

PrefixGPT: Prefix Adder Optimization by a Generative Pre-trained Transformer*

Ruogu Ding¹, Xin Ning¹, Ulf Schlichtmann², Weikang Qian¹

¹Global College, Shanghai Jiao Tong University, China

²Chair of Electronic Design Automation, Technical University of Munich, Germany
rg.ding@sjtu.edu.cn, xin.ning1412@sjtu.edu.cn, ulf.schlichtmann@tum.de, qianwk@sjtu.edu.cn

Abstract

Prefix adders are widely used in compute-intensive applications for their high speed. However, designing optimized prefix adders is challenging due to strict design rules and an exponentially large design space. We introduce PrefixGPT, a generative pre-trained Transformer (GPT) that directly generates optimized prefix adders from scratch. Our approach represents an adder’s topology as a two-dimensional coordinate sequence and applies a legality mask during generation, ensuring every design is valid by construction. PrefixGPT features a customized decoder-only Transformer architecture. The model is first pre-trained on a corpus of randomly synthesized valid prefix adders to learn design rules and then fine-tuned to navigate the design space for optimized design quality. Compared with existing works, PrefixGPT not only finds a new optimal design with a 7.7% improved area-delay product (ADP) but exhibits superior exploration quality, lowering the average ADP by up to 79.1%. This demonstrates the potential of GPT-style models to first master complex hardware design principles and then apply them for more efficient design optimization.

Code — <https://github.com/Mightlaus/PrefixGPT-AAAI26>

1 Introduction

Adders are important arithmetic units used in many applications such as signal processing and artificial intelligence (AI). Among many available designs, prefix adders stand out for their fast speed (Ladner and Fischer 1980). They are typically modeled as prefix graphs with rigid design rules, and their design space is exponentially large. Consequently, the central goal of prefix adder optimization is to navigate this space to find valid designs with an optimized trade-off between circuit area and delay.

Historically, classic designs like the Sklansky (Sklansky 1960), Kogge-Stone (Kogge 1973), and Brent-Kung (Brent and Kung 1982) adders were discovered manually. To explore the design space more systematically, design automation methods were developed. Some heuristic approaches, such as simulated annealing (Moto and Kaneko 2018) and

space pruning (Ma et al. 2018), often struggled to escape suboptimal regions. Recently, AI-based methods using reinforcement learning (RL) (Roy et al. 2021), gradient descent (Song et al. 2024), Monte Carlo tree search (MCTS) (Lai et al. 2024), and large language models (LLMs) (Xiao et al. 2024) have shown greater success.

Despite their achievement, these AI methods often share a common limitation: they frame optimization as a structure refinement task. Whether by adding or deleting nodes (Roy et al. 2021; Lai et al. 2024) or editing textual representations of local connections (Xiao et al. 2024), they start with an existing design and progressively modify it. This paradigm introduces two issues. First, iterative modifications frequently violate the strict design rules of prefix adders. This necessitates an additional and often inconvenient repair step to ensure correctness, which may also introduce extra hardware cost. Second, the performance is highly sensitive to the initial design (e.g., Sklansky). As illustrated in Fig. 1, the effectiveness of two representative methods, PrefixRL (Roy et al. 2021) and ArithTree (Lai et al. 2024), can dramatically change when starting from different manual designs. This reveals a strong *initialization bias* that may prevent these refinement-based methods from discovering high-quality designs.

The recent success of generative pre-trained Transformer (GPT)-style models offers a new perspective (Radford et al. 2018; Touvron et al. 2023). They have showed an exceptional ability in mastering complex natural language and producing high-quality text that simultaneously satisfies structural and logical constraints (Radford et al. 2019; Guo et al. 2025). Prefix adders, with their rigid construction rules, form an even more structured domain than natural language. This motivates our central question: *Can a GPT-style model also be developed to learn the “grammar” of prefix adders and directly generate high-quality designs that are valid by construction?*

We affirmatively answer the above question by proposing PrefixGPT, a customized GPT-style model trained from scratch for generating optimized prefix adders. We formulate the prefix adder design as a sequence generation problem and build a GPT-style model to generate the sequence. We recast the design rules into a dynamic legality mask that is applied during generation, ensuring every adder is valid by construction to avoid any repairs. After pre-trained on a large

*Corresponding author: Weikang Qian.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

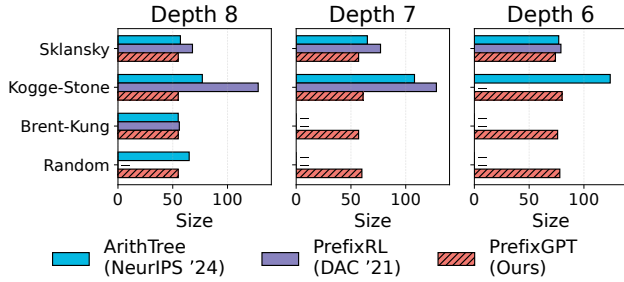


Figure 1: Optimized 32-bit adders discovered by two state-of-the-art methods and PrefixGPT under four initializations (Sklansky, Kogge-Stone, Brent-Kung and Random). Here, depth and size approximate circuit delay and area, respectively; the smaller the better. While the other methods fluctuated or even failed to find a solution (denoted by '-'), PrefixGPT was robust and consistently found better designs across all initializations. A full comparison is in Table 1.

corpus of randomly synthesized prefix adders to learn the design rules, the model is further fine-tuned with an RL objective to generate prefix adders with optimized area and delay. Experiments show that PrefixGPT finds superior designs quickly and robustly. Even under random initialization, its performance outperforms state-of-the-art (SOTA) methods that rely on good initial manual designs, as shown in Fig. 1. Our contributions are summarized as follows:

- We introduce a sequence representation for prefix adders, thereby reformulating their optimization as a sequence generation task that can be addressed by powerful GPT-style models.
- We recast the design rules for prefix adders into a dynamic legality mask that is directly integrated into the generation process, guaranteeing every design is valid by construction and eliminating any repair steps.
- We build PrefixGPT and equip it with a comprehensive pre-training and fine-tuning pipeline, enabling the model to leverage the learned design rules for efficient and robust prefix adder optimization.

2 Preliminaries

2.1 Prefix Adder

Adders are fundamental arithmetic units in digital systems. Given two n -bit binary inputs, $A = (a_{n-1}a_{n-2} \dots a_0)_2$ and $B = (b_{n-1}b_{n-2} \dots b_0)_2$, an n -bit adder computes an n -bit sum $S = (s_{n-1}s_{n-2} \dots s_0)_2$ and a final carry-out c_{n-1} . A prefix adder calculates its sum bit s_i ($1 \leq i \leq n-1$) from a carry-out at bit position $(i-1)$, denoted as c_{i-1} . Therefore, a prefix adder needs to generate n carry-outs c_0, c_1, \dots, c_{n-1} .

The dominant part of a prefix adder is a sub-circuit for generating these carry-outs, which can be abstracted as a *prefix graph* shown in Fig. 2(a). There are two types of nodes in this graph: *input node*, denoted as $\ell_{i:i}$ with $0 \leq i \leq n-1$, located at the top row of the graph, and *merge node*, denoted as $\ell_{j:i}$ with $0 \leq i < j \leq n-1$, located at any remaining row of the graph. For any node $\ell_{j:i}$ ($0 \leq i \leq j \leq n-1$) in

the graph, it outputs a *signal pair*. An input node $\ell_{i:i}$ takes input bits a_i and b_i as its inputs, while a merge node $\ell_{j:i}$ takes two signal pairs from $\ell_{j:k}$ and $\ell_{k-1:i}$ as inputs, where k satisfies that $i < k \leq j$. For example, the merge node $\ell_{4:2}$ in Fig. 2(a) takes two signal pairs from $\ell_{4:4}$ and $\ell_{3:2}$ as inputs. We call the nodes $\ell_{j:k}$ and $\ell_{k-1:i}$ the *more significant parent (MSP)* and the *less significant parent (LSP)*, respectively, of the node $\ell_{j:i}$. In a prefix adder, the carry-out c_i ($0 \leq i \leq n-1$) is a signal from the output of $\ell_{i:0}$. In order to generate all carry-outs, the prefix graph must include n nodes $\ell_{0:0}, \ell_{1:0}, \dots, \ell_{n-1:0}$. For readers interested in the actual circuit implementation of the input and merge nodes, we refer them to the appendix. In summary, a valid prefix graph should satisfy the following three design rules:

1. **Input rule:** It includes nodes $\ell_{0:0}, \ell_{1:1}, \dots, \ell_{n-1:n-1}$.
2. **Output rule:** It includes nodes $\ell_{0:0}, \ell_{1:0}, \dots, \ell_{n-1:0}$.
3. **Merge rule:** For any merge node $\ell_{j:i}$ ($0 \leq i < j \leq n-1$) in the graph, there must exist an integer k such that $i < k \leq j$ and both nodes $\ell_{j:k}$ and $\ell_{k-1:i}$ exist in the graph.

The structure of the prefix graph directly affects the area and delay of an adder. Specifically, the *size* of the graph, which is the total number of merge nodes in the graph, is proportional to the adder area, while the *depth* of the graph, which is the number of levels of the graph, is proportional to the delay of the adder. To design a prefix adder with optimized area and delay translates to finding a prefix graph with optimized size and depth.

2.2 Generative Pre-trained Transformers (GPTs)

GPTs have been successfully used to tackle tasks in diverse fields, including biology (Luo et al. 2022), chemistry (M. Bran et al. 2024), finance (Seshakagari et al. 2025), and programming (Bucaioni et al. 2024). These GPT-style models are typically built upon a decoder-only architecture (Radford et al. 2018) and operate *autoregressively*: given a sequence of input tokens (x_1, \dots, x_t) , the model predicts the next token x_{t+1} . The generated token is then appended to the sequence, forming a new input sequence $(x_1, \dots, x_t, x_{t+1})$ for the subsequent generation step. This process repeats until a terminal condition, such as an end-of-sequence (EOS) token or a maximum length limit, is met.

A GPT-style model is typically trained in two stages. The first stage is *pre-training*, where the model is trained on a large corpus of general-domain text to build a broad understanding of syntax and factual knowledge. The second stage is *fine-tuning*, where it is adapted for a specific downstream task and often uses techniques like RL to align its behavior with the desired objectives. We employ the same two-stage training process for our PrefixGPT.

3 Methods

This section describes the details of PrefixGPT. To facilitate the development of a GPT-style model for prefix adder generation, Sec. 3.1 presents a method that converts a prefix graph into a coordinate sequence. It also shows how to recast the design rules into a dynamic legality mask used in the generation process. A custom decoder-only Transformer

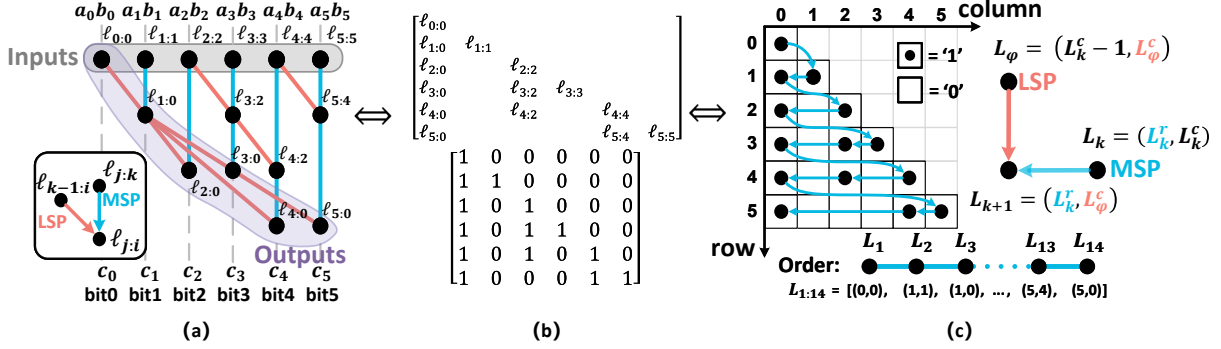


Figure 2: Equivalent representations of a 6-bit prefix adder with size = 8 and depth = 4. (a) Graph. (b) Matrix. (c) Coordinate Sequence. The blue line always connects a node and its MSP, while red line for its LSP.

model (Sec. 3.2) with a full pre-training (Sec. 3.3) and fine-tuning (Sec. 3.4) pipeline is then proposed for optimized adder exploration.

3.1 Sequence Representation of Prefix Adder

To develop a GPT-style model for generating prefix adders, we need to find a suitable sequence representation of the underlying prefix graph. We first convert the graph into a matrix representation proposed in (Roy et al. 2021). As shown in Fig. 2(b), the prefix graph of an n -bit prefix adder can be represented as an $n \times n$ binary lower-triangular matrix, where the entry at row j and column i ($0 \leq i \leq j \leq n - 1$) is a ‘1’ if and only there exists a node $\ell_{j:i}$ in the prefix graph. We call the matrix *prefix matrix*. We record the coordinates of the ‘1’s by scanning the prefix matrix from the top row to the bottom, and within each row, from the diagonal entry to the leftmost. This scan creates a *coordinate sequence* $L_{1:N} = (L_1, \dots, L_N)$, where N is the sequence length and $L_p = (L_p^r, L_p^c)$ ($1 \leq p \leq N$) is the coordinate in the matrix of the p -th ‘1’ in the scan order with L_p^r giving the row coordinate and L_p^c giving the column coordinate. Note that a coordinate $L_p = (L_p^r, L_p^c)$ in the sequence corresponds to a node $\ell_{L_p^r:L_p^c}$ in the prefix graph.

For a GPT-style model, it repeatedly generates the next sequence entry given the current partial sequence until a complete sequence is generated. Thus, to develop a GPT-style model generating the coordinate sequence, it is important to identify the valid choices for the next coordinate L_{k+1} given the current sequence $L_{1:k}$. We distinguish two cases.

Case 1: the last coordinate in the current sequence, i.e., L_k , is at column 0 of the prefix matrix. By the input rule described in Sec. 2.1, all diagonal entries of the prefix matrix are ‘1’. Thus, in this case, by the scanning order, the next ‘1’ we visit is located at the diagonal entry of the next row of L_k (see Fig. 2(c)). Thus, we have $L_{k+1} = (L_k^r + 1, L_k^c + 1)$.

Case 2: L_k is not at column 0 of the prefix matrix. Then, the scan of the current row is not finished until it reaches the ‘1’ guaranteed to exist at column 0 by the output rule (Sec. 2.1). Consequently, the next coordinate L_{k+1} must be in the same row and we have $L_{k+1}^r = L_k^r$, as shown in Fig. 2(c). Moreover, as L_k is not at column 0 of the prefix matrix, L_{k+1} cannot be a diagonal entry. Consequently,

L_{k+1} must correspond to a merge node in the prefix graph. Following (Roy et al. 2021), we restrict the MSP of the merge node corresponding to L_{k+1} to be the node corresponding to L_k . Note that this restriction simplifies the sequence generation task but sacrificing some optimality, and we will leave it as a future work to relax this restriction. By the merge rule in Sec. 2.1, a merge node with its MSP at column L_k^c must have its LSP located at row $(L_k^c - 1)$, which in turn implies that the next column coordinate L_{k+1}^c must be the column of an existing node in row $(L_k^c - 1)$. Formally, we have $L_{k+1}^r = L_k^r$, and $L_{k+1}^c \in \{L_\varphi^c \mid 0 \leq \varphi < k, L_\varphi^r = L_k^c - 1\}$.

Example 1. Consider the example in Fig. 2(c). If the last coordinate in the partial sequence, L_k , is $(1, 0)$, the subsequent coordinate can only be $(L_k^r + 1, L_k^c + 1) = (2, 2)$. Alternatively, suppose $L_k = (3, 3)$. This node acts as the MSP of L_{k+1} , and we can derive that the row coordinate for the LSP of L_{k+1} must be $L_k^c - 1 = 3 - 1 = 2$. Since there are only two nodes at row 2 of the prefix matrix in Fig. 2(c) and their coordinates are $(2, 0)$ and $(2, 2)$, the valid choices for the column coordinate of L_{k+1} are 0 and 2. Hence, the only valid choices for L_{k+1} are $(3, 0)$ and $(3, 2)$.

The above rule enables constructing a *legality mask* that dynamically filters out all invalid choices during each coordinate generation step. Specifically, when sampling the next coordinate L_{k+1} , the mask pre-computes all infeasible options conditioned on $L_{1:k}$ and set their probabilities to zero, thereby preventing invalid sampling. Notably, the legality mask supports a parallel implementation on GPUs. We detail its actual implementation with an example in the appendix.

3.2 Model Architecture

Standard GPT-style models (Radford et al. 2019) typically generate one-dimensional token sequences and lack the ability to directly handle two-dimensional (2D) coordinate data. To overcome this limitation, PrefixGPT introduces two techniques, a *spatial coordinate embedding* and a *two-head Transformer backbone*, which will be introduced next. Fig. 3(a) illustrates the overall architecture.

Spatial Coordinate Embedding The spatial coordinate embedding is designed to effectively encode each node’s

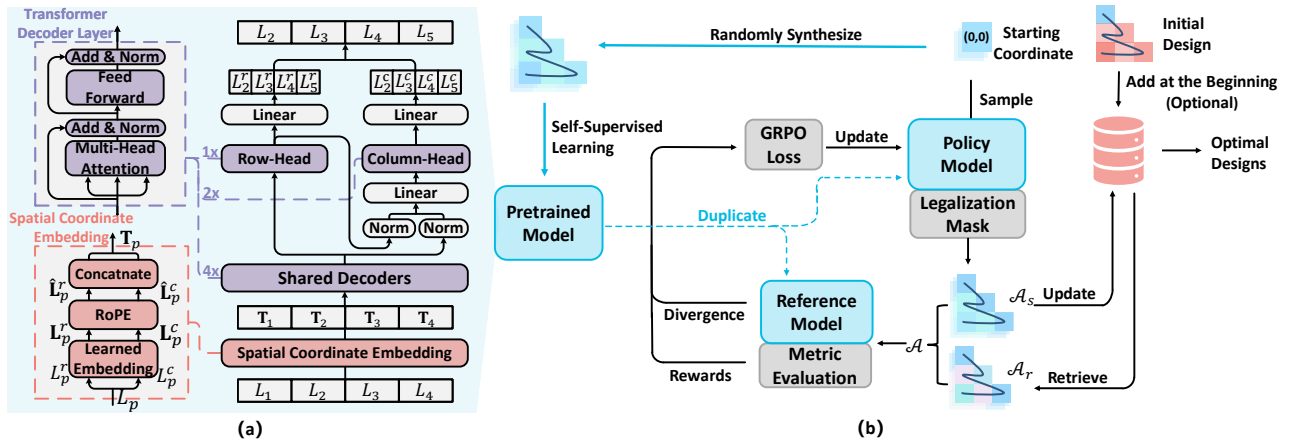


Figure 3: PrefixGPT. (a) Model architecture. (b) Training pipeline: solid blue arrows denote the pre-training phase and the solid black arrows denote fine-tuning phase.

2D coordinates and their relative positional information into unified tokens.

As shown in the bottom-left box in Fig. 3(a), we apply separate learned embedding layers to the row and column indices of each coordinate $L_p = (L_p^r, L_p^c)$, producing two d -dimensional vectors \mathbf{L}_p^r and $\mathbf{L}_p^c \in \mathbb{R}^d$. These are enhanced with the rotary positional embedding (RoPE) technique (Su et al. 2024) to encode relative-position information. Conceptually, RoPE associates each integer position m with an orthogonal rotation matrix $\mathbf{R}(m) \in \mathbb{R}^{d \times d}$, which has the key properties of a rotation: $\mathbf{R}(m)^\top \mathbf{R}(m) = \mathbf{I}$ and, crucially, $\mathbf{R}(m)^\top \mathbf{R}(n) = \mathbf{R}(n - m)$. Applying this to our coordinate embedding $\mathbf{L}_p^{r/c} \in \mathbb{R}^d$,¹ the RoPE-enhanced embedding, denoted as $\hat{\mathbf{L}}_p^{r/c}$, is obtained by a simple matrix-vector product as $\hat{\mathbf{L}}_p^{r/c} = \mathbf{R}(L_p^{r/c}) \mathbf{L}_p^{r/c}$. When computing the inner product between any two RoPE-enhanced embeddings for positions p and q , the rotation matrices combine elegantly:

$$\begin{aligned} \langle \hat{\mathbf{L}}_p^{r/c}, \hat{\mathbf{L}}_q^{r/c} \rangle &= \left(\mathbf{R}(L_p^{r/c}) \mathbf{L}_p^{r/c} \right)^\top \left(\mathbf{R}(L_q^{r/c}) \mathbf{L}_q^{r/c} \right) \\ &= (\mathbf{L}_p^{r/c})^\top \mathbf{R}(L_p^{r/c})^\top \mathbf{R}(L_q^{r/c}) \mathbf{L}_q^{r/c} \\ &= (\mathbf{L}_p^{r/c})^\top \mathbf{R}(L_q^{r/c} - L_p^{r/c}) \mathbf{L}_q^{r/c}, \end{aligned}$$

where $\mathbf{R}(L_q^{r/c} - L_p^{r/c})$ is solely a function of the relative distance between the coordinates. This enables PrefixGPT to effectively capture not only the absolute positions but also the relative distances between merge operations.

Finally, for each coordinate L_p , we concatenate its RoPE-enhanced embeddings to form a fused *coordinate token* $\mathbf{T}_p = \hat{\mathbf{L}}_p^r \parallel \hat{\mathbf{L}}_p^c \in \mathbb{R}^{2d}$. The resulting token sequence $\mathbf{T}_{1:k} = (\mathbf{T}_1, \dots, \mathbf{T}_k)$ is then passed through the two-head Transformer backbone for next coordinate generation.

¹Throughout this paper, we use the superscript r/c as a notational shorthand. It indicates that an expression or operation applies independently to both its row (r) and column (c) counterparts. For example, an entity $X^{r/c}$ represents both X^r and X^c .

Two-Head Transformer Backbone To sequentially predict the next coordinate in the sequence, the Transformer backbone first transforms the spatial coordinate embedding into shared hidden states $\mathbf{H}_{1:k}^{\text{share}}$ through four stacked Transformer decoder layers, called *shared decoder*. These shared states then feed into a *row head* and a *column head*, which predict the row and column coordinates, respectively.

The row head independently refines the shared hidden states to produce the hidden states of the row, $\mathbf{H}_{1:k}^{\text{row}}$, with a Transformer decoder layer as $\mathbf{H}_{1:k}^{\text{row}} = \text{RowHead}(\mathbf{H}_{1:k}^{\text{share}})$. Then, a linear projection is applied to predict the probabilistic distribution of the next row coordinate as follows:

$$P(L_{k+1}^r | L_{1:k}) = \text{Softmax}(\mathbf{W}^{\text{row}} \mathbf{h}_k^{\text{row}}),$$

where $\mathbf{h}_k^{\text{row}}$ is the k -th component of $\mathbf{H}_{1:k}^{\text{row}}$, and we sample the next row coordinate from the obtained distribution.

Since a valid next column coordinate depends on the row coordinate as described in Sec. 3.1, the column head integrates both $\mathbf{H}_{1:k}^{\text{share}}$ and $\mathbf{H}_{1:k}^{\text{row}}$ to generate the hidden states of the column, $\mathbf{H}_{1:k}^{\text{col}}$, through two Transformer decoder layers, which is further processed by a final linear layer to predict the next column coordinate as follows:

$$\begin{aligned} \mathbf{H}_{1:k}^{\text{col}} &= \text{ColumnHead}(\mathbf{W}_p (\mathbf{H}_{1:k}^{\text{share}*} \parallel \mathbf{H}_{1:k}^{\text{row}*})), \\ P(L_{k+1}^c | L_{1:k}) &= \text{Softmax}(\mathbf{W}^{\text{col}} \mathbf{h}_k^{\text{col}}), \end{aligned}$$

where $\mathbf{H}_{1:k}^{\text{share}*}$ and $\mathbf{H}_{1:k}^{\text{row}*}$ denote hidden states after the root mean square normalization (Zhang and Sennrich 2019), and \parallel marks concatenation. The next column coordinate is also sampled from the distribution $P(L_{k+1}^c | L_{1:k})$.

3.3 Pre-training

We first teach the model design rules through pre-training before asking it to generate optimized designs. As illustrated by the solid blue arrows in Fig. 3 (b), we pre-train PrefixGPT using self-supervised learning on a corpus of one million valid coordinate sequences. Each sequence is generated by performing a random walk starting at the same first coordinate $L_1 = (0, 0)$ and proceeding by randomly selecting a

valid next coordinate until the EOS coordinate $(n - 1, 0)$ is reached for an n -bit adder.

During pre-training, PrefixGPT learns to predict the next valid coordinate by minimizing the loss between its predictions and the ground-truth labels. For an input sequence $L_{1:k}$, its loss sums the individual cross-entropy loss from both row and column as follows:

$$\mathcal{L}_{\text{pre}} = -\frac{1}{k} \sum_{p=1}^k (\log P(L_p^r | L_{1:p-1}) + \log P(L_p^c | L_{1:p-1})).$$

The pre-training is both efficient and transferable. The training corpus can be synthesized on a single GPU with roughly one minute under our implementation settings, and once PrefixGPT is pre-trained on a maximal bit-width n , it can be instantly adapted to any m -bit tasks ($m \leq n$), simply by resetting the EOS coordinate from $(n - 1, 0)$ to $(m - 1, 0)$. For this work, PrefixGPT was pre-trained on sequences up to $n = 48$, due to GPU memory constraints.

3.4 Fine-tuning for Optimized Adder Generation

The fine-tuning process is indicated by the solid black arrows in Fig. 3(b). In this stage, we employ RL to align the output of the pre-trained model with our design objective, i.e., an optimized area-delay trade-off, and continuously exploring the solution space to discover higher-quality adders.

Fine-tuning begins by duplicating the pre-trained model into two copies: a trainable policy model π_θ and a frozen reference model π_{ref} . The policy model evolves to discover new designs, and the reference model preserves the pre-training behavior to avoid performance catastrophic collapse (Schulman et al. 2017; Stiennon et al. 2020). At each RL iteration, π_θ samples a group of G valid adders $\mathcal{A}_s = (\alpha_1, \dots, \alpha_G)$ using the legality mask, as shown in Fig. 3(b). Each adder α_i can be represented by a coordinate sequence $L_{1:N_i}$, where N_i is its length. We evaluate each α_i with a scalar reward set as the negative area-delay product (ADP) $r_i = -\text{area}(\alpha_i) \times \text{delay}(\alpha_i)$. We choose this reward since a commonly used target for circuit optimization is minimizing the ADP. In parallel, π_{ref} re-scores the same sequences to compute a Kullback-Leibler (KL) divergence \mathbb{D}_{KL} (Schulman et al. 2017) between the policy model and the reference model. We extend its approximate form for both the row and the column factors:

$$\begin{aligned} \mathbb{D}_{\text{KL}}(L_p) &= \mathbb{D}_{\text{KL}}^r(L_p) + \mathbb{D}_{\text{KL}}^c(L_p), \\ \mathbb{D}_{\text{KL}}^{r/c}(L_p) &= \frac{\pi_\theta^{r/c}(L_p)}{\pi_{\text{ref}}^{r/c}(L_p)} - \log \frac{\pi_\theta^{r/c}(L_p)}{\pi_{\text{ref}}^{r/c}(L_p)} - 1, \\ \pi_\theta^{r/c}(L_p) &= P_\theta(L_p^{r/c} | L_{1:p-1}), \\ \pi_{\text{ref}}^{r/c}(L_p) &= P_{\text{ref}}(L_p^{r/c} | L_{1:p-1}), \end{aligned}$$

where P_θ and P_{ref} represent the prediction probability given by the policy model and reference model, respectively.

We adopt the lightweight group relative policy optimization (GRPO) (Shao et al. 2024) method for efficient training. Compared with proximal policy optimization (Schulman et al. 2017), GRPO prevents the expensive advantage

model by directly calculating the advantage \hat{A}_i with a normalized reward as $\hat{A}_i = \frac{r_i - \mu_{\mathcal{A}}}{\sigma_{\mathcal{A}}}$, where $\mu_{\mathcal{A}}$ and $\sigma_{\mathcal{A}}$ are the mean and the standard deviation, respectively, of the rewards of all adders in a sample group \mathcal{A} . We optimize the policy model by maximizing the following objective function:

$$\mathcal{J}(\theta) = \frac{1}{G} \sum_{i=1}^G \frac{1}{N_i} \sum_{p=1}^{N_i} \left\{ \gamma^p s_\theta(L_p) \hat{A}_i - \beta \mathbb{D}_{\text{KL}}(L_p) \right\}, \quad (1)$$

where $s_\theta(L_p) = \pi_\theta^r(L_p) + \pi_\theta^c(L_p)$, and $\gamma = 0.99$ and $\beta = 0.001$ are constants.

We also propose a *best-design retrieval* mechanism to reuse advantageous designs and reinforce high-quality patterns during training. Specifically, at each iteration, all sampled designs are logged into a database. We retrieve from the global database of all previously generated designs a small subset \mathcal{A}_r of designs with the lowest ADP, where the subset size is limited to at most 10% of the total batch size. The retrieved designs are then merged with the newly sampled batch \mathcal{A}_s to form $\mathcal{A} = \mathcal{A}_s \cup \mathcal{A}_r$ for the GRPO update.

The best-design retrieval also supports the use of good manual designs for initialization. Specifically, we can add good manual designs into the database at the beginning of fine-tuning. Then, those retrieved with high initial quality typically dominate the GRPO advantage calculations, gently steering the policy toward learning their superior architectures.

4 Experiment Results

4.1 Experiment Setup

We implemented PrefixGPT in PyTorch with CUDA 12.6 and ran the experiments on a single NVIDIA RTX 4090 GPU and an AMD EPYC 9374F CPU. The model used an embedding dimension of $d = 128$. We pre-trained the model for 5 epochs on a corpus of 1,000,000 randomly synthesized valid coordinate sequences using the Adam optimizer with a learning rate of 10^{-4} . For fine-tuning, we conducted RL for 200 iterations, and we sampled $G = 512$ valid designs per iteration with a temperature of 0.8. The legality mask was applied through all fine-tuning iterations.

We evaluated PrefixGPT on four adder bit-widths, 16, 24, 32, and 48, and comprehensively compared it with two SOTA methods, PrefixRL (Roy et al. 2021), which uses deep Q-learning, and ArithTree (Lai et al. 2024), an MCTS-based method. To test whether PrefixGPT overcomes initialization bias, we evaluate all methods with four initializations: three manual designs, Sklansky (Sk.), Kogge-Stone (Ko.), and Brent-Kung (Br.), and a random (Ra.) setting. In the random setting, ArithTree and PrefixRL started searching with a randomly generated valid design. In contrast, PrefixGPT began its search with no initial design in its database. To ensure a fair comparison, we evaluated each combination of bit-width and initial design by running 10 independent trials, and each trial was strictly capped at 100,000 candidate design evaluations, resulting in a total of 1,000,000 evaluations per combination for each method.

We also compared PrefixGPT with a generative baseline, PrefixLLM (Xiao et al. 2024). However, since it only scales

to 16-bit adders, and its generation time per design is over 200 seconds, which is excessively long, we used the data from its authors and the original paper for comparison.

4.2 Comparison of Prefix Graph Size

As mentioned in Sec. 2.1, the size and depth of a prefix graph reflect the area and delay, respectively, of the corresponding prefix adder. Given the inherent trade-off between size and depth, following (Lai et al. 2024), we set three strict depth limits for each bit-width and compare the minimum achievable size under these limits. Specifically, for each bit-width, the depth limits are h , $h + 1$, and $h + 2$, where h is the minimum possible depth for an adder of the given bit-width (Snir 1986). As shown in Table 1, PrefixGPT achieves the best size on all 12 pairs of bit-width and depth limit, with a maximum size reduction of 59.1% on 48-bit with Kogge-Stone initialization. Moreover, prior methods are highly sensitive to initialization and may produce much larger or even invalid designs (indicated by ‘-’) when tested under tight depth limits, as even a minor structural modification can easily violate the depth limit. Additionally, for PrefixLLM, the minimum sizes obtained under Ko. initialization of 16-bit adders are 31, 28, and 24 for depths 5, 6, and 7, respectively, which falls behind PrefixGPT.

Bit-Width	Depth	ArithTree Size				PrefixRL Size				PrefixGPT Size			
		Sk.	Ko.	Br.	Ra.	Sk.	Ko.	Br.	Ra.	Sk.	Ko.	Br.	Ra.
16-bit	5	31	45	-	-	-	-	-	-	31	31	31	32
	6	25	25	25	25	26	45	26	-	25	25	25	25
	7	24	24	24	24	25	27	24	30	24	24	24	24
	Average:	27	31	25	25	26	36	25	30	27	27	27	27
24-bit	6	45	78	-	-	88	-	-	-	45	50	50	64
	7	41	49	-	51	83	-	-	42	40	40	42	
	8	39	39	39	40	48	76	-	65	39	39	39	40
	Average:	42	55	39	40	50	82	-	65	42	43	43	49
32-bit	6	77	124	-	-	79	-	-	-	74	80	76	78
	7	65	108	-	77	128	-	-	57	61	57	60	
	8	57	77	55	65	68	128	56	-	55	55	55	55
	Average:	66	103	55	65	75	128	56	-	62	65	63	64
48-bit	7	119	217	-	-	224	-	-	102	94	98	102	
	8	105	191	-	123	222	-	-	99	87	88	100	
	9	94	155	-	120	207	-	-	86	86	86	86	
	Average:	106	188	-	122	218	-	-	96	89	91	96	

Table 1: Prefix graph sizes obtained by various methods. A ‘-’ indicates that the size is not found for a method and is excluded when computing the rounded average.

4.3 Comparison of Area and Delay

To align with industrial design practice, we evaluated the synthesized results by measuring their post-synthesis area and delay. This was done with the open-source logic synthesis tool ABC (Berkeley Logic Synthesis and Verification Group 2008) using Nangate45 Open Cell Library (Nangate Inc. 2008) as the standard cell library.

Fig. 4 plots the ADP distributions, aggregated over all four initializations for each bit-width and each method. At

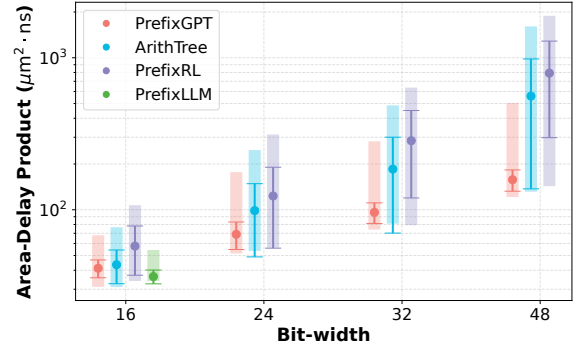


Figure 4: ADP comparison of various methods. In the figure, the vertical strips span the min-max range, while the markers show the mean ADP with ± 1 standard deviation.

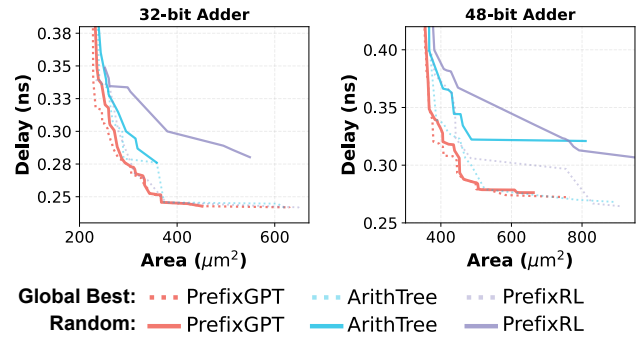


Figure 5: Prefix adder optimization results after ABC synthesis. In each plot, solid lines indicate the Pareto frontiers achieved from random initializations and the dashed lines represent the global Pareto frontier, aggregated from the designs across all Sk., Ko., Br. and Ra. initialization strategies.

16-bit, PrefixGPT achieved an minimum ADP of $31.13 \mu\text{m}^2 \cdot \text{ns}$, better than that of PrefixLLM ($34.01 \mu\text{m}^2 \cdot \text{ns}$). PrefixGPT’s superiority was most evident at higher bit-widths, excelling in both performance and robustness. At 48-bit, it discovered a design with a minimum ADP of $121.3 \mu\text{m}^2 \cdot \text{ns}$, which was a 7.7% and 15.1% improvement over the minima found by ArithTree and PrefixRL, respectively. Averaged over all designs, PrefixGPT achieves a mean ADP of $91.14 \mu\text{m}^2 \cdot \text{ns}$ for 32-bit adders and $185.15 \mu\text{m}^2 \cdot \text{ns}$ for 48-bit adders, representing a 71.9% and 79.1% reduction over the mean ADP of ArithTree. Moreover, PrefixGPT’s designs exhibit substantially greater stability. Its standard deviation in ADP is consistently lower, and at 48-bit, it shows a reduction of over 94% compared to ArithTree (25.2 vs. 422.2).

This robustness is further demonstrated in Fig. 5. Even under random initializations (solid lines), PrefixGPT yields designs that closely match its global Pareto frontier (dashed lines), aggregated over all initialization settings. Remarkably, these random-initialized results often outperform the best global frontiers achieved by the other methods.

PrefixGPT runs very efficiently, taking approximately 7 ms to generate one design for the 32-bit task. We report the

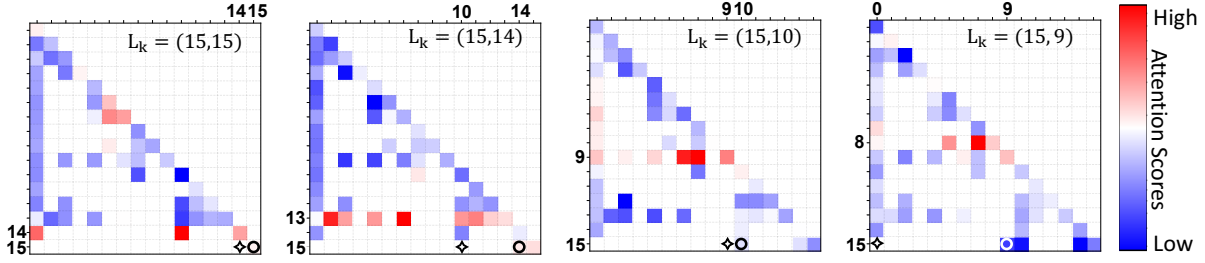


Figure 6: The model’s attention scores for the final four generation steps of an unseen 16-bit adder. To predict the next cell, L_{k+1} (starred), the model correctly focuses its attention on the potential LSPs at row $(L_k^c - 1)$. Note that L_k is denoted in circle.




Figure 7: Ablation studies. Solid lines denote the mean average and the color stripes represent ± 1 standard deviation.

full wall-clock runtime in appendix.

4.4 Investigation of the Effectiveness

To investigate the sources of PrefixGPT’s effectiveness, we repeated the 32-bit random optimization 20 times with different components ablated:

- **In model architecture**, the RoPE block was removed from the model architecture, leaving only the learned embedding block to encode absolute positional information.
- **In model training**, the entire pre-training phase was omitted. Instead, a model with randomly initialized parameters was used for fine-tuning.
- **In fine-tuning**, the KL-divergence regularization term was removed from the objective function shown in Eq. (1), and the best-design retrieval was disabled.

Fig. 7 shows the normalized average reward over the 20 trials. The removal of the pre-training phase and RoPE proved most detrimental, respectively resulting in 31.6% and 16.6% lower final average reward. While ablating the KL-divergence and the best-design retrieval strategy led to a 1.3% decrease in final average reward, their removal had a dramatic impact on stability, increasing the standard deviation from 0.018 to 0.038 (+110%) at the final iteration.

These results suggest that PrefixGPT’s strong performance primarily arises from its pre-training phase, where the model may internalize the fundamental design principles. We tested this hypothesis by removing the legality mask and prompting the pre-trained model to autoregressively generate designs from the start coordinate. As shown in Fig. 8,




Figure 8: Generation legal rate after pre-training with varying sample sizes, evaluated without the legality mask.

the generation legal rate (the proportion of valid designs in each generated batch) rapidly converges to nearly 100% with sufficient training, demonstrating the model’s learned ability to adhere to design rules and generate structurally valid adders. An inspection of the model’s internal mechanism provides further confirmation. As shown in Fig. 6, the attention scores of the column head demonstrated a remarkable alignment with the merge rule described in Sec. 3.1. For example, when $L_k = (15, 14)$, the model’s attention primarily focuses on coordinates in row $L_k^c - 1 = 13$, which are the valid candidates for the LSP of L_{k+1} . This combined evidence, from both external behavior and internal mechanism, suggests that during pre-training, PrefixGPT moves beyond pattern memorization to develop a foundational understanding of the design space, which likely enables its robust and effective search.

5 Conclusion

We introduce PrefixGPT, a generative pre-trained Transformer that presents a new paradigm for prefix adder optimization. Instead of relying on structural refinements, PrefixGPT directly generates optimized prefix-adder topologies from scratch through a sequence-based formulation. Extensive experiments demonstrate that our approach achieves state-of-the-art performance, discovering novel designs with improved area–delay product and substantially reduced variance across initializations. These results show that GPT-style models can effectively handle complex hardware-design constraints and autonomously generate high-quality circuits beyond current heuristic or RL methods.

Acknowledgments

This work is supported by the National Key R&D Program of China under grant number 2021ZD0114701.

References

- Berkeley Logic Synthesis and Verification Group. 2008. ABC: A System for sequential synthesis and verification. <https://people.eecs.berkeley.edu/~alanmi/abc/>. Versioned open-source software.
- Brent; and Kung. 1982. A regular layout for parallel adders. *IEEE Transactions on Computers*, 100(3): 260–264.
- Bucaioni, A.; Ekedahl, H.; Helander, V.; and Nguyen, P. T. 2024. Programming with ChatGPT: How far can we go? *Machine Learning with Applications*, 15: 100526.
- Guo, D.; Yang, D.; Zhang, H.; Song, J.; Zhang, R.; Xu, R.; Zhu, Q.; Ma, S.; Wang, P.; Bi, X.; et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Kogge, P. 1973. A parallel algorithm for the efficient solution of general class of recurrence equations. *IEEE Transactions on Computers*, 22(8): 786–792.
- Ladner, R. E.; and Fischer, M. J. 1980. Parallel prefix computation. *J. ACM*, 27(4): 831–838.
- Lai, Y.; Liu, J.; Pan, D. Z.; and Luo, P. 2024. Scalable and effective arithmetic tree generation for adder and multiplier designs. In *Advances in Neural Information Processing Systems*, volume 37, 139151–139178. Curran Associates, Inc.
- Luo, R.; Sun, L.; Xia, Y.; Qin, T.; Zhang, S.; Poon, H.; and Liu, T.-Y. 2022. BioGPT: Generative pre-trained transformer for biomedical text generation and mining. *Briefings in Bioinformatics*, 23(6): bbac409.
- M. Bran, A.; Cox, S.; Schilter, O.; Baldassari, C.; White, A. D.; and Schwaller, P. 2024. Augmenting large language models with chemistry tools. *Nature Machine Intelligence*, 6(5): 525–535.
- Ma, Y.; Roy, S.; Miao, J.; Chen, J.; and Yu, B. 2018. Cross-layer optimization for high speed adders: A pareto driven machine learning approach. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 38(12): 2298–2311.
- Moto, T.; and Kaneko, M. 2018. Prefix sequence: Optimization of parallel prefix adders using simulated annealing. In *2018 IEEE International Symposium on Circuits and Systems (ISCAS)*, 1–5. IEEE.
- Nangate Inc. 2008. Open Cell Library v2008 10 SP1. <http://www.nangate.com/openlibrary/>.
- Radford, A.; Narasimhan, K.; Salimans, T.; Sutskever, I.; et al. 2018. Improving language understanding by generative pre-training. Technical report, OpenAI.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I.; et al. 2019. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8): 9.
- Roy, R.; Raiman, J.; Kant, N.; Elkin, I.; Kirby, R.; Siu, M.; Oberman, S.; Godil, S.; and Catanzaro, B. 2021. PrefixRL: Optimization of parallel prefix circuits using deep reinforcement learning. In *2021 58th ACM/IEEE Design Automation Conference (DAC)*, 853–858. IEEE.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Seshakagari, H. R. B.; Umashankar, A.; Harikala, T.; Jayasree, L.; and Severance, J. 2025. Dynamic financial sentiment analysis and market forecasting through large language models. *International Journal of Human Computations & Intelligence*, 4(1): 397–410.
- Shao, Z.; Wang, P.; Zhu, Q.; Xu, R.; Song, J.; Bi, X.; Zhang, H.; Zhang, M.; Li, Y.; Wu, Y.; et al. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Sklansky, J. 1960. Conditional-sum addition logic. *IRE Transactions on Electronic computers*, EC-9(2): 226–231.
- Snir, M. 1986. Depth-size trade-offs for parallel prefix computation. *Journal of Algorithms*, 7(2): 185–201.
- Song, J.; Swope, A.; Kirby, R.; Roy, R.; Godil, S.; Raiman, J.; and Catanzaro, B. 2024. CircuitVAE: Efficient and scalable latent circuit optimization. In *Proceedings of the 61st ACM/IEEE Design Automation Conference*, 1–6.
- Stiennon, N.; Ouyang, L.; Wu, J.; Ziegler, D.; Lowe, R.; Voss, C.; Radford, A.; Amodei, D.; and Christiano, P. F. 2020. Learning to summarize with human feedback. *Advances in Neural Information Processing Systems*, 33: 3008–3021.
- Su, J.; Ahmed, M.; Lu, Y.; Pan, S.; Bo, W.; and Liu, Y. 2024. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568: 127063.
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; Rodriguez, A.; Joulin, A.; Grave, E.; and Lample, G. 2023. LLaMA: Open and efficient foundation language models. *arXiv:2302.13971*.
- Xiao, W.; Putrevu, V. S. C.; Hemadri, R. V.; Garg, S.; and Karri, R. 2024. PrefixLLM: LLM-aided prefix circuit design. *arXiv preprint arXiv:2412.02594*.
- Zhang, B.; and Sennrich, R. 2019. Root mean square layer normalization. *Advances in Neural Information Processing Systems*, 32.