ANALOG CIRCUIT TOPOLOGY DESIGN AND SIZING WITH FLOW MATCHING GRAPH LEARNING

Anonymous authors

Paper under double-blind review

ABSTRACT

The soaring demand for electronic devices calls for novel and more efficient analog circuits design. Deep generative models have shown promise in assisting topology, parameter sizing, and layout design process, but existing approaches treat these tasks separately and lack generalizability across diverse problem settings. In this work we introduce a flow matching model for automatic analog circuit design, which achieves high-quality sampling across a variety of topologies and representations. Our model showcases state-of-the-art performance on end-to-end topology design and sizing on the Open Circuit Benchmark (OCB) dataset, and on transistor-level topology generation on the AnalogGenie dataset. Code and models are provided as external supplementary files to this submission.

1 Introduction

The automation of analog circuit design stands as an active area of research, driven both by the demand for increasingly efficient architectures to sustain the growth of the electronics industry and by the intrinsic complexity of the task, which is notoriously more challenging than digital circuit design due to its greater diversity of components. Accordingly, the literature presents a wide range of data-driven approaches aimed at automating one or more steps of the analog design workflow, which traditionally includes topology discovery (Lu et al., 2022; 2023; Poddar et al., 2024), parameter sizing (Wang et al., 2018; 2020; Krylov et al., 2023), and layout prediction (Kunal et al., 2019; Xu et al., 2019; Liu et al., 2025a).

Despite significant progress, several hurdles remain. Many methods exhibit limited generalizability, restricting their applicability to a small set of circuit topologies. Others rely on multiple models trained for different subtasks or require substantial computational resources. The absence of widely adopted benchmarks and open models has also often been cited as a limiting factor for faster advancement in the field. This issue has been partly addressed by the recent release of benchmarks and models targeting topology generation and device sizing (Dong et al., 2023; Gao et al., 2025), enabling more systematic comparisons.

In terms of model architectures, the long-standing paradigm of representing circuits as graphs (Ren et al., 2020; Wang et al., 2020; Hakhamaneshi et al., 2022) now coexists with the recent adoption of Large Language Models (LLM)-based methodologies (Yin et al., 2024; Liu et al., 2024a;b), which harness the exceptional ability of LLMs for sequence modeling to generate circuit design as textual outputs (Chang et al., 2024; Lai et al., 2025). There is however still room for improvement, while the generalizability of these methods beyond pre-defined settings remains an open question.

We argue that graph-based representations of circuits holds untapped potential to address these limitations. In particular, recent progress in generative modeling of graphs using denoising diffusion (Vignac et al., 2022) and flow matching (Eijkelboom et al., 2024; Qin et al., 2025) lets us foresee promising applications for analog circuit design. These models are notorious for their high sample quality (Esser et al., 2024), and can accommodate the conditional generation of multimodal data, opening the door to applications such as circuit completion or parameter sizing within a single architecture. Diffusion models have already proved successful for device sizing (de Azevedo et al., 2025; Eid et al., 2025) and topology discovery (Liu et al., 2025b), but for a limited scope of circuits. Further, no end-to-end framework has been proposed to date that jointly tackles these tasks.

In this work, we introduce a multimodal flow matching model, **CircuitFlow**, for end-to-end generation of analog circuit topology and device sizing. Built on a graph transformer backbone (Ma et al., 2023), it shows remarkable sampling quality and allows a very fine control on the denoising process through a modality-dependent time sampling scheme. We evaluate our approach on established benchmarks: first at the behavioral level on the OCB dataset (Dong et al., 2023), and then at the transistor level on the AnalogGenie dataset (Gao et al., 2025) with a separate model for the prediction of port connectivity, achieving in both cases state-of-the-art improvement in the quality of generated circuits. The main contribution of this work is the extension of discrete flow matching to analog circuit design, yielding a unified framework that jointly addresses topology generation and device sizing across multiple representation levels and circuit complexities This is primarily enabled by a novel dimension- and modality- dependent time sampling scheme, which grants unprecedented flexibility for diverse inference-time applications without the need for additional training. The remainder of this paper is structured as follows. Related work is discussed in Sections 2. Section 3 introduces the theoretical foundations of this work. The proposed approach is detailed in Section 4. Experimental results are presented in Section 5 and future research directions in Section 6.

2 RELATED WORK

2.1 DATA-DRIVEN TOPOLOGY DESIGN AND SIZING OF ANALOG CIRCUITS

Topology generation. Data-driven approaches for topology design include Reinforcement Learning (RL) (Fan et al., 2021; Zhao & Zhang, 2022), Bayesian Optimization (BO) in the continuous latent space of a Variational Auto-Encoder (VAE) (Lu et al., 2022; 2023; Dong et al., 2023; Shen et al., 2024), and retrieval from predefined building blocks (Fayazi et al., 2023; Poddar et al., 2024). RL and BO methods often suffer from slow convergence, while retrieval-based strategies depend heavily on the completeness of predefined architectures and typically lack flexibility. Recent works leverage pre-trained LLMs to generate topologies as text output (Chang et al., 2024; Lai et al., 2025), but have not yet scaled beyond a limited set of circuit types and complexity. AnalogGenie (Gao et al., 2025) demonstrates strong scalability to diverse, transistor-level topologies, but its GPT-based backbone requires extensive data augmentation to enforce permutation invariance over input graph nodes.

Device Sizing. Parameter sizing, whether at the behavior or transistor level, has been widely explored using RL (Wang et al., 2018; Settaluri et al., 2020; Wang et al., 2020; Cao et al., 2022; Gao et al., 2023; Cao et al., 2024), BO (Lyu et al., 2018), supervised learning (Hakhamaneshi et al., 2022; Krylov et al., 2023) or LLMs (Yin et al., 2024; Liu et al., 2024a;b). Some works aim to address both topology design and sizing (Fayazi et al., 2023; Lu et al., 2023; Liu et al., 2025b), but do so in several stages with separate, dedicated models.

2.2 FLOW MATCHING FOR GRAPH GENERATION

Flow matching models (Lipman et al., 2023; Esser et al., 2024; Tong et al., 2024) have emerged as a sample-efficient alternative to diffusion models (Sohl-Dickstein et al., 2015; Ho et al., 2020). They have been extended to discrete state-space (Campbell et al., 2024; Gat et al., 2025) and proved successful on diverse graph generation tasks (Eijkelboom et al., 2024; Qin et al., 2025). To date, flow matching has not been applied to analog circuit design. A handful of methods have explored the use of denoising diffusion for sizing and topology discovery (de Azevedo et al., 2025; Eid et al., 2025; Liu et al., 2025b), but remain restricted to a narrow range of circuit topologies.

3 Preliminaries

3.1 Continous Flow Matching

The objective of continuous flow matching is to learn an approximation function u_t^{θ} of a quantity u_t called a *vector field*, which, given an arbitrary dimension d, is a function of \mathbb{R}^d in itself that describes the instantaneous change of a *flow* x_t with respect to a *time dimension*:

$$u_t(x_t) = \frac{dx_t}{dt}$$
, with $x_t \in \mathbb{R}^d$. (1)

Given time-varying probability distribution p_t , u_t is said to generate the probability path p_t if x_t is a random variable that follows p_t , where the prior p_0 is typically a standard Gaussian or uniform distribution, and p_1 is the unknown data distribution. This is expressed by the continuity equation which links u_t to p_t :

$$\frac{dp_t(x_t)}{dt} = -\operatorname{div}(p_t(x_t)u_t(x_t)),\tag{2}$$

where $\operatorname{div}(f(x)) = \sum_i \frac{\partial f(x)_i}{\partial x_i}$ is the divergence operator. Ideally, for a given p_t that satisfies the continuity equation and a trained approximation function u_t^θ , sampling from the data distribution can be achieved by drawing x_0 from the prior and solving Equation (1) up to t=1. In practice, one cannot express u_t and p_t directly, but may instead define them as the expectations of a *conditional* path and velocity field over the data distribution p_1 (Lipman et al., 2023):

$$p_t(x_t) = \int p_t(x_t|x_1)p_1(x_1)dx_1,$$
(3)

$$u_t(x_t) = \int u_t(x_t|x_1) \frac{p_t(x_t|x_1)p_1(x_1)}{p_t(x_t)} dx_1.$$
(4)

This definition allows two important results. The first corollary is that if $u_t(x_t|x_1)$ generates $p_t(x_t|x_1)$, then u_t generates p_t . Hence it is enough to define a conditional velocity field and probability path that respect the continuity equation, which is a much easier task. Second, the same parameter set θ minimizes the following objectives:

$$\mathcal{L}_{\text{FM}} = \mathbb{E}_{t, x_t \sim p_t} \Big[\|u_t^{\theta}(x_t) - u_t(x_t)\|_2^2 \Big], \text{ and } \mathcal{L}_{\text{CFM}} = \mathbb{E}_{t, x_1, x_t \sim p_t(x_t | x_1)} \Big[\|u_t^{\theta}(x_t) - u_t(x_t | x_1)\|_2^2 \Big].$$
(5)

As \mathcal{L}_{CFM} offers a tractable objective, it is therefore enough to reason in terms of conditional quantities. In summary, if $p_t(x_t|x_1)$ and $u_t(x_t|x_1)$ are chosen adequately such that $u_t(x_t|x_1)$ generates $p_t(x_t|x_1)$, then minimizing Equation (5) amounts to fitting a neural network u_t^{θ} which generates p_t , i.e., which may then be used to sample from p_1 . Lipman et al. (2023) propose to write the conditional probability path as a Gaussian:

$$p_t(x_t|x_1) = \mathcal{N}(x_t; \mu_t(x_1), \sigma_t^2(x_1)I), \tag{6}$$

with $\mu_0(x_1)=0$, $\mu_1(x_1)=x_1$, $\sigma_0(x_1)=1$, and to write the flow x_t as $x_t=\sigma_t(x_1)x_0+\mu_t(x_1)$ where $x_0\sim p_0$. When $\sigma_t\to 0$, and $\mu_t(x_1)=tx_1+(1-t)x_0$, one obtains the well-known rectified flow (Liu, 2022):

$$x_t = tx_1 + (1 - t)x_0, (7)$$

$$u_t(x_t|x_1) = x_1 - x_0 = \frac{x_1 - x_t}{1 - t},$$
 (8)

which can be used in Equation (5) to compute \mathcal{L}_{CFM} . Once trained, the model can be used to draw from p_1 starting from a noise sample $x_0 \sim p_0$ and following denoising Euler steps Δ_t :

$$x_{t+\Delta_t} = x_t + \Delta_t u_t^{\theta}(x_t). \tag{9}$$

3.2 DISCRETE FLOW MATCHING

One approach (Campbell et al., 2024) to modeling discrete data $x_1 \in [1, \dots, S]^D$, where each dimension of x_1 can take S different states, is to consider the whole flow x_t as discrete, and allow state transitions to occur *one dimension at a time*. This translates in the following factorization of $p_{t+\Delta t}(x_{t+\Delta t}|x_t)$:

$$p_{t+\Delta t}(x_{t+\Delta t}|x_t) = \prod_{d} p_{t+\Delta t}(x_{t+\Delta t}^d|x_t).$$
(10)

This allows to define the generative process using a rate matrix $R_t \in \mathbb{R}^{S \times S}$ which characterizes state transition over single dimensions and replaces the velocity field u_t from the continuous setting, as can be seen from the denoising process:

$$x_{t+\Delta t}^d \sim \operatorname{Cat}(\delta\{x_t^d, x_{t+\Delta t}^d\} + \Delta t \, R_t(x_t^d, x_{t+\Delta t}^d)). \tag{11}$$

The rate matrix satisfies $R_t(i,j) \ge 0$ if $i \ne j$ and $R_t(i,i) = -\sum_{j\ne i} R_t(i,j)$. As in the continuous case, R_t must satisfy the continuity equation (known as the Kolmogorov equation in the discrete

case, see Gat et al. (2025)) to ensure that it generates p_t , and that solving Equation (11) amounts to sampling from the data distribution p_1 . Again, it is more convenient to define a *conditional rate* matrix $R_t(x_t, x_{t+\Delta t}|x_1)$ that generates the conditional distribution $p_{t|1}$, and that can be used to recover the marginal rate matrix R_t through the following expectation:

$$R_t(x_t, j) = \mathbb{E}_{p_{1|t}(x_1|x_t)} R_t(x_t, j|x_1). \tag{12}$$

This time however the conditional rate matrix can be computed in closed form, and one instead aims to learn the posterior probabilities $p_{1|t}(x_1|x_t)$. In the multivariate case, each dimension is learned separately, such that the training objective writes:

$$\mathcal{L}_{\text{DFM}} = -\mathbb{E}_{t,x_1,x_t} \sum_{d} \log(p_{1|t}^{\theta}(x_1^d|x_t)), \tag{13}$$

with $p_{1|t}^{\theta}$ the approximation function. Noised vector x_t is sampled per dimension by interpolating between x_1^d and a prior, that we take here equal to the product of marginal probability mass functions over all states, simply written Cat(m), with $m \in \mathbb{R}^S$ (see Appendix A for the computation of m):

$$x_t^d \sim \text{Cat}(t\delta\{x_1^d, x_t^d\} + (1 - t) \times m)).$$
 (14)

Inference is done independently from training, which does not require access to the conditional rate matrix. For the latter, Campbell et al. (2024) introduce the following expression:

$$R_{t}(x_{t}^{d}, x_{t+\Delta t}^{d} = j | x_{1}^{d}) = \frac{\text{ReLU}(\partial_{t} p_{t|1}(x_{t+\Delta t}^{d} = j | x_{1}^{d}) - \partial_{t} p_{t|1}(x_{t}^{d} | x_{1}^{d})}{S.p_{t|1}(x_{t}^{d} | x_{1}^{d})} + R^{\text{DB}}(x_{t+\Delta t}^{d} = j | x_{1}^{d}),$$
(15)

where the *detailed balance* term R^{DB} allows an adjustable level of stochasticity (see Appendix A for the derivation of both terms for our choice of prior distribution). The expectation over x_1^d can be derived in closed form to give the final expression of the marginal rate matrix R_t (see Appendix B):

$$R_{t}(x_{t}^{d}, j) = \frac{\left(1 - m_{j} + m_{x_{t}^{d}}\right)}{S(1 - t)m_{x_{t}^{d}}} p_{1|t}^{\theta}(x_{1}^{d} = j|x_{t})$$

$$+ \frac{\text{ReLU}\left(m_{x_{t}^{d}} - m_{j}\right)}{S(1 - t)m_{x_{t}^{d}}} (1 - p_{1|t}^{\theta}(x_{1}^{d} = j|x_{t}) - p_{1|t}^{\theta}(x_{1}^{d} = x_{t}^{d}|x_{t}))$$

$$+ \eta p_{1|t}^{\theta}(x_{1}^{d} = x_{t}^{d}|x_{t}) + \eta \frac{t + (1 - t)m_{x_{t}^{d}}}{(1 - t)m_{j}} p_{1|t}^{\theta}(x_{1}^{d} = j|x_{t}), \tag{16}$$

where $\eta \in \mathbb{R}^+$ is the tunable noise level.

Multimodal Flows. In the case of a multimodal flow (x_t, y_t) , Campbell et al. (2024) showed that if the noising process $p_{t|1}(x_t, y_t|x_1, y_1)$ factorizes over its variables such that:

$$p_{t|1}(x_t, y_t|x_1, y_1) = \prod_{d}^{D_x} p_{t|1}(x_t^d|x_1^d) \prod_{d}^{D_y} p_{t|1}(y_t^d|y_1^d), \tag{17}$$

then the process composed separately of R_t^x (respectively u_t^x if x_t is continuous) and R_t^y (resp. u_t^y), as defined above, generates the marginal multimodal flow $p_t(x_t, y_t)$. This allows to sample t independently for each variable, which enables a remarkable flexibility of the denoising process.

4 GRAPH FLOW MATCHING FOR ANALOG TOPOLOGY DESIGN AND SIZING

For the remainder of this work, we represent circuits as undirected graphs $\mathcal G$ composed of a set of D_v nodes $\mathcal V$, a set of edges $\mathcal E\subseteq \mathcal V\times \mathcal V$ that connect the nodes, and when applicable, a node feature vector $\mathcal F$ that provides component sizes. Individual node elements and node features are respectively noted v^d and f^d $\forall d \leq D_v$, and take values in $\{1,\ldots,S\}$ and $\mathbb R$, respectively, where S is the total number of node types. Individual edges are noted e^d $\forall d \leq D_e = D_v \times D_v$ and take values in $\{0,1\}$. The objective of this work is to train a multimodal flow matching model to sample from the data distribution $p_1(\mathcal G)$. This section first defines this flow along with its training process. We then describe the architecture of the employed graph transformer model, and give details about the chosen data representation.

4.1 MULTIMODAL FLOW FOR FLEXIBLE CIRCUIT MODELING

The flow we consider here is composed of two discrete variables, node types \mathcal{V}_t and edges \mathcal{E}_t , along with the continuous device sizes \mathcal{F}_t . As before, the noising process will factorize over variables, and as suggested by Campbell et al. (2024) the noising time index will be sampled independently for each variable, yielding respectively t_v , t_e and t_f for node types, edges and features. Here however we go one step further and sample time independently for each dimension, such that $t_v, t_f \in [0,1]^{D_v}$ and $t_e \in [0,1]^{D_e/2}$ (graphs are undirected and only half of the edges are modeled). As we shall see shortly, this makes the sampling process particularly flexible, allowing a whole range of key applications. The noising process thus writes, using the notations $t = (t_v, t_e, t_f)$ and $\mathcal{G}_t = (\mathcal{V}_{t_v}, \mathcal{E}_{t_e}, \mathcal{F}_{t_f})$:

$$p_{t|1}(\mathcal{G}_t|\mathcal{G}_1) = \prod_{d}^{D_v} p_{t_v|1}(v_{t_v}^d|v_1^d) p_{t_f|1}(f_t^d|f_1^d) \prod_{d}^{D_e/2} p_{t_e|1}(e_{t_e}^d|e_1^d), \tag{18}$$

where we omitted the dimension dependency on time for the sake of clarity. Each dimension is noised independently according to Equations (14) and (7) for discrete and continuous variables, respectively, and this can be seen as an extension of the previous multimodal flow where every dimension is a variable on its own. From Proposition 4.2 of Campbell et al. (2024), we know that the following process generates $p_t(\mathcal{G}_t) = \mathbb{E}_{p_1(\mathcal{G}_1)}[p_{t|1}(\mathcal{G}_t|\mathcal{G}_1)]$:

$$R_t(v_{t_v}^d, j) = \mathbb{E}_{p_{1|t}^\theta(v_1^d|\mathcal{G}_t)}[R_t(v_{t_v}^d, j|v_1^d)], \tag{19}$$

$$R_t(e_{t_e}^d, j) = \mathbb{E}_{p_{1|t}^d(e_1^d|\mathcal{G}_t)}[R_t(e_{t_e}^d, j|e_1^d)], \tag{20}$$

$$u_t(f_{t_f}^d) = \mathbb{E}_{p_{1|t}^\theta(f_1^d|\mathcal{G}_t)}[u_t(f_{t_f}^d|f_1^d)], \tag{21}$$

where we abused notations and simply referred to all marginal and conditional rate matrices as R_t and $R_t(.|\mathcal{G}_1)$ to keep notations uncluttered. The denoising process for $v_{t_v}^d$ and $e_{t_e}^d$ is done according to Equation (11), where the two rate matrices $R_t(v_{t_v}^d,j)$ and $R_t(e_{t_e}^d,j)$ are computed in closed form using Equation (16) and their respective marginal prior distributions. The denoising process for features $f_{t_f}^d$ follows Equation (9), using the vector field $u_t^\theta(f_t^d)$ learned by the model.

Loss function. Our generative model is trained by minimizing \mathcal{L}_{CFM} on continuous variables and \mathcal{L}_{DFM} on discrete variables. The overall loss function writes:

$$\mathcal{L} = \mathbb{E}_{t,\mathcal{G}_1,\mathcal{G}_t} \left[-\sum_{d}^{D_v} \log(p_{1|t}^{\theta}(v_1^d|\mathcal{G}_t)) - \sum_{d}^{D_e/2} \log(p_{1|t}^{\theta}(e_1^d|\mathcal{G}_t)) + \sum_{d}^{D_v} \|u_t^{\theta,d}(\mathcal{G}_t) - u_t(f_{t_f}^d|f_1^d)\|_2^2 \right]. \tag{22}$$

Flexible denoising. The main advantage of the proposed framework is that it allows all dimensions, i.e., components or groups of components, to be denoised independently. Very diverse applications are therefore possible with the same model depending on the chosen *time sampling scheme*:

- (a) End-to-end topology design and sizing $-t_v, t_e, t_f : 0 \to 1$;
- (b) Circuit completion / conditional inpainting $-\mathcal{G} = [(\mathcal{V}, \mathcal{E}, \mathcal{F}), (\mathcal{V}', \mathcal{E}', \mathcal{F}')]$ such that $t_v, t_e, t_f = 1$, and $t_{v'}, t_{e'}, t_{f'} : 0 \to 1$;
- (c) Topology-conditional sizing t_v , $t_e = 1$, and $t_f : 0 \to 1$;
- (d) Link prediction $t_v = 1$, and $t_e : 0 \to 1$.

Another noteworthy application is *continued denoising*, where an output graph \mathcal{G} undergoes additional denoising steps if it fails to meet predefined criteria. By fixing t between 0 and 1 independently for each device and edge, one can control the extent to which one part of the circuit must be denoised further or preserved, allowing a very fine-grained supervision. Illustrations of several applications can be found in Figure 1 and Appendix F.

4.2 Network Architecture

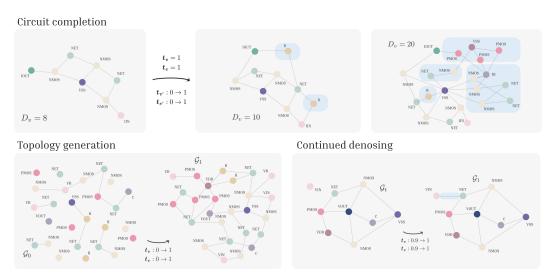


Figure 1: Multiple applications enabled by our framework, here on the AnalogGenie dataset.

We base our denoising network on the graph transformer architecture of Ma et al. (2023) which combines a global receptive field with an expressive random walk structural encoding, achieving strong performance across diverse graph learning tasks. To adapt it to our framework, we introduce two key modifications. First, we incorporate a time-conditioning mechanism, inspired by Diffusion Transformers (Peebles & Xie, 2023), mapping t_v (and respectively t_e and t_f) into multiplicative and additive biases α_v , α_v' , β_v , β_v' , γ_v , and γ_v' , applied at different stages of each transformer layer. Second, we add a dedicated processing path for node features f, mirroring the sequence of operations used for node types v. The graph transformer layer of CircuitFlow is illustrated in Figure 2. Overall, our model is a light architecture of 2.05M parameters.

4.3 Unified Circuit Graph Representation

Topology design and device sizing can be performed at different levels of the circuit representation hierarchy, which mainly includes the behavior level and the transistor level. In the former case, transistors do not appear explicitly, but are included in larger substructures, such as single-stage operational amplifiers (op-amps). Following prior work (Ren et al., 2020; Hakhamaneshi et al., 2022) we represent devices along with circuit and voltage nodes (i.e., nets and ports) as graph nodes. This results in a unified circuit representation that accommodates both representation levels, while conforming to the writing conventions of the SPICE simulation software. Finally, we proceed in a *hierarchical approach* to predict the pin connectivity of transistors. As noted by Gao et al. (2025), this step is necessary to disambiguate topologies which do not specify pin assignments. We therefore train a dedicated model to regress edge prob-

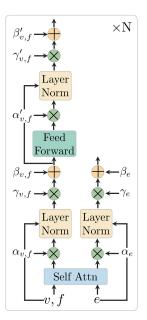


Figure 2: The conditional graph transformer layer of CircuitFlow.

abilities between transistors and their neighboring nodes, conditioned on the output of the topology generation model. This ensures a one-to-one mapping between generated graphs and SPICE netlists across all representation levels. Further details on the circuit representation and the hierarchical model are provided in Appendices C and D.

5 EXPERIMENTS

We evaluates the performance of CircuitFlow on two standard benchmarks for circuit topology generation: OCB (Dong et al., 2023) and AnalogGenie (Gao et al., 2025), which represent circuits at

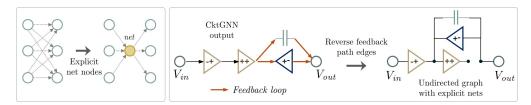


Figure 3: Limitations of the OCB/CktGNN representation. Without explicit nets, graphs are more complex (**left**), and ambiguous (**right**): the output of the second single-stage op-amp (++), on the left side of the right panel, is *not* linked to the input of the third op-amp (+-), which belongs to a *feedback path*: arrows direction should be reversed. The DAG representation wrongly portrays the circuit as valid: it is *open* and therefore invalid, as shown when the feedback path edges are reversed.

distinct abstraction levels. Since OCB also provides device sizes, we additionally assess the model's effectiveness on the sizing task on this dataset. Both datasets however require preprocessing to map circuits into our unified representation, which we describe in the following section. We release the processed datasets together with the code and model weights.

5.1 Datasets and Preprocessing

OCB. The OCB dataset consists of 10,000 DAGs describing up to 3-stage op-amps at the *behavioral level*, split into 9,000 training samples (composed of 3,957, or 44%, unique topologies) and 1,000 test samples. Each graph comes in two versions, one where components are grouped using a predefined set of *subgraphs*, and one decomposed into individual components, the latter including device sizes. To keep the approach general we work at the component level, though several preprocessing steps are more conveniently done at the subgraph level, which requires to map the circuit back to its component and recover the sizes.

The first of these steps is to *identify and revert feedback path edges*: as illustrated in Figure 3 (right), using DAGs leads to *ambiguous topologies*, preventing the identification of feedback path components and misrepresenting their input/output pin assignment. We therefore revert all input and output edges to subgraphs containing feedback op-amps (*gm*- in OCB terminology). Following common practice, we then add *circuit nets* as graph nodes, fully disambiguating device connections. This also has the advantage to simplify the graphs. Finally DAGs are converted into undirected graphs, and we validate the robustness of the whole process by ensuring that all circuits remain simulatable.

AnalogGenie. The AnalogGenie dataset contains 3,350 samples spanning 11 analog circuit types (op-amps, SC-samplers, bandgap references, power converters, etc.). Circuits are represented at the transistor level and are substantially larger than OCB samples (see Table 5 in Appendix E for an exhaustive comparison). Gao et al. (2025) highlight the need to represent transistor pins to disambiguate otherwise similar topologies, but omit net nodes. This has two detrimental consequences: (i) as discussed in the previous section, it creates unnecessary complexity: adding net nodes reduces mean graph density from 0.09 to 0.04; (ii) it breaks the invariance between the pins of symmetric devices (resistors, capacitors, inductors), which must therefore be learned. We therefore preprocess AnalogGenie circuits by adding net nodes, and removing pin nodes for symmetric devices. Resulting dataset statistic can be found in Table 5.

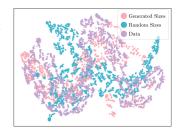


Figure 4: The t-SNE embeddings of circuit properties shows how the model has learned to match data sizes.

5.2 BEHAVIOR-LEVEL TOPOLOGY GENERATION AND SIZING ON THE OCB DATASET

Topology Generation. We first evaluate the quality of circuit topologies generated by CircuitFlow, using several graph learning baselines: D-VAE (Zhang et al., 2019), DAGNN (Thost & Chen, 2021),

Model	V.U.N.↑	Val. sim.↑	Val. graphs↑	Val. circuits↑	Uniqueness ↑	Novelty ↑
DAGNN	_	_	83.1	74.2	_	97.2
PACE	_	_	83.1	75.1	_	97.1
D-VAE	44.5	58.1	67.7	59.5	84.3	94.5
CktGNN	47.9	74.2	85.1	81.4	72.6	93.0
CircuitFlow (Ours)	74.3	92.7	99.4	98.4	85.8	91.1

Table 1: Evaluation of output topology quality across various architectures on OCB. Results that could not be reproduced are reported from Dong et al. (2023). All metrics are expressed in **percentage**, and uniqueness is computed over 10,000 samples.

Simul. out	Data	CircuitFlow	Random feats.
Gain Pm Ugw	1.60 ± 0.73 2.71 ± 1.09 12.59 ± 8.99	1.45 ± 0.70 2.20 ± 1.18 18.50 ± 8.99	1.31 ± 0.68 1.97 ± 1.09 5.55 ± 6.25
$KL(p_{data} .)$	0	0.48	1.15

Table 2: SPICE simulation outputs and KL divergence of circuit performance distributions on OCB. CircuitFlow has learned to produce sizes that maximize all three properties.

PACE (Dong et al., 2022), and CktGNN (Dong et al., 2023). D-VAE and CktGNN were retrained using the code provided by Dong et al. (2023). We apply the same processing to CktGNN's outputs as described above for OCB subgraphs to reflect actual circuit topologies. Metrics from this work are used here, including **Valid graphs** (percentage of connected graphs with input and output nodes) and **Valid circuits** (circuit with a main path composed only of single-stage op-amps). However, these do not capture graphs with nodes of degree one, which indicates open circuits, and ignore feedback paths. This means that some circuits that appear valid are not simulatable. We therefore add a **Valid sim** metric, similarly to Gao et al. (2025), that gives the proportion of circuits simulatable with SPICE with default parameters. Finally we also report the **V.U.N** (Vignac et al., 2022), which gives the fraction of outputs that are simultaneously simulatable, unique, and novel.

Results. Results from Table 1 show the exceptional ability of CircuitFlow to produce novel and valid outputs using as few as 100 denoising Euler steps, improving on CktGNN by 18 points in Valid sim and 26 points in V.U.N, hence establishing a new state-of-the-art on the OCB dataset.

Sizing Experiment. Our multimodal flow matching model is the first architecture capable of jointly generating both circuit topologies and device features, providing an efficient alternative to the common practice of addressing these tasks separately. In this section, we investigate how indirect control over circuit performance can be achieved through our training objective, which maximizes the log-likelihood of output sizes. This aligns the proposed framework with practical analog design, where sizing is typically driven by performance objectives. To illustrate this, we simulate the *gain*, *phase margin* (*pm*) and *unit-gain frequency* (*ugw*) of 10,000 generated circuits, repeat the operation with randomly sampled sizes, and compare with data values. As reported in Table 2, by matching the distribution of dataset sizes, our model has indirectly learned to maximize all three circuit properties. Finally, we note that the output sizes can also constitute highly effective initialization points for classical optimization algorithms, such as BO or genetic algorithm. Figure 4 represents t-SNE embeddings of the simulated quantities, further illustrating the alignment between generated and data distributions.

5.3 Transistor-Level Topology Generation on the AnalogGenie Dataset

This section explores the scalability of CircuitFlow to the more complex and diverse graphs of the AnalogGenie benchmark. Following Gao et al. (2025), we compare with Lamagic Chang et al. (2024) and AnalogCoder (Lai et al., 2025) which, while both trained on distinct datasets and being limited in terms of circuit complexity, represent important milestone works. We train CircuitFlow with our hierarchical approach, using a dedicated model for the regression of pin assignment prob-

abilities. The number of denoising Euler steps was increased to $1{,}000$, which reduces rare but valid outlier topologies and improves pin assignment. Finally, we explore using continued denoising as a post-processing strategy. To this end we select invalid circuits based on predefined but simple rules: disconnected graphs, absence of a VSS node, or node degree inconsistent with device pin number. Those circuits then undergo 10 additional denoising steps starting from t=0.5, ensuring that most of the topology is preserved (examples of this process can be found in Figure 7 in appendix). The whole process is repeated up to 10 times, obtaining significant improvements at a negligible cost .

Results can be found in Table 3. Overall, we achieve a new state-of-the-art on the AnalogGenie dataset over all considered metrics, improving V.U.N over AnalogGenie by 11 to 22 points. Notably, this is done with minimal inductive bias, and a very light preprocessing pipeline.

Model	V.U.N. (%) ↑	Val. sim (%)↑	Uniqueness (%)↑	Novelty (%) ↑	Max Node Number↑
LaMAGIC	_	68.2	_	12.7	4
AnalogCoder	_	57.3	_	8.9	10
AnalogGenie	62.2	73.1	88.5	100	63
CircuitFlow (Ours)	73.7	74.9	98.8	100	71
CircuitFlow + post-pro.	84.4	85.9	98.5	100	71

Table 3: Output topology quality on the AnalogGenie dataset.

6 LIMITATIONS AND FUTURE WORK

CircuitFlow offers a unified framework for the joint generation of topology and device sizes, achieving state-of-the-art quality in generated circuit architectures. While our training strategy already enables indirect control over circuit performances, the practical utility of CircuitFlow for analog circuit design would be further enhanced by direct conditioning mechanisms, potentially fully replacing classical optimization pipelines. This would however require further scaling of the current datasets, along with costly simulation efforts, which is a long-standing bottleneck in the field. At the transistor level, one promising direction is to integrate topology and pin assignment into a single model. In practice, we found the current design to be the most effective balance between sampling quality and model complexity, given the large graph sizes and relatively small scale of the AnalogGenie dataset. Importantly, this constraint is not specific to our approach, and rather reflects a broader domain-wise challenge. Still, our results demonstrate that strong generative models can already be trained effectively on existing open benchmarks.

7 Conclusion

This work introduces CircuitFlow, a flow matching framework for joint analog circuit topology generation and device sizing. By leveraging independent time sampling across dimensions and modalities, the model achieves remarkable inference-time flexibility, enabling applications such as circuit completion, error correction or link prediction. Experiments show that CircuitFlow consistently produces valid, novel, and simulatable circuits, outperforming prior state-of-the-art models across both considered benchmarks. These results demonstrate its scalability from behavioral-level op-amps to diverse transistor-level circuits, while requiring only minimal preprocessing and inductive bias. Finally, this work opens new research directions in generative circuit design, achieving together tasks that were previously treated in isolation. A promising next step is to enable fine-grained control over device sizes, steering generation toward specific performance objectives.

REPRODUCIBILITY STATEMENT

All results reported here are fully reproducible using the provided code and weights, and the precise pipeline will be described on the project's github page. Likewise, we also realease the preprocessing code for both dataset and for the outputs from CktGNN, together with the exact dataset versions used to train our models.

REFERENCES

- Andrew Campbell, Jason Yim, Regina Barzilay, Tom Rainforth, and Tommi Jaakkola. Generative flows on discrete state-spaces: Enabling multimodal flows with applications to protein co-design. arXiv preprint arXiv:2402.04997, 2024.
- Weidong Cao, Mouhacine Benosman, Xuan Zhang, and Rui Ma. Domain knowledge-infused deep learning for automated analog/radio-frequency circuit parameter optimization. In *Proceedings of the 59th ACM/IEEE Design Automation Conference*, pp. 1015–1020, 2022.
- Weidong Cao, Jian Gao, Tianrui Ma, Rui Ma, Mouhacine Benosman, and Xuan Zhang. Rose-opt: Robust and efficient analog circuit parameter optimization with knowledge-infused reinforcement learning. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 2024.
- Chen-Chia Chang, Yikang Shen, Shaoze Fan, Jing Li, Shun Zhang, Ningyuan Cao, Yiran Chen, and Xin Zhang. Lamagic: Language-model-based topology generation for analog integrated circuits. *arXiv preprint arXiv:2407.18269*, 2024.
- Filipe Parrado de Azevedo, Nuno Calado Correia Lourenço, and Ricardo Miguel Ferreira Martins. Comprehensive application of denoising diffusion probabilistic models towards the automation of analog integrated circuit sizing. *Expert Systems with Applications*, pp. 128414, 2025.
- Zehao Dong, Muhan Zhang, Fuhai Li, and Yixin Chen. Pace: A parallelizable computation encoder for directed acyclic graphs. In *International conference on machine learning*, pp. 5360–5377. PMLR, 2022.
- Zehao Dong, Weidong Cao, Muhan Zhang, Dacheng Tao, Yixin Chen, and Xuan Zhang. CktGNN: Circuit graph neural network for electronic design automation. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=NE2911Kq1sp.
- Pedro Eid, Filipe Azevedo, Nuno Lourenço, and Ricardo Martins. Using denoising diffusion probabilistic models to solve the inverse sizing problem of analog integrated circuits. *AEU-International Journal of Electronics and Communications*, 195:155767, 2025.
- Floor Eijkelboom, Grigory Bartosh, Christian Andersson Naesseth, Max Welling, and Jan-Willem van de Meent. Variational flow matching for graph generation. *Advances in Neural Information Processing Systems*, 37:11735–11764, 2024.
- Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first international conference on machine learning*, 2024.
- Shaoze Fan, Ningyuan Cao, Shun Zhang, Jing Li, Xiaoxiao Guo, and Xin Zhang. From specification to topology: Automatic power converter design via reinforcement learning. In 2021 IEEE/ACM International Conference On Computer Aided Design (ICCAD), pp. 1–9. IEEE, 2021.
- Morteza Fayazi, Morteza Tavakoli Taba, Ehsan Afshari, and Ronald Dreslinski. Angel: Fully-automated analog circuit generator using a neural network assisted semi-supervised learning approach. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 70(11):4516–4529, 2023.
- Jian Gao, Weidong Cao, and Xuan Zhang. Rose: Robust analog circuit parameter optimization with sampling-efficient reinforcement learning. In 2023 60th ACM/IEEE Design Automation Conference (DAC), pp. 1–6. IEEE, 2023.
- Jian Gao, Weidong Cao, Junyi Yang, and Xuan Zhang. Analoggenie: A generative engine for automatic discovery of analog circuit topologies. In *The Thirteenth International Conference on Learning Representations*, 2025.
- Itai Gat, Tal Remez, Neta Shaul, Felix Kreuk, Ricky TQ Chen, Gabriel Synnaeve, Yossi Adi, and Yaron Lipman. Discrete flow matching. *Advances in Neural Information Processing Systems*, 37: 133345–133385, 2025.

- Kourosh Hakhamaneshi, Marcel Nassar, Mariano Phielipp, Pieter Abbeel, and Vladimir Stojanovic. Pretraining graph neural networks for few-shot analog circuit modeling and design. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 42(7):2163–2173, 2022.
 - Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
 - Dmitrii Krylov, Pooya Khajeh, Junhan Ouyang, Thomas Reeves, Tongkai Liu, Hiba Ajmal, Hamidreza Aghasi, and Roy Fox. Learning to design analog circuits to meet threshold specifications. In *International Conference on Machine Learning*, pp. 17858–17873. PMLR, 2023.
 - Kishor Kunal, Meghna Madhusudan, Arvind K Sharma, Wenbin Xu, Steven M Burns, Ramesh Harjani, Jiang Hu, Desmond A Kirkpatrick, and Sachin S Sapatnekar. Align: Open-source analog layout automation from the ground up. In *Proceedings of the 56th Annual Design Automation Conference 2019*, pp. 1–4, 2019.
 - Yao Lai, Sungyoung Lee, Guojin Chen, Souradip Poddar, Mengkang Hu, David Z Pan, and Ping Luo. Analogcoder: Analog circuit design via training-free code generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 379–387, 2025.
 - Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *The Eleventh International Conference on Learning Repre*sentations, 2023.
 - Bingyang Liu, Haoyi Zhang, Xiaohan Gao, Zichen Kong, Xiyuan Tang, Yibo Lin, Runsheng Wang, and Ru Huang. Layoutcopilot: An Ilm-powered multi-agent collaborative framework for interactive analog layout design. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 2025a.
 - Chengjie Liu, Weiyu Chen, Anlan Peng, Yuan Du, Li Du, and Jun Yang. Ampagent: An Ilmbased multi-agent system for multi-stage amplifier schematic design from literature for process and performance porting. *arXiv preprint arXiv:2409.14739*, 2024a.
 - Chengjie Liu, Yijiang Liu, Yuan Du, and Li Du. Ladac: Large language model-driven auto-designer for analog circuits. *Authorea Preprints*, 2024b.
 - Chengjie Liu, Jiajia Li, Yabing Feng, Wenhao Huang, Weiyu Chen, Yuan Du, Jun Yang, and Li Du. Diffckt: A diffusion model-based hybrid neural network framework for automatic transistor-level generation of analog circuits. *arXiv preprint arXiv:2507.00444*, 2025b.
 - Qiang Liu. Rectified flow: A marginal preserving approach to optimal transport. *arXiv* preprint *arXiv*:2209.14577, 2022.
 - Jialin Lu, Liangbo Lei, Fan Yang, Li Shang, and Xuan Zeng. Topology optimization of operational amplifier in continuous space via graph embedding. In 2022 Design, Automation & Test in Europe Conference & Exhibition (DATE), pp. 142–147. IEEE, 2022.
 - Jialin Lu, Liangbo Lei, Jiangli Huang, Fan Yang, Li Shang, and Xuan Zeng. Automatic op-amp generation from specification to layout. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 42(12):4378–4390, 2023.
 - Wenlong Lyu, Fan Yang, Changhao Yan, Dian Zhou, and Xuan Zeng. Batch bayesian optimization via multi-objective acquisition ensemble for automated analog circuit design. In *International conference on machine learning*, pp. 3306–3314. PMLR, 2018.
 - Liheng Ma, Chen Lin, Derek Lim, Adriana Romero-Soriano, Puneet K Dokania, Mark Coates, Philip Torr, and Ser-Nam Lim. Graph inductive biases in transformers without message passing. In *International Conference on Machine Learning*, pp. 23321–23337. PMLR, 2023.
 - William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 4195–4205, 2023.

- Souradip Poddar, Ahmet Budak, Linran Zhao, Chen-Hao Hsu, Supriyo Maji, Keren Zhu, Yaoyao
 Jia, and David Z Pan. A data-driven analog circuit synthesizer with automatic topology selection
 and sizing. In 2024 Design, Automation & Test in Europe Conference & Exhibition (DATE), pp.
 1–6. IEEE, 2024.
 - Yiming Qin, Manuel Madeira, Dorina Thanou, and Pascal Frossard. Defog: Discrete flow matching for graph generation. In *Proceedings of the 42nd International Conference on Machine Learning (ICML)*, 2025. URL https://arxiv.org/abs/2410.04263.
 - Haoxing Ren, George F Kokai, Walker J Turner, and Ting-Sheng Ku. Paragraph: Layout parasitics and device parameter prediction using graph neural networks. In 2020 57th ACM/IEEE Design Automation Conference (DAC), pp. 1–6. IEEE, 2020.
 - Keertana Settaluri, Ameer Haj-Ali, Qijing Huang, Kourosh Hakhamaneshi, and Borivoje Nikolic. Autockt: Deep reinforcement learning of analog circuit designs. *arXiv preprint arXiv:2001.01808*, 2020.
 - Jinyi Shen, Fan Yang, Li Shang, Changhao Yan, Zhaori Bi, Dian Zhou, and Xuan Zeng. Atom: An automatic topology synthesis framework for operational amplifiers. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 2024.
 - Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learn-ing*, pp. 2256–2265. pmlr, 2015.
 - Veronika Thost and Jie Chen. Directed acyclic graph neural networks. *arXiv preprint* arXiv:2101.07965, 2021.
 - Alexander Tong, Kilian FATRAS, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrid Rector-Brooks, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models with minibatch optimal transport. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856.
 - Clement Vignac, Igor Krawczuk, Antoine Siraudin, Bohan Wang, Volkan Cevher, and Pascal Frossard. Digress: Discrete denoising diffusion for graph generation. *arXiv preprint arXiv:2209.14734*, 2022.
 - Hanrui Wang, Jiacheng Yang, Hae-Seung Lee, and Song Han. Learning to design circuits. *arXiv* preprint arXiv:1812.02734, 2018.
 - Hanrui Wang, Kuan Wang, Jiacheng Yang, Linxiao Shen, Nan Sun, Hae-Seung Lee, and Song Han. Gcn-rl circuit designer: Transferable transistor sizing with graph neural networks and reinforcement learning. In 2020 57th ACM/IEEE Design Automation Conference (DAC), pp. 1–6. IEEE, 2020.
 - Biying Xu, Keren Zhu, Mingjie Liu, Yibo Lin, Shaolan Li, Xiyuan Tang, Nan Sun, and David Z Pan. Magical: Toward fully automated analog ic layout leveraging human and machine intelligence. In 2019 IEEE/ACM International Conference on Computer-Aided Design (ICCAD), pp. 1–8. IEEE, 2019.
 - Yuxuan Yin, Yu Wang, Boxun Xu, and Peng Li. Ado-llm: Analog design bayesian optimization with in-context learning of large language models. In *Proceedings of the 43rd IEEE/ACM International Conference on Computer-Aided Design*, pp. 1–9, 2024.
 - Muhan Zhang, Shali Jiang, Zhicheng Cui, Roman Garnett, and Yixin Chen. D-vae: A variational autoencoder for directed acyclic graphs. *Advances in neural information processing systems*, 32, 2019.
 - Zhenxin Zhao and Lihong Zhang. Analog integrated circuit topology synthesis with deep reinforcement learning. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 41(12):5138–5151, 2022.

A DERIVATION OF THE CONDITIONAL RATE MATRIX EXPRESSION

In this paper, we consider the prior distribution over node types and edges as the joint marginal distribution over all S states, where S is the number of node types for nodes and 2 for edges. In the following, we write indiscriminately the variable of interest as $x_t \in [1, ..., S]^D$.

We have then:

$$x_0 \sim \text{Cat}(m_1, m_2, \dots, m_S),$$
 (23)

where:

$$\forall i, m_i \in [0, 1] \quad \text{and} \quad m_i = \frac{1}{ND} \sum_{x \in X_{dup}} \sum_d \delta\{x^d, i\}$$
 (24)

Now the x_t^d is sampled from the distribution obtained by an interpolation between the one-hot distribution in x_1^d and the p_0 .

$$x_t^d \sim \text{Cat}(t\delta\{x_1^d, x_t^d\} + (1-t) \times m))$$
 (25)

The derivation of the conditional rate matrix is inspired from the uniform case of Campbell et al. (2024), substituting the uniform distribution with the marginal one. We first need to derive $\partial_t p_{t|1}(x_t^d|x_1^d)$ in order to compute the first term of the conditional rate matrix in Equation (15).

$$\partial_t p_{t|1}(x_t^d | x_1^d) = \partial_t \left(t\delta\{x_t^d, x_1^d\} + (1-t)m_{x_t^d} \right)$$
(26)

$$= \delta\{x_t^d, x_1^d\} - m_{x_t^d} \tag{27}$$

Thus, for $x_t^d \neq j$ (diagonal entries will be computed later):

$$R_t^*(x_t^d, j \mid x_1^d) = \frac{\text{ReLU}\left(\partial_t p_{t|1}(j \mid x_1^d) - \partial_t p_{t|1}(x_t^d \mid x_1^d)\right)}{Sp_{t|1}(x_t^d \mid x_1^d)}$$
(28)

$$= \frac{\text{ReLU}\left(\delta\{j, x_1^d\} - m_j - \delta\{x_t^d, x_1^d\} + m_{x_t^d}\right)}{S\left(t\delta\{x_t^d, x_1^d\} + (1 - t)m_{x_t^d}\right)}.$$
 (29)

This simplifies as:

$$R_t^*(x_t^d, j \mid x_1^d) = \frac{\left(1 - m_j + m_{x_t^d}\right)}{S(1 - t)m_{x_t^d}} \delta\{j^d, x_1^d\} (1 - \delta\{x_t^d, x_1^d\}) + \frac{\text{ReLU}(m_{x_t^d} - m_j)}{S(1 - t)m_{x_t^d}} (1 - \delta\{j, x_1^d\}) (1 - \delta\{x_t^d, x_1^d\}).$$
(30)

We turn now to the derivation of the *detailed balance* term of the conditional rate matrix, which allows to inject stochasticity in the denoising process. As per Campbell et al. (2024), to ensure that the rate matrix still obeys the continuity equation, $R_t^{\rm DB}$ must satisfy the following detailed balance condition:

$$p_{t|1}(i|x_1^d)R_t^{\rm DB}(i,j|x_1^d) = p_{t|1}(j|x_1^d)R_t^{\rm DB}(j,i|x_1^d) \tag{31} \label{eq:31}$$

Following their general recipe for DB rate matrix expression and again considering $i \neq j$, we assume:

$$R_t^{\text{DB}}(i,j|x_1^d) = a_t \delta\{i, x_1^d\} + b_t \delta\{j, x_1^d\}$$
 (32)

Substituting this into Equation (31) yields:

$$(t\delta\{i, x_1^d\} + (1 - t)m_i) (a_t\delta\{i, x_1^d\} + b_t\delta\{j, x_1^d\})$$
 (33)

 $= (t\delta\{j, x_1^d\} + (1-t)m_j) (a_t\delta\{j, x_1^d\} + b_t\delta\{i, x_1^d\}).$

As this must be true for all i and j as long as $i \neq j$, one may fix $i = x_1^d$ to force a simpler relation between a_t and b_t :

$$b_t = a_t \frac{t + (1 - t)m_i}{(1 - t)m_i} \tag{35}$$

(34)

In the following a_t is set to a *noise level* η which can be seen as a re-noising rate of clean data. Finally, the detailed balance conditional rate matrix writes:

$$R_t^{\text{DB}}(i,j|x_1^d) = \eta \delta\{i, x_1^d\} + \eta \frac{t + (1-t)m_i}{(1-t)m_i} \delta\{j, x_1^d\}.$$
 (36)

B MARGINALIZATION OF THE CONDITIONAL RATE MATRIX

The unconditional rate matrix $R_t(x_t^d,j)$ can be computed in closed form by marginalizing the overall conditional rate matrix $R_t(x_t^d,j\mid x_1^d)=R_t^*(x_t^d,j\mid x_1^d)+R_t^{\mathrm{DB}}(x_t^d,j\mid x_1^d)$ over x_1^d , where the probabilities $p_{1\mid t}^\theta(x_1^d\mid x_t)$ are output by the learned flow matching model θ :

$$R_t(x_t^d, j) = \mathbb{E}_{p_{1|t}^\theta(x_1^d \mid x_t)} \left[R_t^*(x_t^d, j \mid x_1^d) + R_t^{\text{DB}}(x_t^d, j \mid x_1^d) \right]$$
(37)

$$= \mathbb{E}_{p_{1|t}^{\theta}(x_{1}^{d}|x_{t})} \left[\frac{\left(1 - m_{j} + m_{x_{t}^{d}}\right)}{S(1 - t)m_{x_{t}^{d}}} \delta\{j, x_{1}^{d}\} (1 - \delta\{x_{t}^{d}, x_{1}^{d}\}) \right. \\ + \frac{\text{ReLU}(m_{x_{t}^{d}} - m_{j})}{S(1 - t)m_{x_{t}^{d}}} (1 - \delta\{j, x_{1}^{d}\}) (1 - \delta\{x_{t}^{d}, x_{1}^{d}\}) \\ + \eta \delta\{x_{t}^{d}, x_{1}^{d}\} + \eta \frac{t + (1 - t)m_{x_{t}^{d}}}{(1 - t)m_{j}} \delta\{j, x_{1}^{d}\} \right].$$

$$(38)$$

Integrating over x_1^d yields the final expression for the marginal rate matrix, wherein each term can be easily computed:

$$R_{t}(x_{t}^{d}, j) = \frac{\left(1 - m_{j} + m_{x_{t}^{d}}\right)}{S(1 - t)m_{x_{t}^{d}}} p_{1|t}^{\theta}(x_{1}^{d} = j|x_{t})$$

$$+ \frac{\text{ReLU}\left(m_{x_{t}^{d}} - m_{j}\right)}{S(1 - t)m_{x_{t}^{d}}} (1 - p_{1|t}^{\theta}(x_{1}^{d} = j|x_{t}) - p_{1|t}^{\theta}(x_{1}^{d} = x_{t}^{d}|x_{t}))$$

$$+ \eta p_{1|t}^{\theta}(x_{1}^{d} = x_{t}^{d}|x_{t}) + \eta \frac{t + (1 - t)m_{x_{t}^{d}}}{(1 - t)m_{j}} p_{1|t}^{\theta}(x_{1}^{d} = j|x_{t})$$
(39)

Diagonal entries of R_t are then obtained following $R_t(i, i) = -\sum_{i \neq i} R_t(i, j)$.

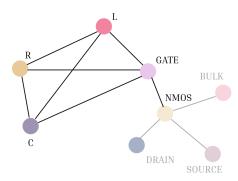
We can finally compute the transition probabilities

$$p_{t+\Delta t|t}(x_{t+\Delta t}^d = j \mid x_t) = \delta\{j, x_t^d\} + R_t(x_t^d, j)\Delta t, \tag{40}$$

which, when done on all dimensions d and all time steps t, finally allows to sample new data points from p_1 .

C GRAPH REPRESENTATION OF A CIRCUIT

The graph-level abstraction we use to represent circuits as graphs is inspired by the SPICE syntax. Just as in a SPICE netlists, there are two main elements: electrical nodes and devices. Electrical nodes include ports (e.g., VDD, VOUT, VIN, ...) and net nodes, which serve as connection points linking two or more pins from different devices. Each device connects to a circuit through its pins, which are linked to electrical nodes. The explicit representation of net nodes (together with transistor pins) therefore allows a 1-to-1 mapping with SPICE netlists.



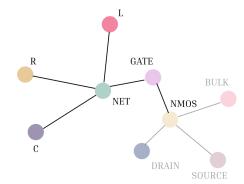
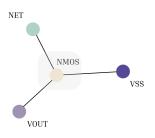


Figure 5: Comparison of circuit graph representations. Left: Direct device-to-device connections create a complex, densely connected graph structure. Right: Introduction of NET nodes as intermediate connection points simplifies the topology by eliminating direct inter-device connections, resulting in a cleaner graph representation.

D PIN-LEVEL PREDICTION

Our two-stage approach for pin-level topology generation involves a first model which generates a circuit topology with device interconnections. Based on the output of the first stage, a second model determines how pins connect to the neighboring components. The process is represented in Figure 6.



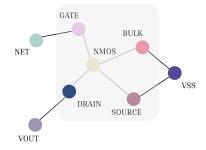


Figure 6: (Left) Device-level circuit, where the NMOS is treated as a single node connected to three nets (NET, VOUT, and VSS). (Right) The corresponding pin-level graph, with explicit pin connections.

The pin assignment model shares the same architecture as the topology generation model, but is trained to regress edge probabilities between a node's pins and the neighboring net nodes. This is done using a classical 2-class cross entropy objective:

$$\mathcal{L}_{PIN} = -\frac{1}{|\mathcal{E}_L|} \sum_{e \in \mathcal{E}_L} y_e \log(\hat{y}_e)$$
(41)

The accuracy of the pin assignment model is reported in Table 4, after 150 training epochs.

Metric	Value (%)
Precision	97.54
Recall	98.87
F1-score	98.2
Accuracy	98.2

Table 4: Test set performance metrics of the pin assignment model.

E DATASETS STATISTICS

Table 5 presents the main features of the datasets used in this paper. AnalogGenie (no pins) corresponds to the dataset version that is used to train the topology generation model, which is completed by a pin assignment model.

Dataset	OCB	AnalogGenie	AnalogGenie (no pins)
# of graphs	10,000	3,350	3,350
% uniqueness	44.0	99.4	99.4
Number of node types	9	81	28
Avg. # of nodes per sample	12	107	38
Avg. # of edges per sample	15	145	60
Avg. density	0.25	0.042	0.11

Table 5: Main dataset statistics.

F ILLUSTRATIONS OF DENOISING APPLICATIONS

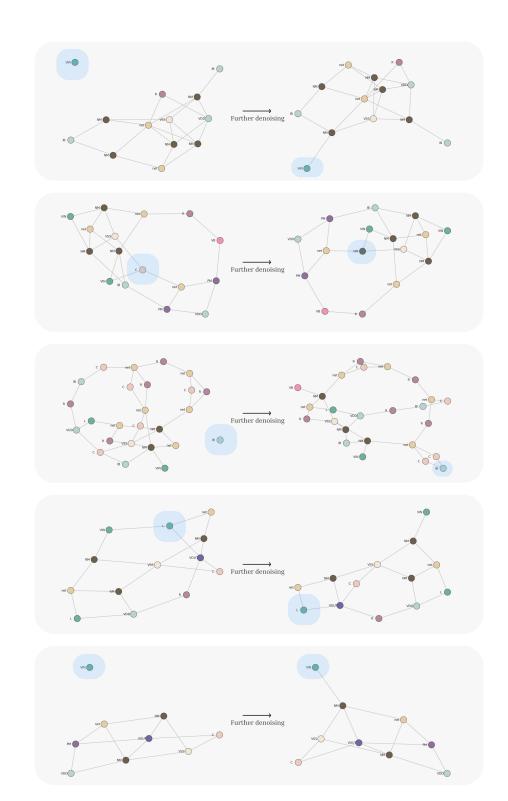


Figure 7: Examples of continued denoising for topology correction on the AnalogGenie dataset.

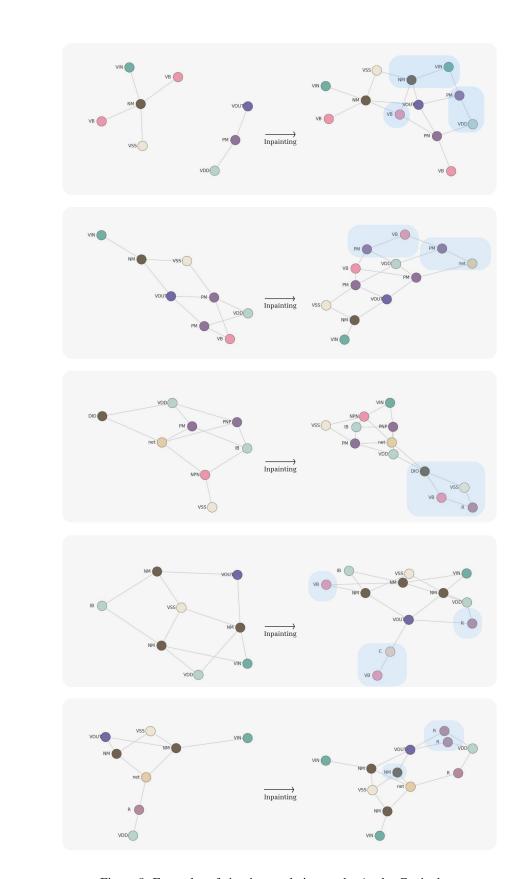


Figure 8: Examples of circuit completion on the AnalogGenie dataset.