CHOP: Clustered Hybrid Optimization for Logic Synthesis With Self-Supervised Prediction

Rongliang Fu, *Graduate Student Member, IEEE*, Ran Zhang, Ziyang Zheng, *Student Member, IEEE*, Zhengyuan Shi, Yuan Pu, Junying Huang, Bei Yu, *Senior Member, IEEE*, Qiang Xu, *Senior Member, IEEE*, and Tsung-Yi Ho, *Fellow, IEEE*

***Abstract*— Hybrid optimization is emerging as a promising approach for enhancing logic synthesis by combining multiple logic optimization methods. This article introduces CHOP, a novel framework that leverages vertex-level feature extraction and optimization-aware clustering to achieve significant improvements in the efficiency and quality of logic synthesis. Unlike existing cluster-level approaches that follow a "partition-then-predict" paradigm, CHOP operates under a "predict-then-partition" approach where the optimization method most appropriate for a given circuit vertex is determined by analyzing its local structural and functional context. To capitalize on this insight, CHOP extracts a subgraph centered on each vertex and calculates its circuit cost, which serves as a vertex feature for the subsequent circuit partitioning process. Furthermore, we introduce a self-supervised prediction model to obtain these vertex features efficiently. Following circuit partitioning, CHOP determines the optimization method for each partition and then iteratively applies these methods to derive the optimized circuit. The experimental results from logic optimization and LUT mapping tasks on the EPFL benchmark demonstrate the effectiveness and efficiency of CHOP.**

***Index Terms*— Circuit partitioning, hybrid logic optimization, logic synthesis, self-supervision model.**

I. Introduction

LOGIC synthesis plays a crucial role in modern electronic design automation (EDA) [1], involving logic optimization and technology mapping. Early efforts focused on simplifying Boolean expressions using algebraic techniques. As the complexity of digital circuits increased, the need for more advanced optimization techniques led to the development of multilevel logic optimization [2]. This approach considers entire networks of logic gates, rather than isolated expressions, thereby enhancing optimization efficacy. The introduction of directed acyclic graphs (DAGs) to represent Boolean networks enabled more effective logic optimization by facilitating the exploration of different structural configurations. Various types of graph structures, such as AND-inverter graphs (AIGs) [3], majority-inverter graphs (MIGs) [4], XOR-AIGs (XAGs) [5], and XOR-MIGs (XMGs) [6], have been developed to effectively capture different logic functions and facilitate algebraic manipulation. For example, Fig. 1 shows the different graph representations for the 1-bit full adder. The AIG representation requires 8 nodes and has a depth of 5, while the MIG representation requires only three nodes and has a depth of 2. The XAG representation requires five nodes and has a depth of 3, while the XMG representation requires only three nodes and has a depth of 2. This comparison underscores the diverse capabilities of different graph structures in optimizing logic functions.

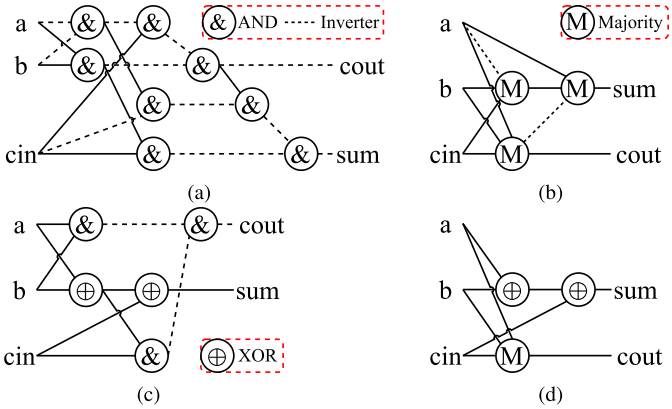


Fig. 1. Graph representations for the 1-bit full adder using (a) AIG with eight nodes and a depth of 5, (b) MIG with three nodes and a depth of 2, (c) XAG with five nodes and a depth of 3, and (d) XMG with three nodes and a depth of 2.

The introduction of these graphs has facilitated the development of numerous logic optimization techniques, such as rewriting [6], [7], [8], [9], resynthesis [10], [11], and resubstitution [12], [13], [14]. Rewriting involves optimizing the number of gates or depth of logic circuits by rearranging or substituting specific Boolean networks within the circuit with their equivalent forms. Resynthesis focuses on the comprehensive reconstruction or reorganization of the original subcircuit into a new logic structure. In contrast, resubstitution involves

Received 7 July 2025; revised 23 October 2025; accepted 17 December 2025. Date of publication 5 January 2026; date of current version 25 August 2026. This work was supported in part by the Research Grants Council of Hong Kong SAR under Grant CUHK14207523, and in part by the National Natural Science Foundation of China under Grant 62302477. This article was recommended by Associate Editor H. Zheng. *(Corresponding author: Junying Huang.)*

Rongliang Fu, Ziyang Zheng, Zhengyuan Shi, Yuan Pu, Bei Yu, Qiang Xu, and Tsung-Yi Ho are with the Department of Computer Science and Engineering, The Chinese University of Hong Kong, Hong Kong, China (e-mail: rlfu@cse.cuhk.edu.hk; zyzheng23@cse.cuhk.edu.hk; zyshi@cse.cuhk.edu.hk; ypu@cse.cuhk.edu.hk; byu@cse.cuhk.edu.hk; qxu@cse.cuhk.edu.hk; tyho@cse.cuhk.edu.hk).

Ran Zhang and Junying Huang are with the SKLP, Institute of Computing Technology, Chinese Academy of Sciences, Beijing 100190, China (e-mail: zhangran23s@ict.ac.cn; huangjunying@ict.ac.cn).

Digital Object Identifier 10.1109/TCAD.2026.3650860

1937-4151 © 2026 IEEE. All rights reserved, including rights for text and data mining, and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission.
See https://www.ieee.org/publications/rights/index.html for more information.
Authorized licensed use limited to: Chinese University of Hong Kong. Downloaded on September 13,2026 at 14:37:34 UTC from IEEE Xplore. Restrictions apply.

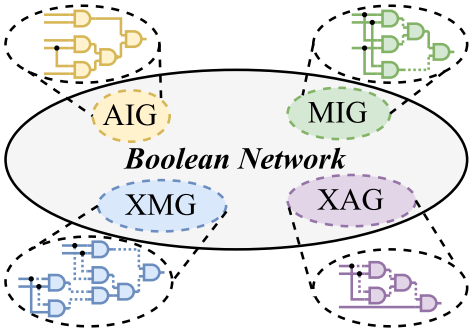


Fig. 2. Hybrid optimization example.

replacing redundant nodes with existing ones, thereby avoiding the introduction of new nodes. However, these methods mainly focus on a single type of graph representation.

Due to the inherent complexity of graph structures, hybrid optimization, a combination of these circuit graph optimization techniques, has the potential to significantly enhance circuit performance [15], [16], [17], [18], as shown in Fig. 2. Luca et al. proposed MixSyn [15], which combines different logic optimization techniques to improve circuit performance. MixSyn first performs XOR optimization for a given Boolean network, and then directly splits the optimized Boolean network into two parts, composed of 1) XOR nodes and 2) remaining nodes, respectively. Finally, it performs AND/OR optimization on the second part, and then merges the resulting network and the first part to obtain the final optimized Boolean network. Thus, MixSyn merely combines XOR and AND/OR optimization techniques, without considering the unique characteristics of the circuit design. Neto et al. [16] proposed LSOracle, which identifies clusters within a circuit and assigns a specific graph optimization technique to each cluster to explore a better circuit design. Although LSOracle achieves promising improvements in the area-delay product (ADP), it has notable shortcomings as follows.

1) It relies on the general graph partition tool, KaHyPar [19], to partition the circuit, without accounting for the unique characteristics of circuit design.
2) It represents Boolean functions using Karnaugh-map images and employs a traditional neural network model to predict the optimization method of each cluster, which may limit prediction performance on complex circuits.

Besides, Pu et al. [17] proposed HeLO, a heterogeneous DAG-based logic optimization framework that combines graph learning and hierarchical clustering. HeLO clusters logic components based on structural and functional similarity and uses a graph neural network (GNN) model to generate topological-functional embeddings for predicting the best-fit DAG type for each subcircuit. However, existing hybrid optimization methods face several fundamental limitations that constrain their effectiveness. First, cluster-level predictions lack the fine-grained vertex-level feature extraction necessary for optimal Boolean network partitioning. Second, the sequential "partition-then-predict" paradigm prevents feedback from prediction results to improve clustering quality. Finally, these methods often rely on general-purpose clustering techniques that fail to capture the intricate relationship between circuit structures and their optimal optimization strategies, leading to suboptimal partitioning decisions.

TABLE I
COMPARISON OF CHOP AND EXISTING HYBRID LOGIC OPTIMIZATION FRAMEWORKS

Aspect	LSOracle [16]	HeLO [17]	CHOP
Prediction Level	Cluster-level	Cluster-level	**Vertex-level**
Prediction Phase	After clustering	After clustering	**Before clustering**
Clustering Method	K-way Hypergraph	Agglomerative	**Hierarchical Clustering**

To address these challenges, this article proposes CHOP, a hybrid optimization framework that incorporates circuit-specific characteristics into the optimization process. As shown in Table I, CHOP differs from traditional logic optimization frameworks. Unlike previous methods that rely on cluster-level prediction, CHOP extracts structural and functional features from each vertex in a Boolean network. These features guide the network partitioning and enable the selection of optimal optimization strategies for each cluster, thereby effectively managing the complexity of modern circuit structures. Overall, this article makes the following contributions.

1) *Paradigm Shift:* A fundamental shift from "partition-then-predict" to "predict-then-partition" paradigm, enabling optimization-aware circuit clustering that naturally aligns with synthesis objectives.
2) *Vertex-Level Prediction:* A novel vertex-level feature extraction method using self-supervised learning that captures fine-grained optimization preferences, surpassing traditional cluster-level approaches.
3) *Hierarchical Clustering:* A circuit-specific hierarchical partitioning algorithm that integrates both structural topology and functional optimization tendencies through adaptive masking strategies.
4) *Optimization Selection:* An intelligent optimization method selection strategy that leverages predicted vertex features to determine the most suitable synthesis technique for each cluster.

Moreover, experimental results on the EPFL benchmark [20] demonstrate the effectiveness and efficiency of our proposed framework compared with the state-of-the-art.

1) For logic optimization, CHOP achieves average node-depth-product (NDP) reductions of 16.07% and 12.39%, with average speeds of 1.92 × and 2.98 ×, compared with the state-of-the-art hybrid optimization methods.
2) For 6-LUT mapping, CHOP achieves average ADP reductions of 17.99% and 36.42%, with average speeds of 1.78 × and 2.48 ×, compared with the state-of-the-art hybrid optimization methods.

The rest of this article is organized as follows. Section II provides some background information. Section III introduces the terminology used and our problem formulation. Section IV introduces our proposed hybrid optimization framework. Section V presents the experimental methods and results. Finally, Section VI summarizes our work.

II. PRELIMINARIES

A. Boolean Network

A Boolean network [21] is a mathematical model that characterizes a system of interconnected binary variables, where

Authorized licensed use limited to: Chinese University of Hong Kong. Downloaded on September 13,2026 at 14:37:34 UTC from IEEE Xplore. Restrictions apply.

each variable can assume one of two possible values: true (1) or false (0). In a Boolean network, the state of each vertex at time $t+1$ is determined by a Boolean function that depends on the states of a subset of the other vertices at time t. Thus, the dynamics of the network are governed by the following equation:

$$x_i(t+1) = f_i\left(x_{i_1}(t), x_{i_2}(t), \ldots, x_{i_k}(t)\right) \tag{1}$$

where $x_i(t)$ is the state of vertex i at time t, f_i is the Boolean function associated with vertex i, and $x_{i_1}(t), x_{i_2}(t), \ldots, x_{i_k}(t)$ are the states of the k input vertices that regulate vertex i. The simplicity of Boolean networks, coupled with their capacity to model complex interactions, makes them a powerful tool for both theoretical analysis and practical applications in diverse domains. Notably, in digital circuit design, Boolean networks facilitate the representation of both combinational and sequential circuits, thereby enhancing our understanding and manipulation of these systems.

B. Logic Representation

In logic synthesis, Boolean networks are typically modeled using DAGs, which provide multiple representation structures. These representations provide structured and efficient ways for describing and optimizing Boolean networks, facilitating the analysis and transformation of Boolean functions. The following four structures are prevalent in representing Boolean networks, each with distinct characteristics and applications.

- AIG [3] is a widely used representation in which the vertices correspond to AND operations, and the edges represent signals or their inversions.
- MIG [4] incorporates majority operations [a 3-input majority is $\mathrm{M}(a,b,c) = (a \wedge b) \vee (a \wedge c) \vee (b \wedge c)$] alongside inverters, allowing for a more compact and expressive representation of certain logic functions.
- XAG [5] merges XOR operations with AND operations and inverters, particularly advantageous for applications involving parity checks and error detection, where XOR operations play a crucial role.
- XMG [6] integrates XOR operations, majority operations, and inverters, suitable for a wide range of applications, including quantum circuit synthesis and cryptographic function implementations.

Overall, as stated in Theorem 1, the XMG has the most expressive power, as it can represent all other structures. The AIG is the simplest structure, while the MIG and XAG are intermediate representations that can capture specific logic functions more efficiently than the AIG.

Theorem 1: AIG $\subset$ MIG / XAG $\subset$ XMG.

Proof: In all graphs, inverters are represented by complemented edge markers.

- AIG $\subset$ XAG Since AIG nodes only have AND operations, an AIG is trivially an XAG where all XOR nodes are unused, i.e., AIG $\subset$ XAG.

Algorithm 1 Leiden Algorithm

Input: A network $G(V, E)$.
Output: Partition of vertices into communities cs.

```
cs ← {{v} | v ∈ V}
while modularity Q improves do
    for v ∈ V do
        cs ← Move v for maximizing Q
    cs ← Identify potential subcommunities in cs
    G(V, E) ← Each c ∈ cs is regarded as a vertex
return cs
```

- AIG $\subset$ MIG / XMG Since an AND operation $f = a \wedge b$ is equivalent to a 3-input majority operation with one input tied to 0: $f = \mathrm{M}(a, b, 0)$, an AIG node is always a special case of a majority node, i.e., AIG $\subset$ MIG / XMG.
- MIG $\subset$ XMG Since MIG nodes only have majority operations, an MIG is trivially an XMG where all XOR nodes are unused, i.e., MIG $\subset$ XMG.
- XAG $\subset$ XMG Since XAG nodes have XOR and AND operations and an AND node is always a special case of a majority node, an XAG is always a special case of an XMG, i.e., XAG $\subset$ XMG.

Therefore, we conclude AIG $\subset$ MIG / XAG $\subset$ XMG. □

C. Community Detection

In complex networks, nodes tend to cluster into relatively dense groups, commonly referred to as communities, whose modular structure is typically unknown in advance. Community detection [22] thus becomes a crucial problem for understanding the structure of large and complex networks. Among the best-known methods for community detection are those based on optimizing modularity [23], a measure of partition quality, such as the Leuven [24] and Leiden [25] algorithms. These methods identify communities by maximizing modularity Q, effectively performing network partitioning tasks. Although the definitions of Q vary slightly, the primary formula is as follows:

$$Q = \frac{1}{2m} \sum_{u,v} \left(A_{u,v} - \frac{k_U . K_{.v}}{2m} \right) \delta(c_u, c_v) \tag{2}$$

where m is the total number of edges in the network, k_v is the degree of vertex v, A is the adjacency matrix of the network, and $\delta(c_u, c_v)$ is an indicator function that equals 1 if vertices u and v belong to the same community, and 0 otherwise.

The Leuven algorithm is characterized by its simplicity and elegance, optimizing the modularity through two fundamental phases: 1) local movement of nodes; and 2) aggregation of the network. In the local movement phase, individual nodes are relocated to the community that yields the largest increase in quality function. Subsequently, in the aggregation phase,

Authorized licensed use limited to: Chinese University of Hong Kong. Downloaded on September 13,2026 at 14:37:34 UTC from IEEE Xplore. Restrictions apply.

an aggregate network is constructed based on the partition obtained in the local movement phase, and each community in this partition is treated as a node in the aggregate network. The two phases are repeated until there is no further increase in the quality function. However, the Leuven algorithm may yield arbitrarily poorly connected communities, even internally disconnected communities in the worst case [25].

To address this problem, the Leiden algorithm [25] introduces a combination of smart local movement, fast local movement, and random neighbor movement to produce partitions in which all communities are well-connected. Also, the Leiden algorithm is faster than the Leuven algorithm and can find superior partitions. Algorithm 1 outlines the process of the Leiden algorithm, which optimizes modularity Q by iteratively refining community partitions. The algorithm begins with an initial partition where each vertex belongs to its own community (line 1). It then iteratively performs the following steps.

Local Moving Phase: Vertices are reassigned to neighboring communities to maximize modularity gain (lines 3–4).

Refinement Phase: The resulting communities are divided into smaller and well-connected subcommunities to ensure connectivity (line 5).

Aggregation Phase: The network is coarsened by treating each refined community as a single vertex (line 6).

This process repeats until convergence (i.e., no further improvement in modularity), producing a high-quality partition with connectivity guarantees (lines 2–6).

During hybrid logic optimization, it is often difficult to identify which vertices should be clustered together into a subcircuit to yield optimal optimization results. Consequently, this article will integrate community detection methods with logic synthesis to enhance circuit clustering. This approach aims to overcome the complexities inherent in modern circuit designs and facilitate the exploration of more effective circuit design strategies.

D. Representation Learning for Circuit

Circuit representation learning has emerged as a pivotal area in EDA, reflecting the broader trend in artificial intelligence toward developing general representations that can be applied to a variety of downstream tasks. Within this field, the DeepGate family of models [26], [27] has pioneered the use of GNNs to encode AIGs, enabling a wide range of EDA applications, such as testability analysis [28], power estimation [29], and SAT solving [30], [31]. Notable advancements include the introduction of Gamora [32] and HOGA [33], which enhance reasoning capabilities by effectively representing both logic gates and their associated cones. In addition, PolarGate [34] addresses the challenges related to functionality representation by leveraging the principles of ambipolar states. Despite the progress made with GNNs, challenges such as over-squashing and over-smoothing persist. To mitigate these issues, the recent work of DeepGate3 [35] utilizes DeepGate2 as a tokenizer and then leverages the global aggregation mechanism of transformers with a connective mask to enhance circuit representation. As a result, DeepGate3 demonstrates enhanced performance across various pretraining tasks at both gate-level and graph-level.

III. Problem Formulation

A. Terminology

A DAG serves as an effective representation of a Boolean network, denoted as $G(V, E)$. In this representation, the set of vertices V typically includes a collection of primary inputs (PIs), primary outputs (POs), and internal nodes within the network. The directed edge set $E \subseteq V \times V$ illustrates the connectivity relationships among the vertices, with the orientation of these edges indicating the direction of the data flow within the circuit. For any vertex $v \in V$, we define $\mathrm{FI}(v)$ and $\mathrm{FO}(v)$ to represent the sets of fan-in and fan-out vertices, respectively. Consequently, for any vertex $u \in \mathrm{FI}(v)$ or $v \in \mathrm{FO}(u)$, there exists a directed edge $(u, v) \in E$. In particular, for any $i \in$ PIs or $o \in$ POs, their respective sets of fan-in and fan-out vertices are empty. Furthermore, the depth of a Boolean network is determined as the maximum logic level among all POs, and the size of the network is the total number of vertices within G. For any vertex v, the logic level ι of a vertex v is recursively defined as $\iota(v) = \max_{u \in \mathrm{FI}(v)} \iota(u) + 1$.

B. Problem Formulation

Hybrid logic optimization involves the application of various circuit graphs and the corresponding logic optimization methods to a Boolean network. Since the circuit design often involves multiple metrics, such as delay and area, the optimization process seeks to achieve a balance among these metrics. This article specifically focuses on optimizing the ADP, a widely used metric in circuit design that effectively balances both delay and area. Hence, the optimization problem in this article can be formulated as follows.

Input: A Boolean network $G(V, E)$.

Output: An optimized circuit $G'(V', E')$.

Constraints: The logic functions remain the same before and after optimization.

Goal:

$$\min \{\mathrm{ADP} = \mathrm{delay} \times \mathrm{area}\} \tag{3}$$

where minimizing the ADP facilitates a more effective balance between delay and area, thereby optimizing the overall performance of the circuit.

This problem poses several substantial challenges. First, it is essential to quantify the closeness between vertices within a Boolean network. Although the relationships among vertices may not be immediately apparent from the circuit's representation, they do objectively exist and can be exploited to enhance the partition of G. We define the closeness between vertices u and v as $\rho_{u,v}$

$$\rho_{u,v} = \begin{cases} 1, & \text{if } (u, v) \in E \\ 0, & \text{otherwise} \end{cases} \tag{4}$$

Authorized licensed use limited to: Chinese University of Hong Kong. Downloaded on September 13,2026 at 14:37:34 UTC from IEEE Xplore. Restrictions apply.

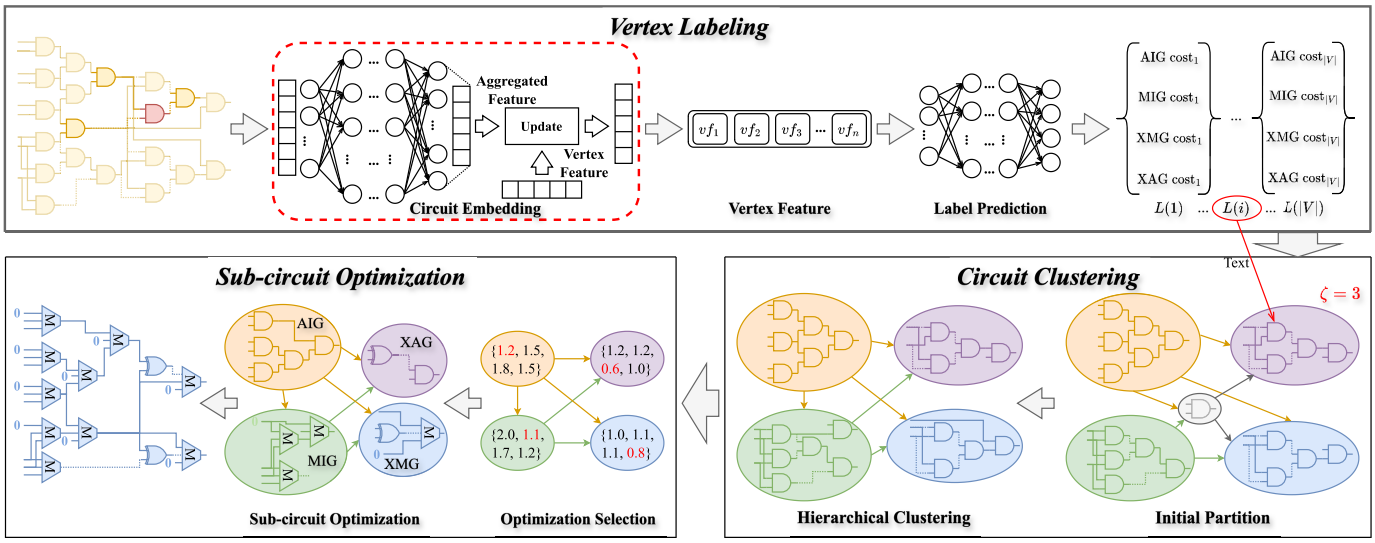


Fig. 3. Flow of the CHOP framework.

where (u, v) denotes a directed edge from u to v. Second, the effective partition of the Boolean network is vital, as the quality of the partition directly influences the outcome of the hybrid optimization process. Finally, executing all optimization methods on each cluster can be exceedingly time-consuming. If we can harness the properties of the cluster to directly identify which circuit representation and corresponding optimization method are most suitable for a given cluster, this predictive capability can substantially reduce the overall runtime of the optimization process. We will address these challenges within the framework of CHOP, thereby enhancing circuit design.

IV. CHOP

To address these limitations and incorporate circuit-specific characteristics into the hybrid optimization process, this section presents CHOP, a hybrid optimization framework that fundamentally shifts from cluster-level to vertex-level analysis. Unlike traditional methods that rely on postpartitioning prediction, CHOP extracts structural and functional features from each vertex in a Boolean network and uses these features to guide optimization-aware partitioning. This paradigm shift enables more precise clustering that naturally aligns with optimization objectives, thereby addressing the core limitations of existing approaches.

As illustrated in Fig. 3, CHOP operates through three main stages: 1) *vertex labeling*, where each vertex is assigned an optimization preference label based on its local structural and functional context; 2) *circuit clustering*, where vertices with similar labels are grouped using a hierarchical clustering algorithm; and 3) *subcircuit optimization*, where each cluster is optimized using its most suitable optimization method, and the results are merged to form the final optimized circuit.

A. Vertex Labeling

Unlike general graphs, Boolean networks exhibit not only complex interconnect topologies but also rich functional semantics, as each vertex corresponds to a specific logic gate or subfunction. This dual complexity poses significant challenges for downstream optimization tasks, which often require detailed, fine-grained knowledge of circuit characteristics at the vertex level. To address these challenges, we propose a learning-based feature extraction method that efficiently encodes both the structural and functional attributes of each vertex within the network. In contrast to traditional approaches that rely on manual feature engineering, which are often time-consuming and difficult to generalize, our method leverages a pretrained neural model to infer expressive representations of circuit elements automatically.

1) Feature Selection: Given that the optimization potential of a circuit vertex is fundamentally determined by its local structural context and functional dependencies, we utilize the features of a subgraph surrounding a vertex to represent its optimization characteristics. This approach is theoretically grounded in the principle that logic optimization techniques operate on localized circuit patterns, where the effectiveness of a specific optimization method depends heavily on the surrounding logic structure. In the context of a Boolean network, we introduce a partial order $\preccurlyeq_k$, defined such that for any two vertices u and v, the relation $u \preccurlyeq_k v$ holds if there exists a path from u to v with a length of at most k. For a specific vertex $v \in V$, we define its expanded subgraph with degree k as follows:

$$g_k(v) = \{u \in V : u \preccurlyeq_k v\} \cup \{u \in V : v \preccurlyeq_k u\}. \quad (5)$$

As an illustrative example, the red vertex located in the upper-left corner of Fig. 3 and its surrounding highlighted vertices collectively construct a subgraph with $k = 1$.

After obtaining the expanded subgraph $g_k(v)$ for vertex v, the next step involves identifying the optimal circuit graph for $g_k(v)$. The ADP values produced through different optimization methods for the expanded subgraph $g_k(v)$ may exhibit similarities. Therefore, relying solely on the circuit graph with minimum cost cannot accurately reflect the characteristics of a vertex. This article denotes the costs of different

Authorized licensed use limited to: Chinese University of Hong Kong. Downloaded on September 13,2026 at 14:37:34 UTC from IEEE Xplore. Restrictions apply.

circuit graphs applied to the expanded subgraph as a vector: $\text{Cost} = [\text{cost}_1, \text{cost}_2, \ldots, \text{cost}_n]$, where cost_i represents the ADP generated using the optimization method for the ith structure of circuit graphs. To facilitate the analysis of different vertex features, we apply softmax normalization to generate the final label vector L for each vertex

$$L = \text{softmax}\left(\text{Cost}\right). \tag{6}$$

The processed subgraph set $S(G)$, together with the label set L, serves as training data for learning the labels of vertices within the Boolean network. These labels subsequently guide both the partition and prediction tasks. In Section IV-A.2, we will introduce the GNN architecture adopted for label prediction.

2) Model Selection: To reduce the time-consuming process of acquiring vertex labels L, an accurate and efficient prediction model is essential to lower the overall computational cost. Our task requires considering both structural and functional aspects of logic synthesis, making models that focus solely on functionality, such as FGNN [36] and PolarGate [34], unsuitable. In addition, the partitioned subcircuits in our dataset vary widely in size, demanding a model that generalizes well across different scales. Since incorporating model predictions adds time overhead, it is crucial to use a neural network with low computational complexity to keep the overall runtime reasonable. Given these criteria, we chose DeepGate, a state-of-the-art transformer-based model that captures both structural and functional features while offering strong generalization capabilities [35]. Note that our framework remains flexible and configurable, allowing alternative models to be integrated as needed.

3) Label Prediction: According to the subgraph expansion method in Section IV-A.1, we obtain a set of subgraphs $S(G)$. For each subgraph $g_k(v)$, we prepare the corresponding label $L(v) = [l_1(v), l_2(v), \ldots, l_n(v)]$. Subsequently, we train the DeepGate model with these subgraphs and their corresponding labels, as illustrated in the upper middle part of Fig. 3. Given a vertex v, to predict its corresponding label $\hat{L}(v)$, we first get the expanded subgraph $g_k(v)$ and then use the DeepGate model to embed the subgraph

$$\text{VF} = \text{DeepGate}\left(g_k\left(v\right)\right). \tag{7}$$

$\text{VF} = \{\text{vf}_1, \text{vf}_2, \ldots, \text{vf}_{|V|}\}$ represents the embedding sequence of the subgraph, where $|V|$ denotes the number of vertices $\{v_i, i = 1, 2, \ldots, |V|\}$ within $g_k(v)$, and vf_{v_i} represents the embedding sequence of vertex v_i. Then, we predict the $\hat{L}(v)$ by concatenating both functional and structural embedding and feeding the embedding into a multilayer perceptron (MLP)

$$\hat{L}\left(v\right) = \text{MLP}\left(\text{VF}\right). \tag{8}$$

Finally, we calculate the loss between the predicted label $\hat{L}(v)$ and the ground truth $L(v)$ after the softmax function

$$\text{loss} = \left\|\text{softmax}\left(L\left(v\right)\right), \text{softmax}\left(\hat{L}\left(v\right)\right)\right\|_1. \tag{9}$$

During the process of using DeepGate to predict the label set L for all vertices in the Boolean network G, each vertex v iteratively aggregates information from its neighbors and updates its own representation. This iterative message-passing mechanism is referred to as circuit embedding. The outcome of this embedding process is formally defined in (7), and its structural location is highlighted by the red dotted region in the upper half of Fig. 3. The resulting vertex representations are subsequently fed into an MLP, as described in (8), to generate the final label predictions. This prediction stage corresponds to the area shown in the upper right corner of Fig. 3.

B. Circuit Clustering

As discussed in Section II-C, conventional partition or clustering techniques often struggle to accurately determine which vertices should be grouped together to form subcircuits that facilitate efficient and effective logic optimization. These methods typically rely on purely structural metrics or heuristic community detection algorithms, which may not align well with the functional and optimization-relevant characteristics of Boolean networks. Consequently, they often generate partitions that fail to capture the true logical boundaries and interaction patterns necessary for high-quality synthesis. To address this limitation, we propose a partition algorithm specifically tailored for logic optimization in Boolean networks. Unlike traditional approaches that pursue general-purpose structural decomposition, our algorithm focuses on identifying subcircuits that are particularly conducive to logic simplification and performance improvement. The algorithm leverages both structural features and optimization potential to guide partition decisions in a manner that aligns with the requirements of synthesis and transformation tasks. The complete methodology, including algorithmic details, heuristics, and implementation considerations, is illustrated in Algorithm 2, which consists of two steps: 1) initial partition (lines 1–15) and 2) hierarchical clustering (lines 16–21).

1) Initial Partition: Given the intricate topological and functional structure of Boolean networks, coupled with the cone-based characteristics of most logic optimization techniques, it is essential to perform a high-quality initial partition before global clustering. Such an initial decomposition not only facilitates more effective downstream optimization but also ensures structural compatibility with cone-centric logic synthesis frameworks.

To this end, we introduce a specialized initial partitioning algorithm designed to identify structurally meaningful subgraphs. The process begins with the calculation of the logic levels ι, as defined in Section III-A, for all vertices in the Boolean network G (line 1). The algorithm then traverses the network G from its POs to effectively facilitate the initial partition. Specifically, all POs are inserted into a priority queue P, which organizes the elements in descending order according to their logic levels (line 2). This sorting strategy ensures that vertices with a greater logical depth are prioritized during cone expansion, thus increasing the potential for optimization. The initial partitioning proceeds iteratively (lines 4–15). At each iteration, the vertex with the highest logic level is extracted from P and designated as the root r of a new cone c (line 5). The cone c is then expanded through a traversal toward the input direction, using a breadth-first search strategy (lines 8–13). This expansion follows the direction opposite to

Authorized licensed use limited to: Chinese University of Hong Kong. Downloaded on September 13,2026 at 14:37:34 UTC from IEEE Xplore. Restrictions apply.

Algorithm 2 Circuit Clustering

Input: A Boolean network $G(V, E)$, labels L, ζ.
Output: Partition results cs.

1 $\iota \leftarrow$ Calculate the logic levels of all vertices
2 $P \leftarrow$ Get POs in G sorted by ι
3 $cs \leftarrow \emptyset$
4 **while** $P \neq \emptyset$ **do**
5 $\quad r \leftarrow \text{argmax}_{r \in P} \iota(r)$, $P \leftarrow P \setminus \{r\}$
6 $\quad c \leftarrow \emptyset$, $q \leftarrow \emptyset$
7 $\quad v \leftarrow r$
8 $\quad$ **while** $\iota(r) - \iota(v) < \zeta$ **do**
9 $\quad\quad$ **if** $\iota(r) - \iota(v) > \zeta/2 \wedge |c| < 2^{\zeta/2-1}$ **then**
10 $\quad\quad\quad$ **break**
11 $\quad\quad c \leftarrow c \cup \{v\}$
12 $\quad\quad q \leftarrow q \cup \text{FI}(v)$ ▷ Descendingly sort FI(v) by ι
13 $\quad\quad v \leftarrow q.\text{pop}()$
14 $\quad P \leftarrow P \cup \{u \mid u \in \text{FI}(v) \setminus c, v \in c\}$
15 $\quad cs \leftarrow cs \cup \{c\}$
16 $\psi \leftarrow 1$
17 **while** $|cs| \geq |L|$ **do**
18 $\quad V' \leftarrow$ Regard each $c \in cs$ as a vertex
19 $\quad E' \leftarrow$ Create edges among V' through mask ψ
20 $\quad cs \leftarrow$ Perform `Leiden` for $G'(V', E')$ with L
21 $\quad \psi \leftarrow \psi + 1$
22 **return** cs

the circuit's signal flow, effectively tracing functional dependencies from outputs to inputs and naturally exposing the underlying logic structure.

Notably, the inclusion of a vertex v into a cone c implies that v is no longer eligible to serve as a candidate root in subsequent cone construction iterations. This constraint underscores the importance of carefully maintaining the quality of each cone, as suboptimal selections in earlier stages may compromise the overall effectiveness of the partition and optimization process. To assess the quality of a cone, we primarily consider the number of vertices it encompasses. A cone with a greater number of vertices typically reflects a denser concentration of logic elements within a constrained range of ι difference. This dense grouping implies a higher likelihood of internal logic redundancy within the cone, which in turn enhances its potential for logic simplification and optimization. Consequently, striking an appropriate balance between cone size and logic optimization potential becomes essential. To achieve this balance, an early stop mechanism is introduced (lines 9 and 10), which serves as a dynamic control for the cone expansion process. Specifically, as the cone expands from its root vertex r, the algorithm monitors both the ι range and the size of the current cone. If the condition $\iota(r) - \iota(v) > \zeta/2$ is met for a newly encountered vertex v, and the current cone size satisfies $|c| < 2^{\zeta/2-1}$, then the expansion is halted prematurely. This stop condition prevents overextending a cone that is too small to yield meaningful optimization benefits. Otherwise, if the condition is not satisfied, all vertices that meet the criterion $\iota(r) - \iota(v) < \zeta$ are incorporated into the cone c.

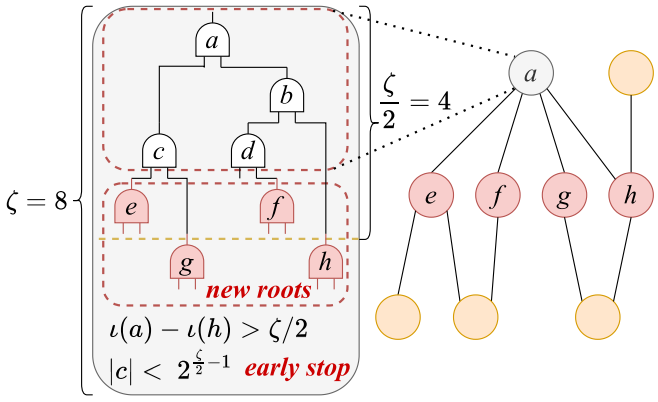


Fig. 4. Example of early stop mechanism.

An illustrative example of this early stop strategy is shown in Fig. 4, where the threshold parameter is set to $\zeta = 8$. During the expansion of a cone rooted at vertex a, the algorithm reaches vertex h and evaluates whether to continue expansion. At this stage, the cone contains four vertices, which is less than the threshold $2^{\zeta/2-1} = 8$. Therefore, the expansion is terminated early, and the vertices e, f, g, and h, which lie beyond the acceptable size depth tradeoff, are added to the candidate set P for potential use as roots in future cone construction. This example highlights how the mechanism effectively maintains both diversity in cone root selection and structural quality in logic clustering.

After expansion, the algorithm updates the priority queue P (line 14) and stores the constructed cone c in the partition set cs for subsequent processing (line 15). The algorithm terminates when the priority queue P is empty. As illustrated in the lower right portion of Fig. 3, this procedure yields a set of five distinct clusters, each representing a localized logic region. These clusters serve as the output of the initial partition stage and lay the foundation for subsequent hierarchical clustering and subcircuit optimization stages.

2) Hierarchical Clustering: Hierarchical clustering begins with the initial partition obtained from Section IV-B.1 and progressively merges these clusters to yield high-quality clustering results (lines 16–21). ψ is a masking that will be utilized in the subsequent hierarchical clustering phase, and it gradually increases with the number of iterations (line 16). Initially, each cluster $c \in$ cs is treated as an individual vertex and added to the vertex set V' (line 18). The edges in the intermediate graph G', denoted as E', are constructed based on the connectivity information from the original graph G, with the process governed by a masking ψ (line 19). Specifically, for any pair of vertices u and v belonging to clusters c_i and c_j, respectively, an edge is established between c_i and c_j if $\rho_{u,v} > 0$, that is, their logic level difference satisfies $\iota(u) - \iota(v) \leq \psi$.

This masking strategy provides two key benefits that enhance the quality and efficiency of the subsequent partition and optimization processes. First, from the perspective of Boolean function decomposition, a significant difference in the ι values between two vertices can be interpreted as an indication of the extraction of common factors. In such

Authorized licensed use limited to: Chinese University of Hong Kong. Downloaded on September 13,2026 at 14:37:34 UTC from IEEE Xplore. Restrictions apply.

scenarios, clustering these vertices within the same subgraph by retaining their connecting edge may obscure opportunities for more effective algebraic simplification. To mitigate this, the algorithm is designed to mask such edges, effectively ignoring their influence during clustering. This selective omission enables the clustering process to better explore the potential optimization. Second, from a timing or path-sensitivity perspective, vertices that reside on the critical path, or more generally on deeper logic paths, are typically more sensitive to interconnect changes. Retaining the edges between such vertices is essential for preserving the original timing characteristics and enabling delay-aware optimizations. The masking strategy, by design, preserves these structurally and temporally important edges, ensuring that clustering does not inadvertently degrade circuit performance. As a result, the clusters produced through this strategy naturally promote depth-aware logic grouping. This not only facilitates more effective optimization in subsequent logic synthesis stages but also contributes to the construction of logic networks with reduced logic depth and improved performance profiles.

In the original graph G, each vertex v is associated with a label $L(v)$, which reflects its cost information for different logic optimization algorithms. In the constructed graph G', this property must be preserved to ensure meaningful guidance during the hierarchical clustering process. Accordingly, for each vertex c in G', its label $L(c)$ is defined as follows:

$$L(c) = \frac{\Sigma_{v \in c} L(v)}{|c|} \tag{10}$$

where $|c|$ refers to the number of vertices in G which are contained in c. For an edge (c_i, c_j) in G', its weight w_{c_i,c_j} is defined as

$$w_{c_i,c_j} = \sum_{u \in c_i} \sum_{v \in c_j} \rho_{u,v} \times (\psi - \iota(u) + \iota(v)) + s_{c_i,c_j} \tag{11}$$

$$s_{c_i,c_j} = \sum_{\lambda=0}^{|L|-1} \mathbb{1}\left(r_{c_{i,\lambda}}, r_{c_{j,\lambda}}\right) \times 2^{|L|-r_{c_{i,\lambda}}} \tag{12}$$

where $r_{c_{i,\lambda}}$ represents the index value of $l_\lambda(c_i)$ in ascending order of $L(c_i)$. The indicator function $\mathbb{1}(r_{c_{i,\lambda}}, r_{c_{j,\lambda}})$ equals 1 if $r_{c_{i,\lambda}} = r_{c_{j,\lambda}}$ and equals 0 otherwise. When two vertices c_i and c_j exhibit high similarity s_{c_i,c_j}, this increases the tendency of these vertices to be grouped into the same cluster.

After constructing the graph G', the Leiden algorithm is applied to cluster the vertices within this graph (line 20). Following each clustering iteration, the masking is updated by $\psi = \psi + 1$ (line 21). Constraining ψ to a relatively small value enables the hierarchical clustering to capture variations across topological levels in a fine-grained way. This process is repeated iteratively until the number of resulting clusters falls below the specified threshold $|L|$. As illustrated in the lower part of Fig. 3, the hierarchical clustering procedure ultimately generates four distinct cluster groups, each visually distinguished by a different color.

The computational complexity analysis of Algorithm 2 reveals its efficiency advantages. During the initial partition stage, each vertex is evaluated once to determine its initial cluster assignment, resulting in a complexity of $O(|V|)$ for vertex processing and $O(|E|)$ for edge traversal (lines 4–15). In the subsequent hierarchical clustering stage, the masking filtering process depends on the number of edges in the Boolean network, yielding $O(|E|)$ complexity (line 19). Since the `Leiden` algorithm operates with linear complexity in the number of edges, this stage contributes $O(|E'|)$ where $|E'| \leq |E|$ due to masking (line 20). Assuming the algorithm performs I hierarchical iterations until convergence (lines 17–21), the overall time complexity can be expressed as $O(I \times |E|)$, where I is typically small in practice (usually $I \leq 10$ for most circuits).

C. Subcircuit Optimization

After circuit clustering, the Boolean network is divided into multiple clusters, each representing a localized subcircuit. To fully exploit the optimization potential of these subcircuits, it is necessary to select appropriate logic optimization strategies tailored to the structural and functional characteristics of each cluster. To this end, we propose a logic optimization selection algorithm that comprises two key components: tendency acquisition and subcircuit optimization. The tendency acquisition module is responsible for analyzing each cluster and inferring its optimization preferences or tendencies, which may reflect a specific Boolean network structure for which the subcircuit is more suitable. Based on this guidance, the subcircuit optimization module then applies the corresponding optimization method to each cluster, ensuring that the selected strategy aligns with the intrinsic logic structure and optimization potential of the subcircuit.

1) Tendency Acquisition: Each cluster $c \in$ cs is associated with a circuit graph optimization method most suitable for it. The tendency $T(c)$ can be estimated based on the label L of the vertices within c, enabling a more targeted and effective optimization strategy, i.e.,

$$T(c) = \text{argmin } L(c). \tag{13}$$

Since $L(c)$ is generated by aggregation of $L(v), v \in c$, choosing the minimum value of $L(c)$ means that most vertices in c prefer to choose the $T(c)$th optimization algorithm. The lower left part of Fig. 3 shows an example, where the optimization method corresponding to the index marked by red digits is selected for the cluster.

2) Subcircuit Optimization: Each cluster is subsequently optimized using its designated logic optimization technique. The optimized clusters are then merged into a complete Boolean network by aligning their PIs and POs based on their original correspondences. As a result, an optimized Boolean network is reconstructed. This process is illustrated in Fig. 3, where the final concatenation of optimized clusters into the output Boolean network is highlighted in blue.

D. Overall Flow of CHOP

In summary, the overall workflow of CHOP is illustrated in Algorithm 3. The process begins by utilizing the pretrained model, as introduced in Section IV-A, to predict the label $L(v)$ for each vertex v within the Boolean network G (line 1). This learning-based approach effectively mitigates the inefficiency

Authorized licensed use limited to: Chinese University of Hong Kong. Downloaded on September 13,2026 at 14:37:34 UTC from IEEE Xplore. Restrictions apply.

Algorithm 3 Overall Flow of CHOP

Input: A Boolean network $G(V, E)$.
Output: An optimized circuit $G'(V', E')$.
1 $L \leftarrow$ Predict the labels of all vertices V
2 $cs \leftarrow$ Perform circuit clustering for G with L
3 **for** $c \in cs$ **do**
4 $\quad o \leftarrow$ Select the optimization method for c
5 $\quad c \leftarrow$ Perform the selected optimization o for c
6 $G'(V', E') \leftarrow$ Merge cs
7 **return** $G'(V', E')$

of conventional feature extraction methods. Subsequently, the network G, along with the resulting set of labels L, is fed to the circuit clustering algorithm detailed in Section IV-B to obtain the clusters cs (line 2). By incorporating both initial partition and hierarchical clustering, the algorithm is capable of generating optimization-oriented subcircuits. Finally, each cluster c undergoes logic optimization using the techniques introduced in Section IV-C (lines 3–5), and the resulting subcircuits are then merged to construct the final optimized circuit $G'(V', E')$ (line 6).

V. Experimental Results

The proposed CHOP framework was implemented in C++ and Python. For evaluation, we employed circuits from EPFL benchmarks [20]. All experiments were conducted on a machine with an Intel[1] Xeon[1] Platinum 8350C processor, an NVIDIA A100 GPU, and 1.5 TB of memory. To evaluate the effectiveness of CHOP, we employed the same four types of logic representations and their corresponding optimization methods as utilized by LSOracle and HeLO. These optimization methods involve rewrite, refactor, balance, and resubstitution, all of which are implemented in the mockturtle [37], an open-source logic network library. The baselines consist of the following two parts.

1) A single logic representation along with its corresponding optimization method, including AIG, MIG, XMG, and XAG.
2) Hybrid optimization frameworks, specifically LSOracle [16] and HeLO [17]. Since we were unable to obtain the prediction model of LSOracle, we performed all the logic optimization methods it provides and selected the best outcome as the final result.

A. Accuracy of Labeling Vertices

We extracted all subgraphs using the method described in Section IV-A.1. During the expansion process, each vertex was expanded with an expansion degree of $k = 3$. Afterward, we applied the four aforementioned logic optimization methods to each subgraph. We evaluated the number of vertices and the logic depth after optimization, which reflect the structure of the subgraph prior to technology mapping. We further assessed the latency and area after performing 6-LUT mapping to characterize the postmapping implementation quality. These evaluation results were then used to compute the labels defined in (6), ensuring that the assigned labels accurately reflect the behavior of each subgraph under different optimization strategies. Finally, we fed these subgraphs and their corresponding labels into the prediction model as training data.

[1]Registered trademark.

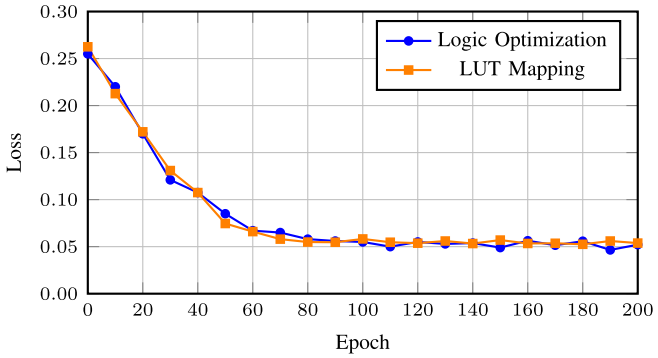


Fig. 5. Training loss curves with logic optimization and LUT mapping as the targets, respectively.

We expanded subgraphs centered at each vertex across a set of circuits from the ISCAS'85 [38] and EPFL [20] benchmarks, and randomly sampled 10 000 subgraphs to construct the training dataset. The model was trained for 200 epochs, with the loss function defined in Section IV-A.3. The training loss curve is illustrated in Fig. 5, where the loss value stabilizes after approximately 80 epochs. Using the index corresponding to the minimum predicted value of L as the evaluation criterion, the model achieves final accuracies of 99.62% and 99.38%, respectively.

B. Setting of ζ

To rigorously evaluate the impact of the parameter ζ, we conducted controlled experiments under two complementary scenarios, designed to analyze its effect across circuits with different logic depths.

For deep logic circuits, we select the `log2` circuit (logic depth: 444, 32 060 nodes before optimization), where ζ is determined by scaling the logic depth with varying percentages. This allows us to examine how different scaling factors influence optimization outcomes. For shallow logic circuits, we use the `voter` circuit (logic depth: 70, 13 758 nodes before optimization), where ζ is evaluated both as a proportional parameter and as fixed values. This dual approach enables a comprehensive analysis of ζ's sensitivity under different depth constraints. The experimental results are illustrated in Figs. 6 and 7.

The analysis reveals distinct optimization patterns across both scenarios. As shown in Fig. 6, setting the scaling percentage to 0.10 for deep circuits uniquely achieves both area and delay reductions below the mean, yielding the lowest ADP among all tested configurations. This indicates that excessively small ζ values underutilize the clustering potential by creating overly fine initial partitions, while excessively large values generate coarse partitions that limit subsequent refinement opportunities.

Authorized licensed use limited to: Chinese University of Hong Kong. Downloaded on September 13,2026 at 14:37:34 UTC from IEEE Xplore. Restrictions apply.

TABLE II
EXPERIMENTAL RESULTS ON THE EPFL BENCHMARK [20] AFTER LOGIC OPTIMIZATION

Circuit	AIG			MIG			XMG			XAG			LSOracle [16]			HeLO [17]			CHOP	
	depth	node	NDP	depth	node	NDP	depth	node	NDP	depth	node	NDP	depth	node	NDP	depth	node	NDP	depth	node
adder	55	1913	4.80	14	1567	1.00	14	1567	1.00	213	1016	9.86	14	1567	1.00	12	2043	1.12	14	1567
arbiter	20	5140	1.09	11	6335	0.74	15	5695	0.90	59	8973	5.59	29	5688	1.74	10	6392	0.68	15	6310
bar	16	3599	1.63	12	3336	1.13	12	3336	1.13	12	3141	1.06	16	3599	1.63	12	3337	1.13	12	2952
cavlc	13	722	1.34	10	702	1.00	15	697	1.49	14	729	1.45	10	702	1.00	10	702	1.00	10	702
ctrl	13	96	1.61	7	172	1.55	8	180	1.86	8	137	1.41	6	124	0.96	6	172	1.33	8	97
dec	3	304	1.00	3	304	1.00	3	304	1.00	3	304	1.00	3	304	1.00	3	304	1.00	3	304
div	1714	122441	2.30	970	94088	1.00	970	94088	1.00	4372	57247	2.74	1917	117962	2.48	743	145573	1.19	970	94088
hyp	16240	320826	1.33	16116	242549	1.00	16116	242549	1.00	20717	197031	1.04	21234	230663	1.25	15964	265710	1.09	16118	242367
i2c	9	1417	0.97	11	1381	1.15	16	1358	1.65	13	1416	1.40	9	1417	0.97	11	1381	1.15	9	1465
int2float	10	258	1.25	8	258	1.00	14	270	1.83	12	254	1.48	8	258	1.00	8	258	1.00	8	258
log2	478	30809	1.50	238	40214	0.97	238	40214	0.97	366	36277	1.35	343	39109	1.36	232	73294	1.73	252	39060
max	80	6909	2.46	47	6418	1.34	47	6418	1.34	187	3739	3.11	47	6418	1.34	54	7812	1.88	60	3749
memctrl	57	79655	1.27	69	51762	1.00	95	47472	1.26	83	56188	1.31	98	54531	1.50	61	56592	0.97	69	51762
multiplier	478	30809	3.16	122	35342	0.93	122	35342	0.93	274	27062	1.59	181	31043	1.21	135	49548	1.44	149	31239
priority	211	734	1.12	127	1101	1.01	126	1101	1.00	166	1107	1.32	126	1101	1.00	126	1101	1.00	126	1101
router	16	443	1.14	13	478	1.00	16	455	1.17	31	291	1.45	13	478	1.00	13	478	1.00	13	478
sin	271	5277	1.75	112	7315	1.00	112	7315	1.00	165	6184	1.25	141	6483	1.12	126	9667	1.49	112	7315
sqrt	6994	19193	1.30	3974	26080	1.00	3974	26080	1.00	5060	17856	0.87	5072	19165	0.94	3942	31486	1.20	3974	26080
square	59	23939	1.35	39	21864	0.81	41	21289	0.83	224	18549	3.96	87	21528	1.78	53	22335	1.13	50	20997
voter	76	9608	1.10	59	14002	1.24	60	13990	1.26	62	11760	1.09	76	10698	1.22	68	15292	1.56	72	9252
Ave. ratio	1.61	1.06	1.67	0.95	1.11	1.04	1.09	1.11	1.18	2.58	0.96	2.22	1.22	1.06	1.27	0.95	1.28	1.20	1.00	1.00

NDP: normalized node-depth product relative to CHOP; depth: logic depth; node: number of nodes.

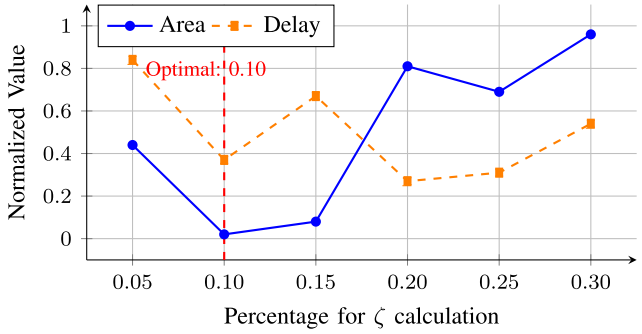


Fig. 6. Parameter sensitivity analysis for deep circuits: Impact of ζ percentage on optimization quality for the `log2` circuit (depth = 444). The optimal percentage of 0.10 minimizes both area and delay deviations, validating our parameter choice $\zeta = \max(\text{depth} \times 0.1, 10)$.

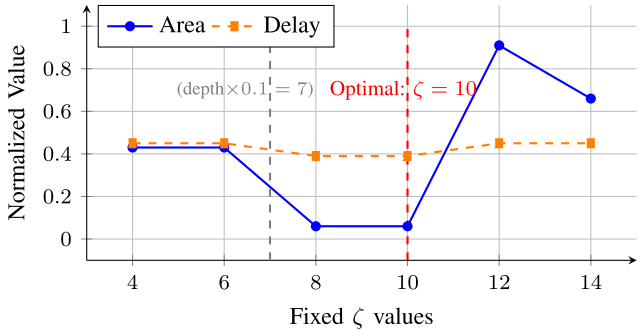


Fig. 7. Parameter sensitivity analysis for shallow circuits: Impact of fixed ζ values on optimization quality for the `voter` circuit (depth=70). The optimal value $\zeta = 10$ provides the best balance, demonstrating the effectiveness of our minimum threshold in the formula $\zeta = \max(\text{depth} \times 0.1, 10)$.

Similar trends emerge for shallow circuits (Fig. 7), where the optimal fixed value $\zeta = 10$ provides superior balance between clustering granularity and optimization effectiveness. To ensure robust performance across diverse circuit depths, we establish a minimum threshold of 10, leading to the final parameter configuration: $\zeta = \max(\text{depth} \times 0.1, 10)$.

C. Results of Logic Optimization

Table II presents the experimental results after logic optimization on the EPFL benchmark, to minimize the NDP, i.e., the product of the number of nodes ("node") and logic depth ("depth"). Notably, the NDP in Table II is the ratio of the NDP of each baseline to the NDP of CHOP. The results indicate that the optimization performance varies significantly across the four single logic representation methods. In most cases, the MIG-based method yields superior performance among these four methods. These findings underscore the limitations of existing logic optimization techniques that focus solely on a single logic representation. In addition, although XMG has the highest expressive power among these four logic representations, the results reveal that XMG-based optimization techniques do not consistently outperform their counterparts. This inconsistency may stem from the heuristic nature of the current optimization algorithms, which may not fully leverage the expressive capabilities of XMG.

Furthermore, hybrid optimization methods also exhibit varying performance outcomes. Notably, in certain cases, the overall performance of the hybrid strategy does not surpass that of applying MIG alone. This phenomenon may result from the relatively strong performance of MIG during the logic optimization phase, while many subgraphs within the hybrid method may have missed the opportunity to adopt MIG as the final optimization method. Overall, CHOP shows the best performance, specifically with average NDP reductions by 29.58%, 1.81%, 10.74%, 33.26%, 16.07%, and 12.39% when compared with AIG, MIG, XMG, XAG, LSOracle, and HeLO, respectively. These significant improvements highlight

Authorized licensed use limited to: Chinese University of Hong Kong. Downloaded on September 13,2026 at 14:37:34 UTC from IEEE Xplore. Restrictions apply.

TABLE III
EXPERIMENTAL RESULTS ON THE EPFL BENCHMARK [20] AFTER 6-LUT MAPPING

Circuit	AIG			MIG			XMG			XAG			LSOracle [16]			HeLO [17]			CHOP	
	delay	area	ADP	delay	area	ADP	delay	area	ADP	delay	area	ADP	delay	area	ADP	delay	area	ADP	delay	area
adder	64	192	2.81	11	397	1.00	14	409	1.31	64	192	2.81	11	397	1.00	9	583	1.20	11	397
arbiter	9	1876	1.21	9	2577	1.67	9	2577	1.67	23	2313	3.82	12	1877	1.62	9	3439	2.22	8	1741
bar	4	512	1.00	4	512	1.00	4	512	1.00	4	512	1.00	4	512	1.00	4	512	1.00	4	512
cavlc	8	114	1.31	6	119	1.03	6	119	1.03	6	119	1.03	8	114	1.31	6	150	1.29	6	116
ctrl	2	28	1.00	2	28	1.00	2	28	1.00	2	28	1.00	2	28	1.00	2	30	1.07	2	28
dec	2	273	1.00	2	273	1.00	2	273	1.00	2	273	1.00	2	273	1.00	2	273	1.00	2	273
div	2005	4213	1.01	776	27491	2.54	776	27491	2.54	1606	25529	4.88	1915	11869	2.71	732	28889	2.52	2010	4180
hyp	8520	49736	1.00	8520	49736	1.00	8520	49736	1.00	8520	49736	1.00	8509	51437	1.03	8651	69321	1.42	8520	49736
i2c	5	342	1.13	5	342	1.13	5	342	1.13	5	342	1.13	5	342	1.13	5	362	1.20	5	302
int2float	4	47	1.00	4	47	1.00	4	47	1.00	4	47	1.00	4	47	1.00	4	56	1.19	4	47
log2	182	7893	1.24	129	9102	1.01	129	9103	1.01	185	8125	1.30	151	8720	1.14	137	22455	2.66	143	8094
max	45	1689	2.49	22	2288	1.65	22	2288	1.65	89	787	2.29	45	1689	2.49	29	2574	2.44	34	899
memctrl	58	12064	1.30	45	11969	1.00	45	11969	1.00	49	11709	1.07	49	12657	1.15	45	11969	1.00	45	11969
multiplier	127	6348	1.77	68	7453	1.12	76	6821	1.14	119	6907	1.81	91	6919	1.39	76	9533	1.59	75	6058
priority	35	225	1.00	33	264	1.11	33	265	1.11	33	265	1.11	35	225	1.00	116	391	5.76	35	225
router	13	53	1.00	13	53	1.00	13	53	1.00	13	53	1.00	13	53	1.00	8	110	1.28	13	53
sin	95	1590	1.27	66	1925	1.07	66	1925	1.07	83	1583	1.10	77	1714	1.11	81	2708	1.84	67	1779
sqrt	2034	4086	1.12	3897	8141	4.29	3897	8141	4.29	3689	7839	3.91	1957	3992	1.06	3701	10664	5.34	1930	3828
square	47	5066	2.11	30	4822	1.28	32	4559	1.29	122	4130	4.46	56	4447	2.21	41	4830	1.75	25	4515
voter	29	2008	1.69	35	2827	2.88	35	2827	2.88	34	3328	3.29	28	1958	1.60	39	3150	3.57	22	1562
Ave. ratio	1.43	1.05	1.37	1.04	1.52	1.44	1.06	1.51	1.46	1.73	1.36	2.00	1.16	1.17	1.35	1.18	1.90	2.07	1.00	1.00

ADP: normalized area-delay product relative to CHOP; delay: circuit delay; area: circuit area.

the importance of the combination of circuit clustering and optimization method selection.

D. Results of LUT Mapping

However, it is worth noting that logic-level optimization metrics alone do not fully capture the practical implications for physical design. This is because different vertex types in a Boolean network not only differ in expressive power but also require varying amounts of physical resources to implement. As a result, logic-level metrics such as gate count and logic depth provide only a partial view of overall design quality. To more accurately assess the practical effectiveness of the optimization methods, Table III presents the results after 6-LUT mapping [39], which offers a closer approximation of the real-world performance in the implementation of integrated circuits. The objective of 6-LUT mapping is to minimize the ADP, i.e., the product of circuit area and delay. Similarly, the ADP values in Table III are normalized for comparison. Notably, consistent with the results in Table II, the MIG-based method achieves the best optimization outcomes among the four single logic representation methods. The difference is that hybrid optimization methods show better results in 6-LUT mapping, especially LSOracle and CHOP.

Overall, CHOP continues to maintain a leading position relative to the baselines. Specifically, CHOP achieves ADP reductions of 19.48%, 17.28%, 18.62%, 30.16%, 17.99%, and 36.42% on average compared with AIG, MIG, XMG, XAG, LSOracle, and HeLO, respectively. In addition, there are two primary reasons why HeLO performs poorly in 6-LUT mapping. First, the framework targets logic optimization, and the input of technology mapping in [17] directly sources from the results of logic optimization. Second, HeLO's circuit partitioning method has significant shortcomings. Specifically, unlike CHOP, HeLO's initial partitioning method extends from POs to PIs without considering the situation of other cones. This makes HeLO's optimization results depend to a certain extent on the order of processing POs during the initial partitioning.

Notably, in circuits where LSOracle demonstrates suboptimal performance, CHOP consistently excels. This discrepancy arises because LSOracle aggregates edges with higher weights for partition, which can sometimes lead to the grouping of vertices with significantly different logic levels, potentially resulting in suboptimal performance in certain circuits. This observation further validates our partition method presented in Section IV-B.2. This disparity further motivates our approach: rather than relying on a single optimization method across all designs, we aim to tailor the optimization strategy to the specific characteristics of each circuit or subgraph cluster. By leveraging structural and functional diversity within the design, our method seeks to achieve a better tradeoff between optimization quality and generalization.

E. Runtime

The runtime comparison between CHOP and two state-of-the-art hybrid optimization methods is presented in Fig. 8. On smaller circuits, CHOP does not exhibit a noticeable speed advantage, primarily because the selected model introduces a fixed runtime overhead that dominates when the circuit size is small. This overhead mainly stems from the cost of subgraph extraction, feature encoding, and inference time of the learning-based components. However, as the circuit size increases, these fixed costs become less significant relative to the overall runtime, and the scalability of CHOP becomes more evident.

Overall, CHOP shows high efficiency, with average speeds of 1.92 × that of LSOracle and 2.98 × that of HeLO for logic optimization, as well as 1.78 × and 2.48 × for LUT mapping.

Authorized licensed use limited to: Chinese University of Hong Kong. Downloaded on September 13,2026 at 14:37:34 UTC from IEEE Xplore. Restrictions apply.

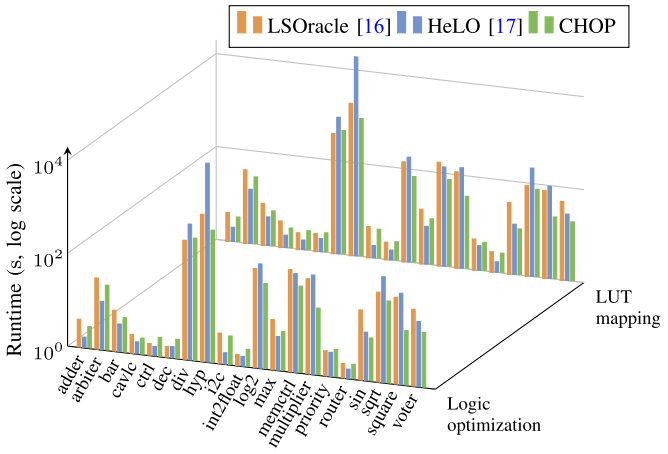


Fig. 8. Comparison of normalized runtime among different optimization methods for logic optimization and 6-LUT mapping, respectively.

This substantial reduction in runtime highlights CHOP's ability to efficiently handle large-scale designs. Furthermore, since CHOP relies on pretrained models for inference rather than costly iterative optimization, its performance is more stable and predictable across different circuit instances. These results collectively demonstrate the runtime efficiency and scalability of CHOP, making it particularly suitable for large and complex circuits where traditional methods may struggle with high computational overhead.

VI. Conclusion

This article proposed CHOP, a novel framework that fundamentally advances hybrid logic synthesis through vertex-level feature extraction and optimization-aware clustering. The key innovation lies in shifting from the traditional "partition-then-predict" paradigm to a "predict-then-partition" approach, enabling more effective alignment between circuit structures and optimization strategies. CHOP's hierarchical clustering algorithm, guided by both structural and functional features, demonstrates superior performance across multiple metrics.

The experimental results validate CHOP's effectiveness. For logic optimization on the EPFL benchmark, CHOP achieved average NDP reductions of 16.07% and 12.39% compared with LSOracle and HeLO, with 1.92 × and 2.98 × speedup, respectively. For 6-LUT mapping, CHOP achieved average ADP reductions of 17.99% and 36.42%, with 1.78 × and 2.48 × speedup, respectively.

Future work will focus on extending CHOP to handle larger industrial circuits, integrating more sophisticated neural architectures for vertex feature extraction, and exploring applications to other EDA tasks such as placement and routing optimization.

Acknowledgment

The research work described in this article was conducted at the JC STEM Laboratory of Intelligent Design Automation.

References

[1] E. Testa, M. Soeken, L. G. Amar, and G. De Micheli, "Logic synthesis for established and emerging computing," *Proc. IEEE*, vol. 107, no. 1, pp. 165–184, Jan. 2019.

[2] R. K. Brayton, R. Rudell, A. Sangiovanni-Vincentelli, and A. R. Wang, "MIS: A multiple-level logic optimization system," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. CAD-6, no. 6, pp. 1062–1081, Nov. 1987.

[3] L. Hellerman, "A catalog of three-variable or-invert and and-invert logical circuits," *IEEE Trans. Electron. Comput.*, vol. EC-12, no. 3, pp. 198–223, Jun. 1963.

[4] L. Amarú, P.-E. Gaillardon, and G. De Micheli, "Majority-inverter graph: A novel data-structure and algorithms for efficient logic optimization," in *Proc. 51st Annu. Design Autom. Conf.*, Aug. 2014.

[5] G. Meuli, M. Soeken, and G. De Micheli, "XOR-AND-inverter graphs for quantum compilation," *Npj Quantum Inf.*, vol. 8, no. 1, Jan. 2022.

[6] W. Haaswijk, M. Soeken, L. Amarù, P.-E. Gaillardon, and G. De Micheli, "A novel basis for logic rewriting," in *Proc. 22nd Asia South Pacific Design Autom. Conf. (ASP-DAC)*, Jan. 2017, pp. 151–156.

[7] A. Mishchenko, S. Chatterjee, and R. Brayton, "DAG-aware AIG rewriting a fresh look at combinational logic synthesis," in *Proc. 43rd Annu. Conf. Design Autom. (DAC)*, 2006, pp. 532–535.

[8] H. Riener, W. Haaswijk, A. Mishchenko, G. De Micheli, and M. Soeken, "On-the-fly and DAG-aware: Rewriting Boolean networks with exact synthesis," in *Proc. Design, Autom. Test Eur. Conf. Exhib. (DATE)*, Mar. 2019, pp. 1649–1654.

[9] A. T. Calvino and G. De Micheli, "Scalable logic rewriting using Don't cares," in *Proc. Design, Autom. Test Eur. Conf. Exhib. (DATE)*, Mar. 2024, pp. 1–6.

[10] S.-Y. Lee and G. D. Micheli, "Heuristic logic resynthesis algorithms at the core of peephole optimization," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 42, no. 11, pp. 3958–3971, Nov. 2023.

[11] L. Amarú et al., "Improvements to Boolean resynthesis," in *Proc. Design, Autom. Test Eur. Conf. Exhib. (DATE)*, Mar. 2018, pp. 755–760.

[12] A. M. R. Brayton and A. Mishchenko, "Scalable logic synthesis using a simple circuit structure," in *Proc. IWLS*, vol. 6, 2006, pp. 15–22.

[13] H. Riener, E. Testa, L. Amaru, M. Soeken, and G. De Micheli, "Size optimization of MIGs with an application to QCA and STMG technologies," in *Proc. 14th IEEE/ACM Int. Symp. Nanosc. Architectures*, Jul. 2018, pp. 157–162.

[14] S.-Y. Lee, H. Riener, A. Mishchenko, R. K. Brayton, and G. De Micheli, "A simulation-guided paradigm for logic synthesis and verification," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 41, no. 8, pp. 2573–2586, Aug. 2022.

[15] L. Amarú, P.-E. Gaillardon, and G. De Micheli, "MIXSyn: An efficient logic synthesis methodology for mixed XOR-AND/OR dominated circuits," in *Proc. 18th Asia South Pacific Design Autom. Conf. (ASP-DAC)*, Jan. 2013, pp. 133–138.

[16] W. L. Neto, M. Austin, S. Temple, L. Amaru, X. Tang, and P.-E. Gaillardon, "LSOracle: A logic synthesis framework driven by artificial intelligence: Invited paper," in *Proc. IEEE/ACM Int. Conf. Comput.-Aided Design (ICCAD)*, Nov. 2019, pp. 1–6.

[17] Y. Pu et al., "HeLO: A he terogeneous l ogic o ptimization framework by hierarchical clustering and graph learning," in *Proc. Int. Symp. Phys. Design*, Mar. 2025, pp. 116–124.

[18] R. Fu et al., "Late breaking results: Hybrid logic optimization with predictive self-supervision," in *Proc. 62nd ACM/IEEE Design Autom. Conf. (DAC)*, Jun. 2025, pp. 1–2.

[19] S. Schlag, T. Heuer, L. Gottesbüren, Y. Akhremtsev, C. Schulz, and P. Sanders, "High-quality hypergraph partitioning," *ACM J. Experim. Algorithmics*, vol. 27, pp. 1–39, Dec. 2022.

[20] L. Amarú, P.-E. Gaillardon, and G. De Micheli, "The EPFL combinational benchmark suite," Tech. Rep., 2015.

[21] S. A. Kauffman, "Metabolic stability and epigenesis in randomly constructed genetic nets," *J. Theor. Biol.*, vol. 22, no. 3, pp. 437–467, Mar. 1969.

[22] S. Fortunato, "Community detection in graphs," *Phys. Rep.*, vol. 486, nos. 3–5, pp. 75–174, Feb. 2010.

[23] M. E. J. Newman and M. Girvan, "Finding and evaluating community structure in networks," *Phys. Rev. E, Stat. Phys. Plasmas Fluids Relat. Interdiscip. Top.*, vol. 69, no. 2, Feb. 2004, Art. no. 026113.

[24] V. D. Blondel, J.-L. Guillaume, R. Lambiotte, and E. Lefebvre, "Fast unfolding of communities in large networks," *J. Stat. Mech., Theory Exp.*, vol. 2008, no. 10, Oct. 2008, Art. no. P10008.

[25] V. A. Traag, L. Waltman, and N. J. van Eck, "From Louvain to leiden: Guaranteeing well-connected communities," *Sci. Rep.*, vol. 9, no. 1, pp. 2045–2322, Mar. 2019.

Authorized licensed use limited to: Chinese University of Hong Kong. Downloaded on September 13,2026 at 14:37:34 UTC from IEEE Xplore. Restrictions apply.

[26] M. Li, S. Khan, Z. Shi, N. Wang, H. Yu, and Q. Xu, "DeepGate: Learning neural representations of logic gates," in *Proc. 59th ACM/IEEE Design Autom. Conf.*, Jul. 2022, pp. 667–672.
[27] Z. Shi et al., "Deepgate2: Functionality-aware circuit representation learning," in *Proc. EEE/ACM Int. Conf. Comput. Aided Design*, Aug. 2023, pp. 1–9.
[28] Z. Shi et al., "DeepTPI: Test point insertion with deep reinforcement learning," in *Proc. IEEE Int. Test Conf. (ITC)*, Sep. 2022, pp. 194–203.
[29] S. Khan, Z. Shi, M. Li, and Q. Xu, "DeepSeq: Deep sequential circuit learning," 2023, *arXiv:2302.13608*.
[30] M. Li, Z. Shi, Q. Lai, S. Khan, S. Cai, and Q. Xu, "On EDA-driven learning for SAT solving," in *Proc. 60th ACM/IEEE Design Autom. Conf. (DAC)*, Jul. 2023, pp. 1–6.
[31] Z. Shi et al., "EDA-driven preprocessing for SAT solving," 2024, *arXiv:2403.19446*.
[32] N. Wu, Y. Li, C. Hao, S. Dai, C. Yu, and Y. Xie, "Gamora: Graph learning based symbolic reasoning for large-scale Boolean networks," in *Proc. 60th ACM/IEEE Design Autom. Conf. (DAC)*, Jul. 2023, pp. 1–6.
[33] C. Deng et al., "Less is more: Hop-wise graph attention for scalable and generalizable learning on circuits," 2024, *arXiv:2403.01317*.
[34] J. Liu, J. Zhai, M. Zhao, Z. Lin, B. Yu, and C. Shi, "PolarGate: Breaking the functionality representation bottleneck of and-inverter graph neural network," in *Proc. 43rd IEEE/ACM Int. Conf. Computer-Aided Design*, Oct. 2024, pp. 1–9.
[35] Z. Shi, Z. Zheng, S. Khan, J. Zhong, M. Li, and Q. Xu, "DeepGate3: Towards scalable circuit representation learning," in *Proc. 43rd IEEE/ACM Int. Conf. Comput.-Aided Design*, Oct. 2024, pp. 1–9.
[36] Z. Wang et al., "Functionality matters in netlist representation learning," in *Proc. 59th ACM/IEEE Design Autom. Conf.*, Jul. 2022, pp. 61–66.
[37] M. Soeken et al., "The EPFL logic synthesis libraries," 2018, *arXiv:1805.05121*.
[38] F. Brglez and H. Fujiwara, "A neutral netlist of 10 combinational benchmark circuits and a target translator in Fortran," Tech. Rep., 1985, pp. 677–692.
[39] G. Liu and Z. Zhang, "A parallelized iterative improvement approach to area optimization for LUT-based technology mapping," in *Proc. ACM/SIGDA Int. Symp. Field-Program. Gate Arrays*, Feb. 2017, pp. 147–156.



Rongliang Fu (Graduate Student Member, IEEE) received the B.S. degree in software engineering from Northwestern Polytechnical University, Xi'an, China, in 2018, and the M.S. degree in computer science and technology from the University of Chinese Academy of Sciences, Beijing, China, in 2021. He is currently pursuing the Ph.D. degree with the Department of Computer Science and Engineering, The Chinese University of Hong Kong, Hong Kong.

His research interests include electronic design automation and computer architecture.



Ran Zhang received the B.S. degree from Northeastern University, Shenyang, China, in 2023. He is currently pursuing the master's degree with the University of Chinese Academy of Sciences (UCAS), under the Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China.

His research interests include electronic design automation and computer architecture.



Ziyang Zheng (Student Member, IEEE) received the bachelor's degree in data science and big data from Harbin Institute of Technology (Shenzhen), Shenzhen, China, in 2023. He is currently pursuing the Ph.D. degree with the CURE Laboratory, The Chinese University of Hong Kong, Hong Kong, where he is advised by Prof. Qiang Xu.

His research interests include multimodal learning, diffusion models, and AI applications in electronic design automation (EDA).



Zhengyuan Shi received the B.Eng. degree from Shandong University, Jinan, China, in 2021. He is currently pursuing the Ph.D. degree with The Chinese University of Hong Kong, Hong Kong.

He has published over 20 papers in leading EDA conferences, including DAC and ICCAD. His research interests include AI for EDA and large circuit models.

Mr. Shi received the Best Paper Nomination Award from DAC 2022 and ASP-DAC 2025. He received the Presidential Scholarship for the B.Eng. degree from Shandong University.



Yuan Pu received the B.S. degree in computer science from The Chinese University of Hong Kong (CUHK), Hong Kong, in 2022, where he is currently pursuing the Ph.D. degree with the Department of Computer Science and Engineering, supervised by Prof. Bei Yu.

His research interests include combinatorial algorithms/AI in EDA.



Junying Huang received the Ph.D. degree in microelectronics and solid-state electronics from the University of Chinese Academy of Sciences, Beijing, China, in 2016.

She is currently an Associate Professor with the High-Throughput Computer Research Center, Institute of Computing Technology, Chinese Academy of Sciences, Beijing. Her research interests include superconductive RSFQ logic, computer architecture, electronic design automation, and hardware security.



Bei Yu (Senior Member, IEEE) received the Ph.D. degree from The University of Texas at Austin, Austin, TX, USA, in 2014.

He is currently a Professor with the Department of Computer Science and Engineering, The Chinese University of Hong Kong, Hong Kong.

Dr. Yu received ten best paper awards from IEEE TSM 2022, DATE 2022, ICCAD 2021 and 2013, ASPDAC 2021 and 2012, ICTAI 2019, *Integration, the VLSI Journal* in 2018, ISPD 2017, SPIE Advanced Lithography Conference 2016, and many other awards, including the DAC Under-40 Innovator Award in 2024, the IEEE CEDA Ernest S. Kuh Early Career Award in 2022, and Hong Kong RGC Research Fellowship Scheme (RFS) Award in 2024. He has served as the TPC Chair for the ACM/IEEE Workshop on Machine Learning for CAD, and on many journal editorial boards and conference committees.



Qiang Xu (Senior Member, IEEE) received the B.E. and M.E. degrees from Beijing University, Beijing, China, in 1997 and 2000, respectively, and the Ph.D. degree from McMaster University, Hamilton, ON, Canada, in 2005.

He is currently a Professor of computer science and engineering with The Chinese University of Hong Kong, Hong Kong. His research interests include electronic design automation, trusted computing, and representation learning.

Dr. Xu is serving as an Associate Editor for IEEE TRANSACTIONS ON COMPUTER-AIDED DESIGN OF INTEGRATED CIRCUITS AND SYSTEMS and *Integration, the VLSI Journal*.



Tsung-Yi Ho (Fellow, IEEE) received the Ph.D. degree in electrical engineering from National Taiwan University, Taipei, Taiwan, in 2005.

He is currently a Professor with the Department of Computer Science and Engineering, The Chinese University of Hong Kong (CUHK), Hong Kong. His research interests include several areas of computing and emerging technologies, especially in the design automation of microfluidic biochips.

Dr. Ho is a Distinguished Member of ACM. He was a recipient of the Best Paper Award from IEEE TRANSACTIONS ON COMPUTER-AIDED DESIGN OF INTEGRATED CIRCUITS AND SYSTEMS in 2015. He serves as the VP Conferences for IEEE CEDA, and the Executive Committee for ASP-DAC and ICCAD.

Authorized licensed use limited to: Chinese University of Hong Kong. Downloaded on September 13,2026 at 14:37:34 UTC from IEEE Xplore. Restrictions apply.