

# GAMORA: Graph Learning based Symbolic Reasoning for Large-Scale Boolean Networks

Nan Wu<sup>1</sup>, Yingjie Li<sup>2</sup>, Cong Hao<sup>3</sup>, Steve Dai<sup>4</sup>, Cunxi Yu<sup>\*2</sup>, Yuan Xie<sup>5</sup>

<sup>1</sup>University of California, Santa Barbara, <sup>2</sup>University of Utah, <sup>3</sup>Georgia Institute of Technology,

<sup>4</sup>NVIDIA, <sup>5</sup>Alibaba DAMO Academy

nanwu@ucsb.edu, yingjie.li@utah.edu, callie.hao@gatech.edu, sdai@nvidia.com, yuanxie@gmail.com

\*correspondence: cunxi.yu@utah.edu

**Abstract**—Reasoning high-level abstractions from bit-blasted Boolean networks (BNs) such as gate-level netlists can significantly benefit functional verification, logic minimization, datapath synthesis, malicious logic identification, etc. Mostly, conventional reasoning approaches leverage structural hashing and functional propagation, suffering from limited scalability and inefficient usage of modern computing power. In response, we propose a novel symbolic reasoning framework exploiting graph neural networks (GNNs) and GPU acceleration to reason high-level functional blocks from gate-level netlists, namely GAMORA, which offers high reasoning performance w.r.t exact reasoning algorithms, strong scalability to BNs with over 33 million nodes, and generalization capability from simple to complex designs. To further demonstrate the capability of GAMORA, we also evaluate its reasoning performance after various technology mapping options, since technology-dependent optimizations are known to make functional reasoning much more challenging. Experimental results show that (1) GAMORA reaches almost 100% and over 97% reasoning accuracy for carry-save-array (CSA) and Booth-encoded multipliers, respectively, with up to six orders of magnitude speedups compared to the state-of-the-art implementation in the ABC framework; (2) GAMORA maintains high reasoning accuracy (>92%) in finding functional modules after complex technology mapping, and we comprehensively analyze the impacts on GAMORA reasoning from technology mapping. GAMORA is available at <https://github.com/Yu-Utah/Gamora>.

## I. INTRODUCTION

Reasoning high-level abstractions (e.g., functional blocks) from bit-blasted Boolean networks (BNs) (e.g., unstructured gate-level netlists) has demonstrated its wide applications in improving functional verification efficiency [7], [16] and identifying malicious logics such as detecting hardware trojan and intellectual property infringement usage [3], [14]. In the era of globalization and democratization of integrated circuit (IC) development and fabrication, such reasoning is expected to bring broader impacts on hardware security, which is at the heart of modern computing systems: more than 40 percent of FPGA/ASIC projects are working under safety-critical development process standards or guidelines [9].

Due to the optimization conducted by RTL synthesis tools, reasoning high-level abstractions such as functional blocks from unstructured or flattened netlists is extremely challenging, since hierarchy and module information is lost during multi-level logic minimization and technology mapping, which is also complicated by functional blocks overlapping and gate sharing. The problem goes further due to the explosion in runtime for large-scale BNs. Conventional reasoning approaches leverage structural analysis and functional propagation. Structural approaches either adopt shape hashing based on circuit topology to find structurally similar wires to form word-level abstractions [15], or rely on reference libraries to map sub-circuits with reference circuits [6]. Functional approaches focus on identifying functionally equivalent gates and wires by cut enumeration [10], [20]. The combination of structural and functional analysis [15], [20], [26] is more prevalent for efficient word-level abstraction and propagation. Despite the achieved success, the performance of these

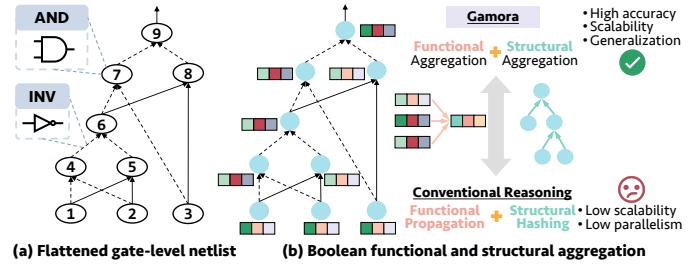


Fig. 1: The inputs to GAMORA are flattened gate-level netlists, with each node as an AND gate and dashed edges as inverters. By encoding Boolean functional information as node features, GAMORA can simultaneously handle functional and structural aggregation, analogous to functional propagation and structural hashing in conventional reasoning but with strong scalability.

conventional approaches is restricted by **limited scalability** and **inefficient utilization of modern computing power**: (1) structural hashing is very time/memory-consuming for large BNs with billions of nodes; (2) functional propagations by symbolic evaluation are solver-ready but extremely expensive, in particular for *bit-blasted* non-linear arithmetic BNs; (3) all these algorithms do not effectively utilize modern computing power due to the difficulty of parallelism.

Recently, we have witnessed the emergence of machine learning (ML) applied for computer systems and electronic design automation (EDA) tasks [23], as an alternative to conventional design solutions. Since circuit netlists or BNs can be easily represented as graphs, graph neural networks (GNNs) are naturally suitable to classify sub-circuit functionality from gate-level netlists [2], analyze impacts of circuit rewriting on functional operator detection [27], and predict boundaries of arithmetic blocks [12].

Motivated by the limitations of conventional approaches and the potentials of GNNs applied on circuit designs, we propose a graph learning-based symbolic reasoning framework to reverse engineer functional blocks from gate-level netlists, namely GAMORA, which has **high reasoning accuracy**, **strong scalability** to BNs with billions of nodes, and **generalization capability** from simple to complex designs. GAMORA employs a multi-task GNN to guarantee reasoning accuracy while simultaneously handling structural and functional information from BNs. Once well trained, GAMORA becomes adept at generalizing to large-scale and complex BNs, leveraging the accelerated inference and parallel processing offered by modern computing systems. We summarize our contributions as follows.

- **Novel multi-task GNN for structure and function fusion.** The message passing mechanisms in GNNs enable simultaneous *Boolean functional and structural aggregation*, corresponding to the symbolic propagation and structural hashing in conventional

reasoning methods, as shown in Figure 1. The multi-task setting allows knowledge sharing across different reasoning sub-tasks to guarantee high reasoning accuracy.

- **Billion-node scalability and parallelism.** We develop domain-specific techniques to compress node features, significantly reducing compute costs. The exploitation of graph learning draws better support from modern computing systems, such as GPU deployment, for scalability to large BNs and parallel execution.
- **Generalization capability.** Unlike many ML-based approaches that are trained with complex designs and infer on simpler ones, GAMORA can easily generalize from simple to complex BNs and handle the reasoning complexity introduced by more advanced designs (such as Booth-encoded multipliers) and technology mapping.
- **Evaluation.** Regarding reasoning performance, GAMORA reaches almost 100% and over 97% reasoning accuracy for carry-save array (CSA) and Booth multipliers, respectively; after technology mapping, the reasoning accuracy is still over 92%. Regarding runtime and scalability, GAMORA can perform reasoning for large BNs with tens of millions of nodes/edges within one second, with a speedup of up to six orders of magnitude compared to the logic synthesis tool ABC [4].

## II. PRELIMINARY AND MOTIVATION

### A. Boolean Networks and And-Inverter Graphs

BNs are well-studied discrete mathematical models with broad applications in chemistry, biology, circuit design, formal verification, etc. For purposes of synthesis and verification, a concise and uniform representation of BNs consisting of inverters and two-input AND-gates, known as and-inverter graphs (AIGs), has found successful use in diverse EDA tasks, since AIGs allow rewriting, simulation, technology mapping, placement, and verification to share the same data structure [18]. In an AIG, each node has at most two incoming edges; a node without incoming edges is a primary input (PI); primary outputs (POs) are denoted by special output nodes; each internal node represents a two-input AND function. Based on De Morgan’s laws, any combinational BN can be converted into an AIG [4] in a fast and scalable manner.

In AIGs, cut enumeration can be used to detect Boolean functions. A feasible cut of node  $n$  is a set of nodes in the transitive fan-in cone of  $n$ , whose truth value assignments completely determine the value of  $n$ . A cut is  $K$ -feasible if there are no more than  $K$  inputs. Figure 2 depicts an example of reasoning XOR functions and full adders from AIGs. In Figure 2(a), the AIG has a 3-feasible cut of node 9 and a 2-feasible cut of node 6; after truth table computation, the functions of node 6 and node 9 are  $IN1 \oplus IN2$  and  $IN1 \oplus IN2 \oplus IN3$ , respectively. Thus, as shown in Figure 2(b), node 6 is an XOR2 function, and node 9 is an XOR3 function. Figure 2(c) shows a full adder bitslice, with the sum as an XOR function and the carry-out as a majority (MAJ) function. By pairing an XOR3 with a MAJ3 with identical inputs, a full adder bitslice can be extracted, which is then aggregated for word-level abstraction.

### B. Word-Level Abstraction

Word-level abstraction significantly reduces the complexity of large-scale BNs by grouping wires into meaningful words and keeping useful information related to control logic, which is widely applied in reasoning functional units from gate-level netlists [15], [20], [26]. Conventional word identification uses *structural shape hashing* and *functional bitslice aggregation*. Structural shape hashing assigns each edge in the BN a shape, which is defined as the directed graph constructed by the backward reachable nodes from this edge within

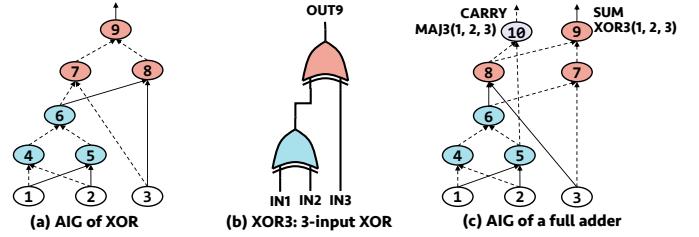


Fig. 2: Netlists of XOR and a full adder. (a) AIG of XOR3 function. (b) XOR3 function:  $OUT9 = XOR3(IN1, IN2, IN3)$ . (c) Full adder with a sum function (i.e., XOR3) and a carry-out function (i.e., MAJ3).

certain depth/steps. Functional bitslice aggregation adopts functional matching to group functionally equivalent nodes and edges by cut enumeration. Typically, structural hashing and functional aggregation are iteratively propagated across neighborhood nodes using symbolic evaluation [15], [20], [26]. However, for large-scale BNs, structural hashing is memory-consuming; functional bitslice aggregation is not efficient due to the requirement of bit-blasting; the computation of symbolic evaluation is also expensive. Motivated by the **limited scalability** and the **difficulty of parallelism**, we propose to exploit **graph learning and GPU acceleration for highly scalable reasoning**.

### C. Graph Neural Network

Since BNs and circuit netlists are naturally represented as graphs, GNNs can be leveraged to classify sub-circuit functionality from gate-level netlists [2], predict the functionality of approximate circuits [5], analyze impacts of circuit rewriting on functional operator detection [27], and predict boundaries of arithmetic blocks [12]. Promising as they are, these approaches focus on graphs with tens of thousands of nodes, and conduct training on complex designs and inference on relatively simpler ones, in which the generalization capability from simple to complex designs is not well examined.

GNNs operate by propagating information along the edges of a given graph. Each node is initialized with a representation, which could be either a direct representation or a learnable embedding obtained from node features. Then, a GNN layer updates each node representation by integrating node representations of both itself and its neighbors in the graph. The propagation along edges extracts structural information from graphs, corresponding to structural shape hashing in conventional reasoning; after encoding Boolean functionality into node features, neighborhood aggregation is analogous to functional aggregation in conventional reasoning. Thus, the inherent message-passing mechanism in GNNs enables simultaneously handling structural and functional information. Motivated by the **analogy between GNN computation and conventional reasoning**, we propose a **multi-task GNN for high-performance reasoning** w.r.t. exact reasoning algorithms, with **strong generalization capability** from simple to complex designs.

## III. PROPOSED APPROACH

### A. Overview

**Problem Formulation.** Figure 3(a) illustrates the overview of GAMORA. The inputs are flattened gate-level netlists in AIG format, without any micro-architectural or RTL information. These AIGs are generated by the logic synthesis tool ABC [4]. The goal is to exploit a multi-task GNN to reason high-level abstractions by performing node-level classification on AIGs, after which functional blocks (e.g., adders) can be extracted based on the annotated AIGs.

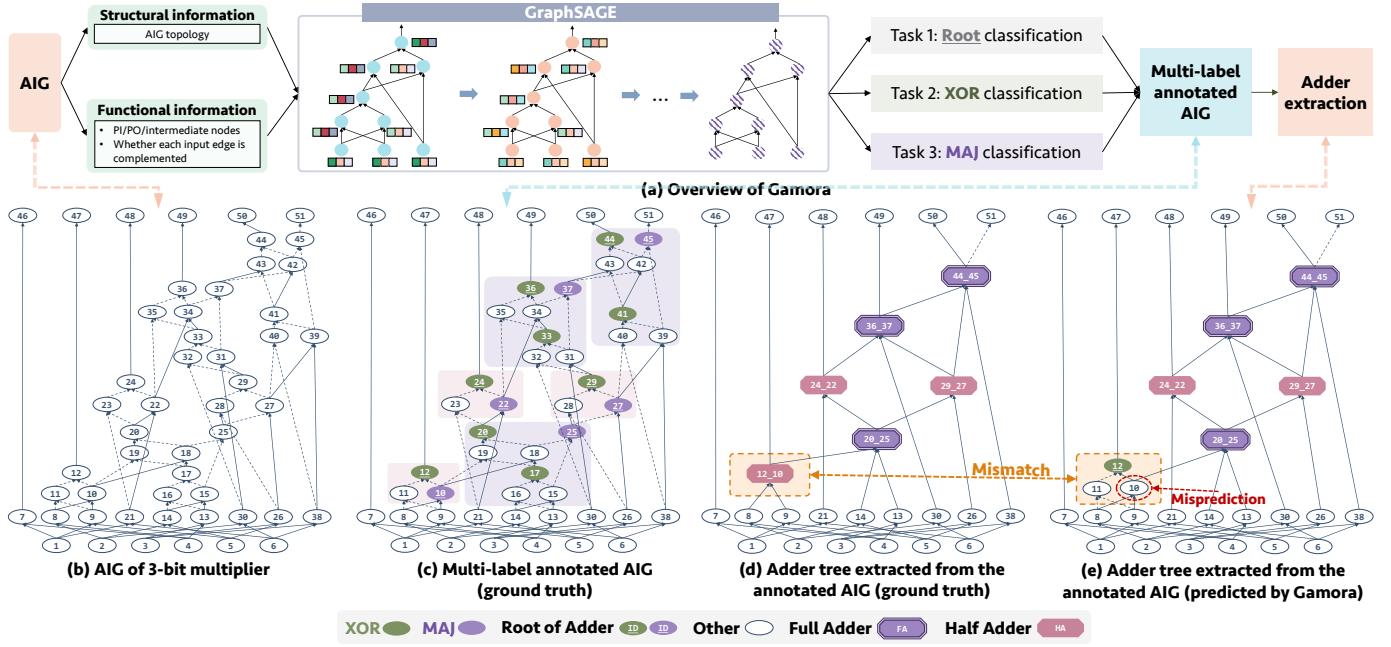


Fig. 3: Overview of GAMORA. (a) GAMORA takes in flattened netlists in AIG format and performs multi-task node classification to reason the Boolean function of each node, after which the adder trees within multiplier netlists can be automatically extracted to improve the efficiency of word-level abstraction. (b) AIG of a 3-bit CSA multiplier after synthesis. (c) Annotated AIG with the Boolean function of each node, using the ground truth provided by ABC. (d) Adder tree extracted based on the *exact* reasoning, including three FAs and three HAs. (e) Adder tree extracted based on the reasoning performed by GAMORA.

**Case Study on Multipliers.** Integer multipliers are indispensable to computationally intensive applications, such as signal processing and cryptography applications. Recent years also witness the strong demand for large integer multipliers in homomorphic encryption [1]. In general, formal multiplier verification is challenging, especially for structurally complex designs such as Booth multipliers [7], [16], [21]. Symbolic computer algebra (SCA) has been successfully employed to verify a variety of integer multipliers [7], [13], [16], [17], [26], which relies heavily on detecting full adders (FAs) and half adders (HAs) in multiplier netlists. The state-of-the-art implementation in ABC framework [26] develops a fast algebraic rewriting approach to extracting adder trees from flattened multiplier netlists by detecting pairs of XOR and MAJ functions, which can handle large bitwidth multipliers (up to 2048-bit) but with extremely long runtime. Thus, targeting integer multipliers, we leverage GNNs to identify XOR and MAJ functions to extract adders from flattened netlists, which improves the efficiency of word-level abstraction from BNs and has strong scalability enabled by GPU acceleration.

#### B. Multi-Task Learning for Boolean Reasoning

Boolean reasoning requires gathering structural and functional information from neighbor nodes, a process that can be imitated by the message-passing mechanism in GNNs. The task of reasoning high-level abstractions from flattened netlists, i.e., pinpointing adders from AIGs, involves a two-step procedure [15], [20], [26]: (1) detecting XOR/MAJ functions to construct adders, and then (2) identifying their boundaries. Therefore, we propose to apply multi-task learning (MTL) for Boolean reasoning to approach its nature, in which knowledge sharing across sub-tasks provides higher reasoning precision. This section details (1) how structural and functional information are fused in node embeddings, (2) how the two-step reasoning is formulated as a multi-task node classification, and (3) the post-processing after performing reasoning on each node in AIGs.

**1) Fusing Structural and Functional Information:** We leverage the message propagation and neighborhood aggregation in GNNs to generate the node embeddings of AIGs that simultaneously fuse structural and functional information. First, the structural information is distilled by passing node embeddings along edges that connect them. Second, the Boolean functional information can be encoded in node features. For each node, there are three node features represented in binary values denoting node types and Boolean functionality. The first node feature indicates whether this node is a PI/PO or intermediate node (i.e., AND gate). The second and the third node features indicate whether each input edge is inverted or not, such that AIGs can be represented as homogeneous graphs without additional edge features. These compressed node features not only encapsulate Boolean functionality of each node but also enable high compute and memory efficiency. Figure 3(b) shows the AIG of a 3-bit CSA multiplier, in which the structural information is presented in the AIG topology and the functional information is encoded in node features. For example, node 1 is a PI with the feature vector [0, 0, 0]; node 7 is an internal node without negation on inputs, so the feature vector is [1, 0, 0]; node 17 has two inputs inverted, with the feature vector [1, 1, 1].

With the emphasis on generalization from simple to complex designs, the specific model employed is GraphSAGE [11]. Given a GraphSAGE model with  $K$  layers, the node embeddings propagated between different layers are computed as follows:

$$h_{\mathcal{N}(v)}^k \leftarrow \text{AGGREGATE}_k(\{h_u^{k-1} \mid \forall u \in \mathcal{N}(v)\}); \quad (1)$$

$$h_v^k \leftarrow \sigma(\mathbf{W}^k \cdot \text{CONCAT}(h_v^{k-1}, h_{\mathcal{N}(v)}^k)).$$

Here,  $\mathcal{N}(v)$  is the immediate neighborhood of node  $v$ ;  $\text{AGGREGATE}_k$  and  $\mathbf{W}^k$  are the aggregation function and the weight matrix for layer  $k$ , respectively, where  $\forall k \in \{1, \dots, K\}$ . After stacking  $K$  layers, the

structural and functional information within  $K$ -hop search depth is fused in the embedding of each node.

2) *Multi-Task Classification*: We identify the Boolean function of each node by multi-task node classification to approach the nature of the problem: there are two steps involved in reasoning functional blocks from unstructured AIGs. The first step detects XOR and MAJ functions from AIGs, which will be used to construct adders. Since each XOR/MAJ function consists of multiple nodes in AIGs, only the root nodes of these functions are labeled as XOR/MAJ with other nodes marked as plain nodes. In addition to the exact XOR/MAJ functions, negation-permutation-negation equivalent functions are also labeled as XOR/MAJ. The second step aims to automatically identify the boundaries of HAs and FAs, and thus we label roots (i.e. the sum and the carry-out functions) and leaves of each adder. Figure 3(c) shows a multi-label annotated AIG of a 3-bit multiplier, using the ground truth provided by ABC. Notably, one node can have multiple labels. For example, node 20 is labeled with XOR and the root of an adder; node 17 is labeled with XOR.

The MTL not only follows the intuition of this two-step reasoning but also exploits divide and conquer, since it is extremely hard for GNNs to reach high prediction accuracy with a single-task multi-label node classification. The employment of MTL enables knowledge sharing across sub-tasks and improves sample efficiency during training, which guarantees high reasoning performance. Specifically, the two-step reasoning is decoupled into three simpler classification tasks using generated node embeddings: *Task 1* classifies the roots and leaves of adders; *Task 2* and *Task 3* detect XOR and MAJ nodes, respectively. We use hard parameter sharing for MTL and the overall loss function  $\mathcal{L}$  is shown below:

$$\mathcal{L} = \alpha \cdot \ell(\hat{y}_1, y_1) + \beta \cdot \ell(\hat{y}_2, y_2) + \gamma \cdot \ell(\hat{y}_3, y_3), \quad (2)$$

in which  $\ell$  is the negative log-likelihood between predictions (i.e.,  $\hat{y}_1$ ,  $\hat{y}_2$ , and  $\hat{y}_3$ ) and the ground truth (i.e.,  $y_1$ ,  $y_2$ , and  $y_3$ ), and  $\alpha$ ,  $\beta$ , and  $\gamma$  are hyper-parameters to adjust the importance of each task. In our implementation,  $\alpha = 0.8$  and  $\beta = \gamma = 1$ .

3) *Adder Tree Extraction from Multi-Labeled Graphs*: After performing the multi-task node classification, we can recognize XOR, MAJ, and root nodes of adders. The XOR and MAJ pairs with identical inputs are matched to construct adders. The conversion from Figure 3(c) to 3(d) depicts the adder tree extraction. In Figure 3(c), the AIG has a set of XOR nodes  $\mathbb{X} = \{12, 17, 20, 24, 29, 33, 36, 41, 44\}$  and a set of MAJ nodes  $\mathbb{M} = \{10, 22, 25, 27, 37, 45\}$ . After removing the nodes that are not marked as adder roots,  $\mathbb{X} = \{12, 20, 24, 29, 36, 44\}$ . Given  $\mathbb{X}$  and  $\mathbb{M}$ , node 12 is  $\text{XOR3}(8, 9, 0)$  and node 10 is  $\text{MAJ3}(8, 9, 0)$ , a three-input XOR/MAJ function with node 8, node 9, and the constant zero as the inputs; node 20 is  $\text{XOR3}(10, 13, 14)$  and node 15 is  $\text{MAJ3}(10, 13, 14)$ ; this matching process continues until all six pairs of XOR and MAJ are generated, which are three FAs and three HAs, as shown in Figure 3(d).

Notably, GAMORA adopts graph learning to mimic the *exact* reasoning. In Figure 3(e), one HA cannot be automatically extracted due to the misprediction of node 10. Our evaluation indicates only several nodes near the least significant bit are always mispredicted due to their shallow neighborhood structure, which has a subtle impact on the efficiency of algebraic rewriting. By fusing structural and functional information into node embeddings and using MTL to approach the reasoning nature, GAMORA is expected to reach as close as possible to the *exact* reasoning precision.

## IV. EXPERIMENT

### A. Experiment Setup

The AIG-based CSA and Booth multipliers are generated by the logic synthesis tool ABC [4], with the ground truth provided by the adder tree extraction command [26]. We consider two technology libraries: (1) the reduced standard-cell library *mcnc.genlib* (with gate input size  $\leq 3$ ) from SIS distribution [19], and (2) ASAP 7nm technologies [24]. The GNN-based framework is implemented in Pytorch Geometric [8]. Two GraphSAGE models are developed for simple and complex design netlists: (1) a shallow 4-layer model with the hidden channel of 32 (for CSA multipliers w/ and w/o simple technology mapping), and (2) a deep 8-layer model with the hidden channel of 80 (for Booth multipliers and after complex technology mapping). The generated node embeddings are passed to a shared linear layer with size of 32 and the ReLU activation function, followed by another linear transformation with softmax for each sub-task to perform node classification. Experiments are performed on a Linux host with AMD EPYC 7742 64-core CPUs and one NVIDIA A100 SXM 40GB GPU. *In general, GAMORA is trained on small bitwidth multipliers (typically less than 32-bit) and evaluated on large bitwidth multipliers (up to 2048-bit).*

### B. Evaluation on Reasoning Performance

We evaluate the reasoning performance from three aspects: (1) how functional and structural information influence the reasoning precision; (2) how design complexity affects model selection and training; (3) how technology mapping complicates the reasoning process and what domain insights can be derived to facilitate more accurate symbolic reasoning on complex BNs.

1) *Reasoning Precision Analysis*: Figure 4 illustrates how the reasoning performance on CSA multipliers is affected by different bitwidth multipliers for training, single/multi-task setting, and the employment of functional information. First, the larger bitwidth multiplier is adopted for training, the higher reasoning precision can be achieved, which typically converges after training with 8-bit multipliers. The main reason is for CSA multipliers, an 8-bit multiplier is able to provide a sufficient variety of structural properties, which can be learned and well generalized to larger multipliers by GAMORA. Second, the multi-task setting conspicuously outperforms the single-task counterpart, indicating that the knowledge sharing across multiple tasks greatly benefits the prediction accuracy of every single task. Third, there is always a boost of accuracy when employing functional information for prediction, since identifying the role of each node relies on not only the surrounding structure but also the function of itself and its neighbors. The synergy of structural and functional information in GAMORA is analogous to the combination of structural hashing and functional propagation in conventional symbolic reasoning.

With the multi-task setting and simultaneously fusing structural and functional attributes, GAMORA achieves almost 100% prediction accuracy in symbolic reasoning for CSA multipliers. It is noted that several nodes near the least significant bit (LSB) are always mispredicted due to their shallow neighborhood structure, as shown in Figure 3(e). This means the HA at LSB cannot be automatically extracted, but can be easily corrected during post-processing.

2) *The Impact of Design Complexity*: We analyze the impact from design complexity by evaluating the reasoning performance on radix-4 Booth-encoded multipliers, as shown in Figure 6. From the model selection aspect, as Booth multipliers generally have more complex structures, deeper models are necessary to characterize neighborhood structures and provide informative node embeddings, thus guaranteeing high prediction accuracy. From the training aspect, larger multipliers

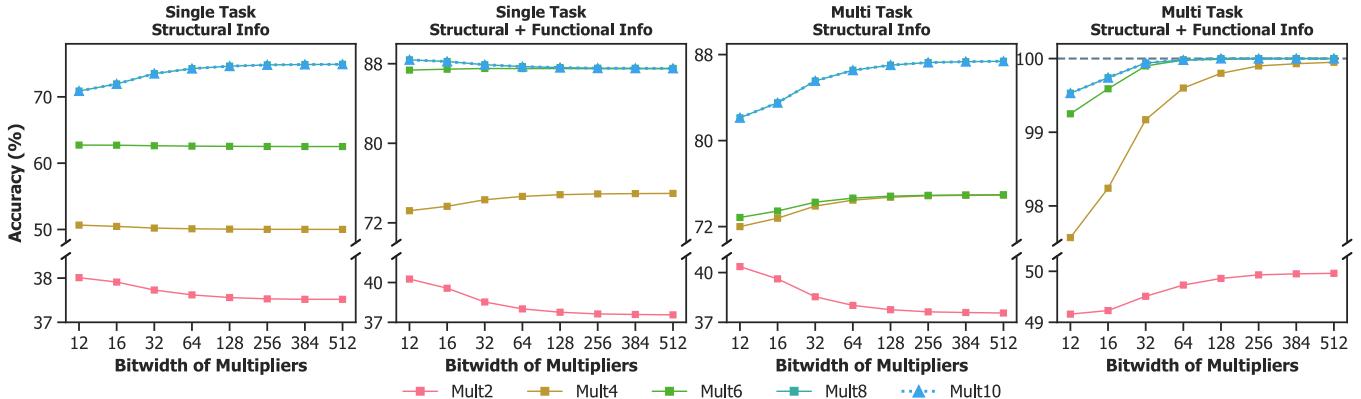


Fig. 4: Sensitivity analysis on CSA multipliers with respect to (1) the bitwidth of multipliers for training (ranging from 2-bit to 10-bit), (2) single/multi-task, and (3) whether employing functional information.

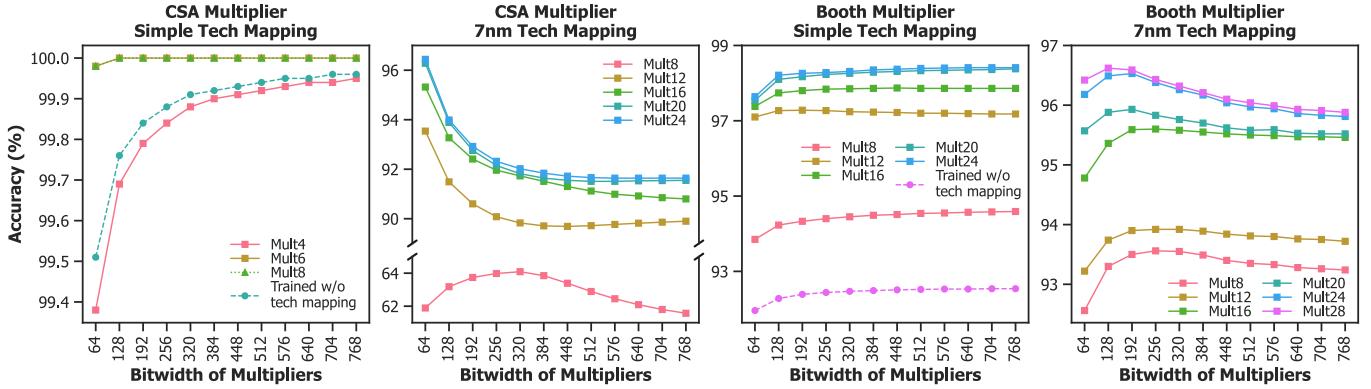


Fig. 5: Evaluation on CSA and Booth multipliers, with simple and complex technology mapping.

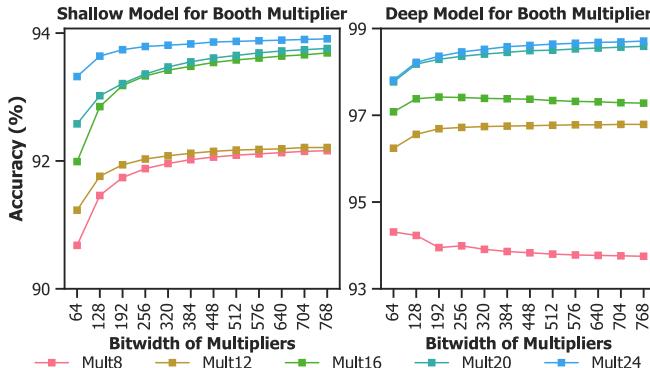


Fig. 6: Evaluation on Booth multipliers with shallow and deep models.

(i.e., up to 24-bit Booth multiplier) are required for training such that adequate variety and representativeness of structural and functional characteristics are exposed to and well captured by GAMORA.

3) *The Impact of Technology Mapping*: It is a known challenge that technology mapping can increase the complexity of formal reasoning on BNs [15], [20], [25]. Thus, we evaluate the performance of GAMORA with respect to different technology mapping options. The multipliers are mapped using the ABC standard-cell mapper (command `map`). Figure 5 depicts the reasoning performance on CSA and Booth multipliers after simple technology mapping [19] and ASAP 7nm technology mapping [24]. Specifically, the ASAP 7nm library contains 161 standard-cell gates, including multi-output cells such as the full adder cell, which significantly increases the complexity and irregularity of post-mapping netlists.

In the simple technology mapping case, the models trained before technology mapping demonstrate good generalization capability, still reaching over 99% and 92% prediction accuracy for CSA and Booth multipliers, respectively; with retraining, comparable reasoning performance to those on original multipliers is achieved with similar sizes of training multipliers. The scenario is fairly different in the case of ASAP 7nm technology mapping, which employs a relatively complex technology library: first, the generalization capability is limited before and after technology mapping; second, the prediction accuracy slightly drops even with retraining; third, it is necessary to use large training multipliers to guarantee performance.

These observations imply several takeaways. First, the more complex technology library is applied, the more difficult it is for learning-based symbolic reasoning, since more complexity is involved both in AIG structures and the functionality of each node. This also implicates attributes related to the technology library should be included in node and edge features. Second, the capability to cope with intricate AIG netlists comes at the expense of more comprehensive training data. One underlying assumption of many supervised ML tasks is the training and testing data should be independent and identically distributed, which is governed by a fundamental principle called empirical risk minimization that provides theoretical performance bounds [22]. Thus, increasing the size of training data can envelop more knowledge of interested statistical properties, ensuring better generalization to testing data.

### C. Runtime and Scalability Analysis

In addition to the high reasoning performance, we demonstrate the superiority of GAMORA by analyzing its runtime and scalability.

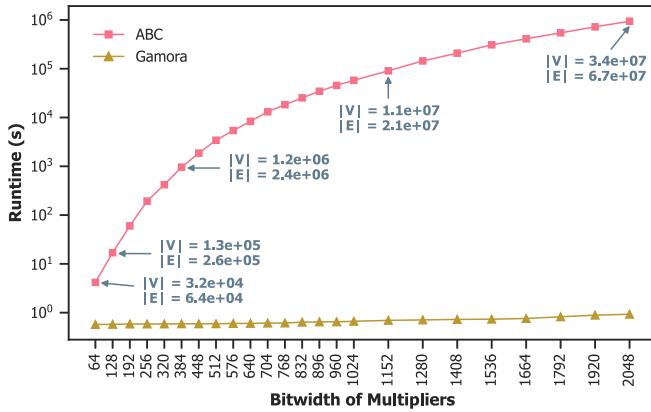


Fig. 7: Runtime comparison between GAMORA and ABC. Note that the number of nodes  $|V|$  and the number of edges  $|E|$  are annotated for scalability analysis.

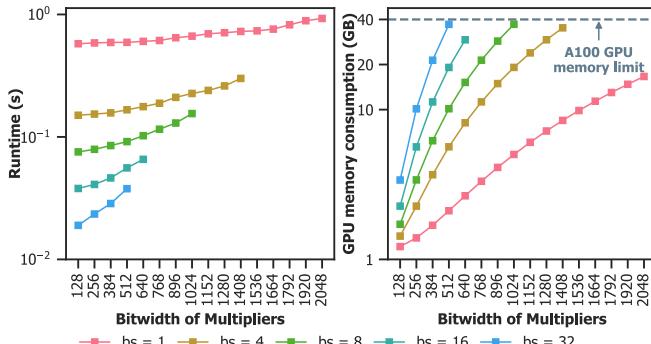


Fig. 8: Average runtime and GPU memory consumption with batched reasoning, where the batch size is denoted as bs. We currently focus on single-GPU implementation.

**Runtime complexity analysis.** Basically, the runtime only relates to the scale of AIGs, i.e., the number of nodes  $|V|$  and the number of edges  $|E|$ . Figure 7 compares the runtime of GAMORA against ABC on CSA multipliers: for large designs such as a 2048-bit CSA multiplier with around 34 million nodes and 67 million edges, GAMORA attains a speedup of up to six orders of magnitude. This shows not only the great efficiency in symbolic reasoning enabled by graph learning but also the scalability to extremely large designs.

**Batched reasoning with single GPU.** Figure 8 shows further acceleration allowed by batched reasoning. Currently, we focus on single GPU implementation, which limits the batch size by the GPU memory, and leave multi-GPU implementation as our future work to support larger batch processing. Even with a single GPU, there already reveal promising results and positive trends benefiting from parallel execution and GPU acceleration.

## V. CONCLUSION

Reasoning high-level abstractions from bit-blasted BNs has benefited functional verification, logic minimization, datapath synthesis, malicious logic identification, etc. In this work, we propose a novel symbolic reasoning framework, GAMORA, which exploits GNNs to imitate structural hashing and functional aggregation in conventional reasoning approaches. Evaluation shows that (1) with the proposed multi-task GNN model, GAMORA offers **high reasoning performance** that reaches almost 100% and over 97% accuracy for CSA and Booth-encoded multipliers, which is still over 92% in finding functional modules after complex technology mapping; (2) with GPU

acceleration on graph learning, GAMORA has **strong scalability** to BNs with over 33 million nodes, with up to **six orders of magnitude speedups** compared to the state-of-the-art implementation in the ABC framework; (3) GAMORA also demonstrates **great generalization capability** from simple to complex designs, such as from small to large bitwidth multipliers, and from before to after technology mapping. GAMORA reveals the great potential of applying GNNs and GPU acceleration to speed up symbolic reasoning.

## VI. ACKNOWLEDGE

This work is supported by National Science Foundation (NSF) under NSF-2047176, NSF-2019336, NSF-2008144, and NSF-2229562 awards.

## REFERENCES

- [1] Abbas Acar et al. A survey on homomorphic encryption schemes: Theory and implementation. *CSUR*, 2018.
- [2] Lila Alrahi et al. Gnn-re: Graph neural networks for reverse engineering of gate-level netlists. *IEEE TCAD*, 2021.
- [3] Ulbert J Botero et al. Hardware trust and assurance through reverse engineering: A tutorial and outlook from image analysis and machine learning perspectives. *ACM JETC*, 2021.
- [4] Robert Brayton and Alan Mishchenko. Abc: An academic industrial-strength verification tool. In *Proc. CAV*. Springer, 2010.
- [5] Tim Bucher et al. Appgnn: Approximation-aware functional reverse engineering using graph neural networks. *arXiv:2208.10868*, 2022.
- [6] Burcin Cakir and Sharad Malik. Reverse engineering digital ics through geometric embedding of circuit graphs. *ACM TODAES*, 2018.
- [7] Maciej Ciesielski et al. Understanding algebraic rewriting for arithmetic circuit verification: a bit-flow model. *IEEE TCAD*, 2019.
- [8] Matthias Fey and Jan E. Lenssen. Fast graph representation learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019.
- [9] Harry Foster. The 2022 wilson research group functional verification study, Accessed: 2022.
- [10] Adria Gascón et al. Template-based circuit understanding. In *Proc. FMCAD*, 2014.
- [11] Will Hamilton et al. Inductive representation learning on large graphs. In *Proc. NeurIPS*, 2017.
- [12] Zhuolun He et al. Graph learning-based arithmetic block identification. In *Proc. ICCAD*, 2021.
- [13] Daniela Kaufmann et al. Verifying large multipliers by combining sat and computer algebra. In *Proc. FMCAD*, 2019.
- [14] Haocheng Li et al. Attacking split manufacturing from a deep learning perspective. In *Proc. DAC*, 2019.
- [15] Wenchao Li et al. Wordrev: Finding word-level structures in a sea of bit-level gates. In *Proc. HOST*, 2013.
- [16] Alireza Mahzoon et al. Revsca: Using reverse engineering to bring light into backward rewriting for big and dirty multipliers. In *Proc. DAC*, 2019.
- [17] Alireza Mahzoon et al. Formal verification of modular multipliers using symbolic computer algebra and boolean satisfiability. In *Proc. DAC*, 2022.
- [18] Alan Mishchenko et al. Dag-aware aig rewriting: A fresh look at combinational logic synthesis. In *Proc. DAC*, 2006.
- [19] Ellen M Sentovich et al. Sis: A system for sequential circuit synthesis. 1992.
- [20] Pramod Subramanyan et al. Reverse engineering digital circuits using structural and functional analyses. *IEEE TETC*, 2013.
- [21] Mertcan Temel and Warren A Hunt. Sound and automated verification of real-world rtl multipliers. In *Proc. FMCAD*, 2021.
- [22] Vladimir Vapnik. Principles of risk minimization for learning theory. *Proc. NeurIPS*, 1991.
- [23] Nan Wu and Yuan Xie. A survey of machine learning for computer architecture and systems. *ACM Comput. Surveys*, 2022.
- [24] Xiaoqing Xu et al. Standard cell library design and optimization methodology for asap7 pdk. In *Proc. ICCAD*, 2017.
- [25] Cunxi Yu et al. Formal verification of arithmetic circuits by function extraction. *IEEE TCAD*, 2016.
- [26] Cunxi Yu et al. Fast algebraic rewriting based on and-inverter graphs. *IEEE TCAD*, 2017.
- [27] Guangwei Zhao and Kaveh Shamsi. Graph neural network based netlist operator detection under circuit rewriting. In *Proc. GLSVLSI*, 2022.