

LongRTL: Graph-Similarity-Guided LLM-driven Long Context RTL Optimization

Yuyang Ye¹, Che-Kuan Shen², Xiangfei Hu³, Yuchen Liu³, Shuo Yin¹, Xufeng Yao¹, Bei Yu^{1,†}, Tsung-Yi Ho¹
¹CUHK ²National Central University ³Southeast University

Abstract—Large Language Models (LLMs) show great promise in RTL code generation and optimization. However, real-world RTL designs are typically long, entangled, and poorly modularized—posing a major challenge due to context-length limitations and lack of structure. To overcome these obstacles, we propose a scalable LLM-based RTL optimization framework guided by graph similarity. Our method introduces three collaborative agents: (1) a Partition Agent that decomposes RTL designs into semantically meaningful AST subtrees, guided by AST graph similarity to reusable design templates; (2) an Optimization Agent that generates RTL submodule code based on partitioned subtrees using multi-modal Retrieval-Augmented Generation (RAG) with both AST and RTL guidance; and (3) a Reconstruction Agent that reassembles optimized submodules based on logic-aware ordering and Graph-RAG prompting, ensuring global functional equivalence. Together, these components enable robust, structure-aware optimization of long-context RTL designs, bridging the gap between toy examples and industrial-scale hardware codebases.

I. INTRODUCTION

Large Language Models (LLMs) have recently demonstrated strong capabilities in hardware design tasks such as RTL code generation [1]–[3], optimization [4], [5], and verification [6], [7]. However, existing LLM-based RTL optimization approaches are largely constrained to short, modular code snippets that neatly fit within the context window of modern LLMs. In contrast, real-world industrial RTL designs are often long, monolithic, and lack clear module boundaries, with complex interdependencies spanning thousands of lines. This disparity fundamentally limits the scalability and practicality of current LLM-driven techniques.

A key challenge lies in the context-length bottleneck: LLMs struggle to process global design semantics and long-range dependencies in oversized RTL designs [8]. Effective LLM-based optimization at scale thus requires a principled workflow to: (1) partition the design into meaningful and optimizable units, (2) perform structure- and function-aware optimization for each part, and (3) reconstruct the design while preserving overall functional equivalence.

To address this, we propose a graph similarity-guided LLM optimization framework for long-context RTL designs. Inspired by how human engineers recognize and reuse known design idioms, our method employs AST-level graph similarity to decompose, optimize, and reconstruct RTL in a modular and semantically consistent manner. As shown in Fig. 1, our framework comprises three collaborative agents: **Partition Agent**. We parse RTL into abstract syntax trees (ASTs) and reason at the behavioral level. A library of high-

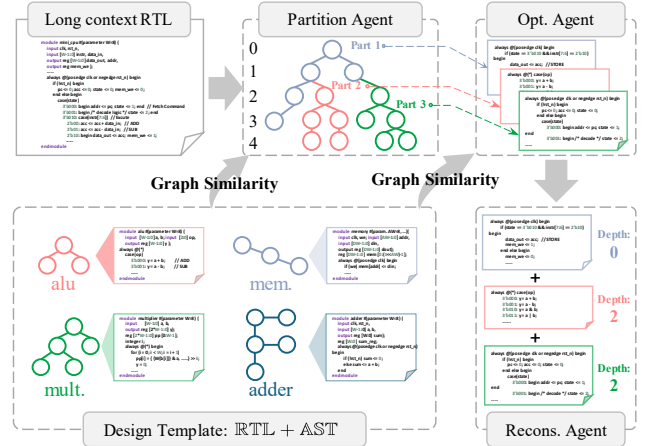


Fig. 1 Overall flow of our work.

quality design templates is built, each paired with its RTL text and AST graph. Given a long RTL input, we compute AST-level graph similarity using a trained GNN, and apply a tree-based dynamic programming algorithm to partition the global AST into K connected subtrees that best match known templates. These subtrees represent reusable design patterns and enable interpretable decomposition—mimicking how engineers reason about large RTL designs.

Optimization Agent. Each partitioned subtree is mapped back to RTL and optimized by an LLM agent equipped with a multi-modal Retrieval-Augmented Generation (RAG) mechanism. Instead of relying solely on RTL text, our framework jointly leverages both RTL code and AST graph structure during retrieval. For each subgraph, we retrieve a small set of RTL+AST template exemplars with similar functionality and structural hierarchy. This targeted retrieval narrows the search space and provides directionally accurate optimization. We further explore multiple rewrites using Monte Carlo Tree Search (MCTS), and select optimized RTL submodules with superior design performance metrics (e.g., delay, area, power).

Reconstruction Agent. Optimized submodules are reassembled into a functionally equivalent design using a logic-aware hierarchical strategy. Crucially, we introduce a graph RAG prompt mechanism that encodes the structural relationship between each partitioned subtree and the original full AST. This enables the LLM to reason about interconnections and avoid logic or signal mismatches. This reconstruction respects logic depth and control hierarchy, ensuring function equivalence and smooth end-to-end synthesis.

Overall, we present the first end-to-end framework that enables practical LLM-based RTL optimization for long-context RTL designs. By combining AST graph partition,

[†]Corresponding author.

This work is partially supported by The Research Grants Council of Hong Kong SAR (No. CUHK14211824 and No. CUHK14210723).

targeted multimodal retrieval, and logic-guided reconstruction, our method significantly extends the capability of LLMs beyond toy examples —paving the way for real-world deployment in RTL code optimization. Our contributions are summarized as follows:

- We propose a graph similarity-guided RTL partitioning approach that segments large-scale designs into semantically meaningful subtrees.
- We develop a multi-modal RAG system that combines RTL text and AST graph as multimodal guidance, enabling structurally grounded and functionally equivalent RTL submodule generation and optimization based on partitioned subtrees.
- We introduce a logic-aware reconstruction strategy that integrates optimized submodules into the RTL design while preserving function equivalence, using a graph-based prompt to connect local and global structure.
- We validate our framework on long-context benchmarks, achieving 100% functional equivalence and significant PPA improvements.

II. PRELIMINARIES

Abstract Syntax Tree (AST). RTL designs can be naturally represented as Abstract Syntax Trees (ASTs), where each node denotes a code construct (e.g., module, always block, if-statement), and edges reflect parent-child syntax relationships. Compared to netlists or control/data-flow graphs, ASTs preserve syntactic structure and semantic hierarchy, making them more suitable for RTL code understanding and transformation—especially when paired with LLMs.

Graph Similarity. Graph similarity measures how structurally and semantically aligned two graphs are. Given two AST graphs \mathbb{G}_1 and \mathbb{G}_2 , a similarity score $S(\mathbb{G}_1, \mathbb{G}_2) \in [0, 1]$ is computed based on topological matching, label consistency, and local subgraph overlap. In our work, graph similarity enables us to match AST subtrees of long context RTL with known functional patterns (e.g., ALUs, FSMs).

III. METHOD

A. Design Templates

To support graph-guided partitioning, retrieval-based optimization, and logic-aware reconstruction, we construct a reusable design template at both RTL and AST levels.

RTL Design Templates: We collect a diverse set of behavioral RTL designs covering common circuit functions: arithmetic units (adder, subtractor, multiplier), data operations (left/right shifter, barrel shifter, bitwise logic), control structures (multiplexer, comparator, priority encoder, decoder), and storage elements (latch, register, counter). Each RTL template is written in synthesizable Verilog and encapsulates a reusable functional unit.

AST Graph Templates: Each RTL template is parsed into an Abstract Syntax Tree (AST), encoded as a graph $\mathbb{G} = (\mathbf{A}, \mathbf{F})$. Here, \mathbf{A} denotes the adjacency matrix encoding the parent-child structure of the AST, and \mathbf{F} is the node feature matrix enables graph similarity computation. Nodes in \mathbb{G} are categorized as: (1) Operation nodes: e.g., Assign, BinOp, If, Cond, Always, etc., capturing control and logic structures; (2) Port nodes: e.g., input, output, wire, reg, inout, representing interface and connectivity.




Fig. 2 Training of AST graph similarity computation.

Each node is annotated with a one-hot encoded feature vector $\mathbf{f}_v \in \mathbf{F}$ based on its type. We develop an AST-based template library $\mathcal{G} = \{\mathbb{G}_1, \dots, \mathbb{G}_N\}$, extracted from RTL design exemplars, to support downstream tasks.

B. AST Graph Similarity Computation

As illustrated in Fig. 2, we measure the graph similarity $S(\mathbb{G}_i, \mathbb{G}_j)$ between two AST graphs \mathbb{G}_i and \mathbb{G}_j via three stages: (1) node-level similarity, (2) graph-level similarity, and (3) contrastive training.

Node-level similarity. In AST graphs, each AST node feature \mathbf{f}_v is encoded using a Graph Convolutional Network (GCN) to propagate local structural information and generate node encoding results \mathbf{h}_v :

$$\mathbf{h}_v = \text{GCN}(\mathbf{f}_v). \quad (1)$$

To compare ASTs at a fine-grained level, we define their node-level similarity as the average cosine similarity over all node pairs based on node encoding results:

$$S_{\text{node}}(\mathbb{G}_i, \mathbb{G}_j) = \frac{1}{|\mathcal{V}_i| \cdot |\mathcal{V}_j|} \sum_{u \in \mathcal{V}_i} \sum_{v \in \mathcal{V}_j} \cos(\mathbf{h}_u, \mathbf{h}_v), \quad (2)$$

where \mathcal{V}_i and \mathcal{V}_j are node sets of AST graphs \mathbb{G}_i and \mathbb{G}_j . \mathbf{h}_u and \mathbf{h}_v are the final node encoding results in \mathbb{G}_i and \mathbb{G}_j after graph learning, respectively. $\cos(\mathbf{h}_u, \mathbf{h}_v)$ denotes the cosine similarity between the two node vectors.

Graph-level similarity. To capture global semantics, we apply attention-based pooling over node encoding results to generate graph embeddings $\mathbf{h}_{\mathbb{G}}$:

$$\alpha_v = \text{sigmoid}(\mathbf{w}^\top \mathbf{h}_v), \quad \mathbf{h}_{\mathbb{G}} = \sum_{v \in \mathcal{V}} \alpha_v \mathbf{h}_v, \quad (3)$$

where \mathbf{w} is a learnable vector. The graph-level similarity between two AST graphs is computed via:

$$S_{\text{graph}}(\mathbb{G}_i, \mathbb{G}_j) = \cos(\mathbf{h}_{\mathbb{G}_i}, \mathbf{h}_{\mathbb{G}_j}), \quad (4)$$

$\cos(\mathbf{h}_{\mathbb{G}_i}, \mathbf{h}_{\mathbb{G}_j})$ denotes the cosine similarity between two graph embeddings. Then, the graph similarity of \mathbb{G}_i and \mathbb{G}_j is computed via:

$$S(\mathbb{G}_i, \mathbb{G}_j) = S_{\text{node}}(\mathbb{G}_i, \mathbb{G}_j) + S_{\text{graph}}(\mathbb{G}_i, \mathbb{G}_j). \quad (5)$$

Contrastive learning. To enforce functional consistency across structurally diverse ASTs, we adopt contrastive learning over the graph learning training. We define a positive pair $(\mathbb{G}_i, \mathbb{G}_j^+)$ if both ASTs correspond to the same design functionality, and define negative pairs $(\mathbb{G}_i, \mathbb{G}_j^-)$ if they belong to different RTL design types.

Based on node-level and graph-level graph similarity, we optimize the following InfoNCE-style [9] contrastive loss:

$$\mathcal{L}_{\text{CL}} = -\log \frac{\exp(S(\mathbb{G}_i, \mathbb{G}_j^+)/\tau)}{\exp(S(\mathbb{G}_i, \mathbb{G}_j^+)/\tau) + \sum_{j=1}^{N^-} \exp(S(\mathbb{G}_i, \mathbb{G}_j^-)/\tau)}, \quad (6)$$

where $S(\mathbb{G}_i, \mathbb{G}_j)$ is the graph similarity between two graphs computed via Equation (5). N_- is the number of negative samples. τ is a temperature hyperparameter (set to 0.1).

This training objective pulls together functionally equivalent AST graphs and pushes apart those with different semantics. Through this contrastive supervision, the learned AST embeddings reflect both structural features and functional alignment. It enhances the reliability of similarity-based matching in downstream partitioning and optimization.

C. Partition Agent

To support scalable optimization of long-context RTL designs, we propose a Tree-based Dynamic Programming (Tree-DP) algorithm that hierarchically partitions a large AST graph \mathbb{G}_{long} into K disjoint, connected subtrees $\{\mathbb{G}_{s_1}, \dots, \mathbb{G}_{s_K}\}$. Tree-DP performs graph similarity-guided partitioning, where each subtree is matched against the reusable design template library \mathcal{G} (described in Section III-A) to ensure functional alignment and structural coherence. It ensures that the resulting subtrees resemble known RTL modules (e.g., adders, shifters, FSMs). This encourages each subtree to exhibit well-formed functional intent, making it more interpretable for LLMs and easier to optimize, regenerate, or replace in downstream tasks. After partitioning, each subtree serves as an LLM-optimizable unit and is passed to the downstream optimization agent. The complete procedure is detailed in Algorithm 1.

Objective. The partitioning process is guided by an AST graph similarity function $S(\cdot, \cdot)$ described in Section III-B. For each candidate subtree \mathbb{G}_{s_k} extracted from \mathbb{G}_{long} , we identify its most similar design pattern from the template library \mathcal{G} , and define the template relevance score as the maximum graph similarity. Our global objective is thus to find a partition configuration that maximizes total similarity:

$$\max_{\{\mathbb{G}_{s_1}, \dots, \mathbb{G}_{s_K}\}} \sum_{k=1}^K \max_{\mathbb{G}_{\text{temp}} \in \mathcal{G}} S(\mathbb{G}_{s_k}, \mathbb{G}_{\text{temp}}). \quad (7)$$

Constraints. To ensure both validity and downstream usability, we impose the following constraints on the resulting partitions: (i) Each subtree \mathbb{G}_{s_k} must be connected and correspond to a syntactically and semantically valid RTL fragment; (ii) The union of all subtrees must exactly cover the full AST: $\bigcup_{k=1}^K \mathbb{G}_{s_k} = \mathbb{G}_{\text{long}}$, with $\mathbb{G}_{s_i} \cap \mathbb{G}_{s_j} = \emptyset$ for all $i \neq j$; (iii) The number of partitions K is either user-specified or adaptively determined based on the design scale and optimization budget.

Workflow. Exploiting the hierarchical structure of ASTs, Tree-DP performs partitioning via four key stages: (1) **Traversal:** A bottom-up traversal is performed over the AST \mathbb{G}_{long} , visiting each node v and recursively computing the best partition scores for subtrees rooted at v under varying partition budgets. (2) **Merge & Score:** At each internal node, we enumerate valid combinations of child subtrees and attempt to merge them into higher-level candidates. Each merged subtree \mathbb{G}_{s_k} is compared against the design template library \mathcal{G} based on graph similarity. (3) **Update:** Among all merge candidates, only the one with the highest cumulative similarity is preserved in the DP table for each partition size k during the bottom-up traversal from leaves to node v . (4) **Backtrack:** After the DP table is populated, Tree-DP traces back from the root of \mathbb{G}_{long} to recover the partition

Algorithm 1: Tree-DP Partitioning for Large AST

Input: Long RTL AST graph \mathbb{G}_{long} , template library \mathcal{G} , target partition count K
Output: Partitioned subtrees $\{\mathbb{G}_{s_1}, \dots, \mathbb{G}_{s_K}\}$

```

1 function BESTMATCHTEMPLATE( $\mathbb{G}_{\text{cand}}, \mathcal{G}$ ):
2   return  $\max_{\mathbb{G}_{\text{temp}} \in \mathcal{G}} S(\mathbb{G}_{\text{cand}}, \mathbb{G}_{\text{temp}})$ ;
3
4 // Graph similarity-driven searching
5 function TREEDP( $v$ ):
6   foreach child  $c$  of  $v$  do
7      $\text{Traverse}$  from leaves to root TREEDP( $c$ )
8
9   for  $k = 1$  to  $K$  do
10    foreach valid  $k$ -way partition of children( $v$ ) do
11      Merge children into  $\mathbb{G}_{\text{cand}}$ ;
12       $s \leftarrow \text{BESTMATCHTEMPLATE}(\mathbb{G}_{\text{cand}}, \mathcal{G})$ ;
13       $S_{\text{total}} \leftarrow s + \sum_{j=1}^k \text{DP}[c_j][k_j]$ ;
14      if  $S_{\text{total}} > \text{DP}[v][k]$  then
15        Update best config  $\text{DP}[v][k] \leftarrow S_{\text{total}}$ ;
16
17 return  $\text{DP}[v]$ 
18
19 Main execution:
20 Initialize  $\text{DP}[v][k] = -\infty$  for all  $v, k$ ;
21 TREEDP(root of  $\mathbb{G}_{\text{long}}$ );
22 Backtrack from  $\text{DP}[\text{root}][K]$  to get  $\{\mathbb{G}_{s_1}, \dots, \mathbb{G}_{s_K}\}$ ;

```

plan $\{\mathbb{G}_{s_1}, \dots, \mathbb{G}_{s_K}\}$, ensuring that the selected subtrees maximize similarity to known design patterns in \mathcal{G} .

D. Optimization Agent

To enable structure-preserving and high-quality RTL optimization of partitioned AST subtrees, we develop an optimization agent based on AST+RTL multimodal retrieval-augmented large language models (LLMs). This agent follows a structured four-stage pipeline: Search, Retrieve, Generate, and Refine, as illustrated in Fig. 3. For each AST subtree \mathbb{G}_{s_i} obtained from the partition agent, the optimization agent searches for AST structurally aligned RTL exemplars, explores diverse rewrite candidates through Monte Carlo Tree Search (MCTS), and iteratively refines outputs based on functional equivalence feedback. Our method achieves faithful optimization of RTL submodules, supporting scalable long-context RTL design optimization.

Search. Given a partitioned AST subgraph \mathbb{G}_{s_i} , we first encode it into a graph embedding using a trained GNN for computing AST-level similarity (Section III-B). We then perform a graph similarity search over the AST template database to retrieve the top- N structurally similar AST designs. This step constructs a high-quality candidate set that reflects similar control and computation patterns, forming a strong foundation for downstream RTL generation.

Retrieve. Each retrieved AST \mathbb{G}_j is paired with its corresponding RTL implementation RTL_j to form a multi-modal exemplar set $\{(\mathbb{G}_j, \text{RTL}_j)\}$. These exemplars—comprising both structural and semantic information—are used to prompt the LLM in a retrieval-augmented generation (RAG) setting. Unlike static prompt engineering, our retrieval dynamically adapts to the structure of \mathbb{G}_{s_i} , enabling query-specific generation that respects the design intent.

Generate. Using the exemplar pairs as prompt, the LLM generates candidate rewrites RTL'_{s_i} for the target partitioned subtree \mathbb{G}_{s_i} . To systematically explore diverse rewrites, we employ Monte Carlo Tree Search (MCTS), where each node in the search tree corresponds to a generated RTL variant. Each expansion samples different prompt configurations

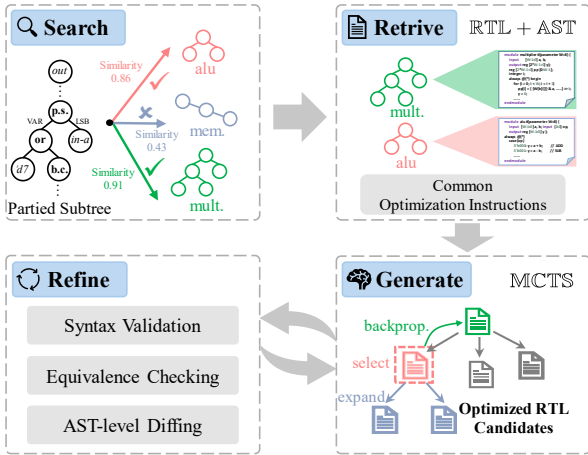


Fig. 3 The optimization agent pipeline.

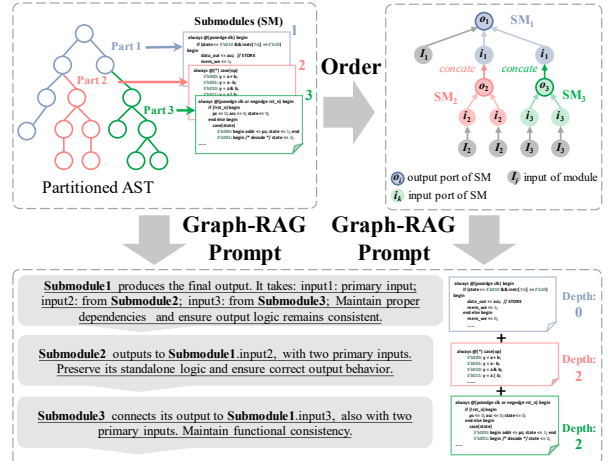


Fig. 4 Graph-RAG prompting via logic-aware ordering.

(e.g., reordering exemplars), and each candidate is evaluated using a hybrid score:

$$\text{Score}(\text{RTL}'_i) = \lambda \cdot S(\mathbb{G}_i, \mathbb{G}'_i) + (1 - \lambda) \cdot \text{PPA}(\text{RTL}'_{s_i}), \quad (8)$$

where $S(\mathbb{G}_i, \mathbb{G}'_i)$ measures the graph-level structural similarity between the original and generated ASTs, and $\text{PPA}(\text{RTL}'_{s_i})$ denotes a performance-aware proxy metric reflecting the RTL's quality. The PPA metric is defined as the normalized average of delay, area, and power of generated RTL. It can be reweighted by the designer to prioritize specific objectives, such as low power or high speed.

Refine. Once MCTS completes—either by rollout limit or convergence, we select the top-scoring candidate as the optimized RTL. The selected RTL undergoes a multi-stage verification pipeline: (1) Syntax validation via Verilog compiler checks; (2) Functional simulation under benchmark inputs to ensure functional equivalence; (3) AST-level diffing to capture unintended logic drift. If any issues are detected, the agent reuses the MCTS procedure with refined prompts until a valid and performant RTL variant is found.

E. Reconstruction Agent

After modular optimization, the reconstruction agent reassembles optimized RTL submodules RTL_{s_i} of partitioned AST subtrees into a functional equivalent design. To support this process, we employ a logic-aware ordering strategy and a Graph-RAG prompting framework, enabling the LLM to generate the functional equivalent top-level RTL with accurate connectivity and control flow, as shown in Fig. 4.

Order. Each optimized RTL submodule RTL_{s_i} corresponds to a partitioned AST subtree \mathbb{G}_{s_i} . To preserve functional semantics and signal dependencies, we sort the submodules based on their logical depth and connectivity within the large AST \mathbb{G}_{long} of the original long context RTL. This facilitates correct instantiation and signal propagation.

Graph-RAG Prompting. We construct a hierarchical Graph-RAG prompt that jointly encodes structural, semantic, and hierarchical information. Specifically, the prompt includes: (i) the set of partitioned AST subtrees \mathbb{G}_{s_i} and their corresponding RTLs RTL_{s_i} ; (ii) the full large AST \mathbb{G}_{long} to retain global module hierarchy and I/O interconnection patterns; and (iii) the original long context RTL for semantic grounding. This fusion enables the LLM to reason over hierarchical control flow, structural alignment, and naming

consistency, producing a `top.v`-style RTL that integrates all submodules coherently.

Validate. We apply the same validation pipeline used in the optimization agent. The generated RTL is checked for syntax errors using Verilog compilers, and functionally verified via equivalence checking. We also perform AST-level diffing to detect structural drift. If inconsistencies are found, the model is re-prompted with updated context to regenerate the erroneous segments. This iterative refinement ensures that the final top-level RTL design is functionally equivalent and structurally faithful to the original RTL design.

IV. EXPERIMENTAL RESULTS

To evaluate our proposed framework, we implement a complete RTL optimization pipeline. The main components are as follows: For AST parsing and preprocessing, we use `PyVerilog` [10] to parse input RTL code into abstract syntax trees (ASTs). For the partition agent, we implement a Tree-based Dynamic Programming (Tree-DP) algorithm in Python, supporting configurable partition budgets and similarity thresholds. Depending on design scale, the partition count K is adaptively chosen from 5, 10, 15, 20. For the optimization and reconstruction agents, we build both upon `ChatGPT 4o` [11] via OpenAI API. The optimization follows the four-stage RAG-based workflow described in Section III-D, while reconstruction leverages a graph-RAG flow with logic-depth-aware prompting, as in Section III-E. The graph similarity module uses a trained Graph Convolutional Network (GCN) [12] to guide both retrieval and partitioning. For PPA Evaluation, all RTL designs are synthesized using Synopsys Design Compiler [13] under the ASAP 7nm technology node [14]. We report post-synthesis delay, area, and power, and compute a normalized PPA score as a weighted average. The weights are user-configurable to support various optimization priorities (e.g., timing or power). For verification, we adopt a multi-stage correctness check to ensure functional and structural equivalence. First, syntax is validated via `iVerilog` [15]. Then, we perform simulation-based functional checking with Python-generated testbenches. Additionally, we employ: (i) `ABC` [16] for SAT-based combinational equivalence checking, and (ii) `egg` [17] for efficient arithmetic reasoning.

TABLE I Benchmark characteristics of single-module long-context RTL designs. Delay (ps), Area (μm^2), and Power (mW) are reported by Synopsys Design Compiler.

Des.	Lines	Nodes	Wires	Delay	Area	Power
add64	195	917	1306	47.6	173.9	39.9
mult32	88	2119	2873	617.2	230.2	194.3
comput	204	1520	2265	952.5	141.6	116.8
traffic	90	66	69	9.6	59.7	1.4
alu	111	909	1110	1555.6	79.1	49.3
radix	300	300	439	269.7	29.9	6.7
asyn	147	560	601	122.7	85.2	13.5
accu	77	66	111	61.2	9.8	3
fpu_pre	235	663	881	40.3	65.0	46.9
fpu_post	638	1692	2119	3009.4	138.5	102.5
mc_sel	201	380	464	299.1	49.9	28.6
mc_rf	166	193	256	126.9	32.8	6.6
buffer	268	774	847	54.2	103.7	24.2
eth	325	527	739	55.3	87.9	12.7
AVE.	217	763	1005	515.8	92.0	46.2

TABLE II Benchmark characteristics of multi-module long-context RTL designs.

Des.	Lines	Nodes	Wires	Delay	Area	Power
ac97	2030	8214	8243	89.9	1173.2	30.2
aes	2195	7043	7953	105.7	922.9	14.6
crca	1034	3812	3855	105.3	340.5	3.1
ecg	1525	66872	67385	305.6	6636.4	41.1
AVE.	1696	21485	21859	151.6	2268.3	22.2

A. Benchmark and Baseline

Design Templates: We construct a high-quality template library from the RTLRewriter suite, curated with experienced Verilog engineers. They are based on industrial design practices and over 50 scholarly works on RTL transformation. Each template includes:

- **RTL code** (RTL): optimized Verilog with favorable PPA (synthesized via Synopsys Design Compiler [13]);
- **AST graph** (G): JSON-serialized syntax tree capturing structural and logical semantics.

The library spans 55 templates across 21 RTL functions, including arithmetic datapaths, memory controllers, finite-state machines (FSMs), MUX-based control.

Long Context Designs: There are two types of large RTLs:

- **Single-Module RTL:** Flat, monolithic designs without module boundaries. They pose major challenges for LLMs due to context overflow and entangled logic.
- **Multi-Module RTL:** Modular SoC-style designs with declared interfaces, yet still exceed context limits and contain non-trivial inter-module dependencies.

TABLE I and TABLE II summarize their scale and PPA metrics, reflecting realistic industry-grade complexity.

Baselines: We compare against five strong baselines:

- **General LLMs:** GPT-4o and Gemini, representing state-of-the-art LLMs without RTL-specific tuning.
- **RTL-Specific LLMs:** RTLCoder and RTLRewriter, both designed for Verilog generation and optimization.
- **Compiler-Based Optimizers:** Yosys+Egg, an open-source compiler, albeit inferior to Synopsys Design Compiler [13].

All LLM-based methods use the same prompt format and context constraints for fair comparison.

TABLE III Functional equivalence pass rate results.

Des.	Y+E	GPT4o	Gemini	Coder	Rewriter	Ours
add64	✓	×	✓	×	✓	✓
mult32	✓	✓	✓	✓	✓	✓
comput	✓	×	×	×	×	✓
traffic	✓	✓	✓	✓	✓	✓
alu	✓	✓	✓	✓	✓	✓
radix	✓	×	×	×	×	✓
asyn	✓	✓	×	✓	✓	✓
accu	✓	✓	✓	✓	✓	✓
fpu_pre	✓	×	×	×	×	✓
fpu_post	✓	×	×	×	×	✓
mc_sel	✓	×	×	×	×	✓
mc_rf	✓	✓	×	✓	✓	✓
buffer	✓	×	×	×	×	✓
eth	✓	×	×	×	×	✓
AVE.	100%	42.9%	35.7%	42.9%	50.0%	100%

B. Performance

Performance on Single-Module Long Context RTL.

This setting poses the core challenge for LLM-based RTL optimization: large monolithic designs (often >200 lines) without clear modular boundaries, causing severe context overflow and logic entanglement. As shown in TABLE III, our method achieves 100% functional equivalence across all test cases, thanks to AST-based partitioning guided by design templates. The reconstruction agent preserves signal consistency through logic-depth-aware ordering and graph-RAG prompting. In contrast, prior methods succeed on less than 50% of the cases, and completely fail when RTL line counts exceed 200—limiting their practical utility on real-world designs. As shown in TABLE IV, our method not only ensures full correctness but also achieves about 25% average PPA improvement over original RTLs, surpassing all baselines in both quality and reliability.

Performance on Multi-Module Long Context RTL.

We further validate our framework on SoC-style RTLs with explicit modules but long overall context and inter-module dependencies. Here, our retrieval-then-optimize strategy remains effective: AST similarity-based retrieval yields structurally relevant exemplars, and MCTS-guided rewriting selects candidates with better PPA. According to TABLE V, our method achieves up to 5% PPA improvement over RTLRewriter, the strongest existing baseline, and about 15% average improvement compared to original RTLs.

Runtime Efficiency. As shown in Fig. 5, we report average runtime performance across both benchmarks. The end-to-end average runtime is approximately 360s for single-module cases and 700s for multi-module ones. Notably, LLM inference is relatively efficient; the majority of the runtime overhead comes from logic synthesis (to evaluate PPA using Synopsys Design Compiler [13] and simulation-based verification (to ensure functional equivalence using ABC [16])). Our modular design enables parallel execution across submodules, allowing the framework to scale efficiently for industrial RTL use cases.

C. Ablation Studies


To evaluate each component’s contribution, we conduct ablations on three aspects: (1) graph similarity-driven partitioning, (2) multimodal retrieval-augmented optimization, and (3) logic-aware reconstruction. As shown in Fig. 6, results are measured by functional equivalence pass rate and normalized PPA metrics.

TABLE IV Comparisons of baseline approaches on single-module long context RTL benchmarks. ‘‘RAT.’’ represents the ratio of optimized RTL to original RTL on PPA metrics (Delay (ps), Area (μm^2), Power (mW)), as synthesized and reported by Synopsys Design Compiler. Because Yosys [18] + Egg [17] performs no substantial logic optimization, the synthesized RTL shows no PPA improvement.


Des.	Yosys [18] + Egg [17]			GPT4o [11]			Gemini [19]			RTLcoder [1]			RTLrewriter [4]			Ours		
	Delay	Area	Power	Delay	Area	Power	Delay	Area	Power	Delay	Area	Power	Delay	Area	Power	Delay	Area	Power
add64	45.3	192.8	41.3	/	/	/	47.5	162.2	39.4	/	/	/	41.8	147.3	34.7	39.2	133.3	32.2
mult32	669.7	259.0	220.1	547.2	194.6	164.7	552.8	198.0	172.5	567.6	216.2	183.2	499.7	178.0	149.7	463.9	158.8	145.9
comput	1047.0	157.6	129.7	/	/	/	/	/	/	/	/	/	/	/	/	740.6	110.9	84.1
traffic	10.2	66.7	1.6	8.5	53.3	1.3	9.0	53.3	1.3	8.8	56.0	1.3	7.8	47.7	1.2	7.4	44.7	1.1
alu	1753.6	84.1	52.5	1413.1	66.6	43.2	1462.9	67.0	44.9	1450.7	74.0	46.6	1320.1	61.0	39.8	1122.9	56.1	39.6
radix	285.0	33.6	7.4	/	/	/	/	/	/	/	/	/	/	/	/	189.2	22.6	5.0
asyn	135.2	97.5	14.4	108.9	70.9	12.2	/	/	/	112.7	78.9	12.6	100.5	66.2	11.4	93.7	61.1	10.5
accu	66.5	10.8	3.3	56.9	8.7	2.6	56.0	9.0	2.7	56.0	9.0	2.7	51.3	8.2	2.3	47.4	7.5	2.2
fpu_pre	44.6	70.4	51.4	/	/	/	/	/	/	/	/	/	/	/	/	31.1	42.6	33.4
fpu_post	3174.5	149.1	115.0	/	/	/	/	/	/	/	/	/	/	/	/	2104.2	104.8	80.8
mc_sel	334.6	57.2	32.6	/	/	/	/	/	/	/	/	/	/	/	/	231.7	35.2	21.3
mc_rf	139.6	34.7	7.5	106.4	28.1	5.4	/	/	/	118.4	30.8	6.3	94.2	25.6	5.0	89.0	24.4	4.7
buffer	58.5	117.6	27.7	/	/	/	/	/	/	/	/	/	/	/	/	42.5	76.7	17.1
eth	60.7	96.6	13.5	/	/	/	/	/	/	/	/	/	/	/	/	40.5	65.3	10.1
AVE.	558.9	102.0	51.3	/	/	/	/	/	/	/	/	/	/	/	/	373.1	67.2	35.1
RAT.	1.08	1.11	1.10	0.89	0.86	0.87	0.94	0.89	0.93	0.92	0.93	0.94	0.82	0.80	0.81	0.72	0.73	0.76

TABLE V Comparisons of baseline approaches on multi-module long context RTL benchmarks.

Des.	Yosys [18] + Egg [17]			GPT4o [11]			Gemini [19]			RTLcoder [1]			RTLrewriter [4]			Ours		
	Delay	Area	Power	Delay	Area	Power	Delay	Area	Power	Delay	Area	Power	Delay	Area	Power	Delay	Area	Power
ac97	107.5	1311.3	33.8	85.1	1124.8	27.7	80.5	1040.5	28.1	82.7	1123.4	28.1	77.9	1078.9	26.6	72.2	1027.5	25.0
aes	122.3	999.0	16.0	97.3	881.4	13.4	93.7	819.7	13.5	97.9	885.4	13.8	88.0	847.4	12.8	82.5	807.1	12.3
crc4	124.8	376.4	3.4	97.6	325.6	2.8	92.9	307.6	2.9	97.5	329.2	2.9	86.9	308.6	2.6	83.6	299.6	2.6
ecg	355.2	7435.7	46.8	286.9	6362.7	37.6	269.1	5988.9	38.9	276.9	6377.3	38.8	260.1	6055.0	36.7	246.7	5826.2	33.3
AVE.	177.5	2530.6	25.0	141.7	2173.6	20.4	134.0	2039.2	20.8	138.8	2178.8	20.9	128.2	2072.5	19.7	121.3	1990.1	18.3
RAT.	1.17	1.12	1.13	0.93	0.96	0.92	0.88	0.90	0.94	0.92	0.96	0.94	0.85	0.91	0.88	0.80	0.88	0.82



(a) On Single-Module RTLs.



(b) On Multi-Module RTLs.

Fig. 5 Average runtime breakdown of our work.

(1) **W/o Partition: Random Subtree Splitting.** Instead of using Tree-DP with template-guided graph similarity, we randomly split the long-context AST \mathbb{G}_{long} into fixed-size subgraphs. Despite using the same optimization agent, the functional pass rate drops by 44.4%, and delay/area increase due to poorly defined module boundaries.

(2) **W/o Optimization: No Multimodal Retrieval Guidance.** We remove RTL+AST-based retrieval and directly prompt the LLM with common RTL optimization instructions like RTLrewriter [4]. Lacking structural grounding, the generated RTL becomes generic and misaligned, leading to noticeable PPA degradation.

(3) **W/o Reconstruction: Naive Concatenation.** We skip logic-aware ordering and graph prompting during reassembly, instead concatenating submodules. This results in broken control/data paths and I/O mismatches, reducing functional correctness by 11.1% and impacting downstream PPA.

V. CONCLUSION

This work presents a complete framework for long-context RTL optimization, comprising three collaborative agents: (1) a graph similarity-guided partition agent that decomposes large ASTs into semantically meaningful subtrees; (2) a

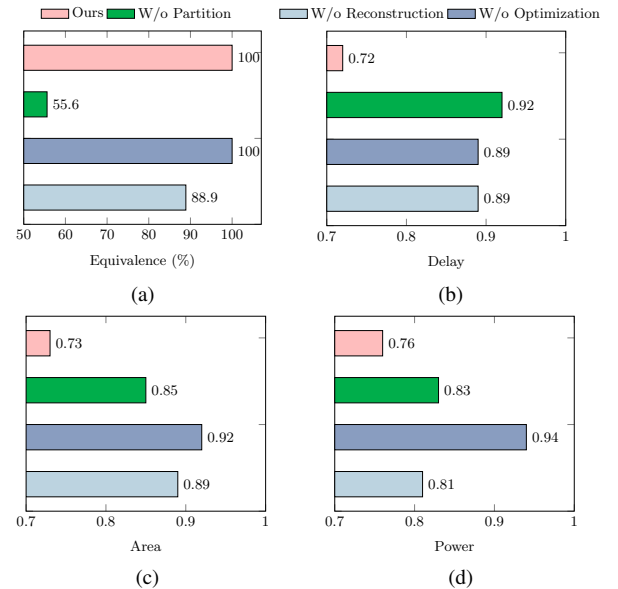


Fig. 6 Comparison among different schemes by (a) functional equivalence pass rate (%), (b) delay, (c) area, and (d) power. All these values are normalized by original RTL design results reported via Synopsys Design Compiler.

multimodal optimization agent that jointly leverages AST structure and RTL semantics for retrieval-augmented generation with MCTS-based candidate selection; (3) a logic-aware reconstruction agent that ensures functional equivalence via graph-RAG prompting and logic-aware reassembly. Experiments on challenging RTL designs confirm 100% functional equivalence and consistent PPA improvements over state-of-the-art baselines, demonstrating the framework’s scalability and real-world applicability.

REFERENCES

- [1] S. Liu, W. Fang, Y. Lu, J. Wang, Q. Zhang, H. Zhang, and Z. Xie, "Rtlcoder: Fully open-source and efficient llm-assisted rtl code generation technique," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 2024.
- [2] S. Thakur, B. Ahmad, H. Pearce, B. Tan, B. Dolan-Gavitt, R. Karri, and S. Garg, "Verigen: A large language model for verilog code generation," *ACM Transactions on Design Automation of Electronic Systems*, vol. 29, no. 3, pp. 1–31, 2024.
- [3] P. Zehua, H. Zhen, M. Yuan, Y. Huang, and B. Yu, "Betterv: Controlled verilog generation with discriminative guidance," in *Forty-first International Conference on Machine Learning*, 2024.
- [4] X. Yao, Y. Wang, X. Li, Y. Lian, R. Chen, L. Chen, M. Yuan, H. Xu, and B. Yu, "Rtlrewriter: Methodologies for large models aided rtl code optimization," in *Proceedings of the 43rd IEEE/ACM International Conference on Computer-Aided Design*, 2024, pp. 1–7.
- [5] Y. Wang, W. Ye, P. Guo, Y. He, Z. Wang, B. Tian, S. He, G. Sun, Z. Shen, S. Chen *et al.*, "Symrtlo: Enhancing rtl code optimization with llms and neuron-inspired symbolic reasoning," *arXiv preprint arXiv:2504.10369*, 2025.
- [6] Y. Hu, J. Ye, K. Xu, J. Sun, S. Zhang, X. Jiao, D. Pan, J. Zhou, N. Wang, W. Shan *et al.*, "Uvllm: An automated universal rtl verification framework using llms," *arXiv preprint arXiv:2411.16238*, 2024.
- [7] X. Wang, G.-W. Wan, S.-Z. Wong, L. Zhang, T. Liu, Q. Tian, and J. Ye, "Chatcpu: An agile cpu design and verification platform with llm," in *Proceedings of the 61st ACM/IEEE Design Automation Conference*, 2024, pp. 1–6.
- [8] C. Deng, Y.-D. Tsai, G.-T. Liu, Z. Yu, and H. Ren, "Scalertl: Scaling llms with reasoning data and test-time compute for accurate rtl code generation," *arXiv preprint arXiv:2506.05566*, 2025.
- [9] A. v. d. Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," *arXiv preprint arXiv:1807.03748*, 2018.
- [10] S. Takamaeda-Yamazaki, "Pyverilog: A python-based hardware design processing toolkit for verilog hdl," in *International Symposium on Applied Reconfigurable Computing*. Springer, 2015, pp. 451–460.
- [11] A. Hurst, A. Lerer, A. P. Goucher, A. Perelman, A. Ramesh, A. Clark, A. Ostrow, A. Welihinda, A. Hayes, A. Radford *et al.*, "Gpt-4o system card," *arXiv preprint arXiv:2410.21276*, 2024.
- [12] T. Kipf, "Semi-supervised classification with graph convolutional networks," *arXiv preprint arXiv:1609.02907*, 2016.
- [13] Synopsys, "Design Compiler User Guide," <https://www.synopsys.com/implementation-and-signoff/rtl-synthesis-test/dc-ultra.html>, 2024.
- [14] L. T. Clark, V. Vashishtha, L. Shifren, A. Gujja, S. Sinha, B. Cline, C. Ramamurthy, and G. Yeric, "Asap7: A 7-nm finfet predictive process design kit," *Microelectronics Journal*, vol. 53, pp. 105–115, 2016.
- [15] S. Williams and M. Baxter, "Icarus verilog: open-source verilog more than a year later," *Linux Journal*, vol. 2002, no. 99, p. 3, 2002.
- [16] R. Brayton and A. Mishchenko, "Abc: An academic industrial-strength verification tool," in *International Conference on Computer Aided Verification*. Springer, 2010, pp. 24–40.
- [17] S. Coward, G. A. Constantinides, and T. Drane, "Automatic datapath optimization using e-graphs," in *2022 IEEE 29th Symposium on Computer Arithmetic (ARITH)*. IEEE, 2022, pp. 43–50.
- [18] C. Wolf, J. Glaser, and J. Kepler, "Yosys-a free verilog synthesis suite," in *Proceedings of the 21st Austrian Workshop on Microelectronics (Austrochip)*, vol. 97, 2013.
- [19] Google, "Gemini," <https://gemini.google.com>, 2025.