

BDD2Seq: Enabling Scalable Reversible-Circuit Synthesis via Graph-to-Sequence Learning

Mingkai Miao¹, Jianheng Tang², Guangyu Hu^{2*}, Hongce Zhang^{1,2*}

¹Hong Kong University of Science and Technology (Guangzhou), Guangzhou, Guangdong, China

²Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong
mmiao815@connect.hkust-gz.edu.cn, {ghuae,jtangbf}@connect.ust.hk, hongcezh@ust.hk

Abstract

Binary Decision Diagrams (BDDs) are instrumental in many electronic design automation (EDA) tasks thanks to their compact representation of Boolean functions. In BDD-based reversible-circuit synthesis, which is critical for quantum computing, the chosen variable ordering governs the number of BDD nodes and thus the key metrics of resource consumption, such as Quantum Cost. Because finding an optimal variable ordering for BDDs is an NP-complete problem, existing heuristics often degrade as circuit complexity grows. We introduce `BDD2Seq`, a graph-to-sequence framework that couples a Graph Neural Network encoder with a Pointer-Network decoder and Diverse Beam Search to predict high-quality orderings. By treating the circuit netlist as a graph, `BDD2Seq` learns structural dependencies that conventional heuristics overlooked, yielding smaller BDDs and faster synthesis. Extensive experiments on three public benchmarks show that `BDD2Seq` achieves around 1.4 times lower Quantum Cost and 3.7 times faster synthesis than modern heuristic algorithms. To the best of our knowledge, this is the first work to tackle the variable-ordering problem in BDD-based reversible-circuit synthesis with a graph-based generative model and diversity-promoting decoding.

Code — <https://github.com/tracymiao111/BDD2Seq>

Extended version — <https://arxiv.org/pdf/2511.08315>

Introduction

Binary Decision Diagrams (BDDs) (Akers 1978), a compact and efficient graph representation of Boolean functions, play a significant role in electronic design automation (EDA) tasks such as logic optimization, circuit synthesis and formal verification (Hachtel and Somenzi 2005). BDDs structurally represent Boolean functions through decision nodes that branch according to Boolean variables, ultimately terminating in nodes representing constant Boolean values. Reduced Ordered Binary Decision Diagrams (ROBDDs) are among the most widely adopted BDD variants due to their canonicity, ensured by adhering to a predefined variable ordering and eliminating redundant nodes through merging isomorphic subgraphs. Consequently, for a given variable ordering,

the ROBDD representation of a Boolean function is unique and compact (Bryant 1986). In the following context, we will indistinguishably use the term BDD to refer to ROBDD.

BDDs have notable applications in reversible circuit synthesis, a critical area in quantum computing and low-power design (Wille and Drechsler 2009). BDD-based synthesis methods generally provide scalability and resource efficiency advantages, especially for handling large Boolean functions (Kerntopf, Perkowski, and Podlaski 2012). Crucially, resource metrics such as Quantum Cost and gate count in reversible circuits directly correlate with the node count of their BDD representations (Soeken et al. 2016; Wille and Drechsler 2009); therefore, minimizing BDD sizes significantly improves these metrics. However, achieving compact BDD representations strongly depends on selecting an effective variable ordering; a poor ordering can dramatically increase BDD size and compromise efficiency.

As will be detailed in the next section, existing methods for deciding variable ordering often struggle to balance computational efficiency and solution quality, exhibiting inconsistent performance especially for large-scale circuits. Due to these limitations, there is significant motivation to explore alternative methods capable of learning from past experiences and structural information inherent in circuits. Recent advances in machine learning, especially techniques that leverage structural and sequential information, provide promising pathways to address these challenges. Its success in EDA tasks, ranging from circuit design optimization to hardware formal verification, demonstrates its potential to address longstanding challenges in this domain (Huang et al. 2021; Guo et al. 2023). Specifically, Graph Neural Networks (GNNs) have emerged as a powerful tool for modeling and analyzing circuit structures as they can inherently capture the relational and structural information in graphs and circuits can naturally be modeled as nodes and edges of graphs (Li et al. 2022; He et al. 2021; Zhu et al. 2022).

Since GNNs do not inherently support sequential predictions, we integrate a generative Pointer Network decoder (Vinyals, Fortunato, and Jaitly 2015), motivated by its demonstrated success in combinatorial optimization tasks (Ma et al. 2019). Unlike traditional deterministic methods, auto-regressive decoders produce probability distributions over candidate sequences. Leveraging this advantage, we employ Diverse Beam Search (Freitag and Al-

*Corresponding authors are Hongce Zhang and Guangyu Hu.
Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Onaizan 2017; Vijayakumar et al. 2018) to simultaneously explore multiple promising candidate sequences, thus improving solution diversity and enhancing the likelihood of identifying higher-quality variable orderings.

In this work, we present a comprehensive framework BDD2Seq for optimizing BDD variable ordering through the integration of graph-based learning and advanced decoding strategies. Our contributions include:

- **BLIF Format Circuit-to-Graph Representation:** We propose a novel method to embed circuit descriptions (in BLIF format) into graph-based representations. This circuit-to-graph embedding facilitates the application of GNN, enabling it to capture intricate structural and relational information inherent in circuit designs.
- **NLP-Inspired Sequence Generation with Diversity.** We incorporate a Pointer-Network decoder with Diverse Beam Search, which—unlike the conventional NLP beam search—retains multiple mutually dissimilar candidate orderings throughout decoding. This diversity markedly raises the chance of discovering near-optimal variable permutations in a single pass.

Extensive experiments on reversible circuit synthesis tasks validate the effectiveness of our proposed framework, demonstrating substantial and consistent improvements over traditional methods in critical metrics, particularly Quantum Cost (QC). Notably, our approach maintains significantly lower computational complexity growth as circuit size increases, addressing a key limitation of conventional heuristics.


Background

Binary Decision Diagram

Binary Decision Diagrams (BDDs) are graph-based structures representing Boolean functions through decision nodes arranged according to Boolean variables. Reduced Ordered Binary Decision Diagrams (ROBDDs) refine BDDs by enforcing a fixed variable ordering and removing redundant nodes, achieving a canonical and compact representation.

The size of a BDD, indicated by the node count, significantly impacts computational efficiency and resource utilization. Consequently, optimizing BDD size is essential, as it directly correlates with performance metrics such as Quantum Cost, gate count, and computational resources in reversible synthesis tasks (Wille and Drechsler 2009; Soeken et al. 2016). One critical factor affecting BDD size is the variable ordering, which, if poorly chosen, can result in exponential growth of the BDD representation and significant slowdown of synthesis. This variability underscores the importance of selecting an effective variable ordering to maintain manageable BDD sizes. To illustrate, **Figure 1** shows the BDD representations of the same Boolean function $f = (x_0x_1 \vee x_2x_3 \vee x_4x_5)$ using two different variable orderings. The ordering $(x_0, x_1, x_2, x_3, x_4, x_5)$ yields a compact BDD structure with total node count 8, whereas the ordering $(x_0, x_3, x_1, x_4, x_2, x_5)$ leads to significant node expansion with total node count 16. This example clearly demonstrates the sensitivity of BDD node count to variable

ordering and shows the importance of selecting efficient orderings to optimize performance in EDA applications.



$ordering_1 = (x_0, x_1, x_2, x_3, x_4, x_5)$ $ordering_2 = (x_0, x_3, x_1, x_4, x_2, x_5)$

Figure 1: $f = (x_0x_1 \vee x_2x_3 \vee x_4x_5)$ with different orderings

However, identifying an optimal variable ordering is an NP-complete problem (Bollig and Wegener 1996), necessitating efficient algorithms to provide near-optimal solutions. Existing approaches typically include heuristic and exact algorithms, each with distinct characteristics. Heuristic methods, such as sifting (Rudell 1993) and its variants (e.g., symmetric sifting (Scholl et al. 1999) and group sifting (Panda and Somenzi 1995)), iteratively relocate variables to positions that locally minimize BDD size, prioritizing computational efficiency. Genetic algorithms (GAs) (Drechsler, Becker, and Göckel 1996) explore the search space more broadly by evolving candidate orderings through selection, crossover, and mutation. Linear methods (Gunther and Drechsler 1999) utilize linear transformations for node-size optimization, providing rapid results with simplicity. On the other hand, exact algorithms systematically navigate the search space using techniques like dynamic programming (Friedman and Supowit 1990; Stergiou 2011) to guarantee optimal solutions.

Despite these various approaches, significant challenges persist: heuristic methods frequently converge prematurely, missing globally optimal solutions, and suffer inconsistent performance across diverse circuits, particularly for more complex Boolean functions. Conversely, exact algorithms ensure optimal solutions but are limited by their exponential computational complexity, restricting practical use to small-scale circuits. A hybrid optimizer melding a steady-state genetic algorithm with several contemporary swarm-intelligence variants was introduced (Awad, Hawash, and Abdalhaq 2023), translating continuous search dynamics into BDD variable permutations and delivering almost linear scalability in both size reduction and runtime. However, its reported effectiveness is validated on only seven benchmark functions, leaving the generality of those gains open to further validation. By framing Boolean functions as hypergraphs and training a 3-hypergraph model, the method

proposed in (Xu et al. 2018a) predicts near-optimal variable orderings with reduced computation time; nevertheless, its BDD node-size reduction remains inferior to both the GA and the Linear approach.

To further illustrate, in **Table 1** we present a comparison of the BDD size and reordering time on three representative circuits from ISCAS85 (Brglez and Fujiwara 1985) and LGSynth91 (Yang 1991): `c880`, `dalu`, and `rot`, using the sifting algorithm, known for its efficiency, and the genetic algorithm, GA, typically recognized for its better performance. For `c880` and `dalu`, sifting is indeed quick but offers limited node-size reduction, while GA achieves improved reduction at the cost of significantly longer runtime. However, the results for the `rot` circuit reveal a critical issue: GA not only becomes excessively time-consuming but also severely underperforms in node-size reduction compared to sifting. This observation underscores a fundamental limitation in existing heuristic algorithms: they lack the ability to learn from the intrinsic structural properties of diverse circuits, leading to inconsistent performance.

Circuit	PI/PO [†]	Gates	Alg. [‡]	Time(s)	Nodes
c880	60/26	383	Sifting	0.74	11437
			GA	208.03	6127
dalu	75/16	1697	Sifting	39.48	898
			GA	63.45	772
rot	135/107	691	Sifting	2.73	9652
			GA	240.21	573369

[†] PI: *primary input*; PO: *primary output*.

[‡] Alg. denotes Algorithms.

Table 1: Sifting/GA for BDD reordering on 3 samples

Consequently, these limitations strongly motivate developing data-driven methods capable of effectively balancing computational efficiency and consistent performance of BDD size optimization.

Reversible Circuit Synthesis

A reversible circuit is the realization of reversible logic (Toffoli 1980), where each output uniquely maps back to a specific input, ensuring that the computational process is inherently invertible. Such one-to-one mappings significantly minimize energy dissipation and power consumption, making reversible circuits particularly attractive for energy-efficient and low-power circuit designs (Landauer 1961). Quantum circuits represent a prominent subclass of reversible circuits (Chiribella, Yang, and Renner 2021). Particularly, arithmetic quantum circuits like quantum adders play a critical role as core subroutines in numerous significant quantum algorithms, such as Shor’s factoring algorithm. Efficient quantum arithmetic circuits directly influence the feasibility and performance of quantum algorithms by reducing gate complexity, circuit depth, and resource demands, making their optimization crucial for practical quantum computing implementations (Gidney and Ekerå 2021).

To efficiently synthesize such arithmetic circuits and other reversible circuits derived from classical Boolean functions, prior works have proposed BDD-based approaches for reversible circuit synthesis (Wille and Drechsler 2009), which leverage the compactness and canonicity of BDDs to generate efficient reversible circuits, where each BDD node directly maps to a sequence of Toffoli and CNOT gates. Compared with alternative synthesis methods, BDD-based approaches typically achieve superior performance in critical metrics such as Gate Count and Quantum Cost (QC), which directly represent the resource requirements of quantum circuit implementations (Kerntopf, Perkowski, and Podlaski 2012).

While BDD-based reversible circuit synthesis offers a highly efficient method for generating reversible circuits, the challenge remains to be minimizing the number of BDD nodes. Fewer BDD nodes directly lead to a smaller, more efficient circuit with fewer gates (Soeken et al. 2016), which ultimately reduces Quantum Cost. **Table 2** and additional results in the extended version explicitly illustrate this impact by comparing reversible circuits synthesized using a natural (unoptimized) ordering with those produced by optimized variable orderings. The results clearly highlight the substantial benefits achievable through effective variable ordering optimization.

Metric	BDD w/o Reorder	w/ GA Reorder
Gates	18	13
Lines	10	9
Quantum Costs	54	37
Transistor Costs	200	144

Table 2: Metrics for reversible C17 circuit with/without GA

Graph Learning for EDA Tasks

Graph learning has recently attracted substantial interest within the field of EDA, primarily because many EDA tasks—including netlist optimization, circuit layout, and timing analysis—naturally map onto graph-based representations (Ma et al. 2020). Unlike traditional heuristic or analytical methods, graph learning explicitly leverages structural information encoded within circuit graphs, enabling models to efficiently capture relational patterns and dependencies. This capability positions graph learning methods as promising candidates for addressing complex circuit-related optimization problems, which are challenging for conventional analytical or heuristic approaches.

Pointer Network with Diverse Beam Search

A Pointer Network generates permutations by pointing to input positions, while Beam Search keeps the k best partial sequences at each step (Vinyals, Fortunato, and Jaitly 2015; Freitag and Al-Onaizan 2017). The pair already excels on permutation problems such as Traveling Salesman Problem (TSP). Because model likelihood and real-world EDA objectives (e.g., Quantum Cost) do not always coincide, we

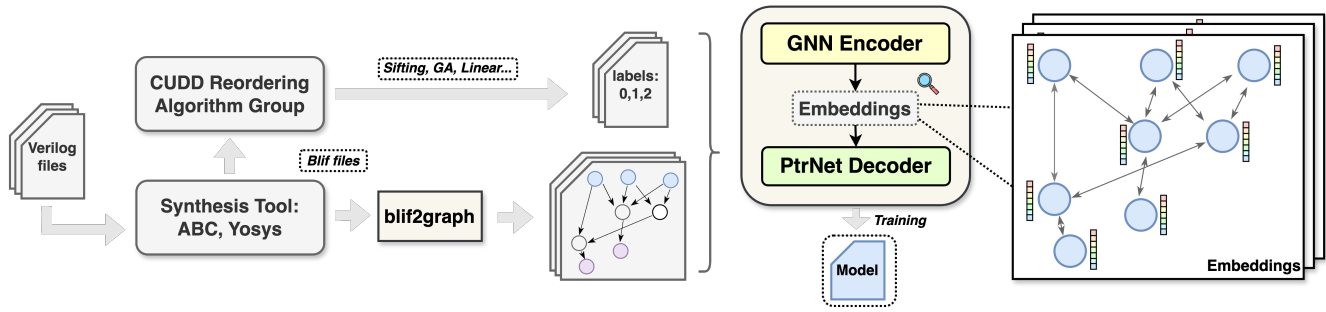


Figure 2: Overall training phase

couple the decoder with Diverse Beam Search (Vijayakumar et al. 2018) to retain several mutually distinct candidates, greatly increasing the chance that at least one aligns with the downstream metric.

Methodology

Overview of Framework

In this section, we will describe the functionality of BDD2Seq in two phases. The first is the training phase, which outlines the steps of data preprocessing, label generation, graph construction, and model training. The second is the inference phase, focused on the reversible circuit synthesis task, where we apply the trained model with Diverse Beam Search strategy to generate the predicted variable ordering, which is then used to perform reversible circuit synthesis.

Training Phase. **Figure 2** overviews the overall training phase. (i) *Label generation.* Each Verilog design is first mapped to a gate-level BLIF netlist with ABC or Yosys (Brayton and Mishchenko 2010; Wolf 2016). We then run several CUDD ordering heuristics (Sifting, GA, Linear, etc.) (Somenzi 2009) and retain the ordering that minimises the BDD node count; this “best-known” sequence serves as the supervisory label. (ii) *Graph construction.* The netlist is converted by BLIF2Graph (details given later) into a directed graph whose node features include the gate truth table and a few basic topological descriptors, such as fan-in, fan-out, and level depth, providing lightweight structural context. (iii) *Model training.* A GNN encoder embeds the graph; a Pointer-Network decoder autoregressively outputs the variable order. Parameters are learned end-to-end by minimizing a weighted negative log-likelihood:

$$\mathcal{L} = \frac{1}{B} \sum_{b=1}^B \frac{\sum_{t=1}^T -\log p_{\theta}(y_{b,t}) w_t m_{b,t}}{\sum_{t=1}^T m_{b,t}},$$

where w_t emphasizes early positions and $m_{b,t}$ masks padding tokens. Model parameters are optimized via backpropagation to minimize this loss.

Inference Phase with Reversible Circuit Synthesis.

Once the model has been trained, the inference process begins by taking new circuit files, followed by the same graph embedding step as described previously.

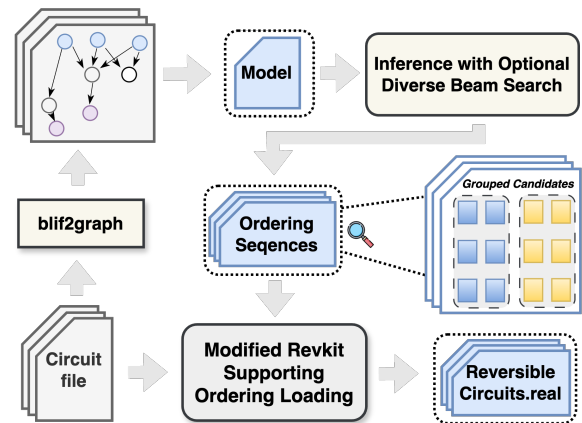


Figure 3: Inference phase with reversible circuit synthesis

Shown as **Figure 3**, the core inference process involves leveraging the trained GNN encoder and Pointer Network decoder, combined with a Diverse Beam Search strategy, to predict high-quality variable ordering sequences. Diverse Beam Search enables the model to explore multiple candidate sequences, significantly increasing the likelihood of identifying an optimal or near-optimal variable ordering.

Once generated, the predicted ordering sequence is directly provided to a modified version of the Revkit synthesis framework (Soeken et al. 2010), a widely used reversible circuit synthesis toolkit, which we adapted to accept externally provided variable orderings. With this predicted ordering, Revkit constructs the circuit BDD and then synthesizes the reversible circuit. The synthesized circuit is subsequently outputted as a *.real* file, explicitly representing the optimized reversible circuit leveraging the variable ordering predicted by our model.

Circuit Embedding - BLIF2Graph

Prior GNN-based work (Li et al. 2023) converts circuits into And Inverter Graph (AIG/AIGER) form before building the graph representation, a format that contains only AND gates and inverters and thus offers limited functional diversity; the lack of native support for BLIF netlists, where each gate is specified by a full truth table, motivates our richer




Figure 4: BLIF2Graph: BLIF to graph with embeddings

BLIF2Graph embedding. As illustrated in **Figure 4**, our proposed embedding method transforms circuit descriptions provided in BLIF format into graph representations suitable for processing by GNNs. In the BLIF format, circuits are defined through inputs, outputs, and logic gates (using `.names` as statement), where each logic gate’s behavior is specified by a compact Boolean relationship between inputs and outputs.

A key innovation in our embedding method is the **truth-table encoding**, a scheme designed to enable the GNN to reason about the circuit’s logic function, not just its topology. It compactly encodes the logical function of each gate into a uniform-length vector embedded directly as node features, allowing the GNN to accurately capture functional dependencies among nodes. **Algorithm 1** outlines this encoding process concisely.

To enrich the representation, we concatenate three lightweight structural scalars per node—topological rank, depth (longest PI-to-node path), and fan-in / fan-out—so the model can distinguish sources from sinks and recognize hubs without incurring extra message-passing steps.

Algorithm 1: Truth Table Embedding for BLIF Gates

Require: BLIF file B , predefined max vector length L
Ensure: Uniform truth-table embeddings V_g for gate g

- 1: **for each** gate $g \in B$ **do**
- 2: Extract inputs $\{x_1, x_2, \dots, x_n\}$ from gate definition
- 3: Compute 2^n input combinations lexicographically:

$$C_g = [(0 \dots 0), (0 \dots 01), \dots, (1 \dots 1)]$$

- 4: Construct truth-table embedding vector V_g :

$$V_g[i] = \begin{cases} 1 & \text{if } C_i \text{ yields } 1, \\ 0 & \text{otherwise} \end{cases}, \quad \forall C_i \in C_g$$

- 5: Pad V_g to length L :

$$V_g \leftarrow [V_g; \mathbf{0}_{L-2^n}]$$

- 6: **end for**
 - 7: **return** $\{V_g\}$ embeddings for all gates
-

Integrating these structural attributes alongside logic-based truth-table embeddings enables the GNN to simultaneously model the functional and relational complexities

inherent in circuit structures, thereby facilitating effective learning for optimized BDD variable ordering. We ablate graph construction by comparing BLIF2Graph with prior AIG-to-Graph (Li et al. 2023) under identical settings; full details and results are in the extended version.

Variable Ordering Prediction

With the circuit embeddings and GNN-based encoding presented, the next step is to decode these learned representations into variable ordering sequences. This decoding process translates graph embeddings into optimized variable permutations, directly determining the efficiency of the resulting reversible circuits.

Pointer-Network Decoding Loop. Since the variable ordering problem requires producing a permutation of the input variables, a Pointer Network (Vinyals, Fortunato, and Jaitly 2015) is a natural choice. Rather than selecting tokens from a predefined vocabulary, the decoder points to nodes in the input graph, which guarantees that every output sequence is a valid permutation. **Figure 5** sketches the decoding loop. Decoding begins with a special `<start>` token. At each step an LSTM (Hochreiter and Schmidhuber 1997) receives the previously selected variable (or the ground-truth variable provided by teacher forcing (Williams and Zipser 1989) during initial training), attends to all node embeddings, masks out nodes that have already been chosen, and converts the masked attention scores into log probabilities. The decoder samples the next variable from this distribution; during inference it may instead choose the top k candidates via Diverse Beam Search. The loop repeats until the `<end>` token is emitted, after which the collected variables constitute the final ordering passed to the synthesis backend.

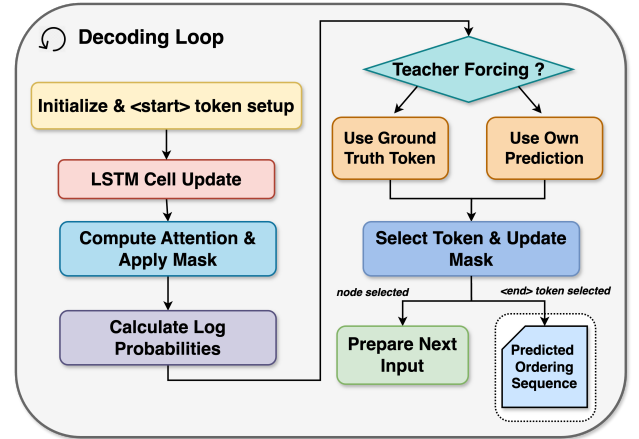


Figure 5: Pointer Network decoding loop

Diverse Beam Search. The likelihood produced by our model serves as a useful proxy, yet it is not perfectly correlated with the ultimate target, Quantum Cost (QC); relying solely on the single most-probable ordering can therefore be suboptimal and calls for a more diversified decoding

strategy. To raise the chance of emitting a low-QC ordering in one forward pass, we embed *Diverse Beam Search* (Vijayakumar et al. 2018) into the decoder.

In contrast to standard beam search, Diverse Beam Search divides the beams into distinct groups, penalizing token selections repeated across these groups to encourage diversity. Formally, this diversity is enforced by applying a penalty to the raw attention scores (logits) before normalization:

$$\text{atnScores}[token] \leftarrow (1 - \alpha) \times \text{atnScores}[token], \quad (1)$$

where $\alpha (0 \leq \alpha \leq 1)$ controls the intensity of diversity enforcement. The penalized scores are then normalized using a log-softmax operation to yield consistent log-probabilities.

Algorithm 2 details this algorithm. At each decoding step, candidate beams are partitioned into n groups, each containing $\frac{m}{n}$ sequences. Attention scores for each token are computed and penalized if previously selected by an earlier group within the current step. The highest-scoring tokens are appended, and the top sequences per group are retained for subsequent steps. This iterative process continues until all variables are ordered, producing diverse, high-quality variable ordering sequences.

By systematically exploring multiple diverse candidate orderings, our method enhances the likelihood of identifying high-quality variable orderings for the subsequent reversible circuit synthesis.

Algorithm 2: Diverse Beam Search for BDD Variable Ordering Prediction

Require: Beam width m , groups n , penalty α , variables X

Ensure: Set of diverse variable ordering sequences

```

1: Initialize groups  $\{G_1, \dots, G_n\}$ , each with  $\frac{m}{n}$  beams
2: for each decoding step  $t = 1, 2, \dots, |X|$  do
3:   for each group  $G_i, i = 1, \dots, n$  do
4:     for each candidate beam in group  $G_i$  do
5:       Compute raw  $\text{atnScores}$  for available tokens
6:       for token previously selected in  $G_j, j < i$ 
7:         end for
8:       Normalize  $\text{atnScores}$  to log probabilities:
9:          $\mathbf{p} \leftarrow \text{log\_softmax}(\text{atnScores})$ 
9:       Append token with highest log-probability
10:      end for
11:      Retain top  $\frac{m}{n}$  sequences per group  $G_i$ 
12:    end for
13:  end for
14: return Final sequences from all groups

```

Experiments

Experiment Setting

Having detailed our methodology, we now proceed to empirically validate the performance of our method on BDD-based reversible circuit synthesis.

We begin by describing the dataset preparation. We construct our dataset by data augmentation techniques including circuit decomposition, random signal negation (Xu et al. 2018a), and fuzzing via `aigfuzz`. The first two techniques take circuits from the LGSynth91 benchmark (Yang 1991) as a reference. The resulted dataset comprises 5241 circuit variants, which are split into training, validation, and testing sets at a ratio of 7:2:1. As for evaluation, we take the circuits from ISCAS85 (Brglez and Fujiwara 1985), LGSynth91, and Revlib (Wille et al. 2008). Although the variants in training phase use LGSynth91 as reference, the original LGSynth91 circuits are unseen and are only used for evaluation purpose. Furthermore, Revlib and ISCAS85 are completely independent benchmarks.

All the experiments were conducted on a platform running Ubuntu 20.04.6 LTS, equipped with 2 NVIDIA GeForce RTX 3090 GPUs and dual Intel® Xeon® Platinum 8375C CPUs at 2.90 GHz.

Graph Encoder Selection

To identify an effective graph encoder for predicting BDD variable orderings, we conducted a comparative evaluation of four prominent GNN architectures: Graph Convolutional Network (GCN) (Kipf and Welling 2017), Graph Isomorphism Network (GIN) (Xu et al. 2018b), GraphSAGE (Hamilton, Ying, and Leskovec 2018), and Graph Attention Network (GAT) (Veličković et al. 2018). Ranking fidelity was measured with Kendall’s τ and Spearman’s ρ ,

$$\tau = \frac{C - D}{\frac{1}{2}n(n-1)}, \quad \rho = 1 - \frac{6 \sum_i d_i^2}{n(n^2 - 1)},$$

where C and D count concordant and discordant pairs, d_i is the rank difference for element i , and n is the sequence length. As shown in the extended version, GAT achieves the highest test-set correlation ($\tau = 0.6182, \rho = 0.6748$) and is thus adopted as our default encoder. Given the task difficulty, with some instances involving permutations of exceeding 200 variables, achieving these scores constitutes a strong rank correlation.

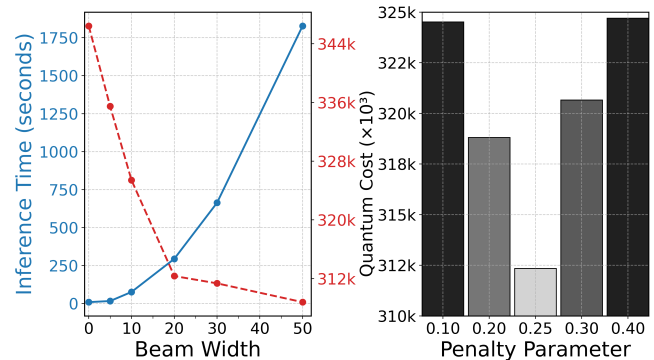


Figure 6: Ablation experiment on Diverse Beam Search

Evaluation on Reversible Circuit Synthesis

We now evaluate the practical impact on reversible circuit synthesis. Experiments are conducted on 148 combinational

Metric	Quality (50 [*])	Balance (20 [*])	Efficiency	SIFT	SYMM [†]	GROUP [†]	GA	LINEAR
Gates	96643	98437	108181	144112	135112	161648	2547934	134377
Lines	28962	29462	31968	41213	38876	45602	686839	38744
Quantum Cost	307367	312345	346425	471792	439452	532588	8628394	437213
Transistors	1164728	1184520	1309936	1776136	1657200	2001736	32346744	1648568
Time (s)	2076.2	521.1	156.2	576.7	630.2	600.4	735.8	581.0

^{*} The numbers in parentheses indicate beam width, i.e., Quality (50) means beam width = 50 for Quality mode.

[†] SYMM: symmetric sifting; GROUP: group sifting.

Table 3: Reversible circuit synthesis metrics comparison using the evaluation dataset

Circuit [‡]	PI/PO	BDD2Seq (QC [*])		Heuristics (QC [*])				
		Balance	Efficiency	SIFT	SYMM [†]	GROUP [†]	GA	LINEAR
dc1_142 ^R	4/7	160	168	160	186	186	186	181
c17 ^L	5/2	37	49	49	49	37	37	49
bw_116 ^R	5/28	924	937	943	937	943	937	943
con1_136 ^R	7/2	88	103	96	95	96	95	96
inc_170 ^R	7/9	579	579	579	592	592	592	621
alu2_96 ^R	10/6	1218	1415	1436	1299	1366	1298	1376
cm151a ^L	12/2	70	70	92	92	92	92	92
add6_92 ^R	12/7	118	566	499	118	474	118	499
alu4_98 ^R	14/8	4455	4934	7222	4334	7623	4403	5583
t481_208 ^R	16/1	139	140	152	152	152	152	152
pm1 ^L	16/13	234	273	273	244	267	244	261
vda ^L	17/39	4224	4224	4477	4477	4460	4169	4470
mux ^L	21/1	135	135	170	170	170	170	170
cm150a_128 ^R	21/1	136	136	186	186	186	186	186
frg1_160 ^R	28/3	598	629	747	747	827	653	747
c880 ^L	60/26	37802	39992	61131	61131	78157	35122	61216
dal ^L	75/16	5297	5652	31145	31145	47170	5235	14330
x4 ^L	94/71	2662	2782	4459	4459	4334	3006	4459
apex5_104 ^R	117/88	9852	10697	10349	9467	10227	10283	10349
rot ^L	135/107	38453	48344	78639	78375	110887	7339159	70456
frg2 ^L	143/139	7189	7189	12468	12361	12154	9391	12111
pair ^L	173/137	20754	20799	46544	46444	49917	1039245	47818
Total		135079	149813	261816	256912	330317	8454773	236165

^{*} Metric: Quantum Cost (QC).

[†] SYMM: symmetric sifting; GROUP: group sifting.

[‡] Superscripts denote datasets: ^L ISCAS85, ^L LGSynth91, ^R Revlib.

[§] See the full tables in the extended version for detailed results.

Table 4: Quantum Cost comparison: BDD2Seq vs. traditional algorithms

circuits from ISCAS85, LGSynth91, and Revlib.

Ablation Study on Diverse Beam Search. To quantify the benefit of adapting Diverse Beam Search and to identify cost-effective settings, we ablate the two hyper-parameters: beam width m and diversity penalty α as shown in **Figure 6**. All other components of BDD2Seq are held fixed. Using $m/2$ diversity groups, increasing m from 0 to 20 cuts Quantum Cost monotonically, confirming that Diverse Beam Search indeed surfaces better orderings. Beyond $m=20$ the quality gain saturates while runtime rises sharply, so we

adopt $m=20$ as the best cost-speed trade-off. With $m=20$, sweeping α shows a bowl-shaped curve: weak penalties ($\alpha < 0.2$) yield insufficient sequence variety, whereas overly strong penalties ($\alpha > 0.3$) scatter probability mass onto low-quality beams. The minimum Quantum Cost occurs at $\alpha=0.25$, which balances exploration and exploitation. These empirically chosen values ($m=20, \alpha=0.25$) define the **Balance** mode used later. The ablation confirms that Diverse Beam Search is indispensable for our performance gains and that its impact can be tuned without prohibitive runtime overhead.

Circuit‡	PI/PO	BDD2Seq (s)		Heuristics (s)				
		Balance	Efficiency	SIFT	SYMM†	GROUP†	GA	LINEAR
dc1_142 ^R	4/7	0.07	0.01	0.01	0.05	0.06	0.01	0.05
c17 ^I	5/2	0.08	0.01	0.07	0.05	0.06	0.01	0.01
bw_116 ^R	5/28	0.08	0.02	0.01	0.09	0.07	0.01	0.06
con1_136 ^R	7/2	0.13	0.02	0.01	0.06	0.05	0.01	0.05
inc_170 ^R	7/9	0.13	0.01	0.01	0.08	0.07	0.01	0.06
alu2_96 ^R	10/6	0.23	0.03	0.01	0.15	0.08	0.03	0.08
cm151a ^L	12/2	0.29	0.02	0.06	0.05	0.05	0.02	0.01
add6_92 ^R	12/7	0.29	0.01	0.01	0.12	0.08	0.02	0.08
alu4_98 ^R	14/8	0.44	0.09	0.03	0.65	0.20	0.18	0.17
t481_208 ^R	16/1	0.47	0.03	0.01	0.11	0.08	0.04	0.08
pm1 ^L	16/13	0.46	0.02	0.06	0.05	0.06	0.02	0.01
vda ^L	17/39	0.56	0.08	0.14	0.14	0.15	0.19	0.03
mux ^L	21/1	1.03	0.39	0.24	0.23	0.26	0.51	0.43
cm150a_128 ^R	21/1	0.84	0.24	0.25	0.25	0.19	0.33	0.37
frg1_160 ^R	28/3	1.34	0.03	0.01	0.25	0.08	0.20	0.08
c880 ^I	60/26	26.22	15.52	47.71	44.43	45.20	121.45	16.86
dalul ^L	75/16	67.32	28.17	182.95	212.63	202.55	270.63	431.41
x4 ^L	94/71	11.84	0.11	0.19	0.20	0.20	0.64	0.04
apex5_104 ^R	117/88	18.27	0.20	0.08	3.62	0.41	1.47	0.42
rot ^L	135/107	47.58	15.69	61.42	60.18	61.68	139.15	20.52
frg2 ^L	143/139	27.54	0.35	1.19	1.20	1.18	2.41	0.25
pair ^L	173/137	39.87	1.06	4.86	4.81	4.91	18.63	1.03
Total		245.08	62.15	299.32	329.39	317.67	555.94	472.03

‡ Superscripts denote datasets: ^I ISCAS85, ^L LGSynth91, ^R Revlib.

† SYMM: symmetric sifting; GROUP: group sifting.

§ Synthesis time reported as 0.01s represents Revkit’s measurement limit ($\leq 0.01s$).

Table 5: Time consumption (seconds): BDD2Seq vs. traditional algorithms§

Overall Reversible-Synthesis Comparison. Thanks to the tunable beam width in Diverse Beam Search, BDD2Seq can be deployed in three practical modes. **Efficiency** (greedy decoding) sacrifices some optimality for minimum latency; **Balance** fixes the beam width at 20 and targets a cost–speed sweet spot; **Quality** widens the beam to 50 to chase the lowest Quantum Cost.

Table 3 shows that, regardless of modes, BDD2Seq beats classical heuristics on every metric. The Balance configuration is usually preferable: it trims Quantum Cost by $1.4\times$ – $27.6\times$ while keeping runtime moderate. When turnaround time dominates, the Efficiency mode is still 3.7 – $4.7\times$ faster than heuristic re-ordering yet continues to lower cost. At the other extreme, the Quality mode pushes cost even lower, albeit with proportionally longer runtime.

Detailed Evaluation on Benchmark Circuits. **Table 4** and **Table 5** contrast BDD2Seq with mainstream heuristics over 22 benchmarks spanning from 4 to 173 primary inputs, reflecting the complexity of the search space. Across the board, BDD2Seq delivers the lowest Quantum Cost, with the advantage widening on large designs such as *c880*, *dalul*, and *rot*. In addition, runtime grows far more slowly than that of heuristic re-ordering, making our approach both cheaper and faster as circuit complexity rises.

Between the two operating modes, **Balance** (beam width 20) secures the best Quantum Cost on 19/22 circuits while incurring only a modest runtime increase over the ultra-fast **Efficiency** mode. In addition, although GA and Linear heuristics were reported to be two of the strongest variable–ordering techniques (Xu et al. 2018a), BDD2Seq surpasses both, delivering lower Quantum Cost and shorter synthesis time.

Conclusion

In this work, we introduced BDD2Seq, a novel deep learning approach to address the problem of optimal variable ordering in BDDs specifically for reversible circuit synthesis. From graph to sequence learning with enhanced searching strategy, our method effectively captures circuit structural and functional information and efficiently explores promising variable permutations.

Extensive experimental results confirm that BDD2Seq significantly outperforms traditional heuristic methods, achieving notable reductions in Quantum Cost and substantially faster synthesis time, particularly in circuits with large numbers of primary inputs. While its three operating modes let users trade runtime for additional cost savings when needed.

Acknowledgments

We thank Yusheng Zhao and Mingfei Yu for their helpful discussions and support for this work.

References

- Akers. 1978. Binary decision diagrams. *IEEE Transactions on computers*, 100(6): 509–516.
- Awad, A.; Hawash, A.; and Abdalhaq, B. 2023. A Genetic Algorithm (GA) and Swarm-Based Binary Decision Diagram (BDD) Reordering Optimizer Reinforced With Recent Operators. *IEEE Transactions on Evolutionary Computation*, 27(3): 535–549.
- Bollig, B.; and Wegener, I. 1996. Improving the variable ordering of OBDDs is NP-complete. *IEEE Transactions on Computers*, 45(9): 993–1002.
- Brayton, R.; and Mishchenko, A. 2010. ABC: an academic industrial-strength verification tool. In *Proceedings of the 22nd International Conference on Computer Aided Verification, CAV'10*, 24–40. Berlin, Heidelberg: Springer-Verlag. ISBN 364214294X.
- Brglez, F.; and Fujiwara, H. 1985. A neutral netlist of 10 combinational benchmark circuits and a target translator. In *Fortran. ISCAS'85*.
- Bryant, R. E. 1986. Graph-based algorithms for boolean function manipulation. *Computers, IEEE Transactions on*, 100(8): 677–691.
- Chiribella, G.; Yang, Y.; and Renner, R. 2021. Fundamental energy requirement of reversible quantum operations. *Physical Review X*, 11(2): 021014.
- Drechsler, R.; Becker, B.; and Göckel, N. 1996. Genetic algorithm for variable ordering of OBDDs. *IEE Proceedings-Computers and Digital Techniques*, 143(6): 364–368.
- Freitag, M.; and Al-Onaizan, Y. 2017. Beam search strategies for neural machine translation. *arXiv preprint arXiv:1702.01806*.
- Friedman, S.; and Supowit, K. 1990. Finding the optimal variable ordering for binary decision diagrams. *IEEE Transactions on Computers*, 39(5): 710–713.
- Gidney, C.; and Ekerå, M. 2021. How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits. *Quantum*, 5: 433.
- Gunther, W.; and Drechsler, R. 1999. Minimization of BDDs using linear transformations based on evolutionary techniques. In *1999 IEEE International Symposium on Circuits and Systems (ISCAS)*, volume 1, 387–390 vol.1.
- Guo, W.; Zhen, H.-L.; Li, X.; Luo, W.; Yuan, M.; Jin, Y.; and Yan, J. 2023. Machine learning methods in solving the boolean satisfiability problem. *Machine Intelligence Research*, 20(5): 640–655.
- Hachtel, G. D.; and Somenzi, F. 2005. *Logic synthesis and verification algorithms*. Springer Science & Business Media.
- Hamilton, W. L.; Ying, R.; and Leskovec, J. 2018. Inductive Representation Learning on Large Graphs. *arXiv:1706.02216*.
- He, Z.; Wang, Z.; Bai, C.; Yang, H.; and Yu, B. 2021. Graph learning-based arithmetic block identification. In *2021 IEEE/ACM International Conference On Computer Aided Design (ICCAD)*, 1–8. IEEE.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long Short-Term Memory. *Neural Computation*, 9: 1735–1780.
- Huang, G.; Hu, J.; He, Y.; Liu, J.; Ma, M.; Shen, Z.; Wu, J.; Xu, Y.; Zhang, H.; Zhong, K.; et al. 2021. Machine learning for electronic design automation: A survey. *ACM Transactions on Design Automation of Electronic Systems (TODAES)*, 26(5): 1–46.
- Kerntopf, P.; Perkowski, M.; and Podlaski, K. 2012. Synthesis of reversible circuits: A view on the state-of-the-art. In *2012 12th IEEE International Conference on Nanotechnology (IEEE-NANO)*, 1–6.
- Kipf, T. N.; and Welling, M. 2017. Semi-Supervised Classification with Graph Convolutional Networks. *arXiv:1609.02907*.
- Landauer, R. 1961. Irreversibility and Heat Generation in the Computing Process. *IBM Journal of Research and Development*, 5(3): 183–191.
- Li, M.; Khan, S.; Shi, Z.; Wang, N.; Yu, H.; and Xu, Q. 2022. Deepgate: Learning neural representations of logic gates. In *Proceedings of the 59th ACM/IEEE Design Automation Conference*, 667–672.
- Li, Y.; Liu, M.; Mishchenko, A.; and Yu, C. 2023. Verilog-to-PyG – A Framework for Graph Learning and Augmentation on RTL Designs. *arXiv:2311.05722*.
- Ma, Q.; Ge, S.; He, D.; Thaker, D.; and Drori, I. 2019. Combinatorial optimization by graph pointer networks and hierarchical reinforcement learning. *arXiv preprint arXiv:1911.04936*.
- Ma, Y.; He, Z.; Li, W.; Zhang, L.; and Yu, B. 2020. Understanding graphs in EDA: From shallow to deep learning. In *Proceedings of the 2020 international symposium on physical design*, 119–126.
- Panda, S.; and Somenzi, F. 1995. Who are the variables in your neighborhood. In *Proceedings of the 1995 IEEE/ACM International Conference on Computer-Aided Design, ICCAD '95*, 74–77. USA: IEEE Computer Society. ISBN 0818672137.
- Rudell, R. 1993. Dynamic variable ordering for ordered binary decision diagrams. In *Proceedings of 1993 International Conference on Computer Aided Design (ICCAD)*, 42–47.
- Scholl, C.; Moller, D.; Molitor, P.; and Drechsler, R. 1999. BDD minimization using symmetries. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 18(2): 81–100.
- Soeken, M.; Frehse, S.; Wille, R.; and Drechsler, R. 2010. RevKit: A Toolkit for Reversible Circuit Design. volume 18.
- Soeken, M.; Tague, L.; Dueck, G. W.; and Drechsler, R. 2016. Ancilla-free synthesis of large reversible functions using binary decision diagrams. *Journal of Symbolic Computation*, 73: 1–26.

- Somenzi, F. 2009. CUDD: CU decision diagram package-release 2.4.0. *University of Colorado at Boulder*, 21.
- Stergiou, S. 2011. Implicit permutation enumeration networks and binary decision diagrams reordering. In *Proceedings of the 48th Design Automation Conference, DAC '11*, 615–620. New York, NY, USA: Association for Computing Machinery. ISBN 9781450306362.
- Toffoli, T. 1980. Reversible computing. In *International colloquium on automata, languages, and programming*, 632–644. Springer.
- Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Liò, P.; and Bengio, Y. 2018. Graph Attention Networks. arXiv:1710.10903.
- Vijayakumar, A. K.; Cogswell, M.; Selvaraju, R. R.; Sun, Q.; Lee, S.; Crandall, D.; and Batra, D. 2018. Diverse Beam Search: Decoding Diverse Solutions from Neural Sequence Models. arXiv:1610.02424.
- Vinyals, O.; Fortunato, M.; and Jaitly, N. 2015. Pointer networks. *Advances in neural information processing systems*, 28.
- Wille, R.; and Drechsler, R. 2009. BDD-based synthesis of reversible logic for large functions. In *2009 46th ACM/IEEE Design Automation Conference*, 270–275.
- Wille, R.; Große, D.; Teuber, L.; Dueck, G. W.; and Drechsler, R. 2008. RevLib: An Online Resource for Reversible Functions and Reversible Circuits. In *38th International Symposium on Multiple Valued Logic (ismvl 2008)*, 220–225.
- Williams, R. J.; and Zipser, D. 1989. A learning algorithm for continually running fully recurrent neural networks. *Neural computation*, 1(2): 270–280.
- Wolf, C. 2016. Yosys open synthesis suite.
- Xu, F.; He, F.; Xie, E.; and Li, L. 2018a. Fast OBDD reordering using neural message passing on hypergraph. *arXiv preprint arXiv:1811.02178*.
- Xu, K.; Hu, W.; Leskovec, J.; and Jegelka, S. 2018b. How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826*.
- Yang, S. 1991. *Logic synthesis and optimization benchmarks user guide: version 3.0*. Citeseer.
- Zhu, K.; Chen, H.; Turner, W. J.; Kokai, G. F.; Wei, P.-H.; Pan, D. Z.; and Ren, H. 2022. Tag: Learning circuit spatial embedding from layouts. In *Proceedings of the 41st IEEE/ACM International Conference on Computer-Aided Design*, 1–9.