hardware and software, we base ourselves on an analogy between binary decision diagrams (BDDs) and 1T1M crossbars.

Lastly, in previous works [12], [16], little or no attention is made to the underlying architecture. Many of these rely on a simple computing architecture consisting of multiple crossbars connected to a bus. Unfortunately, such bus-crossbar architecture is not energy- and latency-efficient. In our work, we target staircase architectures where a staircase is a hardwired collection of crossbars. The main idea is that the bus utilization will be reduced, which translates into energy and latency improvements of the overall computing system.

In this article, we propose a new computing paradigm called path-based in-memory computing. The paradigm is capable of evaluating Boolean functions using 1T1M crossbar arrays. We also propose a framework called PATH to automatically map computation to 1T1M crossbars or PATH systems with staircase structures. The main innovations of this article are summarized, as follows.

- 1) A new computing paradigm, called path-based in-memory computing, is introduced. The paradigm executes Boolean functions fast and efficiently using only READ operations instead of using slow and energy-consuming WRITE operations.
- 2) We introduce a framework, called PATH, to synthesize Boolean functions into a single crossbar design. The PATH framework exploits an analogy between BDDs and 1T1M crossbars to map Boolean functions into crossbar designs. A BDD with $|V|$ nodes and $|E|$ edges can be mapped into a crossbar of dimensions $|V| \times |E|$.
- 3) We further introduce an equivalent bipartite graph data structure for the BDD. By means of node merging, this bipartite graph can be further compressed into a smaller, equivalent graph. This compression results in an area reduction of 16%.
- 4) When the crossbar design becomes too large to be physically realizable, the PATH framework provides a partitioning algorithm to map Boolean functions to staircase structures. The objective is to minimize the bus utilization by minimizing the hardware resources in terms of the number of staircases.
- 5) The experimental evaluation is performed on ten circuits from the Revlib benchmark suite, eight control circuits from the EPFL benchmark suite, and eight ISCAS85

benchmarks. Compared with the state-of-the-art in-memory paradigm COMPACT [16], PATH improves energy and latency with $1006\times$ and $10\times$ on average.

The remainder of this article is organized as follows: preliminaries are provided in Section II. The PATH paradigm is introduced in Section III. In Section IV, we review some state-of-the-art digital in-memory computing paradigms. Problem formulation and a high-level overview of the PATH framework are given in Section V. The crossbar-level synthesis framework is detailed in Section VI, and the partitioning algorithm is provided in Section VII. Optimization steps are introduced in Section VIII. The experimental evaluation is performed in Section IX. This article is concluded in Section X.

II. PRELIMINARIES

A. Binary Decision Diagrams

A BDD is a graph representation of a Boolean function. The directed acyclic graph (DAG) consists of internal decision nodes and two leaf (terminal) nodes. The terminal nodes represent the output "0" and "1", respectively. The internal decision nodes are assigned a Boolean variable, and each internal decision node has a positive and negative output edge. The positive edge corresponds to the positive literal, and the negative edge corresponds to the negative literal. A BDD is evaluated by traversing the graph from the root nodes to one of the leaf nodes based on an instance of the Boolean variables. BDDs commonly refer to reduced order binary decision diagrams (ROBDDs) where nodes and edge have been eliminated to reduce the size of the representation [35]. When a BDD is used to represent a multioutput function, the BDD will have a separate root node for each output of the Boolean function [36].

B. Memristor Crossbar Arrays

In this section, we will review one-transistor one-memristor (1T1M) crossbars [37]. A model for a 1T1M crossbar is illustrated in Fig. 2(c). A 1T1M crossbar array consists of wordlines, bitlines, and selectorlines. Each wordline is connected to each bitline using a series-connected memristor and access transistor. The vertically aligned access transistors share a single selectorline. Both the memristors and the access transistors act functionally as switches that can be turned ON and OFF. The switch corresponding to a memristor is ON (or OFF) based on if the memristor is programmed to have low (or high) resistance. The switch corresponding to the access transistor is turned ON (or OFF) based on if the selectorline is charged (or discharged, depending on the type of transistor).

C. In-Memory Computing Architecture

Traditionally, bus architectures have been leveraged for in-memory computing [38], [39], [40]. In this computing architecture, the crossbars are connected to a bus. An example of a bus architecture with six crossbars is illustrated in Fig. 1(a). However, in our work, we target a PATH system with staircase structures. A staircase structure is a collection of crossbars that have hardwired interconnections. In Fig. 1(b), we illustrate a staircase architecture of six staircases where each staircase consists of five hardwired crossbars.




Fig. 1. Comparison of a traditional bus architecture and a staircase architecture. Each staircase is a collection of hardwired crossbars. (a) Bus architecture. (b) Staircase architecture.

III. PATH-BASED COMPUTING

PATH aims to evaluate Boolean functions using in-memory computing. An example of the flow for the synthesis and evaluation of PATH is shown in Fig. 2. The flow for PATH consists of a one-time slow and expensive compilation phase and a fast and efficient execution phase. The input to the compilation phase is a Boolean function specified in a hardware descriptive language (Verilog, VHDL), which is shown in Fig. 2(a). The input is first synthesized into an abstract crossbar design $\mathcal{D}$, which is shown in Fig. 2(b). The 1T1M crossbar design specifies the state of each nonvolatile memory device (0/1) and the Boolean variable assigned to each selectorline. Here, the Boolean literals a , b , and c are assigned to the first, second, and third selectorline, respectively. The input and output assignment to the wordlines are also specified. Next, the memory devices within a nanoscale crossbar are programmed ON (LRS) or OFF (HRS), which is shown in Fig. 2(c). The state of the devices are programmed to LRS or HRS by applying a voltage with appropriate polarity and magnitude [41]. We use a write-and-verify scheme to ensure the correct programming [42].

In the execution phase, an instance of Boolean variables is provided to the selectorlines. The selectorlines control the switches represented by the access transistors. The state of the switches controlled by the memory devices are also shown in Fig. 2(d). Next, an input voltage is applied to the top-most wordline and an output voltage is measured across a resistor connected to the bottom-most wordline. If the output voltage is high, the Boolean function evaluates to true. Otherwise, the function evaluates to false. For the input instance $(a, b, c) = (1, 1, 0)$, the function evaluates to true because there exists a path from the input to the output, as illustrated in Fig. 2(e). In contrast, the function evaluates to false for the input instance $(1, 0, 0)$. Observe that the memristors must not be reprogrammed to evaluate the same Boolean function for different input vectors. In Fig. 3, a more detailed example is shown for the Boolean function $f = a \vee \neg b$. More specifically, we show the state of the crossbar for all four input vectors. Again, the crossbar must not be reprogrammed for different input vectors.

The one-time compilation phase is both slow and expensive. Mainly, due to the expensive WRITE operations used to program the platform. On the other hand, the cost is amortized across each execution of the Boolean function. The execution phase is fast and efficient because it only involves charging/decharging the selectorlines and performing READ

operations. The advantageous properties compared with other in-memory paradigms comes from the *novel* use of the access transistors. No previous paradigms have used the access transistors to perform logic.

IV. COMPARISON OF DIGITAL IN-MEMORY COMPUTING PARADIGMS

In this section, we compare and review some of the most prevalent state-of-the-art digital in-memory computing paradigms and our proposed path-based in-memory computing paradigm PATH. These digital in-memory computing paradigms are IMPLY [45], MAGIC [11], MAJORITY [13], and FLOW [44]. In Fig. 4, we illustrate the steps during execution for each of the logic styles to evaluate the Boolean function $f = (a \wedge b) \vee \neg c$ for the input vector $a = 1, b = 1, c = 1$. For consistency, we employ the following basic operations among these logic styles: 1) READ; 2) WRITE; and 3) COPY (READ+WRITE). The definitions are provided in the legend at the top of Fig. 4.

A. IMPLY

IMPLY logic is based on the Boolean operation material implication (IMP) [43], [45]. The IMP operation $P \rightarrow Q$ can be realized in hardware using two memristors P and Q . By applying voltages over the memristors P and Q , the result is obtained in the memristor Q . Thus, IMPLY logic is destructive in terms of its inputs [43]. Further, extensive design automation tools for IMPLY-based in-memory computing have not been developed, usually requiring manual labor to design circuits [46]. One of the few automation tools for IMPLY is described in [47]. In Fig. 4(a), we observe that IMPLY requires many intermediate steps of READ and WRITE operations to realize the Boolean function f . The required IMP operations are also provided in Fig. 4(a).

B. MAGIC

The MAGIC logic style [11] is based on the Boolean operation NOR, and can be considered the successor of IMPLY. The NOR operation can be realized using three memristors. The NOT operation is a NOR operation where one input is always “1”. In contrast with IMPLY, MAGIC is not destructive for its inputs [11] when applying the appropriate voltage. Further, there is an additional memristor for the output to be realized. Over the years, many papers have been proposed using the MAGIC logic style [12], [46], [48]. In Fig. 4(b), we show the steps to realize the Boolean function f using READ, WRITE, and COPY operations.

C. MAJORITY

The majority operation is a Boolean function that evaluates to true when half or more of its inputs evaluate true. For in-memory computing, the majority operation with three inputs is primarily interesting due to its one-to-one correspondence with a single memristor [13]. We define the majority operation as $Z' = M(X, \neg Y, Z) = (X \wedge Z) \vee (\neg Y \wedge Z) \vee (X \wedge \neg Y)$. Then let X and Y be the inputs to the two terminals of the memristor, and let Z be the resistive state of the memristor. By applying the appropriate voltages to the inputs and programming the memristor to the appropriate resistive state, the majority




Fig. 2. Flow for the synthesis and evaluation of Boolean functions for PATH. (a) Program in Verilog code. (b) Abstract crossbar design obtained through synthesis. (c) Physical crossbar with the nonvolatile memory devices programmed and Boolean variables assigned to the selectorlines. (d) State of the switches (open/closed) with respect to the state of the nonvolatile memory devices (ON/OFF) and the instance $(a, b, c) = (1, 1, 0)$ of the Boolean variables. (e) Boolean function f evaluates to 1 because there is a path from the input to the output.




Fig. 3. Execution of all four input vectors on a crossbar for the Boolean function $f = a \vee \neg b$. Observe how the state of the memristors does not change for different input vectors, but only the state of the access transistors changes using the selector lines.

function can be executed in-situ. The resulting value Z' is then stored as a resistive value in the memristor [49]. Several synthesis methods have been proposed in recent years [14], [49], [50], [51], [52], many of which rely on majority inverter graphs (MIGs) as data structure. In Fig. 4(c), we illustrate the steps using READ and WRITE operations for the majority logic style.

D. FLOW

FLOW (flow-based computing) is a digital in-memory computing paradigm which relies on the absence/presence of electrical current to perform its computations [16], [44], [53]. Initially, the input variables, their negations, and the Boolean truth values (0/1) are assigned to the memristors. Program execution consists of two steps. In the first step, the memristors are programmed to their resistive states (0 for high, 1 for low), as shown in the first step of Fig. 4(d). In the second step, a high-input voltage is applied to the input wordline (V_{in}), and the Boolean function is *read* out as follows: if there is a path from the input wordline to the output wordline through memristors in a low-resistive state, then the Boolean function evaluates to true. Otherwise, the Boolean function evaluates to false.

E. PATH

In our proposed PATH, the program execution solely relies on READ operations, and the application of an input voltage to perform computations. WRITE operations are only performed once during the previous step, i.e., the compilation phase, for a given Boolean function. In the first step of Fig. 4(e), we show the crossbar design. The memristors are in their resistive states, which were programmed only once prior, during 1T1M crossbar reconfiguration (see previous Section III). During program execution there are no WRITE operations, and thus

these resistive states do not change. However, the selectorlines are charged accordingly to open/close access transistors. The memristors are programmed to their resistive states only once for a given Boolean function. Given another input vector $a = 0, b = 1, c = 0$, the crossbar design in Fig. 4(e), would remain, and is thus invariant to the input vector. Then, in the second step, the evaluation is *read* out by sensing the presence/absence of electrical current. The crossbar invariancy makes PATH a strong contender for repeated computations.

V. PATH FRAMEWORK

First, we outline the problem formulations in Section V-A. Then, we give a high-level overview the PATH framework in Section V-B.

A. Problem Formulation

Our overall objective is to synthesize a Boolean function ϕ into a PATH system. We approach this larger problem by solving two smaller problems, as follows.

- 1) *Problem I:* We propose a synthesis method to construct a crossbar design $\mathcal{D}$ for a Boolean function ϕ . The algorithm is based on an analogy between a BDD for the function ϕ and a 1T1M crossbar. We further improve the synthesis method by transforming the BDD into an equivalent graph-based data structure such that we can reduce its graph size by merging nodes. This transformation results in smaller crossbar designs, and subsequently power and latency improvements.
- 2) *Problem II:* Based on the analogy of Problem I, we propose a synthesis method to construct a topology $\mathcal{T}$ for a PATH system of staircase structures S_j . A staircase structure S_j is an ordered set of crossbars X_i . Between each X_i and X_{i+1} , there are hardwired intercrossbar connections from the wordlines of crossbar X_i to the selectorlines of crossbar X_{i+1} .

B. Overview of Framework

In this section, we give an outline of the PATH framework. An overview of the synthesis flow is shown in Fig. 5. First, in Section VI, we discuss synthesis for a single crossbar. Then, in Section VII, we discuss the staircase partitioning. For the crossbar synthesis, we introduce an algorithm to construct a single crossbar design $\mathcal{D}$ based on an analogy between a bipartite graph, derived from a BDD for the Boolean function(s),




Fig. 4. Comparison of the last step of the compilation phase and the steps during the execution phase for digital in-memory logic styles (a) IMPLY [43], (b) MAGIC [11], (c) MAJORITY [13], (d) FLOW [44], and the proposed (e) PATH. For a given Boolean function, the compilation phase is only performed once and the execution phase is performed many times. For the compilation phase, either a series of instructions or a crossbar is provided. The Boolean function $f = (a \wedge b) \vee \neg c$ is computed for the input vector $a = 1, b = 1, c = 1$. Intermediate variables are denoted by a dollar sign (\$). For each logic style, the operations READ, WRITE, and COPY (= READ + WRITE) are provided, as defined in the legend at the top of the figure. The WRITE operation writes the required resistive state to a memristor or executes the primary Boolean function for IMPLY/MAGIC/MAJORITY by applying the required input voltages. The READ operation reads the resistive state of a memristor or evaluates the Boolean function for FLOW/PATH. The COPY operation performs both a READ and WRITE operation. Observe how PATH does not require any WRITE operations during the execution phase.

and a 1T1M crossbar. This algorithm consists of four steps: 1) graph preprocessing; 2) graph transformation; 3) node merging; and 4) crossbar realization. Then, in Section VII, we introduce a partitioning algorithm. The main steps for the partitioning are the graph partitioning, and the realization of staircase intra- and interconnections. The framework is illustrated with an example in Fig. 6.

VI. CROSSBAR SYNTHESIS

The input to the framework is a BDD, and the output is a crossbar design. The BDD is obtained using CUDD [54] which is subsequently pruned into a graph G . The details are provided in Section VI-A. In the next step, the graph transformation step in Section VI-B, the pruned graph G is converted into a bipartite graph B . This graph is then compressed into a new bipartite graph B' in Section VI-C. The last step is the

crossbar synthesis step, which constructs a crossbar design D for the given bipartite graph B' . The details are provided in Section VI-D.

A. Graph Preprocessing

The input to the graph preprocessing step is a BDD. In Fig. 6(a), a multioutput BDD for a full adder is provided. The Boolean functions are $c_{\text{out}} = (a_0 \wedge b_0) \vee (a_0 \wedge c_{\text{in}}) \vee (b_0 \wedge c_{\text{in}})$ and $s_0 = a_0 \oplus b_0 \oplus c_{\text{in}}$, respectively. The graph preprocessing involves removing the zero output node and all the edges connected to the zero terminal node. The zero terminal node can be removed because it corresponds to $\neg c_{\text{out}}$ and $\neg s_0$. The one terminal node will be connected to the input, which we label in . The edges in the BDD are labeled with their respective decision variables. The positive (negative) edge connected to node with the decision variable x_i will be labeled x_i ($\neg x_i$).




Fig. 5. Overview of the PATH framework, including the crossbar synthesis and staircase partitioning.

Finally, we reverse the edges and we label the nodes from 1 to $|V|$ where $|V|$ is the number of nodes. The resulting graph of the BDD in Fig. 6(a) is shown in Fig. 6(b).

B. Graph Transformation

In this step, the resulting pruned graph is converted into a directed bipartite graph. This graph transformation is introduced as an intermediary data structure for the node merging step. Let $G = (V, E)$ be the pruned graph where V is a set of nodes and E is a set of edges, and let $B = (U_1, U_2, F)$ be a bipartite graph where U_1 and U_2 are sets of nodes and F is a set of edges. The sets U_1 and U_2 are disjoint and independent [55], and F is a new set of edges between nodes from U_1 and U_2 . Let $v \in V$ correspond to a node $u_1 \in U_1$, and let $e \in E$ correspond to a node $u_2 \in U_2$. For each node $v \in V$, we introduce a node $u_1 \in U_1$. For each edge in the BDD, we introduce a node with two edges. More specifically, for an edge $e = (v_1, v_2, l) \in E$ where $v_1 \in V$, $v_2 \in V$ and l is a literal, we create a new node $u_2 = (u_1^1, u_1^2, l) \in U_2$ where u_1^1 is the image of v_1 and u_1^2 is the image of v_2 . Then, we realize the connections between nodes and edges by introducing two new edges in F for each node $u_2 \in U_2$ such that $F = \{(u_1^1, u_2), (u_2, u_1^2) \mid u_2 = (u_1^1, u_1^2, l), u_2 \in U_2\}$. An example of the transformation of the pruned graph G into a bipartite graph B is illustrated in Fig. 6(c). Note that we represent the nodes in U_2 with their literals l instead of the triple (u_1^1, u_1^2, l) for clarity.

C. Node Merging

In the bipartite graph, we observe that a node $u_1 \in U_1$ may have outgoing edges to more than one node $u_2 \in U_2$ with the same literal l . For example, in Fig. 6(c), we observe that node 2 in the bipartite graph B has two outgoing edges to two distinct nodes with both label $\neg b$. In this section, we propose to merge such nodes with the same label into a single node.

More formally, let $B = (U_1, U_2, F)$ be the bipartite graph and let $u_1 \in U_1$ be a node with outgoing edges to nodes $u_2^i = (u_1, u_i, l)$ and $u_2^j = (u_1, u_j, l)$ where $i \neq j$, and $u_1, u_i, u_j \in$

U_1 . Then we define a mapping $B = (U_1, U_2, F) \rightarrow B' = (U_1', U_2', F')$ as follows:

$$U_1 \rightarrow U_1' : u_1 \mapsto u_1$$

$$U_2 \rightarrow U_2' : u_2 = (u_1, u_i, l) \mapsto u_2' = (u_1, l)$$

$$F \rightarrow F' : f = (u_1, (u_1, u_i, l)) \mapsto f' = (u_1, (u_1, l)) \text{ and}$$

$$f = ((u_1, u_i, l), u_1) \mapsto f' = ((u_1, l), u_1).$$

Based on the aforementioned mapping function, we can merge the two nodes with label $\neg b$ into one node such that we obtain a compressed bipartite graph B' as illustrated in Fig. 6(d). This operation is valid due to that the nodes $u_2 \in U_2$ represent literals l , and the edges between $u_1 \in U_1$ and $u_2 \in U_2$ represent conjunctions between u_1 and u_2 . Thus, for two such edges (u_1, u_i) and (u_1, u_j) , we have the following: $u_1 \wedge u_i = u_1 \wedge l = u_1 \wedge u_j$.

D. Crossbar Realization

The outlined crossbar realization is based on an analogy between the bipartite graph $B' = (U_1', U_2', F')$ and 1T1M crossbars. The nodes $u_1 \in U_1'$ correspond to wordlines and the nodes $u_2 \in U_2'$ correspond to bitline-selectorline pairs. The path-based paradigm is based on creating paths by turning on and off connections in the crossbar design. The connections correspond with the edges $f \in F'$, which are realized using the memristors. The crossbar mapping consists of a node assignment step and an edge assignment step.

1) *Node Assignment*: The node assignment involves assigning the nodes $u_1 \in U_1'$ to the wordlines of the crossbar design $\mathcal{D}$ and the nodes $u_2 \in U_2'$ to the bitline-selectorline pairs of crossbar design $\mathcal{D}$.

2) *Edge Assignment*: Next, for each edge $f = (u_1, u_2)$ or $f = (u_2, u_1)$, $u_1 \in U_1, u_2 \in U_2, f \in F$, we program the corresponding memristor at the intersection of wordline u_1 and selectorline u_2 to a low-resistive state (ON). Further, the input and output are assigned to the respective wordlines. The resulting crossbar design $\mathcal{D}$ for the Boolean functions f_1 and f_2 is shown in Fig. 6(e).

VII. PARTITIONING FOR STAIRCASE STRUCTURES

In this section, we propose a partitioning algorithm to synthesize the Boolean function ϕ into a topology $\mathcal{T}$ of staircase structures. A topology is a DAG of staircase structures with potentially multiple edges between different staircase structures where each staircase structure is an ordered set of crossbars with intercrossbar connections between two consecutive crossbars. An overview of the partitioning scheme is illustrated in Fig. 7.

The input of the partitioning algorithm is a bipartite graph $B = (U_1, U_2, F)$ and the output is a topology $\mathcal{T}$ of staircase structures. The bipartite graph B is obtained by means of the preprocessing steps described in Sections VI-A and VI-B. The idea of the partitioning scheme is that the given bipartite graph B is partitioned into smaller bipartite graphs $B_i = (U_{1,i}, U_{2,i}, F_i)$, $|U_{1,i} + U_{2,i}| \leq |U_1 + U_2|$. For each B_i , a crossbar design $\mathcal{D}_i$ is constructed, which is part of a staircase structure. Unfortunately, it is not straightforward to partition the graph B into B_i such that the size of B_i is maximized while meeting the dimensions of crossbar X_i . The partitioning makes that intermediate evaluations must be propagated to other




Fig. 6. Example of the synthesis flow (a) Multioutput BDD for a full adder with Boolean functions c_{out} and s_0 . (b) Nodes of the BDD are relabeled and the edges are assigned a Boolean literal based on their shape (positive literal for solid edges and negative literal for dashed edges). Further, the negative terminal node 0 is removed. (c) Bipartite graph $B = (U_1, U_2, F)$ is constructed from the pruned graph G . (d) Bipartite graph B is compressed into an equivalent bipartite graph $B' = (U'_1, U'_2, F')$ using node merging. (e) Crossbar design $\mathcal{D}$ is constructed with dimensions $|U'_1| \times |U'_2|$ where each node $u_1 \in U'_1$ is assigned to a wordline and each node $u_2 \in U'_2$ is assigned to a bitline-selectorline pair. For each edge $f = (u_1, u_2)$ or $f = (u_2, u_1)$, $u_1 \in U'_1, u_2 \in U'_2$, a memristor is programmed to a low-resistive state (ON) to realize the connections.




Fig. 7. High-level overview of the partitioning scheme. (a) Input of the partitioning scheme is the compressed bipartite graph B' and the user-defined parameter $T = 2$. (b) Using the graph partitioning algorithm explained in Section VII-A, the bipartite graph is decomposed into smaller bipartite subgraphs B_i . (c) Individual subgraphs B_i are each synthesized into crossbar designs $\mathcal{D}_i$ using the crossbar realization in Section VI-D. (d) Finally, a staircase topology $\mathcal{T}$ is constructed by realizing the intrastaircase and interstaircase connections, as explained in Section VII-B.

crossbars and/or staircases. Further, only the first crossbar X_1 in a staircase structure is connected to the bus, which brings that the intermediate results and literals can only be fed to this first crossbar. To address these constraints, we propose the following: a user-defined parameter defines the maximum dimensions which may be used to synthesize a bipartite graph B_i . Here, we assume the number of wordlines and the number of bitline-selectorline pairs is equal for a crossbar. An algorithm to construct such topology $\mathcal{T}$ is provided in Section VII-A. Next, staircase intra- and interconnections must be realized for the aforementioned constraints. These are discussed in Section VII-B.

A. Graph Partitioning

In Algorithm 1, we provide the first part of the partitioning scheme. We are given a bipartite graph $B = (U_1, U_2, F)$ as input, and a user-defined threshold T_i for the amount of logic that will be placed within each crossbar X_i . The output of the algorithm is a topology $\mathcal{T}$ of staircase structures S where each

S is an ordered set of crossbars X_i such that X_i precedes X_{i+1} . The partitioning algorithm has two auxiliary variables $V_{i,1}$ and $V_{i,2}$ which will contain the nodes assigned to the wordlines and selectorlines, respectively. The nodes that are assigned to $V_{i,1}$ are in U_1 , and the nodes that are assigned to $V_{i,2}$ are in U_2 .

The algorithm iterates in a topological sort over the nodes $u_2 \in U_2$. In each iteration, node u_2 is assigned to a crossbar, together with its neighboring nodes. Recall that the nodes u_2 are the edges $e \in E$ in our original graph $G = (V, E)$. When assigning a node $u_2 \in U_2$ to a crossbar X_i , we want each neighboring node u_1 to be assigned to X_i as well. This is due to that u_2 represents an edge $e = (v_1, v_2) \in E$ between two nodes $v_1, v_2 \in V$. Thus, we want both its endpoints to be present in the crossbar X_i .

When assigning a node u_2 to the wordlines of a crossbar $V_{i,2}$, we must not exceed the logic threshold T_i we have set. Similar for its neighboring nodes u_1 when assigning to the selectorlines $V_{i,1}$ (condition of if statement on line 7). If the condition fails, we create a bipartite subgraph $B_i =$

Algorithm 1 Partitioning Algorithm for Staircase Structures

Input: $B = (U_1, U_2, F)$, $T = \{T_0, \dots, T_L\}$
Output: $\mathcal{T}$ // Set of staircase designs

```

1: function TOPOLOGICALSTAIRCASEPARTITIONING( $B, T$ )
2:    $i = 1, V_{i,1} = \emptyset, V_{i,2} = \emptyset, S \leftarrow \emptyset$ 
3:    $\mathcal{T} \leftarrow \emptyset$ 
4:   for  $u_2 \in \text{TOPOLOGICALSORT}(U_2)$  do
5:      $V'_{i,1} \leftarrow V_{i,1} \cup \{u_1 | f = (u_1, u_2) \vee f = (u_2, u_1), \forall f \in F\}$ 
6:      $V'_{i,2} \leftarrow V_{i,2} \cup \{u_2\}$ 
7:     if  $|V'_{i,1}| \leq T_i \wedge |V'_{i,2}| \leq T_i$  then
8:        $V_{i,1} \leftarrow V'_{i,1}$ 
9:        $V_{i,2} \leftarrow V'_{i,2}$ 
10:    else
11:       $F_i \leftarrow \{f | \exists u_1 \in V_{i,1}, u_2 \in V_{i,2}:$ 
12:         $f = (u_1, u_2) \vee f = (u_2, u_1), f \in F\}$ 
13:       $B_i \leftarrow (V_{i,1}, V_{i,2}, F_i)$  // Create bipartite subgraph
14:       $S \leftarrow S \cup \{B_i\}$ 
15:       $i \leftarrow i + 1$ 
16:       $V_{i,1} = \{u_1 | f = (u_1, u_2) \vee f = (u_2, u_1), \forall f \in F\}$ 
17:       $V_{i,2} = \{u_2\}$ 
18:      if  $|S| = L$  then
19:         $\mathcal{T} \leftarrow \mathcal{T} \cup \{S\}$ 
20:         $i \leftarrow i - 1, S \leftarrow \emptyset$ 
21:      end if
22:    end if
23:  end for
24: return  $\mathcal{T}$ 
25: end function

```

($V_{i,1}, V_{i,2}, F_i$) (lines 11 and 12), and we add B_i to the current staircase S (line 13). When the current staircase S has reached its maximum depth L (line 17), then we will add the current staircase S to the topology $\mathcal{T}$ (line 18), and we create a new staircase S (line 19). The algorithm stops when all nodes $u_2 \in U_2$ have been processed.

In Fig. 7(a), we take the compressed bipartite graph B' and the user-defined parameter $T = 2$ as input for the partitioning algorithm. In Fig. 7(b), we illustrate the partitioning of the bipartite graph into multiple subgraphs B_i . Each subgraph B_i is delineated by a dashed line, and all its nodes have the same number i . These subgraphs are subsequently synthesized into crossbar designs $\mathcal{D}_i$, as explained in Section VI-D and grouped into staircase structures.

B. Realization of Staircase Intra- and Interconnections

While the algorithm partitions the bipartite graph B into bipartite subgraphs B_i , which are mapped to crossbars, the hardware architecture imposes additional constraints on the design. We have identified three staircase intra- and interconnections that must be made to realize the crossbar mapping to a partitioning over staircase structures: 1) edge preparation; 2) node propagation; and 3) literal propagation. In Fig. 7(c), we take the crossbar designs $\mathcal{D}_i$ as input. The output is a topology $\mathcal{T}$ of staircases by realizing the staircase intra- and interconnections, as illustrated in Fig. 7(d).

1) *Edge Preparation:* For each crossbar X_i , $i > 1$, the selectorlines are connected to the wordlines of previous crossbar X_{i-1} . In the mapping algorithm in Section VI-D, the nodes $u_2 \in U_2$ are assigned to the selectorlines. This entails that the nodes must be *prepared* in crossbar X_{i-1} . In Fig. 8(a), is illustrated how the nodes $u_{2,j}$ for crossbar X_1 (in black) are prepared in crossbar X_0 (in gray).




Fig. 8. Example of the intraconnections. (a) Edge prep. (b) Node prop. (c) Literal prop.

Algorithm 2 Binary Search Over the Threshold

Input: $B = (U_1, U_2, F)$, D
Output: $\mathcal{T}$ // Set of staircase designs

```

1: function BINARYSEARCH( $T$ )
2:    $low \leftarrow 0, high \leftarrow D$ 
3:    $\mathcal{T} \leftarrow \emptyset$ 
4:    $T \leftarrow \lfloor (low + high)/2 \rfloor$ 
5:   while  $low \neq high$  do
6:      $\mathcal{T}' \leftarrow \text{TOPOLOGICALSTAIRCASEPARTITIONING}(B, T)$ 
7:     if  $\mathcal{T}' \neq \emptyset$  then // Solution found
8:        $\mathcal{T} \leftarrow \mathcal{T}'$ 
9:        $low \leftarrow T$  // Increase lower bound
10:    else // No solution found
11:       $high \leftarrow T$  // Decrease upper bound
12:    end if
13:     $T \leftarrow \lfloor (low + high)/2 \rfloor$ 
14:  end while
15:  return  $\mathcal{T}$ 
16: end function

```

2) *Node Propagation:* A node $u_1 \in U_1$ may appear in multiple crossbars X_i among multiple staircases S_j . From the structure of a pruned graph G , we know that each node $v \in V$ has at most two outgoing edges. At some point, the node will be *realized*, i.e., its two outgoing edges have been assigned. Let that point be denoted as X_r . From this point X_r forward, any other occurrence of v is to realize incoming edges of v . When v occurs at some later point in the same staircase X_i , $i > r$, we must propagate v to that crossbar X_i . This is illustrated in Fig. 8(b) where node v is realized in crossbar X_1 and propagated to crossbar X_4 . The intermediary nodes v' are in gray, and the start and end point are in black. Similarly, when v occurs in multiple staircases, then node v must be propagated from its point of realization X_r to all other staircases.

3) *Literal Propagation:* A literal l may appear in a crossbar X_i , $i \geq 2$. For each such literal l , we must propagate the literal up to layer X_{i-2} . For example, in Fig. 8(c), the literal l appears in crossbar X_4 , and is thus propagated from the first crossbar X_1 to the last crossbar X_4 .

VIII. ADDITIONAL OPTIMIZATION

In this section, we introduce two additional optimization steps. In Section VIII-A, we propose a method to improve the search over the user-defined threshold T , and in Section VIII-B, we propose an optimization to have a more fine-grained exploration of the search space.

A. Partitioning Search

The partitioning algorithm presented in Section VII requires a user-defined parameter T , which is a threshold for the amount of logic that will be placed in a crossbar X . As this

variable is unknown in advance, we propose a binary search over T . Let all crossbars X_i in a staircase structure have the same dimensions $D \times D$. In Algorithm 2, we provide the binary search algorithm for the topological staircase partitioning. The input is the bipartite graph $B = (U_1, U_2, F)$, and the dimensions D of the crossbars. The output is a topology $\mathcal{T}$. The idea is that when for a given threshold T , no solution can be found, we decrease the threshold T . Potentially, no solution is found due to the intra- and interconnections explained in Section VII-B. The node propagations, literal propagations, and edge preparations may result in that a crossbar exceeds its dimensions while constructing, and consequently the partitioning algorithm fails to find a solution for the given constraints. In the other case, when for a given $\mathcal{T}$, a solution can be found, we retain this solution and attempt to find a better solution by increasing the threshold $\mathcal{T}$.

B. Node Splitting

In this section, we provide an optimization step to improve the overall synthesis. Due to the node merging optimization laid out in previous Section VI-C, the node degree for all nodes $u_2 \in U_2$ may increase. The partitioning algorithm in Algorithm 1 assigns such nodes u_2 and its neighboring nodes $u_1 \in U_1$ to a single crossbar in a staircase. When the node degree of u_2 , $\delta(u_2)$, is greater than the logic threshold T , such node cannot be assigned to a crossbar. A solution would be to increase the threshold T , but this brings with it that there is less room for node propagations, literal propagations, and edge preparations. Hence, there is a fine balance which must be sought between the threshold T and the node degree $\delta(u_2)$. Therefore, we propose to split nodes $u_2 \in U_2$ for which $\delta(u_2) > T$ into two nodes u_2^1 and u_2^2 .

In Algorithm 3, we present an algorithm to cope with such nodes. The algorithm can be used in combination with Algorithm 2. More specifically, line 6 in Algorithm 2 can be replaced with $\mathcal{T}' \leftarrow \text{SPLITWRAPPER}(B, T)$.

Algorithm 3 consists of two parts: 1) $\text{SPLITWRAPPER}(B, T)$ and 2) an auxiliary function $\text{SPLITNODE}(B, T)$. The former continues to split nodes $u_2^* \in U_2$ with maximum degree $\delta(u_2^*)$ as long as a node in B is changed (line 16). We use the auxiliary function $\text{SPLITNODE}(B, T)$ to perform this operation. On line 2, we seek such a node $u_2^* \in U_2$ with maximum node degree $\delta(u_2^*)$. When this node degree is smaller than the threshold, we do not need to split. Hence, we return our current bipartite graph B (lines 3–5). Otherwise, we create a new bipartite graph B' where u_2^* is replaced by two new nodes u_2^1 and u_2^2 such that its number of edges is equal, or differs at most by one edge (lines 6–12).

IX. EXPERIMENTAL EVALUATION

The experiments are conducted on a machine with 20 Intel Core i9-9900X and 128GB RAM. The framework is implemented in Python 3.8 and the source code is publicly available on GitHub.¹ We use the ABC [56] binding for CUDD [54] to construct the BDDs with dynamic variable reordering based on symmetric sifting [57]. In Table II, an overview is provided of ten benchmarks from the Revlib benchmark suite [58], eight control benchmarks

Algorithm 3 Node Splitting Algorithm

Input: $B = (U_1, U_2, F)$, T
Output: $\mathcal{T}$ // Set of staircase designs

```

1: function SPLITNODE( $B, T$ )
2:    $u_2^* \leftarrow \arg \max \delta(u_2), \forall u_2 \in U_2$ 
3:   if  $\delta(u_2^*) \leq T$  then
4:     return  $B$ 
5:   end if
6:    $X \leftarrow \{u_1 | (u_1, u_2^*) \in F\}$ 
7:    $Y \leftarrow \{u_1 | (u_2^*, u_1) \in F\}$ 
8:    $U_1' \leftarrow U_1$ 
9:    $U_2' \leftarrow U_2 \setminus \{u_2^*\} \cup \{u_2^1, u_2^2\}$ 
10:   $F' \leftarrow \{(u_1, u_2^1) | u_1 \in X\} \cup \{(u_1, u_2^2) | u_1 \in X\} \cup$ 
        $\{(u_2^1, u_1), (u_2^2, u_1) | u_1 \in Y, u_2^1 \in U_2, u_2^2 \in U_2,$ 
        $Y_1 \subseteq Y, Y_2 \subseteq Y, |Y_1| - |Y_2| \leq 1\}$ 
11:   $B' \leftarrow (U_1', U_2', F')$ 
12:  return  $B'$ 
13: end function

14: function SPLITWRAPPER( $B, T$ )
15:   $B' \leftarrow \emptyset, \mathcal{T} \leftarrow \emptyset$ 
16:  while  $B' \neq B$  do
17:     $\mathcal{T} \leftarrow \text{TOPOLOGICALSTAIRCASEPARTITIONING}(B, T)$ 
18:    if  $\mathcal{T} \neq \emptyset$  then
19:      return  $\mathcal{T}$ 
20:    else
21:       $B' \leftarrow \text{SPLITNODE}(B, T)$ 
22:    end if
23:  end while
24:  return  $\mathcal{T}$ 
25: end function

```




Fig. 9. High-level overview of the architecture. (a) Bank consists of multiple tiles T . (b) Each tile T contains multiple staircases S . The topology of the staircases is according to a H-Tree. (c) Each staircase S contains a series of crossbars X .

from the EPFL benchmark suite [59], and eight ISCAS85 benchmarks [60]. We report the number of inputs, outputs for each benchmark, as well as the number of nodes and edges for the respective BDD.

We evaluate the PATH systems by building an architectural model. In Fig. 9, we illustrate the high-level architecture. The architecture consists of several tiles T on a bank, as illustrated in Fig. 9(a). Each tile T has a H-Tree of staircases S as topology [61], [62], as shown in Fig. 9(b). Four staircases with an I/O of 128 bits are connected to a Wide-I/O bus. Each staircase S is a series of crossbars X , as illustrated in Fig. 9(c).

In our experimental evaluation, we will compare the performance of the proposed PATH framework with COMPACT [16], ArC [50], and CONTRA [12]. The performance is compared in terms of energy, latency, and area. The parameters for the comparisons are given below. To evaluate our architecture, we set the power consumption for the bus and 128×128 crossbar to 13 mW, 0.3 mW [38]. For

¹<https://github.com/sventhijssen/path>

TABLE II

OVERVIEW OF TEN REVLIB BENCHMARKS, EIGHT EPFL CONTROL BENCHMARKS, AND EIGHT ISCAS85 BENCHMARKS. FOR EACH BENCHMARK, THE NUMBER OF INPUTS AND OUTPUTS IS GIVEN. FOR THEIR RESPECTIVE BDDs, THE NUMBER OF NODES AND EDGES IS GIVEN

Benchmark	Benchmark		BDD		Pruned	
	Inputs (num)	Outputs (num)	Nodes (num)	Edges (num)	Nodes (num)	Edges (num)
Revlib						
in0	15	11	385	766	384	680
apex2	39	3	567	1130	566	1042
spla	16	46	594	1184	593	864
pdv	16	40	621	1238	620	887
misex3	14	14	674	1344	673	1094
tial	14	8	897	1790	896	1717
apex4	9	19	990	1976	989	1874
cps	24	109	1080	2156	1079	1633
apex5	117	88	1259	2514	1258	2387
seq	41	35	1302	2600	1301	2041
EPFL						
arbiter	256	129	25109	50214	25108	49758
cavlc	10	11	436	868	435	776
ctrl	7	26	89	174	88	128
dec	8	256	512	1020	511	510
i2c	147	142	1204	2404	1203	1936
int2float	11	7	159	314	158	301
priority	128	8	772	1540	771	1539
router	60	30	219	434	218	379
ISCAS85						
c432	36	7	1291	2578	1290	2463
c499	41	32	111146	222164	111114	212466
c880	60	26	5776	11448	5750	11151
c1355	41	32	111146	222164	111114	212466
c1908	33	25	30605	61110	30580	57308
c2670	233	140	8249	15940	8109	14621
c5315	178	123	15454	30416	15331	27477
c7552	207	108	33983	67534	33875	65400
Normalized			1.00	1.00	1.00	0.85

our design, we have a 4-channel 128-bit Wide-IO bus with a rate of 400 MHz [63]. The area for the respective components are $0.2 \mu\text{m}^2$, 15.7 mm^2 [38]. For COMPACT, we extrapolate the area. The latency for the bus and crossbar components are 15 ns and 100 ns, respectively.

In Section IX-A, we first evaluate the crossbar synthesis. Then, in Section IX-B, we evaluate the proposed PATH framework, including the proposed node merging and partitioning algorithm. Finally, in Section IX-C, we make a comparison of the PATH framework with other digital in-memory computing paradigms.

A. Evaluation of Crossbar Synthesis

In this section, we will evaluate the crossbar synthesis. For the evaluation, we do not impose any restrictions on the crossbar dimensions, such that the number of wordlines (rows), and the number of bitline-selectorline pairs (columns) can be infinitely large. We evaluate the crossbar synthesis first without and then with the proposed node merging. In Table III, we provide the number of nodes and edges for the pruned graph G , as well as the hardware resources for both approaches. For the synthesis without node merging, we observe that the number of rows and the number of columns correspond to the number of nodes and edges of the pruned graph, respectively. This is due to the analogy between BDDs and 1T1M crossbars. Next, we report the number of rows and columns for the approach with node merging. We observe that the number of columns (selectorline-bitlines pairs) reduces by 16% on average, resulting in an area reduction of 16% on average.

TABLE III

NUMBER OF NODES AND EDGES FOR THE PRUNED GRAPH G , AND THE CORRESPONDING HARDWARE RESOURCES FOR A CROSSBAR DESIGN USING THE UNCOMPRESSED AND COMPRESSED BIPARTITE GRAPH

Benchmark	Pruned graph		Without node merging		With node merging	
	Nodes (num)	Edges (num)	Rows (num)	Cols (num)	Rows (num)	Cols (num)
Revlib						
in0	384	680	384	680	384	565
apex2	566	1042	566	1042	566	879
spla	593	864	593	864	593	767
pdv	620	887	620	887	620	750
misex3	673	1094	673	1094	673	849
tial	896	1717	896	1717	896	1143
apex4	989	1874	989	1874	989	1157
cps	1079	1633	1079	1633	1079	1248
apex5	1258	2387	1258	2387	1258	2132
seq	1301	2041	1301	2041	1301	1560
EPFL						
arbiter	25108	49758	25108	49758	25108	41441
cavlc	435	776	435	776	435	530
ctrl	88	128	88	128	88	100
dec	511	510	511	510	511	510
i2c	1203	1936	1203	1936	1203	1837
int2float	158	301	158	301	158	265
priority	771	1539	771	1539	771	1539
router	218	379	218	379	218	351
ISCAS85						
c432	1290	2463	1290	2463	1290	1929
c499	111114	212466	111114	212466	111114	198936
c880	5750	11151	5750	11151	5750	8540
c1355	111114	212466	111114	212466	111114	198936
c1908	30580	57308	30580	57308	30580	53100
c2670	8109	14621	8111	14622	8111	13337
c5315	15331	27477	15331	27477	15331	23850
c7552	33875	65400	33875	65400	33875	53708
Normalized avg.	1.00	1.00	1.00	1.00	1.00	0.84

From this, we conclude that it is advantageous to work with the compressed bipartite graph B' , and we will use this graph in subsequent sections. Thus, a BDD with $|V|$ nodes and $|E|$ edges can be synthesized into a crossbar of dimensions $|V| \times |E|$, which is an upper bound. Empirically, we conclude that on average a BDD with $|V|$ nodes and $|E|$ edges can be synthesized into a crossbar of dimensions $|V| \times 0.84|E|$.

B. Evaluation of the PATH Framework

In this section, we evaluate the PATH framework. In our first experiment, we evaluate the hardware resources for varying staircase depth L , i.e., the number of crossbars in a staircase structure. These hardware resources are the number of staircases, the number of staircase interconnections, and the critical path length. In Table IV, we give an overview of these hardware resources as well as the synthesis time for varying staircase depths $L \in \{1, 2, 4, 6\}$.

We observe that the number of required staircases decreases when the staircase depth L increases, with a reduction of 24% on average for a staircase structure of six layers compared with a single crossbar. For example, for the benchmark arbiter of the EPFL benchmark suite, the number of staircases reduces from 889 for $L = 1$ to 691 for $L = 6$. The number of interconnections may increase or decrease, depending on the benchmark. For example, for arbiter, the number of interconnections increases from 49,973 to for $L = 1$ to 51,035 for $L = 6$. This is due to that the logic threshold tends to be lower for larger staircase structures, requiring more node splits, and consequently more node propagations. However, for the majority of the benchmarks (17 out of 26), the number of intercrossbar connections decreases, with a reduction of 8%

TABLE IV
COMPARISON OF THE HARDWARE RESOURCES AND SYNTHESIS TIME (T) FOR VARYING PATH ARCHITECTURES. THE HARDWARE RESOURCES ARE EXPRESSED IN TERMS OF NUMBER OF STAIRCASES (S), NUMBER OF INTERCONNECTIONS (I), AND CRITICAL PATH LENGTH (C)

Benchmark	L = 1				L = 2				L = 4				L = 6			
	S (num)	I (num)	C (num)	T (min)	S (num)	I (num)	C (num)	T (min)	S (num)	I (num)	C (num)	T (min)	S (num)	I (num)	C (num)	T (min)
Revlib																
in0	11	547	11	0.1	10	527	10	0.2	9	526	9	0.4	9	529	9	0.4
apex2	16	767	16	0.1	13	706	13	0.2	11	630	11	0.2	10	579	10	0.4
spla	14	694	14	0.2	12	666	12	0.3	9	513	9	0.2	7	430	7	0.2
pdc	14	645	14	0.2	12	622	12	0.2	9	526	9	0.2	8	457	8	0.2
misex3	16	835	15	0.2	15	839	15	0.2	13	819	13	0.3	12	814	12	0.4
tial	27	1422	23	0.2	23	1404	21	0.4	21	1429	19	0.3	21	1414	19	0.8
apex4	31	1683	31	0.2	27	1693	26	0.5	25	1686	24	1.0	25	1698	24	1.0
cps	27	1393	25	0.4	24	1388	23	0.5	20	1368	20	0.4	20	1364	20	0.7
apex5	47	2077	27	0.4	36	2060	26	0.5	34	2102	28	0.9	33	2110	27	0.9
seq	37	1826	25	0.4	31	1821	27	0.6	28	1811	22	1.1	28	1831	27	0.9
EPFL																
arbitr	889	49973	302	8.6	762	50348	314	12.7	717	50546	311	17.8	691	51035	307	15.5
cavlc	11	610	11	0.0	11	627	11	0.0	9	593	9	0.0	9	593	9	0.1
ctrl	1	0	1	0.0	1	0	1	0.0	1	0	1	0.0	1	0	1	0.0
dec	6	192	4	0.9	6	202	4	0.0	5	206	4	0.1	5	196	4	0.1
i2c	32	1616	25	0.4	30	1664	27	1.6	28	1639	25	0.3	29	1647	25	0.5
int2float	4	146	4	0.0	3	115	3	0.0	3	106	3	0.0	2	67	2	0.0
priority	19	495	18	0.1	15	425	14	0.3	14	426	14	0.0	16	481	16	4.3
router	4	87	4	0.0	4	95	4	0.1	4	99	4	0.0	4	97	4	0.2
ISCAS85																
c432	40	2121	40	0.3	36	2086	36	0.5	33	2071	33	1.6	32	2049	32	1.0
c499	3592	160724	108	26.4	3212	163939	101	32.7	3020	165236	91	106.9	2884	161093	88	82.0
c880	189	8004	43	1.5	167	7931	43	4.5	155	7750	42	4.9	150	7666	43	4.8
c1355	3592	160724	108	26.1	3212	163939	101	32.0	3020	165236	91	112.1	2884	161093	88	80.4
c1908	939	39903	50	8.7	835	39959	42	9.5	765	40000	45	31.3	731	38959	43	24.9
c2670	352	8576	45	2.7	314	8235	41	5.1	295	7914	40	8.1	272	7199	41	11.3
c5315	393	10699	25	3.9	353	10114	22	7.0	293	8694	20	7.8	277	7834	19	10.6
c7552	1032	35470	108	7.7	894	34416	87	14.6	779	32417	82	21.2	694	30593	78	365.2
Normalized	1.00	1.00	1.00	1.00	0.89	0.98	0.92	1.72	0.80	0.95	0.85	2.27	0.76	0.92	0.83	5.66




Fig. 10. Number of staircases in terms of BDD size for eight ISCAS85 benchmarks.

on average for six layers compared with a single crossbar. For example, for benchmark cavlc of the Revlib benchmark suite, the number of staircase interconnections decreases from 610 for $L = 1$ to 566 for $L = 6$. Finally, we observe that the critical path length reduces by 17% on average for $L = 6$ compared with $L = 1$. The reduction of the number of staircases brings with it that the critical path length decreases. This is because the critical path length is at most the number of staircases, and the number of staircases for $L = 6$ is lower than the number of staircases for $L = 1$. From these results, we conclude it is best to utilize a PATH system with larger staircase structures.

Next, we make an analysis of the hardware utilization in terms of the intermediate data structure. More specifically, in Fig. 10, we show the number of required staircases in terms of the number of BDD nodes for different staircase depths. The crossbar dimensions are 128×128 , and the BDDs are




Fig. 11. Number of (a) Staircases and (b) Interconnections for the EPFL benchmark arbitrator for varying dimensions D .




Fig. 12. Percentage for each of the components in a staircase topology using the PATH framework. These components include logic, node propagation, edge preparation, and literal propagation. The percentages are the averages over the ten Revlib, eight EPFL benchmarks, and eight ISCAS85 benchmarks.

collected from the eight ISCAS85 benchmarks. From this figure, we clearly observe there is a linear trend between these two dimensions. Further, we observe at first glance that the number of required staircases decreases for increasing




Fig. 13. Normalized energy, latency, and area for PATH, COMPACT [16], ArC [50], and CONTRA [12] for ten Revlib benchmarks, eight EPFL benchmarks, and eight ISCAS85 benchmarks. (a) Energy consumption. (b) Latency. (c) Area.

staircase depth (the line for $L = 2$ lies higher than for $L = 4$, and $L = 4$ lies higher than for $L = 6$). This corresponds with the results in Table IV. For $L = 2$, $L = 4$, and $L = 6$, the trendline is described by the following equations, respectively: $0.0285x - 0.0126$, $0.0268x - 0.02782$, and $0.0255x - 0.3139$ where x is the number of BDD nodes.

Now, we will evaluate the PATH framework in terms of the crossbar dimensions. We evaluate on the benchmark arbiter using a staircase depth of six crossbars. In Fig. 11(a), we show the trendline for the number of staircases in function of the crossbar dimensions. We observe that for increasing crossbar dimensions, the number of staircases decreases. This is expected as there is more room in a crossbar for both logic and node propagations. In Fig. 11(b), we observe that the number of interconnections decreases as the crossbar dimensions increase. This is also expected as more logic can be realized within a single staircase, and thus less interstaircase communication is required.

In Section VII-B, we have highlighted that the partitioning method requires some intra- and interconnections in order to be a functional computing paradigm. We make an analysis of the components that constitute to the overall synthesis using partitioning. These components are logic, edge preparation, node propagation, and literal propagation. This analysis may give further insight in the synthesis method with the objective to improve any future work on our proposed framework.

In Fig. 12, we show the percentages for each of the components for varying number of layers L where $L \in \{1, 2, 4, 6\}$. We observe that the percentage of logic, which is defined by the threshold T , decreases for increasing number of layers L . This is due to that the number of node propagations increases when the number of layers L increases, which can also be seen in the figure. As mentioned earlier in Section VIII-B, there is a fine balance between the threshold T and the node degree $\delta(u_2)$. When the node degree decreases, the number of node propagations increase, and vice versa.

C. Comparison With Other Digital In-Memory Computing Paradigms

In this section, we make a comparison of the PATH paradigm with other digital in-memory computing paradigms. More specifically, we compare with COMPACT [16], ArC [50], and CONTRA [12]. COMPACT is the state-of-the-art synthesis method for flow-based computing, ArC for MAJORITY, and CONTRA is the state-of-the-art MAGIC-based general-purpose synthesis method. No comparison is provided with IMPLY-based logic [9], [43] as recent papers have shown that IMPLY-based logic is inferior to MAGIC-based logic [46], [64].

In Fig. 13, a detailed comparison is given for the normalized energy consumption, latency, and area for all benchmarks (except spla and pdc as CONTRA failed to generate a result) using PATH, COMPACT, and CONTRA. For ArC, only the ISCAS85 benchmarks are reported based on the results in [12]. Note that the latency only reflects the execution time, i.e., the runtime to evaluate Boolean functions, and does not include the synthesis time as reported in Table IV nor the crossbar programming. Compared with PATH, COMPACT requires approximately $1006\times$ more energy, and has approximately $10\times$ longer latency on average. The advantageous performance mainly stems from that COMPACT is a flow-based computing framework where the devices are continuously switched for each evaluation, resulting in many expensive (in terms of energy and latency) WRITE operations. No partitioning scheme exists for COMPACT, so we extrapolated the crossbar size of a 128×128 crossbar to the required dimensions. For COMPACT, some benchmarks require more area than PATH, so we have truncated the plot at unity for clarity (e.g., arbiter has $4\times$ area). On average, COMPACT requires $5.8\times$ of the area of PATH. Further, we observe that CONTRA consumes approximately $2166\times$ more energy and is approximately $15\times$ slower than PATH on average. Similarly to previous argument, CONTRA is much less energy-efficient and slower than PATH

due to the large number of write operations. The path-based paradigm only utilizes WRITE operations in the compilation phase, which is amortized across many function evaluations. On average, CONTRA requires only 2% of the area of PATH. Lastly, ArC requires on average $175.96\times$ more energy than PATH and is $8.30\times$ slower than PATH due to the many WRITE operations.

X. SUMMARY AND FUTURE WORK

In this article, we have introduced a new READ-based in-memory computing paradigm, called PATH, by leveraging access transistors to perform logic. We have introduced a framework, called PATH, to automatically synthesize Boolean circuits to PATH systems. The PATH framework relies on an analogy between bipartite graphs and 1T1M crossbars. The bipartite graphs are derived from BDDs, and serve as an intermediate data representation. Further, we have introduced an optimization technique wherein these bipartite graphs are compressed, resulting in an area reduction of 16%. Finally, we have introduced a partitioning algorithm to map Boolean functions to a topology of staircase structures, where a staircase structure is an ordered set of crossbars, which have hardwired connections between them. By introducing staircases, the bus utilization diminishes, which results in high energy and latency improvements. Our experimental results demonstrate that the paradigm is orders of magnitude faster than state-of-the-art in-memory computing paradigms with energy improvements of $1006\times$ on average. The latency improvements are $10\times$ on average. For future work, we envision that leveraging alternative intermediate data structures may improve the overall synthesis method. Further, alternative or orthogonal approaches to our proposed partitioning algorithm are an interesting trajectory for further research.

REFERENCES

- [1] D. Reinsel, J. Gantz, and J. Rydning, *The Digitization of the World From Edge to Core*, Int. Data Corp., Framingham, MA, USA, 2018, p. 16.
- [2] S. Li, L. D. Xu, and S. Zhao, "The Internet of Things: A survey," *Inf. Syst. Front.*, vol. 17, pp. 243–259, Apr. 2015.
- [3] I. Tomkos, D. Klonidis, E. Pikasis, and S. Theodoridis, "Toward the 6G network era: Opportunities and challenges," *IT Prof.*, vol. 22, no. 1, pp. 34–38, 2020.
- [4] W. Liu, Z. Wang, X. Liu, N. Zeng, Y. Liu, and F. E. Alsaadi, "A survey of deep neural network architectures and their applications," *Neurocomputing*, vol. 234, pp. 11–26, Apr. 2017.
- [5] R. Bommasani et al., "On the opportunities and risks of foundation models," 2021, *arXiv:2108.07258*.
- [6] J. Backus, "Can programming be liberated from the von Neumann style?" *Commun. ACM*, vol. 21, no. 8, pp. 613–641, 1978.
- [7] T. N. Theis and H.-S. P. Wong, "The end of Moore's law: A new beginning for information technology," *Comput. Sci. Eng.*, vol. 19, no. 2, pp. 41–50, Mar./Apr. 2017.
- [8] H. Esmailzadeh et al., "Dark silicon and the end of multicore scaling," in *Proc. ISCA*, 2011, pp. 365–376.
- [9] E. Lehtonen, J. Poikonen, and M. Laiho, "Implication logic synthesis methods for memristors," in *Proc. ISCAS'12*, 2012, pp. 2441–2444.
- [10] S. Shirinzadeh, M. Soeken, and R. Drechsler, "Multi-objective BDD optimization for RRAM circuit design," in *Proc. IEEE DDECS*, 2016, pp. 1–6.
- [11] S. Kvatinisky et al., "MAGIC—Memristor-aided logic," *IEEE Trans. Circuits Syst. II, Exp. Briefs*, vol. 61, no. 11, pp. 895–899, Nov. 2014.
- [12] D. Bhattacharjee et al., "CONTRA: Area-constrained technology mapping framework for memristive memory processing unit," in *Proc. ICCAD*, 2020, pp. 1–9.
- [13] P.-E. Gaillardon et al., "The programmable logic-in-memory (PLiM) computer," in *Proc. Design, Autom. Test Europe Conf. Exhibit. (DATE)*, 2016, pp. 427–432.
- [14] S. Shirinzadeh, M. Soeken, P.-E. Gaillardon, and R. Drechsler, "Fast logic synthesis for RRAM-based in-memory computing using majority-inverter graphs," in *Proc. Design, Autom. Test Europe Conf. Exhibit. (DATE)* 2016, pp. 948–953.
- [15] A. Velasquez and S. Jha, "Automated synthesis of crossbars for nanoscale computing using formal methods," in *Proc. NANOARCH*, 2015, pp. 130–136.
- [16] S. Thijssen, S. K. Jha, and R. Ewetz, "COMPACT: Flow-based computing on nanoscale crossbars with minimal semiperimeter," in *Proc. DATE*, 2021, pp. 232–237.
- [17] R. Ronen et al., "The bitlet model: A parameterized analytical model to compare PIM and CPU systems," *ACM J. Emerg. Technol. Comput. Syst.*, vol. 18, no. 2, pp. 1–29, 2022.
- [18] H. Akinaga and H. Shima, "Resistive random access memory (ReRAM) based on metal oxides," *Proc. IEEE*, vol. 98, no. 12, pp. 2237–2251, Dec. 2010.
- [19] G. W. Burr et al., "Phase change memory technology," *J. Vac. Sci. Technol. B*, vol. 28, no. 2, pp. 223–262, 2010.
- [20] Y. Huai et al., "Spin-transfer torque MRAM (STT-MRAM): Challenges and prospects," *AAPPS Bull.*, vol. 18, no. 6, pp. 33–40, 2008.
- [21] M. Hu et al., "Memristor-based analog computation and neural network classification with a dot product engine," *Adv. Mater.*, vol. 30, no. 9, 2018, Art. no. 1705914.
- [22] M. Imani, S. Gupta, Y. Kim, and T. Rosing, "FloatPIM: In-memory acceleration of deep neural network training with high precision," in *Proc. 46th Int. Symp. Comput. Archit.*, 2019, pp. 802–815.
- [23] M. R. H. Rashed, S. K. Jha, and R. Ewetz, "Hybrid analog-digital in-memory computing," in *Proc. IEEE/ACM Int. Conf. Comput. Aided Design (ICCAD)*, 2021, pp. 1–9.
- [24] D. Ielmini and H.-S. P. Wong, "In-memory computing with resistive switching devices," *Nat Electron.*, vol. 1, no. 6, pp. 333–343, 2018.
- [25] O. Mutlu, S. Ghose, J. Gómez-Luna, and R. Ausavarungnirun, "A modern primer on processing in memory," in *Emerging Computing: From Devices to Systems: Looking Beyond Moore and Von Neumann*. Singapore: Springer, 2022, pp. 171–243.
- [26] W. Haensch et al., "Compute in-memory with non-volatile elements for neural networks: A review from a co-design perspective," *Adv. Mater.*, vol. 35, no. 37, 2022, Art. no. 2204944.
- [27] C. Xu et al., "Overcoming the challenges of crossbar resistive memory architectures," in *Proc. HPCA*, 2015, pp. 476–488.
- [28] Z. Liao, J. Fu, and J. Wang, "Ameliorate performance of memristor-based ANNs in edge computing," *IEEE Trans. Comput.*, vol. 70, no. 8, pp. 1299–1310, Aug. 2021.
- [29] Z. Zhao et al., "Hardware-software co-design of slimmed optical neural networks," in *Proc. 24th Asia South Pacific Design Autom. Conf.*, 2019, pp. 705–710.
- [30] C. Feng et al., "A compact butterfly-style silicon photonic-electronic neural chip for hardware-efficient deep learning," *ACS Photon.*, vol. 9, no. 12, pp. 3906–3916, 2022.
- [31] F. Arute et al., "Quantum supremacy using a programmable superconducting processor," *Nature*, vol. 574, no. 7779, pp. 505–510, 2019.
- [32] G. Li, A. Wu, Y. Shi, A. Javadi-Abhari, Y. Ding, and Y. Xie, "On the co-design of quantum software and hardware," in *Proc. 8th Annu. ACM Int. Conf. Nanoscale Comput. Commun.*, 2021, pp. 1–7.
- [33] H. Kim, Y. Kim, S. Ryu, and J.-J. Kim, "Algorithm/hardware co-design for in-memory neural network computing with minimal peripheral circuit overhead," in *Proc. 57th ACM/IEEE Design Autom. Conf. (DAC)*, 2020, pp. 1–6.
- [34] M. R. H. Rashed, S. Thijssen, S. K. Jha, F. Yao, and R. Ewetz, "Stream: Towards read-based in-memory computing for streaming based data processing," in *Proc. 27th Asia South Pacific Design Autom. Conf. (ASP-DAC)*, 2022, pp. 690–695.
- [35] K. S. Brace, R. L. Rudell, and R. E. Bryant, "Efficient implementation of a BDD package," in *Proc. DAC*, 1990, pp. 40–45.
- [36] S.-I. Minato et al., "Shared binary decision diagram with attributed edges for efficient Boolean function manipulation," in *Proc. DAC*, 1990, pp. 52–57.
- [37] M. Wang et al., "A selector device based on graphene-oxide heterostructures for memristor crossbar applications," *Appl. Phys. A, Solids Surf.*, vol. 120, no. 2, pp. 403–407, 2015.
- [38] A. Shafiee et al., "ISAAC: A convolutional neural network accelerator with in-situ analog arithmetic in crossbars," *ACM SIGARCH Comput. Archit. News*, vol. 44, no. 3, pp. 14–26, 2016.
- [39] A. Ranjan, S. Jain, J. R. Stevens, D. Das, B. Kaul, and A. Raghunathan, "X-MANN: A crossbar based architecture for memory augmented neural networks," in *Proc. 56th Annu. Design Autom. Conf.*, 2019, pp. 1–6.

- [40] Z. Zhu et al., "MNSIM 2.0: A behavior-level modeling tool for memristor-based neuromorphic computing systems," in *Proc. Great Lakes Symp. VLSI*, 2020, pp. 83–88.
- [41] J. Solanki, K. Beckmann, J. Pelton, N. Cady, and M. Liehr, "Effect of resistance variability in vector matrix multiplication operations of 1T1R ReRAM crossbar arrays using an embedded test platform," in *Proc. IEEE 32nd Microelectron. Design Test Symp. (MDTS)*, 2023, pp. 1–5.
- [42] M. Liehr, J. Hazra, K. Beckmann, S. Rafiq, and N. Cady, "Impact of switching variability of 65nm CMOS integrated hafnium dioxide-based ReRAM devices on distinct level operations," in *Proc. IEEE Int. Integr. Rel. Workshop (IIRW)*, 2020, pp. 1–4.
- [43] S. Kvatinisky, G. Satat, N. Wald, E. G. Friedman, A. Kolodny, and U. C. Weiser, "Memristor-based material implication (IMPLY) logic: Design principles and methodologies," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 22, no. 10, pp. 2054–2066, Oct. 2014.
- [44] D. E. Rodriguez, J. E. Van Nostrand, S. Jha, and A. Velasquez, "Computation of Boolean formulas using sneak paths in crossbar computing," U.S. Patent 9 319 047, 2016.
- [45] J. Borghetti, G. S. Snider, P. J. Kuekes, J. J. Yang, D. R. Stewart, and R. S. Williams, "'Memristive' switches enable 'stateful' logic operations via material implication," *Nature*, vol. 464, no. 7290, pp. 873–876, 2010.
- [46] R. B. Hur et al., "SIMPLE MAGIC: Synthesis and in-memory mapping of logic execution for memristor-aided logic," in *Proc. ICCAD*, 2017, pp. 225–232.
- [47] S. Chakraborti, P. V. Chowdhary, K. Datta, and I. Sengupta, "BDD based synthesis of boolean functions using memristors," in *Proc. 9th Int. Design Test Symp. (IDT)*, 2014, pp. 136–141.
- [48] R. Ben-Hur et al., "SIMPLER MAGIC: Synthesis and mapping of in-memory logic executed in a single row to improve throughput," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 39, no. 10, pp. 2434–2447, Oct. 2020.
- [49] S. Shirinzadeh, M. Soeken, P.-E. Gaillardon, and R. Drechsler, "Logic synthesis for RRAM-based in-memory computing," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 37, no. 7, pp. 1422–1435, Jul. 2018.
- [50] D. Bhattacharjee, R. Devadoss, and A. Chattopadhyay, "ReVAMP: ReRAM based VLIW architecture for in-memory computing," in *Proc. Design, Autom. Test Europe Conf. Exhibit. (DATE)*, 2017, pp. 782–787.
- [51] D. Bhattacharjee, L. Amafu, and A. Chattopadhyay, "Technology-aware logic synthesis for ReRAM based in-memory computing," in *Proc. Design, Autom. Test Europe Conf. Exhibit. (DATE)*, 2018, pp. 1435–1440.
- [52] D. Bhattacharjee, Y. Tavva, A. Easwaran, and A. Chattopadhyay, "Crossbar-constrained technology mapping for ReRAM based in-memory computing," *IEEE Trans. Comput.*, vol. 69, no. 5, pp. 734–748, May 2020.
- [53] D. Chakraborty and S. Jha, "Automated synthesis of compact crossbars for sneak-path based in-memory computing," in *Proc. DATE*, 2017, pp. 770–775.
- [54] F. Somenzi, *CUDD: CU Decision Diagram Package—Release 2.4.0*, Univ. Colorado, Boulder, CO, USA, 2012.
- [55] D. B. West et al., *Introduction to Graph Theory*, vol. 2. Upper Saddle River, NJ, USA: Prentice Hall, 2001.
- [56] R. E. Bryant, "Symbolic Boolean manipulation with ordered binary-decision diagrams," *ACM Comput. Surveys*, vol. 24, no. 3, pp. 293–318, 1992.
- [57] S. Panda, F. Somenzi, and B. F. Plessier, "Symmetry detection and dynamic variable ordering of decision diagrams," in *Proc. IEEE/ACM Int. Conf. Comput.-Aided Design*, 1994, pp. 628–631.
- [58] R. Wille et al., "RevLib: An online resource for reversible functions and reversible circuits," in *Proc. ISMVL*, 2008, pp. 220–225.
- [59] L. Amarú, P.-E. Gaillardon, and G. De Micheli, "The EPFL combinational benchmark suite," in *Proc. 24th Int. Workshop Logic Synth. (IwLS)*, 2015, pp. 1–5.
- [60] F. Brglez, P. Pownall, and R. Hum, "Accelerated ATPG and fault grading via testability analysis," in *Proc. IEEE Int. Symp. Circuits Syst.*, 1985, pp. 695–698.
- [61] S. Gudaparthi, S. Narayanan, R. Balasubramonian, E. Giacomini, H. Kambalasubramanyam, and P.-E. Gaillardon, "Wire-aware architecture and dataflow for CNN accelerators," in *Proc. 52nd Annu. IEEE/ACM Int. Symp. Microarchit.*, 2019, pp. 1–13.
- [62] G. Krishnan, S. K. Mandal, C. Chakrabarti, J.-S. Seo, U. Y. Ogras, and Y. Cao, "Interconnect-aware area and energy optimization for in-memory acceleration of DNNs," *IEEE Design Test*, vol. 37, no. 6, pp. 79–87, Dec. 2020.

- [63] T. Zhang, C. Xu, K. Chen, G. Sun, and Y. Xie, "3D-SWIFT: A high-performance 3D-stacked wide IO DRAM," in *Proc. 24th Great Lakes Symp. VLSI*, 2014, pp. 51–56.
- [64] P. L. Thangkhiew, R. Gharipinde, and K. Datta, "Efficient mapping of Boolean functions to memristor crossbar using MAGIC NOR gates," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 65, no. 8, pp. 2466–2476, Aug. 2018.



Sven Thijssen received the bachelor's degree in informatics from KU Leuven, Leuven, Belgium, in 2018, and the master's degree in computer science from the University of Central Florida (UCF), Orlando, FL, USA, in 2021, where he is currently pursuing the Ph.D. degree in computer science.

His research interests are in-memory computing and beyond von Neumann computing.

Mr. Thijssen has received the ORCGS Doctoral Fellowship from UCF in 2020 and in 2021 he has received a Best Paper Nomination at DATE.



Muhammad Rashedul Haq Rashed received the bachelor's degree in electrical and electronics engineering from Bangladesh University of Engineering and Technology, Dhaka, India, in 2015. He is currently pursuing the Ph.D. degree in computer engineering with the University of Central Florida, Orlando, FL, USA.

His research interests include EDA for emerging computing paradigms, computer-aided design, and artificial intelligence.

Mr. Rashed has received a Best Paper Nomination at ICCAD 2022.



Sumit Kumar Jha (Senior Member, IEEE) received the B.Tech. degree (Hons.) in computer science and engineering from the Indian Institute of Technology Kharagpur, Kharagpur, India, in 2004, and the Ph.D. degree in computer science from Carnegie Mellon University, Pittsburgh, PA, USA, in 2010.

He is a Professor of Computer Science with the Florida International University, Miami, FL, USA. He has worked on Research and Development Problems with Microsoft Research, India; General Motors; INRIA France; and the Air Force Research Lab Information Directorate. His research has been supported by the National Science Foundation, DARPA, the Office of Naval Research, the Air Force Office of Scientific Research, the Oak Ridge National Laboratory, the Royal Bank of Canada, the Florida Center for Cybersecurity, the Air Force Research Laboratory, and the National Nuclear Security Administration.

Dr. Jha is a recipient of the IEEE Orlando Engineering Educator Excellence Award. He was awarded the prestigious Air Force Young Investigator Award and his research has led to four Best Paper Awards. He is a Full Member of the Sigma Xi.



Rickard Ewetz (Member, IEEE) received the M.S. degree in applied physics and electrical engineering from Linköpings Universitet, Linköping, Sweden, in 2011, and the Ph.D. degree in electrical and computer engineering from Purdue University, West Lafayette, IN, USA, in 2016.

He is currently an Assistant Professor with the Electrical and Computer Engineering Department, University of Central Florida, Orlando, FL, USA. His research interests include physical design and computer-aided design for in-memory computing

using emerging technologies.

Dr. Ewetz has Best Paper Nominations from ASP-DAC 2019 and DATE 2021.