

# Think with Self-Decoupling and Self-Verification: Automated RTL Design with Backtrack-ToT

Zhiteng Chao<sup>1</sup>, Yonghao Wang<sup>1†</sup>, Xinyu Zhang<sup>1,2</sup>, Jiaxin Zhou<sup>3</sup>, Tenghui Hua<sup>1,2</sup>, Husheng Han<sup>1</sup>, Tianmeng Yang<sup>4</sup>, Jianan Mu<sup>1†</sup>, Bei Yu<sup>5</sup>, Rui Zhang<sup>1</sup>, Jing Ye<sup>1,2</sup> and Huawei Li<sup>1,2</sup>

<sup>1</sup>State Key Lab of Processors, Institute of Computing Technology, CAS, Beijing, China

<sup>2</sup>University of Chinese Academy of Sciences, Beijing, China

<sup>3</sup>Beijing Normal University, China; <sup>4</sup>Peking University, China; <sup>5</sup>Chinese University of Hong Kong, China

{chaozhiteng, wangyonghao22s, zhangxinyu23s}@ict.ac.cn, zhoujiaxinmail.bnu.edu.cn,

{huatenghui24s, hanhusheng}@ict.ac.cn, youngtimmy@pku.edu.cn, mujianan@ict.ac.cn,

byucse.cuhk.edu.hk, {zhangrui, yejing, lihuawei}@ict.ac.cn

**Abstract**—Large language models (LLMs) hold promise for automating integrated circuit (IC) engineering using register transfer level (RTL) hardware description languages (HDLs) like Verilog. However, challenges remain in ensuring the quality of Verilog generation. Complex designs often fail in a single generation due to the lack of targeted decoupling strategies, and evaluating the correctness of decoupled sub-tasks remains difficult. While the chain-of-thought (CoT) method is commonly used to improve LLM reasoning, it has been largely ineffective in automating IC design workflows, requiring manual intervention. The key issue is controlling CoT reasoning direction and step granularity, which do not align with expert RTL design knowledge. This paper introduces VeriBToT, a specialized LLM reasoning paradigm for automated Verilog generation. By integrating Top-down and design-for-verification (DFV) approaches, VeriBToT achieves self-decoupling and self-verification of intermediate steps, constructing a Backtrack Tree of Thought with formal operators. Compared to traditional CoT paradigms, our approach enhances Verilog generation while optimizing token costs through flexible modularity, hierarchy, and reusability.

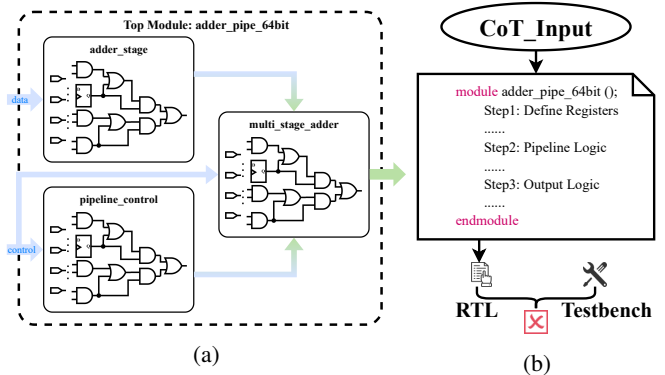
**Index Terms**—Backtrack Tree, ToT, DFV, Top-down

## I. INTRODUCTION

Digital circuits lie at the core of modern computing systems, yet their design remains a demanding process. Engineers rely on register transfer level (RTL) hardware description languages (HDLs) to translate functional requirements into logic gate combinations, which are then verified and mapped into transistor-level designs for fabrication. Although the downstream tasks of verification and transistor mapping have been largely automated, the initial writing of HDL code still demands significant human effort. This deficiency represents a pressing bottleneck in chip design workflows, inflating costs and undermining efficiency [1].

Recent advances in artificial intelligence, notably in large language models (LLMs) [2], present an encouraging avenue for mitigating this challenge. In fields like natural language processing (NLP), LLMs have achieved remarkable success in software [3] and hardware code generation [4]–[12]. This triumph underscores their potential to automate and streamline complex workflows.

† Corresponding Authors



## Top-Down Modular Design Flow with Submodule-Level Verify

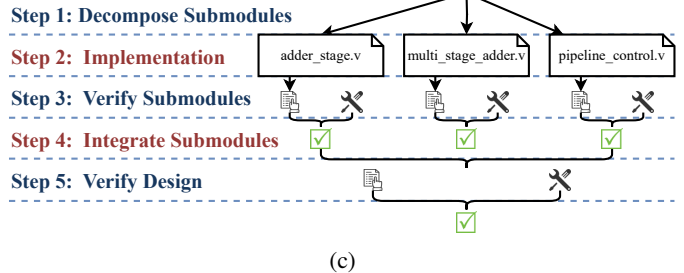


Fig. 1: The modular concurrency in circuit design is illustrated in (a), (b) presents common scenarios for the CoT paradigm in generating an RTL module’s Verilog code step by step, (c) depicts the human circuit design paradigm.

A pivotal factor enabling the success of LLMs in software code generation is the “chain-of-thought” (CoT) reasoning paradigm [13]–[15], which decomposes intricate problems into smaller, sequential steps. By enhancing the accuracy and logical consistency of generated outputs, CoT significantly advances automated code generation. However, directly applying CoT methods to the domain of automated RTL design has yielded limited results as sequence reasoning process is vulnerable and hard to capture hierarchical dependencies of complex submodules.

We conducted a task-specific analysis to better understand

the challenges of naive CoT in automated RTL design. Circuit design is fundamentally different from describing sequential functional behaviors in programming languages like C or Python. As illustrated in Fig. 1a, all functional modules in RTL design are implemented as concrete hardware circuits and operate in parallel to achieve the desired functionality. This parallelism introduces numerous interdependencies between modules. For LLMs, generating RTL requires considering a significantly broader range of contextual information across multiple RTL modules, as well as achieving precise cycle-level design coupling. Consequently, reasoning and decision-making in RTL design are considerably more complex than in software code generation, which typically involves understanding only a limited range of and coarse grained preceding contexts. When naive CoT is applied to enhance LLMs for generating RTL as shown in Fig. 1b, it is often observed that the generated results are uncontrollable, producing unnecessary and meaningless content, without improving the generation capability of LLMs.

The RTL design process by which human engineers write code is illustrated in Fig. 1c and exhibits two key characteristics from an engineering design perspective. Firstly, engineers consider factors such as circuit complexity and adopt a Top-down approach [16] to partition the design into modules. This ensures that each module remains highly cohesive internally while maintaining low coupling with other modules, thereby reducing both the difficulty of developing individual modules and the complexity of integrating them. Secondly, during the coding process, engineers verify each module’s functionality before merging them, iterating layer by layer [17]. This approach significantly lowers the overall debugging complexity. Together, these two practices facilitate efficient iteration and integration in modular design.

In this paper, we introduce VeriBToT, a novel Backtrack-Tree-of-Thought framework for modular RTL design, inspired by human IC design methods. Our approach divides the LLM reasoning process into two cognitive pathways: self-decoupling for partition refinement, and self-verification for ensuring the correctness of generated code. To implement these pathways, we design five prompt generation operators based on human IC design interactions, integrated into a Backtrack Tree of Thought. Experiments on open-source benchmarks show that VeriBToT enhances code correctness compared to existing CoT methods. The main contributions of this work are as follows:

- 1) We propose a novel Backtrack-Tree-of-Thought framework that integrates IC design methodologies (e.g., Top-down and DFV), achieving refined control over the thought process for LLM automatic RTL generation.
- 2) By leveraging the modular and hierarchical principles of the Top-down methodology, we decouple the ToT thought step nodes into corresponding RTL submodules within the IC design. The leaf nodes, representing frontier thought steps, are empowered to self-decouple their leaf-module designs into lower-level submodules.
- 3) By leveraging Top-down and DFV methodologies, we decouple common self-refinement operations found in

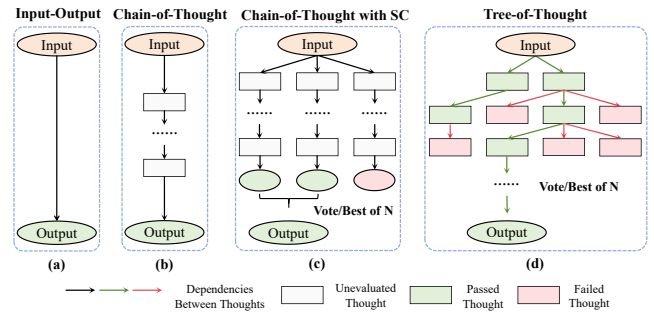


Fig. 2: The comparison between traditional CoT paradigms.

traditional prompt engineering [18], achieving modular verification and design-verification co-iteration in thought process, granting the intermediate thought steps the capability of self-verification.

- 4) Extensive experimental results show that the proposed VeriBToT significantly improves code correctness, with a significant improvement on the benchmark compared to existing CoT methods.

## II. MOTIVATION AND RELATED WORK

CoT-based prompt engineering [13] is used to improve reasoning and problem-solving performance for LLMs. The core idea is to explicitly guide the model through a multi-step reasoning by breaking down a complex task into smaller, sequential steps. This approach mimics human thought processes. Numerous variants have been developed around CoT, such as CoT-SC [19], tree of thoughts (ToT) [20], [21], and graph of thoughts (GoT) [22]. The comparison of different CoT reasoning paradigms is illustrated in the Fig. 2. The traditional CoT method offers some degree of assistance in the field of code generation [15], [23]–[26].

However, there are obstacles to enhancing RTL generation capabilities by directly applying any traditional CoT. Current CoT approaches lack the modular design and verification practices essential for hardware code development. These methods follow a linear, unidirectional flow, where large functional modules are decomposed without reconsidering the reasonableness of the partitioning. This often leads to overly complex modules or excessive inter-module interactions, making it difficult for LLMs to generate correct and integrable code. Moreover, CoT lacks single step verification, relying instead on functional validation after completing the entire code. This approach results in unclear feedback, making it hard to pinpoint and address specific issues. The lack of fine-grained control in current CoT methods also can lead to a significant increase in token consumption. As demonstrated by traditional ToT in Fig. 2d, both the tree depth and the number of child nodes for each node are hyperparameters, and setting these values too large can incur substantial computational costs. Consequently, the absence of modular thinking and iterative verification are key challenges that limit CoT’s effectiveness in RTL generation.

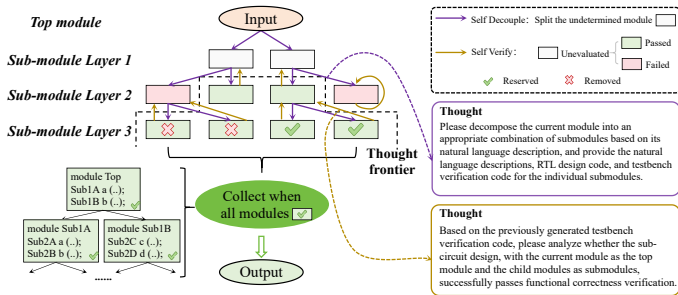


Fig. 3: The motivation of VeriBToT: each node in the diagram represents a thought step, which is limited to a module within the hierarchical RTL design (including natural language descriptions, RTL design Verilog, and testbench Verilog).

Inspired by human-centric IC design methodologies (Top-down, DFV), we introduce a dedicated ToT paradigm, VeriBToT, for automatic RTL generation by LLMs as shown in Fig. 3. The Top-down [16] approach is a hierarchical design methodology that starts with system-level requirements and progressively refines the design to lower abstraction levels, promoting early validation, design reuse, and team collaboration. Design-for-verification (DFV) [17] enhances the verifiability of designs by incorporating verification-friendly features early in the development process, aiming to reduce verification complexity and improve efficiency through modularization and better test planning. The process begins at inputting the natural language description of top-level module design. Then, self-decoupling thought prompt is designed for hierarchical sub-module decomposition to minimize inter-module dependencies, thereby enabling each sub-module to be designed and verified independently without affecting the others. Each thought step generates a complete design flow for a sub-module, including natural language descriptions, RTL design Verilog, and testbench Verilog. For instance, a top-level module *Top* may be decomposed into multi-layer lower level sub-modules (e.g., *Sub1A* could contain *Sub2A* and *Sub2B*, and *Sub1B* could contain *Sub2C* and *Sub2D*). When to invoke the self-decoupling thought requires the LLM to self-assess the complexity of the current design.

The self-verify thought prompt is designed to verify the correctness of the RTL code generated for each sub-module. Based on the already generated testbench, LLMs assess whether the RTL code for the sub-design rooted at the current node is functionally correct. Starting from the leaf node, the process iterates upward. If the LLMs determine that the implementation at the current node is incorrect, the existing sub-design is discarded and rewritten.

### III. THE PROPOSED VERIBTOT FRAMEWORK

#### A. Overview

We designed VeriBToT, a specialized LLM ToT reasoning paradigm for automated RTL design, featuring thought-step self-decoupling and self-verification capabilities. As shown in Fig. 3, the core of VeriBToT lies in how to formally define these two cognitive pathways. Inspired by the practical IC

design process, we define five specific operations that structure the entire reasoning process.

#### B. Operator Design

We have designed five operators to perform VeriBToT. These operators structure the reasoning process into a tree of thought, and the interactions between the operators can be abstracted as collaboration among different design personnel in an IC company design workflow, with the entire process adhering to the principles of Top-down and DFV design.

- 1) *Branch Generator*: Under a correct Top-down design module decomposition plan created by the product manager, RTL designers and testers will break down a complex design into appropriate submodules and separately design the RTL and testbench for each submodule.
- 2) *Node Evaluator*: Testers follow the module decomposition plan designed by the product manager to create separate test plans for different levels of modules according to their specifications. Once the RTL designs for all submodules under the current subtree are completed, the tester uses the pre-designed testbench to verify whether the RTL functionality of the current subsystem is correct.
- 3) *Node Rethinker*: When a single bottom-level module designed by an engineer fails through the testbench, the tester requests the design engineer to immediately redesign the RTL of this module.
- 4) *Backtrack Executor*: An engineer responsible for a middleware module composed of lower-level modules discovers during the actual development process that the product manager's module decomposition plan contains errors, making it impossible to produce middleware that meets overall design requirements. The engineer requests that the product manager revise the module decomposition plan to determine whether the next step should involve redesigning based on the current middleware or modifying the design of the layer above.
- 5) *Code Aggregator*: The product manager ultimately assembles all the validated hierarchical modules into a complete design for approval.

By implementing these operators, each intermediate thought step can achieve self-decoupling and self-validation, enhancing the overall design ability of LLMs.

#### C. Reasoning Process

We define several notions used in this paper based on concepts from previous methods to facilitate the formalization of the problem. We use **lowercase letters**  $x, y, z$  to represent **language sequences**. Here,  $x$  denotes the original natural language description of the RTL design, serving as the input for the entire problem;  $y$  represents the final answer to the problem, including the complete Verilog code as the output; and  $z_1, \dots, z_n$  denotes a **chain of thoughts** that bridges  $x$  and  $y$ , where each  $z_i$  is a intermediate thought step that guides the process of RTL generation. In the **ToT paradigm**, LLMs are permitted to explore multiple reasoning paths across the nodes of a tree, where an intermediate **state** in the tree is represented

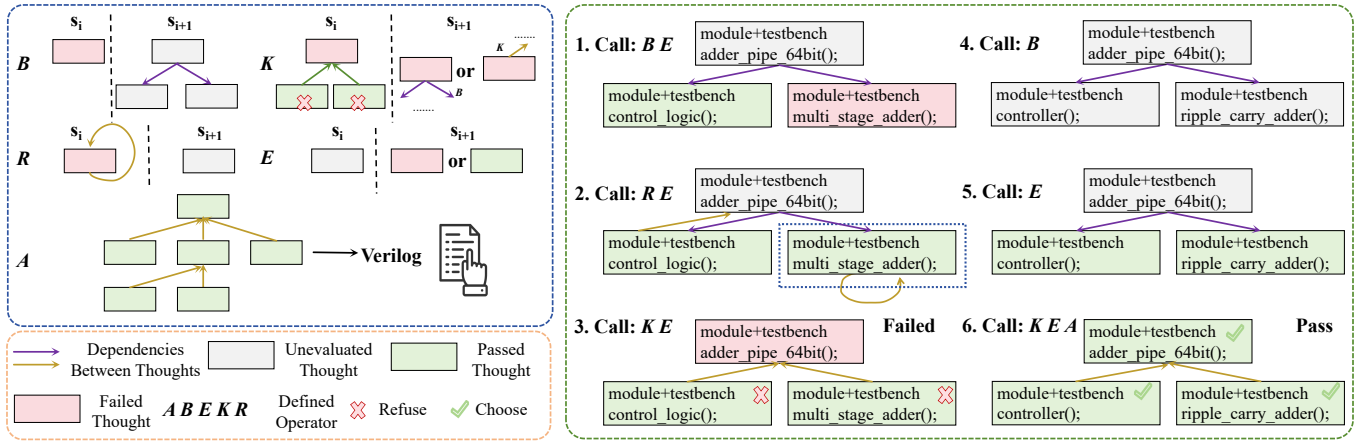


Fig. 4: The left side of the figure illustrates a schematic representation of how pre-defined operators control the thought process. The right side depicts how a 64-bit adder is generated through the thought process of VeriBToT, achieved through the combination of these operators; to differentiate among the five operators we have defined, we employ the following abbreviations: *Branch Generator* is represented by  $B$ , *Node Evaluator* by  $E$ , *Node Rethinker* by  $R$ , *Backtrack Executor* by  $K$ , and *Code Aggregator* by  $A$ .

as  $s = [x, z_1, \dots, z_i]$ , a **node of thought** is represented as  $n$  sampled from the current state  $s$ , and a collection of all states for ToT is denoted by  $S = [s_1, \dots, s_n]$ .

The formal representation of the five operators and a complete example of the VeriBToT reasoning process are illustrated in Fig. 4. To formalize the discussion of the exploration process of VeriBToT, we use **uppercase letters**:  $L$  to denote natural language,  $D$  to denote RTL design, and  $V$  to denote testbench for verification. From the perspective of the content of language sequences, each intermediate thought step  $z_i$  can be split into three parts: the RTL design language description  $z_i^L$ , the Verilog code for the current module RTL design  $z_i^D$ , and the testbench for the current design  $z_i^V$ . The entire process can be broadly categorized into the following key components, organized according to the depth-first search (DFS) traversal method of the Backtrack Tree algorithm:

**Tree Initialization.** The process begins with a root node  $s_0 = x$ ,  $n_0 = x$ . The input sequence  $x = [x^L, x^V]$  represents a natural language description of the overall RTL design and the testbench used for final functionality testing.

**Branching.** For a newly generated leaf node  $n = [n^L, n^D, n^V]$ , VeriBToT calls *Node Evaluator* for self-verifying  $n^D$  by  $n^V$ , both of which are generated at the same step for  $n$ . If the LLM determines that the validation is successful, it proceeds to the next node. If the validation fails again, the node faces two choices. If the LLM determines that the current design is still relatively complex, it decides to invoke the *Branch Generator*. The LLM, using a “branching prompt,” decomposes the current leaf node’s design into several new leaf nodes, each of which describes the complete design requirements  $n^L$ , implementation  $n^D$ , and validation  $n^V$  for a bottom-level module. If the LLM considers the current design to be simple enough, it decides to invoke *Backtrack Executor*, which we will discuss next.

**Backtracking.** The invocation of the *Backtrack Executor* can occur not only after the failure of leaf node generation mentioned above but also when all the leaf nodes of any subtree pass self-validation through the *Node Evaluator*. At this point, VeriBToT begins backtracking towards the root node. For each non-leaf node encountered during this process, the *Node Rethinker* is called based on the submodules of this node that have already passed validation, generating new  $[n^D, n^V]$ , followed by self-validation. During this process, if any non-leaf node fails validation again, backtracking occurs at that node. At this moment, two choices arise. If the LLM determines that there are issues with the submodule divisions of the current node, it uses a “removing prompt” to clear all sub-nodes in the subtree rooted at this node and restarts branching from the backtracked node. Conversely, if the LLM believes that the design of the current node itself is sound but that the task decomposition of the parent node’s higher-level module is erroneous, it will continue backtracking to the parent node, removing the backtracked node and its sibling nodes. The LLM will then evaluate in the same manner iteratively until it backtracks to a node where the submodule division is valid and recommences branching.

**Tree Finality.** When the evaluator believes that the entire RTL design has been correctly completed, the *Code Aggregator* operation is executed to return the final RTL Verilog output  $y$ , and EDA tools are invoked to verify whether it passes the top-level testbench  $x^V$ .

Fig. 4 illustrates how a practical RTL design example is developed using the VeriBToT framework. For a 64-bit adder, referred to as *adder\_pipe\_64bit*, the LLM initially cannot complete a correct design within a single Verilog module. As a result, it decomposes this into two submodules: the control module *control\_logic* and the multi-stage adder module *multi\_stage\_adder*. After both submodules are designed and validated through self-validation, the LLM backtracks to the

Reasoning Paradigm	VerilogEval-Human				RTLML			
	DeepSeek		ChatGPT-4		DeepSeek		ChatGPT-4	
	<i>pass@1</i>	<i>pass@5</i>	<i>pass@1</i>	<i>pass@5</i>	<i>pass@1</i>	<i>pass@5</i>	<i>pass@1</i>	<i>pass@5</i>
IO	0.27	0.36	0.34	0.45	0.34	0.46	0.42	0.56
CoT	0.29	0.37	0.33	0.47	0.30	0.40	0.38	0.50
CoT-SC	0.27	0.37	0.38	0.50	0.24	0.38	0.40	0.46
ToT	0.26	0.39	0.35	0.48	0.28	0.48	0.38	0.48
VeriBToT	0.32	0.44	0.43	0.58	0.42	0.56	0.48	0.62
VeriBToT-	0.29	0.36	0.36	0.52	0.26	0.44	0.41	0.57
Domain Specific	Thakur		RTLCoder		Thakur		RTLCoder	
	<i>pass@1</i>	<i>pass@5</i>	<i>pass@1</i>	<i>pass@5</i>	<i>pass@1</i>	<i>pass@5</i>	<i>pass@1</i>	<i>pass@5</i>
IO	30.3	43.9	36.7	45.5	14.9	24.1	39.8	48.3

Domain specific models lack the ability to follow instructions based on different reasoning paradigms; the light green indicates that the current result surpasses one domain-specific fine-tuned model, while the dark green indicates that it surpasses two.

TABLE I: Reasoning paradigm evaluation on the full benchmark about functional *Pass@1* and *Pass@5*.

Design	IO			CoT			CoT-SC			ToT			VeriBToT		
	Syn.	Fun.	Tok.	Syn.	Fun.	Tok.	Syn.	Fun.	Tok.	Syn.	Fun.	Tok.	Syn.	Fun.	Tok.
adder_16bit	3	3	1.11	5	3	1.35	4	3	2.62	4	3	1.39	5	5	1.88
adder_pipe_64bit	5	1	0.67	5	1	0.82	4	2	1.47	5	3	1.94	5	2	2.34
multi_pipe_8bit	0	0	0.94	0	0	1.26	0	0	2.03	3	0	1.86	4	0	1.41
multi_16bit	1	1	1.1	3	0	1.02	1	0	1.92	4	2	1.45	4	1	1.58
barrel_shifter	5	0	0.62	5	0	0.67	4	0	1.65	4	1	1.41	5	1	2.12
width_8to16	5	0	0.53	2	1	1.01	2	1	0.97	5	1	0.99	4	3	1.62
calendar	5	0	0.52	5	0	0.98	2	0	1.38	5	0	1.1	5	5	1.45
freq_div	5	0	0.74	5	1	1.16	5	2	1.08	3	0	1.78	5	4	1.44
dff8r	3	0	0.43	4	4	0.65	0	0	1.41	3	0	1.23	5	5	1.21
fsm2s	0	0	0.64	0	0	0.78	0	0	1.52	0	0	1.10	2	1	1.82
fsm3comb	0	0	0.63	1	1	0.90	1	1	1.68	0	0	1.34	5	3	2.56
gates4	5	1	0.42	5	4	0.65	4	3	1.36	5	4	1.24	5	5	1.49
popcount3	0	0	0.57	3	1	0.76	0	0	1.25	1	0	1.77	2	2	1.77
m2014_q4i	0	0	0.13	0	0	0.35	0	0	0.82	3	2	1.06	4	4	0.42
ringer	0	0	0.39	0	0	0.71	0	0	0.91	0	0	1.06	4	3	2.53
timer	0	0	0.47	2	2	0.68	0	0	1.29	2	1	1.15	4	3	1.77

TABLE II: The performance of the Deepseek-Coder-V2 in generating Verilog for different hard cases under various reasoning paradigms is measured by the number of successful attempts out of five trials (*#pass@5*), assessing syntax correctness (Syn.) and functional correctness (Fun.); the average token consumption (Tok., the unit is thousand, *k*) is used to evaluate the resource expenditure of the reasoning paradigms, **Bold values** indicate results from our method that are the best.

Design	IO			CoT			CoT-SC			ToT			VeriBToT		
	Syn.	Func.	Tok.	Syn.	Func.	Tok.	Syn.	Func.	Tok.	Syn.	Func.	Tok.	Syn.	Func.	Tok.
adder_16bit	4	3	0.85	4	4	1.88	5	5	3.04	4	4	2.73	5	5	1.76
adder_32bit	5	3	1.37	3	3	1.93	3	3	2.26	2	0	2.19	5	2	2.01
adder_pipe_64bit	5	1	1.07	3	0	1.80	5	1	2.62	4	0	2.65	5	4	2.37
multi_pipe_8bit	5	0	1.97	3	0	1.37	5	0	2.51	2	0	2.58	4	0	0.57
barrel_shifter	1	0	0.74	5	0	1.2	5	0	1.54	5	0	2.3	5	0	2.62
width_8to16	5	3	0.85	5	4	1.45	4	3	1.61	5	3	1.78	5	3	2.32
freq_div	5	1	0.89	5	4	1.70	5	3	3.12	5	4	2.05	5	5	1.99
sequence_detector	0	0	1.03	0	0	1.67	1	0	2.13	1	0	2.93	5	2	1.84
dff8r	5	0	0.43	0	0	0.90	0	0	1.46	0	0	1.38	3	3	1.65
fsm2s	0	0	0.70	4	2	1.00	1	1	1.75	0	0	1.52	4	3	1.77
fsm3comb	0	0	0.66	0	0	0.98	0	0	1.71	3	3	1.81	5	5	1.95
gates4	5	3	0.42	5	5	0.80	5	5	1.42	3	3	1.33	5	5	1.82
popcount3	5	0	0.45	3	0	0.70	4	1	1.27	1	1	1.40	4	2	1.75
lemmings1	0	0	0.88	2	2	1.38	0	0	1.85	0	0	1.51	5	3	1.77
rule90	3	3	0.76	2	1	1.09	2	2	1.54	3	1	1.49	5	5	1.65
vector3	3	1	0.50	1	1	0.95	4	2	1.65	2	0	1.70	5	3	1.75

TABLE III: The performance of the ChatGPT-4 in generating Verilog for different hard cases under various reasoning paradigms, the various metrics are the same as those in TABLE II.

root node and recognizes that they cannot produce a functional adder. It then backtracks further to redefine the submodules, introducing a control module *controller* and a new adder submodule *ripple\_carry\_adder*. Using the testbench, the LLM verifies that the revised design is correct, after which the final

RTL code is collected and returned.

#### IV. EXPERIMENTS

Our proposed method is compared with standard IO prompting and several common CoT prompting paradigms. To better illustrate the versatility of our approach, we consider

four models. Two of them are general-purpose commercial LLMs with strong code generation capabilities: ChatGPT-4 and DeepSeek-Coder-V2 [27]. The other two are domain-specific open-source LLMs fine-tuned for Verilog generation: RTLCoder [6] and Thakur [28].<sup>1</sup> We utilized two NL2V (natural language to Verilog code) benchmarks, RTLLM [10] and VerilogEval-Human [8], for our experiments. VerilogEval-Machine, due to its focus on circuit component connections and truth tables, is not considered as it lacks the flexibility for module decomposition. The key performance indicators used for performance measurement are syntactic correctness, functional correctness, the number of tokens consumed, first pass rate ( $Pass@1$ ) five pass rate ( $Pass@5$ ), and the number of successful attempts in the five ( $\#Pass@5$ ).

### A. Full Benchmark Experimental Results

TABLE I presents the experimental results for  $pass@1$  and  $pass@5$  across the entire dataset of two common academic benchmarks. The results indicate that our method effectively enhances model performance, both with the state-of-the-art model ChatGPT-4 and the open-source model Deepseek. We did not conduct experiments with different inference paradigms for the domain-specific fine-tuned models (Thakur and RTLCoder), as we observed a significant decline in their general instruction-following ability after fine-tuning. These models were unable to comprehend and adhere to the reasoning paradigms for generation.

### B. Ablation Study

We removed the prompt for generating testbenches and self-validation based on them from the original VeriBToT, leaving the LLM to reflect only on the Verilog generated for each node and determine whether backtracking is necessary, labeled as VeriBToT- in TABLE I. The experimental results show a significant performance drop, highlighting the effectiveness of the self-validation backtracking mechanism we proposed. ToT can be seen as an ablation of VeriBToT-, where the backtracking mechanism is removed and the model is forced to perform reasoning up to a predetermined depth. As another ablation experiment, the results show that ToT performs worse than VeriBToT-, proving the effectiveness of the backtracking mechanism. This also partially demonstrates that excessive reasoning in LLMs may not necessarily improve performance [29], particularly when CoTs leads to performance degradation compared to IO in the case of RTLLM.

### C. Hard Case Experimental Results

We conducted a detailed case analysis on the function and syntax pass rates ( $\#pass@5$ ) described in the RTLCoder paper [6], specifically focusing on the relatively low performance in hard cases. Hard cases are defined as those circuits for which the direct IO prompt failed in the functional  $pass@1$  experiment of TABLE I. Due to space limitations, we only present 8 cases each for RTLLM and VerilogEval-Human. The results, as shown in Tables II and III, demonstrate that




Fig. 5: The token consumption comparison between different CoT paradigms for RTL generation.

our method encourages deeper model reasoning, leading to improvements in both syntactic and functional correctness compared to IO and other CoT modes. Moreover, token consumption does not increase significantly; in some instances, it is even lower than that of ToT and CoT-SC.

### D. Token Consumption Evaluation

To assess VeriBToT's token overhead, we analyzed token usage across reasoning paradigms, as shown in Fig. 5. The results indicate that although VeriBToT increases prompt token consumption, the tokens generated in context remain comparable to those of CoT-SC and ToT. Given the additional context introduced—such as testbench code and natural language submodule descriptions—the overall increase is relatively small. This demonstrates the efficiency of the backtracking mechanism in controlling token usage.

## V. CONCLUSION

In this paper, we present VeriBToT, a novel CoT architecture for automating RTL design with LLMs. Inspired by top-down IC design and Design for Verification (DFV), VeriBToT enables self-decoupling and self-verification of thought nodes, represented through a BackTrack Tree. This improves the model's understanding of Verilog design workflows, and our experiments show the broad applicability of the approach. We believe VeriBToT can serve as a unified reasoning paradigm for large-scale circuit generation with LLMs.

### ACKNOWLEDGMENT

This paper is supported in part by the Chinese Academy of Sciences under grant No. XDB0660103, and in part by the National Natural Science Foundation of China (NSFC) under grant No. 62090024.

<sup>1</sup>Experiments were conducted between March 1–20, 2025.

## REFERENCES

- [1] L. Chen, Y. Chen, Z. Chu, W. Fang, T.-Y. Ho, Y. Huang, S. Khan, M. Li, X. Li, Y. Liang *et al.*, “The dawn of ai-native eda: Promises and challenges of large circuit models,” *arXiv preprint arXiv:2403.07257*, 2024.
- [2] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat *et al.*, “Gpt-4 technical report,” *arXiv preprint arXiv:2303.08774*, 2023.
- [3] X. Hou, Y. Zhao, Y. Liu, Z. Yang, K. Wang, L. Li, X. Luo, D. Lo, J. Grundy, and H. Wang, “Large language models for software engineering: A systematic literature review,” *ACM Transactions on Software Engineering and Methodology*, 2023.
- [4] K. Chang, K. Wang, N. Yang, Y. Wang, D. Jin, W. Zhu, Z. Chen, C. Li, H. Yan, Y. Zhou *et al.*, “Data is all you need: Finetuning llms for chip design via an automated design-data augmentation framework,” in *Proceedings of the 61st ACM/IEEE Design Automation Conference*, 2024, pp. 1–6.
- [5] 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*, ser. DAC ’24. New York, NY, USA: Association for Computing Machinery, 2024. [Online]. Available: <https://doi.org/10.1145/3649329.3658493>
- [6] S. Liu, W. Fang, Y. Lu, J. Wang, Q. Zhang, H. Zhang, and Z. Xie, “RTLCode: Fully Open-Source and Efficient LLM-Assisted RTL Code Generation Technique,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, pp. 1–1, 2024.
- [7] M. DeLorenzo, A. B. Chowdhury, V. Gohil, S. Thakur, R. Karri, S. Garg, and J. Rajendran, “Make every move count: Llm-based high-quality rtl code generation using mcts,” *arXiv preprint arXiv:2402.03289*, 2024.
- [8] M. Liu, N. Pinckney, B. Khailany, and H. Ren, “Verilogeval: Evaluating large language models for verilog code generation,” in *2023 IEEE/ACM International Conference on Computer Aided Design (ICCAD)*. IEEE, 2023, pp. 1–8.
- [9] S. Thakur, B. Ahmad, Z. Fan, H. Pearce, B. Tan, R. Karri, B. Dolan-Gavitt, and S. Garg, “Benchmarking Large Language Models for Automated Verilog RTL Code Generation,” in *2023 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2023, pp. 1–6.
- [10] Y. Lu, S. Liu, Q. Zhang, and Z. Xie, “Rtllm: An open-source benchmark for design rtl generation with large language model,” in *2024 29th Asia and South Pacific Design Automation Conference (ASP-DAC)*. IEEE, 2024, pp. 722–727.
- [11] Y. Zhao, D. Huang, C. Li, P. Jin, Z. Nan, T. Ma, L. Qi, Y. Pan, Z. Zhang, R. Zhang *et al.*, “Codev: Empowering llms for verilog generation through multi-level summarization,” *arXiv preprint arXiv:2407.10424*, 2024.
- [12] 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 Trans. Des. Autom. Electron. Syst.*, vol. 29, no. 3, Apr. 2024. [Online]. Available: <https://doi.org/10.1145/3643681>
- [13] J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. Chi, Q. V. Le, D. Zhou *et al.*, “Chain-of-thought prompting elicits reasoning in large language models,” *Advances in neural information processing systems*, vol. 35, pp. 24 824–24 837, 2022.
- [14] J. Li, G. Li, Y. Li, and Z. Jin, “Structured Chain-of-Thought Prompting for Code Generation,” *ACM Trans. Softw. Eng. Methodol.*, Aug. 2024, just Accepted. [Online]. Available: <https://doi.org/10.1145/3690635>
- [15] H. Le, H. Chen, A. Saha, A. Gokul, D. Sahoo, and S. Joty, “Codechain: Towards modular code generation through chain of self-revisions with representative sub-modules,” *arXiv preprint arXiv:2310.08992*, 2023.
- [16] H. Kaeslin, *Top-down digital VLSI design: from architectures to gate-level circuits and FPGAs*. Morgan Kaufmann, 2014.
- [17] F. Sforza, L. Battu, M. Brunelli, A. Castelnuovo, and M. Magnaghi, “A ”design for verification” methodology,” in *Proceedings of the IEEE 2001. 2nd International Symposium on Quality Electronic Design*, 2001, pp. 50–55.
- [18] A. Madaan, N. Tandon, P. Gupta, S. Hallinan, L. Gao, S. Wiegreffe, U. Alon, N. Dziri, S. Prabhume, Y. Yang *et al.*, “Self-refine: Iterative refinement with self-feedback,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [19] X. Wang, J. Wei, D. Schuurmans, Q. Le, E. Chi, S. Narang, A. Chowdhery, and D. Zhou, “Self-consistency improves chain of thought reasoning in language models,” *arXiv preprint arXiv:2203.11171*, 2022.
- [20] S. Yao, D. Yu, J. Zhao, I. Shafran, T. Griffiths, Y. Cao, and K. Narasimhan, “Tree of thoughts: Deliberate problem solving with large language models,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [21] J. Long, “Large language model guided tree-of-thought,” *arXiv preprint arXiv:2305.08291*, 2023.
- [22] M. Besta, N. Blach, A. Kubicek, R. Gerstenberger, M. Podstawski, L. Gianinazzi, J. Gajda, T. Lehmann, H. Niewiadomski, P. Nyczyk *et al.*, “Graph of thoughts: Solving elaborate problems with large language models,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 16, 2024, pp. 17 682–17 690.
- [23] D. Huang, Q. Bu, and H. Cui, “Codecot and beyond: Learning to program and test like a developer,” *arXiv preprint arXiv:2308.08784*, 2023.
- [24] H. Wang, Z. Liu, S. Wang, G. Cui, N. Ding, Z. Liu, and G. Yu, “INTERVENOR: Prompting the Coding Ability of Large Language Models with the Interactive Chain of Repair,” in *Findings of the Association for Computational Linguistics ACL 2024*, 2024, pp. 2081–2107.
- [25] Z. Ni, Y. Li, N. Yang, D. Shen, P. Lv, and D. Dong, “Tree-of-code: A tree-structured exploring framework for end-to-end code generation and execution in complex task handling,” *arXiv preprint arXiv:2412.15305*, 2024.
- [26] Y. Zhou, R. Chen, X. Li, J. Wang, Z. Fang, B. Wang, W. Bai, Q. Cao, and L. Wang, “Vtot: Automatic verilog generation via llms with tree of thoughts prompting,” in *2025 Design, Automation & Test in Europe Conference (DATE)*, 2025, pp. 1–2.
- [27] Q. Zhu, D. Guo, Z. Shao, D. Yang, P. Wang, R. Xu, Y. Wu, Y. Li, H. Gao, S. Ma *et al.*, “Deepseek-coder-v2: Breaking the barrier of closed-source models in code intelligence,” *arXiv preprint arXiv:2406.11931*, 2024.
- [28] S. Thakur, B. Ahmad, Z. Fan, H. Pearce, B. Tan, R. Karri, B. Dolan-Gavitt, and S. Garg, “Benchmarking Large Language Models for Automated Verilog RTL Code Generation,” 2022. [Online]. Available: <https://arxiv.org/abs/2212.11140>
- [29] Z. Wu, L. Qiu, A. Ross, E. Akyürek, B. Chen, B. Wang, N. Kim, J. Andreas, and Y. Kim, “Reasoning or reciting? exploring the capabilities and limitations of language models through counterfactual tasks,” 2024. [Online]. Available: <https://arxiv.org/abs/2307.02477>