

# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 PARETOROUTER: VLSI GLOBAL ROUTING WITH MULTI-OBJECTIVE OPTIMIZATION

Anonymous authors

Paper under double-blind review

## ABSTRACT

Global routing (GR) has been a central task in modern chip design. Many efforts, either ML-based or heuristic, have been proposed that seek to optimize specific business goals, such as overflow (OF) and wirelength (WL) of generated routes. Notably, recent end-to-end neural routers have demonstrated significant speed advantages in optimizing wirelength, yet they struggle to achieve efficiency in reducing overflow. In fact, a good trade-off between the above two metrics has not been achieved, especially when overall efficiency is pursued, as existing ML-based methods often optimize only a single metric. To bridge this gap for more practical industry applications, we propose a flow-matching-based router for GR, called ParetoRouter, which achieves trade-offs between WL and OF, generating highly connected routes at high speed and quality. In the training phase, two differential metric-oriented routing results are utilized to build the training datasets. They are leveraged to design an ‘Average Flow’ between initial pins and final routings. A Pareto sampling method, based on the Das-Dennis method, is also devised to achieve trade-offs between OF and WL in the inference phase. Extensive experimental results show that it achieves SOTA performance on the **overflow reduction** with **less superfluous routes** across all benchmarks with **x10** times speedup over the peer SOTA ML-based method.

## 1 INTRODUCTION

Global routing (GR) (McMurchie et al., 1995; Cheng et al., 2022; Liao et al., 2020) has become one of the most intricate and time-consuming phases in modern VLSI design flows (Kramer & Van Leeuwen, 1984), among the other stages such as logic synthesis (Neto et al., 2021), floorplanning (Li et al., 2022a), and placement (Hao et al., 2021; Shi et al., 2023). With netlists containing millions to billions of nets, global routers interconnect pins, aiming to minimize wirelength and avoid overflow under limited routing resources. In fact, even the simplified ‘2-pin’ case, i.e., connecting each net with exactly two pins under specified constraints, is NP-complete (Paulus et al., 2021).

Classical works on heuristics (Cho et al., 2007; Kastner et al., 2002) often require continuous updates and improvements by human experts to route greedily. To mitigate the reliance on manual effort and enhance design automation and quality, learning-based approaches have been introduced. Notably, as shown in Table 1, most existing ML-based works often rely on generative models or classical solvers (Li et al., 2022b; Yan et al., 2018; Li et al., 2021; 2024; Feng & Feng, 2025) to predict Steiner points (Hwang, 1979) and heuristics post-processing or deep reinforcement learning (DRL) (Mahboubi et al., 2021) to route the discrete predicted points, e.g. NeuralSteiner (Liu et al., 2024) and Hubrouter (Du et al., 2023). DSBRouter (Shi et al., 2025) leverages Diffusion Schrödinger Bridge (DSB) (De Bortoli et al., 2021) to design an end-to-end ML-based router (e.g., from discrete initial pins to connected routes), but it is hard to train and suffers a great generation time. In addition, most ML-based methods optimize only a single metric: for example, Hubrouter primarily targets wirelength, whereas NeuralSteiner and DSBRouter focus on overflow, making it challenging to achieve trade-offs between these two metrics.

Motivated by accelerated diffusion sampling paradigms (e.g., Consistency Models (Song et al., 2023), CM, Flow Matching (Lipman et al., 2023), FM), can we design an end-to-end global router that optimizes both overflow and wirelength simultaneously, generating high-quality connected routes while significantly reducing generation time?

054  
055  
056  
057  
058  
059  
060  
Table 1: Summary of existing machine learning-based solvers for GR.  
061

METHOD	TYPE	END-TO-END	MOO SUPPORT
Hubrouter (Du et al., 2023)	GENERATIVE + RL	✗	✗
NeuralSteiner (Liu et al., 2024)	CNN + POST-PROCESSING	✗	✗
DSBRouter (Shi et al., 2025)	GENERATIVE	✓	✗
ParetoRouter (ours)	GENERATIVE	✓	✓

061 Thus, in this paper, we propose ParetoRouter, which is an FM-based global router that integrates  
 062 a carefully designed yet simple ‘Average Flow’ (AF) to fulfill the diversity of predicted routes,  
 063 facilitating on-demand sampling of routing results. Simply put, during training, ParetoRouter utilizes  
 064 the combination of two differential metric-oriented connected routes as supervisory signals, enabling  
 065 the model to predict the AF in latent space. The design of AF is detailed in Sec 4.1. The FM model  
 066 and the DSB module in DSBRouter serve similar functions.

067 However, as demonstrated by the experimental results in Appendix A.2.2 and A.6.2, FM achieves  
 068 comparable performance to DSB at a lower training cost. In the sampling phase, a fast one-step  
 069 Pareto sampling method based on Das-Dennis approach (Das & Dennis, 1998) is designed to realize  
 070 a controllable generation of routes, which means that ParetoRouter can generate connected routes  
 071 that are more inclined towards overflow (or wirelength) to approximate Pareto Front (PF). This  
 072 sampling module is detailed in Sec. 4.2. Extensive experiments show that ParetoRouter can not only  
 073 generate routes with diversity, but also significantly reduce the generation time and generate nearly  
 074 state-of-the-art (SOTA) routing results under rationally specified parameters. **This paper contributes**  
 075 **as follows:**

076 1) To the best of our knowledge, the proposed ParetoRouter is the first ML-based multi-objective  
 077 global router that produces preference-compliant routing solutions and explicitly explores the Pareto  
 078 front between wirelength (WL) and overflow (OF).

079 2) ParetoRouter incorporates a fast Das–Dennis–based Pareto sampling scheme to approximate Pareto  
 080 Front in GR, which substantially reduces solution-generation time by up to 90% compared with  
 081 DSBRouter (Shi et al., 2025). It enables on-demand OF-oriented or WL-oriented routing results,  
 082 compared to SOTA ML-based methods Du et al. (2023); Liu et al. (2024).

083 3) Experimental results demonstrate that ParetoRouter achieves SOTA OF performance across all  
 084 evaluated benchmarks, reducing OF by an average of 68% and markedly decreasing superfluous  
 085 routing on large-scale nets compared to the SOTA ML-based DSBRouter.

## 087 2 PRELIMINARIES AND PROBLEM DEFINITION

### 089 2.1 OFFLINE MULTI-OBJECTIVE OPTIMIZATION

091 Offline multi-objective optimization (MOO) aims to simultaneously minimize multiple objectives  
 092 using an offline dataset  $\mathcal{D}$  of designs and their corresponding labels. Let the design space be  $\mathcal{X} \subseteq \mathbb{R}^d$ ,  
 093 where  $d$  denotes the dimensionality of the design. The goal of MOO is to identify solutions that  
 094 achieve the best trade-offs among conflicting objectives. Formally, the multi-objective optimization  
 095 problem is defined as:

$$096 \text{Find } \mathbf{x}^* \in \mathcal{X} \text{ such that there is no } \mathbf{x} \in \mathcal{X} \text{ with } f(\mathbf{x}) \prec f(\mathbf{x}^*), \quad (1)$$

098 where  $f : \mathcal{X} \rightarrow \mathbb{R}^m$  is a vector-valued map of  $m$  objective functions, and  $\prec$  denotes Pareto dominance.  
 099 A solution  $\mathbf{x}$  *Pareto dominates* another solution  $\mathbf{x}^*$  (denoted  $f(\mathbf{x}) \prec f(\mathbf{x}^*)$ ) if:

$$100 \quad \forall i \in \{1, \dots, m\}, f_i(\mathbf{x}) \leq f_i(\mathbf{x}^*) \quad \text{and} \quad \exists j \in \{1, \dots, m\} \text{ such that } f_j(\mathbf{x}) < f_j(\mathbf{x}^*). \quad (2)$$

103 In other words,  $\mathbf{x}$  is no worse than  $\mathbf{x}^*$  on every objective and strictly better on at least one. A solution  
 104  $\mathbf{x}^*$  is *Pareto optimal* if there is no  $\mathbf{x} \subseteq \mathcal{X}$  that Pareto dominates  $\mathbf{x}^*$ . The set of all Pareto-optimal  
 105 solutions is the *Pareto set* ( $PS$ ). The corresponding set of objective vectors,  $\{f(\mathbf{x}) \mid \mathbf{x} \in PS\}$ , is  
 106 known as the *Pareto front* ( $PF$ ).

107 The goal of MOO is to compute a set of solutions that closely approximates the PF, providing a  
 108 comprehensive representation of the best attainable trade-offs among the objectives.

108  
109 2.2 FLOW MATCHING

110 Flow matching (FM) (Lipman et al., 2023) is an advanced generative modeling framework that  
 111 has demonstrated superior effectiveness and efficiency compared to other models (Ho et al., 2020;  
 112 Song et al., 2021b;a). At its core is a conditional probability path  $p_t(\mathbf{x} \mid \mathbf{x}_0)$  for  $t \in [0, 1]$ ,  
 113 which evolves from the initial distribution  $p_0(\mathbf{x} \mid \mathbf{x}_0) = q(\mathbf{x}_0)$  to an approximate Dirac delta  
 114  $p_1(\mathbf{x} \mid \mathbf{x}_0) \approx \delta(\mathbf{x} - \mathbf{x}_0)$ . This evolution is conditioned on a specific point  $\mathbf{x}_0$  drawn from  $q(\mathbf{x}_0)$  and  
 115 is governed by the conditional vector field  $u_t(\mathbf{x} \mid \mathbf{x}_0)$ . A neural network with parameters  $\theta$  is trained  
 to learn the marginal vector field  $v(\mathbf{x}, t)$ :

$$116 \hat{v}(\mathbf{x}, t; \theta) \approx v(\mathbf{x}, t) = \mathbb{E}_{\mathbf{x}_0 \sim p_t(\mathbf{x} \mid \mathbf{x}_0)} [u_t(\mathbf{x} \mid \mathbf{x}_0)]. \quad (3)$$

118 The modeled vector field  $\hat{v}(\mathbf{x}, t; \theta)$  serves as a neural Ordinary Differential Equation (ODE) that  
 119 guides the transport from  $q(\mathbf{x}_0)$  to  $p_{data}(\mathbf{x}_1)$ . Normally, FM begins with a sampled noise  $\mathbf{x}_0$  from  
 120  $q(\mathbf{x}_0)$  (Pooladian et al., 2023). Then a linear interpolation with the uncorrupted data  $\mathbf{x}_1$  is constructed:  
 121

$$122 \mathbf{x} \mid \mathbf{x}_0, t = (1 - t) \cdot \mathbf{x}_0 + t \cdot \mathbf{x}_1, \mathbf{x}_0 \sim q(\mathbf{x}_0). \quad (4)$$

124 The conditional vector field is readily derived as  $u_t(\mathbf{x} \mid \mathbf{x}_0) = (\mathbf{x}_0 - \mathbf{x}) / (1 - t)$ . Equivalently,  
 125  $u_t(\mathbf{x} \mid \mathbf{x}_0) = \mathbf{x}_1 - \mathbf{x}_0$ . The following objective is used to minimize the conditional FM:

$$126 \mathbb{E}_{t, p_{data}(\mathbf{x}_0), q(\mathbf{x}_1)} \| \hat{v}(\mathbf{x}, t; \theta) - (\mathbf{x}_1 - \mathbf{x}_0) \|^2. \quad (5)$$

128 After training, samples are generated by integrating the neural ODE driven by the learned vector field  
 129  $\hat{v}(\mathbf{x}, t; \theta)$ .

## 131 2.3 GLOBAL ROUTING VIA FM

132 Typically, given a physical chip and a netlist, a chip canvas is created along with several nets (as shown  
 133 in Fig. 1 (a)), where each net includes pins placed at fixed positions determined by the placement  
 134 of macros and standard cells. The primary goal of global routing (GR) is to establish connections  
 135 for all required pins while simultaneously minimizing the routing WL and OF. Existing ML-based  
 136 methods (Li et al., 2024; Du et al., 2023; Liu et al., 2024) generally approach GR as a two-phase task  
 137 (refer to Fig. 1(a), (b), and (c)): first, the Steiner points (pins) are predicted, and then post-processing  
 138 algorithms are applied to connect the initial pins with the predicted Steiner points. Some approaches,  
 139 such as DSBRouter (Shi et al., 2025), aim to create an end-to-end pipeline (see Fig. 1(a) and (d)).  
 140 However, the DSB used in DSBRouter suffers from efficiency problems, and it is unable to sample  
 141 routes that satisfy multi-objective optimization (MOO).

142 In contrast, ParetoRouter, depicted in Fig. 1(e), leverages a flow-based model (FM) to learn the  
 143 distribution  $p_\theta(\mathbf{x}_1 \mid \mathbf{x}_0)$  of routes  $\mathbf{x}_1$  with high diversity for a given instance  $\mathbf{x}_0$ . Design of the AF  
 144 will be discussed in Sec. 4.1. For sampling, we design a one-step, Das-Dennis-based Pareto sampling  
 145 method, which accelerates the generation process and leads to routing results approximating the  
 146 Pareto Front. The design of the sampling method is detailed in Sec. 4.2.

## 147 3 RELATED WORKS

148 Due to page limits, we leave partial related works to Appendix A.1.

149 **The Task of Global Routing.** Owing to the complexity of VLSI routing, the circuit layout is  
 150 partitioned into rectangular regions, called global cells (GCells) (Cho et al., 2007). Global routing  
 151 is modeled as a grid graph  $G = (V, E)$ , where each GCell corresponds to a vertex ( $v \in V$ ) and  
 152 adjacent GCells are connected by an edge ( $e \in E$ ) representing their shared boundary. Modern chip  
 153 designs employ two or more metal layers for routing, where each layer is assigned either a horizontal  
 154 or vertical direction, yielding a two-dimensional grid abstraction. For each net, the global router  
 155 assigns a connected subset of GCells—linked via multiple edges—to interconnect all pins, typically  
 156 producing a rectilinear Steiner tree (RST) (Chu & Wong, 2005). The Hanan grid (Hanan, 1966) and  
 157 the escape graph (Ganley & Cohoon, 1994) are often exploited to construct a shortest rectilinear  
 158 Steiner minimum tree (RSMT) while avoiding obstacles (Liu et al., 2012), by treating intersection  
 159 points in these graphs as candidate Steiner points.

160 **Classical Global Router.** Global routing is a combinatorial optimization problem that can be  
 161 formulated as a 0–1 integer linear program and solved with a general-purpose solver. In practice,

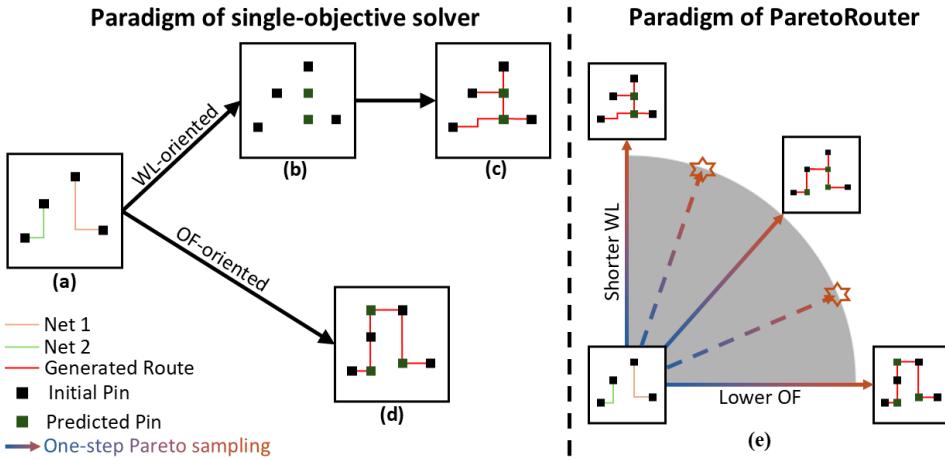


Figure 1: **Difference between ParetoRouter and other solvers.** (a): initial pins and nets. (b): predicted pins and initial pins. (c) and (d): Generated routes via post-processing algorithms. (e): One-step Pareto sampling within ParetoRouter.

classical routers decompose the task into two stages aimed at congestion management: Steiner topology generation and rip-up-and-reroute (RRR). The former commonly relies on FLUTE (Chu & Wong, 2005), which uses lookup tables to build near-optimal Steiner trees in terms of WL for each net, but it is oblivious to congestion. During congestion mitigation, many routers perform edge shifting to move routing demand out of congested regions (Chu & Wong, 2005). NTHU-Route 2.0 (Chang et al., 2008) further introduces a history-based cost that accumulates past congestion and dynamically adjusts routing costs, improving overall quality and efficiency. NCTU-GR 2.0 (Liu et al., 2013) adopts an SMT-aware routing scheme to achieve shorter WL. However, as design complexity and scale grow, these procedures become increasingly time-consuming. Consequently, accelerating congestion resolution with deep learning-based techniques can improve the overall performance of global routing.

**Learning-based Router.** A growing body of works investigates learning to optimize WL and the use of neural networks for global routing, including generating pin-connection orders (Liao et al., 2020), routing segments (Cheng et al., 2022), and customized hub points for rectilinear Steiner trees (RSTs) (Du et al., 2023; Li et al., 2024; Feng & Feng, 2025). The primary practical challenges, however, lie in handling large-scale nets to mitigate overflow (OF) and maintain short WL under limited routing resources. In such settings, judicious detours are essential for relieving congestion, because the WL-minimal RST, e.g., that in Fig. 1(c) produced by Hubrouter (Du et al., 2023) or NeuralSteiner (Liu et al., 2024), may be infeasible in practice. DSBRouter (Shi et al., 2025) can produce low-OF solutions (Fig. 1(d)), but often introduces excessive redundant routing and incurs prohibitive generation time on large-scale nets. PatLabor (Chen et al., 2025b) considers the MOO constraint, but its focus is on balancing WL and RSMT construction delay. Moreover, global routers must be able to generalize to unseen circuit distributions. To meet these challenges, we propose *ParetoRouter*, which enables explicit trade-offs between OF and WL, yielding on-demand routing solutions for nets of arbitrary scale.

#### 4 FRAMEWORK OF PARETOROUTER

For ML-based global routing (GR) solvers, the two-stage paradigm of first predicting Steiner points and then performing routing is intuitive and straightforward. However, it exhibits several limitations: (i) the neural networks used to predict Steiner points are typically trained in a supervised manner, making them brittle under distribution shift and prone to inflating the time complexity of downstream post-processing due to prediction noise; (ii) relying on a single classical solver to provide training targets restricts the diversity of Steiner points available to the model; and (iii) post-processing algorithms are often focused primarily on WL (or OF) minimization, which makes it challenging to satisfy multi-objective optimization (MOO) constraints.

To address these issues, this section introduces ParetoRouter, which integrates MOO into GR to handle constraints explicitly and perform GR in an end-to-end manner. The next two sections first present AF, a training-time mechanism that leverages routing results from multiple classical solvers

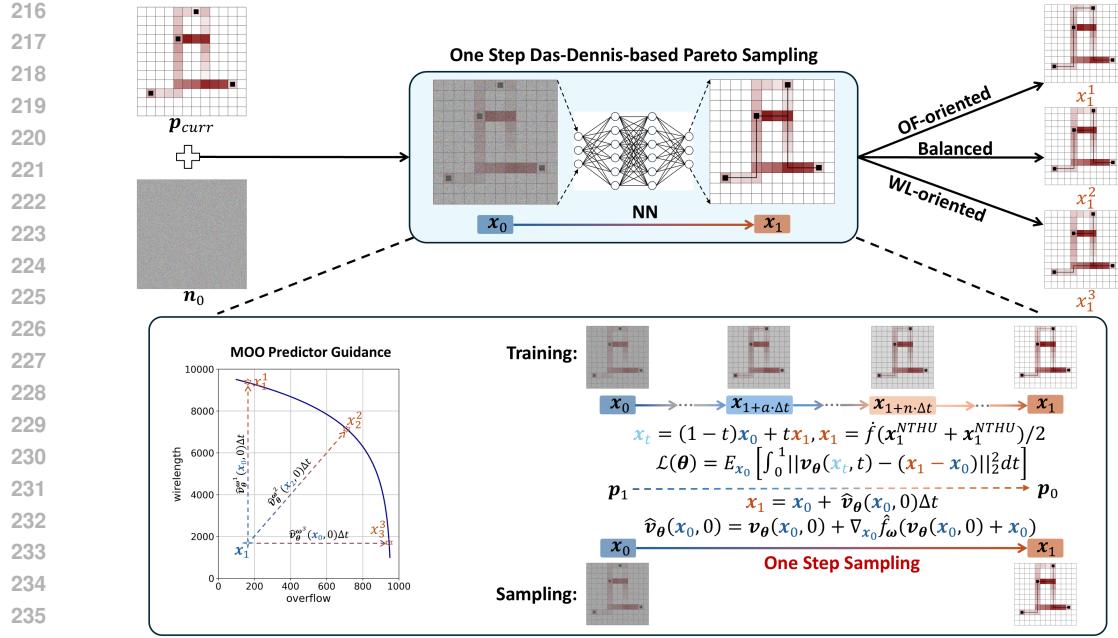


Figure 2: For training, ParetoRouter applies the average of two different solvers’ routing results (NTHURouter and NCTU-GR) to construct the supervisory signals  $x_0$  to enable the NN module to predict more potential routes. While in the sampling phase, weights  $\omega = \{\omega^i\}$  subdivide the objective space into equal partitions. Each weight  $\omega^i \in \omega$  maps to a sampled routing result  $x_0^i$ .

to enable the model to predict a diverse set of routes. Then, during sampling, we introduce a one-step Pareto sampling pipeline based on the Das–Dennis method to initialize multiple weight matrices, thereby guiding the rapid generation of MOO-compliant routes under different weightings in a single step. Fig. 2 depicts the pipeline of proposed ParetoRouter.

#### 4.1 AVERAGE FLOW

Diffusion solvers for GR are not something new; both Hubrouter (Du et al., 2023) and DSBRouter (Shi et al., 2025) try to fulfill the potential of diffusion models for GR. Hubrouter leverages GAN (Goodfellow et al., 2014), DDPM (Ho et al., 2020), and VAE (Kingma & Welling, 2013) to sample predicted ‘hubs’ from simulated Gaussian noise. Anchoring the generative process to stationary noise to enable standard sampling paradigms is intuitive but exhibiting a disconnect between the diffusion trajectory and the structured nature of solution spaces, i.e., the generative process operates in the Gaussian noise space for pin prediction rather than the solution (routes), which causing the restriction on both the controllability of the intermediate states within the generation and the exploitation of prior pin knowledge and need of post-processing algorithms for connected routes. In contrast, DSBRouter (Shi et al., 2025) designs a post-processing-free paradigm that better aligns with heuristic search dynamics, but it is hard to train the backbone and takes a great time for generation due to the inherent constraints of DSB (De Bortoli et al., 2021). Besides, all the ML-based solvers discussed above optimize either WL or OF; none of them can generate routes according to the realistic needs. Consequently, ParetoRouter introduces AF as:

$$\mathbb{E}_{t, p_{data}(\mathbf{x}_1), q(\mathbf{x}_0)} \|\hat{v}(\mathbf{x}, t; \theta) - (\dot{f}(\mathbf{x}_1^{NTHU} + \mathbf{x}_1^{NCTU})/2 - \mathbf{x}_0) \|^2, \quad (6)$$

where  $\mathbf{x}_1^{NTHU}$  and  $\mathbf{x}_1^{NCTU}$  represent the routing results of NTHURouter (Chang et al., 2008) and NCTU-GR (Liu et al., 2013), respectively.  $\dot{f}$  is a scale function to scale the differential routes of these two solvers.  $\mathbf{x}_0$  is the noisy pins with congestion map. We introduce Gaussian noise  $\mathbf{n}_0$  to corrupt the initial pins  $\mathbf{p}_{curr}$ , resulting in an initial sample  $\mathbf{x}_0 = f_n(\mathbf{n}_0 + \mathbf{p}_{curr})$  drawn from  $q(\mathbf{x}_0)$ , where  $f_n$  represents the normalization function. Reasons for noise corruption will be discussed in Appendix A.2.1. Given two flows, starting with the same  $\mathbf{x}_0$ , ending in  $\mathbf{x}_1^{NTHU}$  and  $\mathbf{x}_1^{NCTU}$  respectively, the loss design is like predicting the flow in between. This is why we name it Average Flow. This design is straightforward but has been proven to be very effective in the experiments, which will be discussed in Sec. 5.

270 4.2 ONE STEP DAS-DENNIS-BASED PARETO SAMPLING  
271

272 This section first elucidates the concept of gradient guidance under multi-objective optimization  
273 (MOO) constraints within the ParetoRouter sampling pipeline, along with the design. Then, the  
274 detailed formulation of a weighted score distribution induced by the Das–Dennis method is introduced  
275 to fulfill guided Pareto sampling.

276 **Sampling with Gradient Guidance.** Classifier guidance was initially introduced to steer sample  
277 generation toward designated image categories (Dhariwal & Nichol, 2021). This concept has since  
278 been extended to regression contexts, where it is employed to guide molecular generation (Lee et al.,  
279 2023; Jian et al., 2024; Chen et al., 2025a). Building upon Lemma 1 from Zheng et al. (2023), we  
280 derive the formulation of gradient guidance within the framework of flow matching as follows:

$$281 \tilde{v}(\mathbf{x}, t; \boldsymbol{\theta}) = \hat{v}(\mathbf{x}, t; \boldsymbol{\theta}) + \rho \cdot \nabla_{\mathbf{x}_t} \log p(\mathbf{s} \mid h(\mathbf{x}_t, t)), \quad (7)$$

282 where  $p(\mathbf{s} \mid \mathbf{x}_t, t)$  denotes the distribution of predicted routing score and  $\mathbf{s}$  denotes the computed  
283 properties through classifier function  $h$  whose implementation will be detailed in Sec. A.2.2. More  
284 details of Eq.(7) can be found in Appendix A.3. In implementation, as we assume a one-step  
285 sampling procedure, where the NN module within the ParetoRouter framework is designed to predict  
286 the difference between  $\mathbf{x}_1$  and  $\mathbf{x}_0$ . Consequently, we adopt the following formulation to better align  
287 with this one-step sampling scheme:

$$288 \tilde{v}(\mathbf{x}, t; \boldsymbol{\theta}) = \hat{v}(\mathbf{x}, t; \boldsymbol{\theta}) + \rho \cdot \nabla_{\mathbf{x}} \log p(\mathbf{s} \mid h(\hat{v}(\mathbf{x}, t; \boldsymbol{\theta}) + \mathbf{x}_0, t)). \quad (8)$$

290 **Weighted Score Distribution.** Preceding ML-based works like DSBRouter primarily address the  
291 generation of samples that satisfy a single score  $s$ . In contrast, our proposed ParetoRouter is designed  
292 to optimize two properties, namely OF and WL simultaneously, represented as  $\mathbf{s} = [\hat{f}_1(\mathbf{x}_t), \hat{f}_2(\mathbf{x}_t)]$ .  
293 To address the increased complexity inherent in this multi-objective setting, we decompose the overall  
294 generation task into a series of weighted single-objective subproblems. Specifically, we introduce a  
295 weight vector  $\boldsymbol{\omega} = [\omega_1, \omega_2]$ , where each  $\omega_i > 0$  and  $\sum_{i=1}^m \gamma_i \omega_i = 1$ . The resulting weighted score is:

$$296 \hat{f}_{\boldsymbol{\omega}}(\mathbf{x}_t) = -h \left( \sum_{i=1}^n \gamma_i \hat{f}_i(\mathbf{x}_t) \omega_i \right). \quad (9)$$

300 Here,  $\hat{f}_i$  denotes the predicted score of the  $i$ th objective for  $\mathbf{x}_t$ . Given that the scales of OF and WL  
301 differ, a scaling factor  $\gamma$  is introduced to ensure compatibility. The negative sign reflects that the  
302 objective is minimization. Following the approach in Lee et al. (2023), we define the weighted score  
303 distribution:

$$304 p(\mathbf{s} \mid \hat{v}(\mathbf{x}, t; \boldsymbol{\theta}) + \mathbf{x}_0, t) = e^{\hat{f}_{\boldsymbol{\omega}}(\hat{v}(\mathbf{x}, t; \boldsymbol{\theta}) + \mathbf{x}_0)} / Z, \quad (10)$$

306 where  $Z$  is the normalization constant. By incorporating this formulation into Eq.8, we arrive at the  
307 following predictor update:

$$308 \tilde{v}(\mathbf{x}, t; \boldsymbol{\theta}) = \hat{v}(\mathbf{x}, t; \boldsymbol{\theta}) + \rho \cdot \nabla_{\mathbf{x}} \hat{f}_{\boldsymbol{\omega}}(\hat{v}(\mathbf{x}, t; \boldsymbol{\theta}) + \mathbf{x}_0). \quad (11)$$

310 This vector field  $\tilde{v}(\mathbf{x}, t; \boldsymbol{\theta})$  effectively guides the sampling process toward regions in the input space  
311 that satisfy the desired multi-objective properties encoded by the weighted distribution. To achieve  
312 comprehensive coverage of the Pareto objective space, we employ the Das-Dennis approach (Das  
313 & Dennis, 1998), which partitions the objective space uniformly to generate a diverse set of weight  
314 vectors  $\boldsymbol{\omega}$ . Each weight vector corresponds to a distinct sampled route, thereby facilitating exploration  
315 across the entire trade-off front. The sampling step is performed using the Euler method (Van Kampen,  
316 1976), formulated as:

$$317 \hat{\mathbf{x}}_j = \mathbf{x}_t + \tilde{v}(\mathbf{x}, t; \boldsymbol{\theta}) \Delta t, \quad (12)$$

318 where  $j = t + \Delta t$  represents the next time step. In the ParetoRouter framework, we set  $t = 0$  and  
319  $\Delta t = 1$ . The complete procedure is summarized in Algorithm 1.

321 5 EXPERIMENTS  
322

323 In this section, we empirically compare our proposed ParetoRouter with other ML-based and classical  
324 solvers on ISPD benchmarks.

---

**324 Algorithm 1 One-Step Pareto Sampling**


---

**325 Input:** Dataset  $\mathcal{D}$ , epochs  $T$ ;  
**326 Output:** Generated connected router set  $\mathcal{R}$ ;  
**327**  
1: Train the vector field  $\hat{v}(\mathbf{x}, t; \theta)$  of FM using loss from Eq.(6) on  $\mathcal{D}$ .  
2: Generate weight vectors  $\{\omega^i\}_{i=1}^N$  using the Das-Dennis method.  
3: Initialize  $\mathbf{x}_1$  with scaled Gaussian noise  $\mathbf{n}$  using two classical GR solvers.  
4: **for**  $t = 1$  to  $T$  **do**  
5:     Set  $\Delta t = 1$   
6:     Calculate the score distribution using Eq.(10).  
7:     Calculate the guided vector field  $\tilde{v}(\mathbf{x}, t; \theta)$  using Eq.(11).  
8:     Derive sampled routes (Pareto Front) with Eq.(12).  
9: **end for**  
10: **return**  $\mathcal{R}$

---

**337**  
**338 Table 2: Crrt, WLR and generation time on ISPD07 benchmarks.** Note GAN-HubRouter can  
**339** not directly produces a fully connected result, while DSBRouter consumes a great generation time.  
**340** Compared with them, our ParetoRouter achieves 100% correctness rate with a small generation time.

METRIC	CASE	HUBROUTER (GAN) (DU ET AL., 2023)	DSBROUTER (SHI ET AL., 2025)	PARETOROUTER (OURS)
<b>Correctness Rate</b>	SMALL-4	$0.48 \pm 0.004$	<b><math>1.000 \pm 0.000</math></b>	<b><math>1.000 \pm 0.000</math></b>
	SMALL	$0.12 \pm 0.001$	<b><math>1.000 \pm 0.000</math></b>	<b><math>1.000 \pm 0.000</math></b>
	LARGE-4	$0.000 \pm 0.000$	<b><math>1.000 \pm 0.000</math></b>	<b><math>1.000 \pm 0.000</math></b>
	LARGE	$0.000 \pm 0.000$	<b><math>1.000 \pm 0.000</math></b>	<b><math>1.000 \pm 0.000</math></b>
<b>Wirelength Ratio</b>	SMALL-4	<b><math>1.012 \pm 0.011</math></b>	$1.015 \pm 0.000$	$1.016 \pm 0.000$
	SMALL	$1.002 \pm 0.001$	<b><math>1.001 \pm 0.000</math></b>	$1.002 \pm 0.000$
	LARGE-4	$1.004 \pm 0.021$	<b><math>1.001 \pm 0.000</math></b>	$1.002 \pm 0.000$
	LARGE	<b><math>1.001 \pm 0.000</math></b>	$1.002 \pm 0.000$	$1.003 \pm 0.000$
<b>Generation Time (GPU Sec)</b>	SMALL-4	<b><math>5.88 \pm 0.11</math></b>	$2643 \pm 3.11$	$7.66 \pm 0.30$
	SMALL	$7.15 \pm 0.09$	$2671 \pm 1.68$	<b><math>8.92 \pm 0.06</math></b>
	LARGE-4	<b><math>6.00 \pm 0.07</math></b>	$2687 \pm 3.30$	$8.27 \pm 0.19$
	LARGE	<b><math>7.82 \pm 0.10</math></b>	$2571 \pm 2.24$	$12.33 \pm 0.11$

**353**  
**354 5.1 SETTINGS**

**355**  
**356** For evaluation, we conduct experiments on both ISPD07 (newblue04–newblue07 and  
**357** adaptec01–adaptec05) and ISPD98 (ibm01–ibm06) benchmarks (Alpert, 1998). For both benchmarks,  
**358** we use WL, OF, and generation time as the primary evaluation criteria. Our proposed ParetoRouter is  
**359** compared with three classical routing algorithms — GeoSteiner (Juhl et al., 2018), Labyrinth (Kastner  
**360** et al., 2002), FIUTE (Wong & Chu, 2008) and ES (Chu & Wong, 2005) — as well as three SOTA  
**361** ML-based methods: Hubrouter (Du et al., 2023), NeuralSteiner (Liu et al., 2024) and DSBRouter (Shi  
**362** et al., 2025). We also stdy the Correctness Rate (Crrt), Wirelength Ratio (WLR) (Cheng et al., 2022)  
**363** and Generation Time on newblue04–newblue07 as DSBRouter does.

**364**  
**365** It is worth noting that four SOTA solvers (Liu et al., 2024; Feng & Feng, 2025; Chen et al., 2025b; Li  
**366** et al., 2024) are either tailored to benchmarks with different standards (Liang et al., 2024; Dolgov  
**367** et al., 2019) or have not been publicly released (Liu et al., 2024; Chen et al., 2025b; Feng & Feng,  
**368** 2025). Because NeuralSteiner evaluates on the same benchmarks as the two open-source solvers  
**369** (Hubrouter and DSBRouter), we report its results as provided in the original paper to enable a  
**370** fair comparison. Further details on the experimental benchmarks and additional supplementary  
**371** experiments are given in Appendix A.5.2; A.6.1.

**372**  
**373 5.2 CRRT AND WLR ON PARTIAL ISPD07 BENCHMARKS**

**374**  
**375** Following prior ML-based studies Shi et al. (2025), we evaluate ParetoRouter against existing  
**376** ML-based solvers on the ISPD07 benchmarks (newblue04–newblue07), partitioned into four cate-  
**377** gories—small, small-4, large, and large-4—consistent with earlier works. For simplicity, we report  
**378** results only for GAN-Hubrouter, as other Hubrouter variants yield less competitive results. We  
**379** also cancel RL-based post-processing for GAN-Hubrouter to ensure a fair comparison (Shi et al.,  
**380** 2025). As shown in Table 2, ParetoRouter and DSBRouter, as two end-to-end solvers, achieve  
**381** routing solutions with 100% connectivity, whereas Hubrouter without post-processing maintains

378 **Table 3: Wirelength (WL) & overflow (OF) on ISPD98:** classical global routing (GeoSteiner,  
 379 Labyrinth, Flute+RES) and ML-based methods (Hubrouter, NeuralSteiner, DSBRouter).

380 METRICS	381 MODEL	382 IBM01	383 IBM02	384 IBM03	385 IBM04	386 IBM05	387 IBM06
388 WL	GEOSTEINER	<b>60142</b>	<b>165863</b>	<b>145678</b>	<b>162734</b>	<b>409709</b>	<b>275868</b>
	LABYRINTH	75909	201286	187345	195856	420581	341618
	FLUTE+ES*	61492	169251	146287	167547	411936	280477
	HR-VAE	64703 $\pm$ 1498	176492 $\pm$ 6830	159968 $\pm$ 3281	179895 $\pm$ 5274	434942 $\pm$ 2916	301144 $\pm$ 5832
	HR-DPM	66464 $\pm$ 1586	190588 $\pm$ 2337	168454 $\pm$ 2486	183696 $\pm$ 1736	475820 $\pm$ 5516	320423 $\pm$ 2958
	HR-GAN	61056 $\pm$ 151	167545 $\pm$ 236	147050 $\pm$ 208	164298 $\pm$ 326	411857 $\pm$ 472	277977 $\pm$ 514
	NEURALSTEINER*	61735	170405	148036	166648	415684	283727
	DSBROUTER	61435	174016	152862	163942	420464	342349
388 OF	PARETOROUTER (OURS)	63386	174896	155993	171859	482016	307696
	GEOSTEINER	3342	7399	3944	7420	401	8033
	LABYRINTH	<b>292</b>	384	122	1124	<b>0</b>	502
	FLUTE+ES*	3100	7121	3699	6889	317	7821
	HR-VAE	4721 $\pm$ 667	9919 $\pm$ 801	7311 $\pm$ 692	10433 $\pm$ 1299	909 $\pm$ 106	14103 $\pm$ 1684
	HR-DPM	4933 $\pm$ 700	14117 $\pm$ 1309	9344 $\pm$ 818	11471 $\pm$ 871	2390 $\pm$ 126	17229 $\pm$ 1500
	HR-GAN	3491 $\pm$ 64	7481 $\pm$ 31	4010 $\pm$ 42	7551 $\pm$ 22	419 $\pm$ 7	8039 $\pm$ 12
	NEURALSTEINER	2200	3800	2100	2700	18	2833
396 TIME	DSBROUTER	1430	0	4	10	0	11858
	PARETOROUTER (OURS)	1051	<b>0</b>	<b>0</b>	<b>0</b>	<b>0</b>	<b>0</b>
	GEOSTEINER	<b>3.08</b>	6.91	6.80	<b>9.07</b>	<b>7.72</b>	<b>7.66</b>
	LABYRINTH *	7.11	11.08	11.61	42.03	12.70	21.02
	FLUTE+ES *	3.14	<b>4.90</b>	<b>5.88</b>	15.49	7.88	14.11
	HR-VAE	9.66 $\pm$ 0.08	9.69 $\pm$ 0.04	10.19 $\pm$ 0.06	12.93 $\pm$ 0.07	14.58 $\pm$ 0.00	17.28 $\pm$ 0.16
	HR-DPM	1796.09 $\pm$ 38.68	2772.29 $\pm$ 16.83	2936.52 $\pm$ 21.23	3865.21 $\pm$ 25.07	4369.47 $\pm$ 22.56	4965.08 $\pm$ 121.46
	HR-GAN	41.02 $\pm$ 0.51	46.58 $\pm$ 0.56	52.04 $\pm$ 2.35	67.31 $\pm$ 3.51	72.28 $\pm$ 3.72	88.02 $\pm$ 4.45
396	NEURALSTEINER *	27.18	34.79	46.24	50.37	75.99	70.32
	DSBROUTER	4991	5667	8418	10745	11313	11858
	PARETOROUTER (OURS)	42.37	61.96	68.09	91.20	131.80	121.76

\* EXPERIMENTAL RESULTS CITED FROM RAW MANUSCRIPTS.

only about 30% connectivity on average. ParetoRouter preserves WLR performance comparable to both Hubrouter and DSBRouter, while delivering generation time on par with Hubrouter and superior to DSBRouter. Unless otherwise noted, all reported ParetoRouter results were obtained with  $\omega_1 = \omega_2 = 0.5$ .

### 5.3 OF AND WLR ON REAL-WORLD BENCHMARKS

Noting that, as ParetoRouter produces multiple results to satisfy the MOO constraints and OF is relatively important than WL (Liu et al., 2024), we present the WL and OF of the sampled routes that reduce OF the most.

**Routing Results on ISPD98.** Table 3 shows the WL, OF and generation time for all tested methods on ISPD98 benchmarks. For OF, ParetoRouter achieves the most OF reduction. Compared with the SOTA ML-based OF-oriented DSBRouter, ParetoRouter significantly reduces the total OF with a reduction of 36.06% on `ibm01` and 100% on `ibm05` and only uses an average 1/10 generation time of DSBRouter. In terms of wirelength, ParetoRouter does not incur too much loss, with the least 5.24% on `ibm01` and the most 15.00% on `ibm05` compared with GeoSteiner. For generation time, ParetoRouter remains at the same level as NeuralSteiner, but there is still a gap compared to VAE-based Hubrouter and other classical methods.

**Routing Results on ISPD07.** Table 4 shows the WL, OF and generation time for selected tested methods on ISPD07 benchmarks. We keep the GAN-based Hubrouter and skip other variants of Hubrouter, as the GAN variant gets the best outcome. With the size of nets increasing, ParetoRouter and DSBRouter get the most OF reduction compared to all other methods across all tested benchmarks. But, compared to DSBRouter, ParetoRouter does not incur much increase in WL. For WL and generation time, ParetoRouter shows a similar performance compared to ISPD98.

### 5.4 ABLATION STUDY

Series of ablation studies are conducted to study the effectiveness of classifier gradient guidance, proposed loss function, NN module, as well as classifier guidance module.

**Role of Loss Function, NN Module, and Classifier Guidance Module.** To assess the roles of these three components within the ParetoRouter framework, we perform ablations separately. For the loss function, we remove the  $\mathbf{x}_1^{NTHU}$  term. For the NN module and the classifier-gradient guidance module, we remove the corresponding components from the model architecture. We report results for OF and

Table 5: OF & WL w/ varying ablated components on `ibm01`.

Loss	OF	WL	Time
w/o $\mathbf{x}_1^{NCTU}$	1211	63771	41.11
w/o $\mathbf{x}_1^{NTHU}$	1379	63529	42.56
w/o NN	1565	63450	27.07
w/o Guidance	-	-	-
ParetoRouter	<b>1051</b>	<b>63386</b>	42.37

432 Table 4: **Wirelength (WL) & overflow (OF) on ISPD07.** Comparison of 2 selected classical global  
 433 routing (GeoSteiner, Flute+RES) and 3 ML-based methods (Hubrouter, NeuralSteiner, DSBRouter).

METRICS	MODEL	ADAPTEC01	ADAPTEC02	ADAPTEC03	ADAPTEC04	ADAPTEC05
WL	GEOSTEINER	<b>3389601</b>	<b>3209172</b>	<b>9330748</b>	<b>8865643</b>	<b>9784471</b>
	NCTU-GR	3623718	3331725	9598156	9087206	10560933
	NTHROUTER*	5344000	5229000	13101000	12169000	15538000
	FLUTE+ES*	3418461	3235803	9417934	8896007	9886249
	HR-GAN	3407033	3229110	9355980	8888775	9832110
	NEURALSTEINER*	3438717	3247429	9459117	9003952	9915795
	DSBROUTER	12299050	10072054	29478326	24276147	-
OF	PARETOROUTER(OURS)	3522091	3386722	9502199	9021491	10288329
	GEOSTEINER	35945	53848	142254	45050	102300
	NCTU-GR	<b>0</b>	<b>0</b>	<b>0</b>	<b>0</b>	<b>0</b>
	NTHROUTER*	<b>0</b>	<b>0</b>	<b>0</b>	<b>0</b>	<b>0</b>
	FLUTE+ES*	32518	50947	137104	42306	957704
	HR-GAN	35441	53652	142131	45230	102108
	NEURALSTEINER*	82	255	728	97	431
TIME	DSBROUTER	<b>0</b>	<b>0</b>	<b>0</b>	<b>0</b>	-
	PARETOROUTER(OURS)	<b>0</b>	<b>0</b>	<b>0</b>	<b>0</b>	<b>0</b>
	GEOSTEINER	92.70	123.00	371.02	311.19	320.07
	NCTU-GR	<b>8.96</b>	8.31	26.67	21.00	27.45
	NTHROUTER*	10.0	<b>2.10</b>	<b>10.90</b>	<b>2.60</b>	<b>23.00</b>
	FLUTE+ES *	118.48	187.03	396.51	376.72	360.68
	HR-GAN	593.02	780.44	1324.81	1387.01	1384.96
TIME	NEURALSTEINER*	347.20	461.35	1351.91	1138.66	1106.54
	DSBROUTER	65624	119353	115438	125589	-
	PARETOROUTER(OURS)	422.20	469.51	1561.30	1418.09	1500.11

\* EXPERIMENTAL RESULTS CITED FROM RAW MANUSCRIPTS.

WL on `ibm01`. As ParetoRouter cannot guarantee connectivity of the generated routes without the guidance module, we thus don't report results of the NN-ablated ParetoRouter. Table 5 shows that both WL and OF are affected by all three components. When  $x_1^{NTHU}$  is ablated, both OF and WL increase compared to the complete ParetoRouter. When the guidance module is ablated, OF increases markedly, whereas WL decreases, which is reasonable since Geostiner emphasizes WL optimization. Taken together, these results demonstrate the effectiveness of the proposed components within the ParetoRouter framework.

**Influence of  $\omega$ .** To evaluate whether the proposed Das–Dennis–based sampling can effectively manage the Pareto Front, we examine the generated weightings  $\gamma_i \cdot \omega_i$ . The performance variation observed across different weightings underscores the effectiveness of the proposed sampling scheme in controlling the Pareto front. Specifically, we generate 10 uniformly spaced weightings to approximate the Pareto front in GR. For illustration, we select four distinct  $\omega$  for ParetoRouter and conduct experiments on `ibm01`, as `ibm01` exhibits greater variability in OF. As shown in Table 6, as the ratio  $\gamma_1 \cdot \omega_1$  increases from 0 to 0.2, OF reaches its minimum at  $\gamma_1 \cdot \omega_1 = 0.2$  and then increases as  $\gamma_1 \cdot \omega_1$  continues to grow, whereas WL increases monotonically. This pattern suggests that generating more routes (i.e., larger WL) can exacerbate congestion along existing routes. These observations also provide indirect evidence for the effectiveness of the proposed guided sampling in managing the Pareto Front.

## 6 CONCLUSION AND OUTLOOK

In this paper, we introduce *ParetoRouter*, an end-to-end ML-based global router. Equipped with a simple AF loss and a classifier-gradient guidance module subject to multi-objective optimization (MOO) constraints, ParetoRouter can generate either OF-oriented or WL-oriented routes with 100% connectivity for previously unseen large-scale nets in a single step. Experimental results show that ParetoRouter reduces overflow by an average of 68% while incurring only a modest wirelength penalty—particularly on large-scale nets, thereby narrowing the gap between ML-based solvers and practical chip-design applications.

**Limitations and Future Work.** However, the NN module can still predict superfluous routes that are incorporated into the final routing solution, increasing WL. Moreover, compared with some ML-based solvers (e.g., GAN-Hubrouter and NeuralSteiner), ParetoRouter exhibits slightly longer generation times, which we attribute to the additional gradient computations. In future work, we will focus on optimizing generation time and further reducing WL.

Table 6: OF & WL w/ varying  $\gamma_i \cdot \omega_i$  on `ibm01`.

$\gamma_1 \cdot \omega_1$	$\gamma_2 \cdot \omega_1$	OF	WL
0	1	1505	<b>63317</b>
0.2	0.8	<b>1051</b>	63386
0.5	0.5	1493	63563
0.8	0.2	1671	63847
1	0	1682	63956

486 ETHICS STATEMENT  
487488 This paper aims to advance the state of the art in machine learning and artificial intelligence for  
489 electronic design automation (AI4EDA). This paper does not present immediate, direct negative  
490 social impacts. While the research may entail various societal implications, we do not identify any  
491 that warrant specific emphasis in this paper.  
492493 REPRODUCIBILITY STATEMENT  
494495 To ensure the reproducibility of our research, description of our methodology is detailed comprehensively,  
496 so as the implementation and experimental setups. All experimental results in the paper  
497 are reproducible, and the implementation code of ParetoRouter/code for reproducing experimental  
498 results will be fully open sourced on Github upon publication of this paper.  
499500 LLM USAGE STATEMENT  
501502 The contribution of Large Language Models (LLMs) in the work presented in this article is limited to:  
503 1. polishing the given written statements; 2. reviewing the syntax of the written sentences. We declare  
504 that no experimental results, core implementation of our search, core scientific ideas, experimental  
505 designs, or conclusions have been generated or modified by LLMs. The LLM we used is GPT-5,  
506 owned by OpenAI, and no other LLMs were utilized. All authors have reviewed the final version of  
507 the manuscript and take full responsibility for its content and originality.  
508509 REFERENCES  
510

511 Charles J Alpert. The ispd98 circuit benchmark suite. In *Proceedings of the 1998 international*  
512 *symposium on Physical design*, pp. 80–85, 1998.

513 Yen-Jung Chang, Yu-Ting Lee, and Ting-Chi Wang. Nthu-route 2.0: A fast and stable global router.  
514 In *2008 IEEE/ACM International Conference on Computer-Aided Design*, pp. 338–343. IEEE,  
515 2008.

516 Can Chen, Karla-Luise Herpoldt, Chenchao Zhao, Zichen Wang, Marcus Collins, Shang Shang,  
517 and Ron Benson. Affinityflow: Guided flows for antibody affinity maturation. *arXiv preprint*  
518 *arXiv:2502.10365*, 2025a.

519 Can Sam Chen, Christopher Beckham, Zixuan Liu, Xue Steve Liu, and Chris Pal. Parallel-mentoring  
520 for offline model-based optimization. *Advances in Neural Information Processing Systems*, 36:  
521 76619–76636, 2023.

522 Can Sam Chen, Christopher Beckham, Zixuan Liu, Xue Liu, and Christopher Pal. Robust guided  
523 diffusion for offline black-box optimization. *arXiv preprint arXiv:2410.00983*, 2024.

524 Zhao Chen, Vijay Badrinarayanan, Chen-Yu Lee, and Andrew Rabinovich. Gradnorm: Gradient  
525 normalization for adaptive loss balancing in deep multitask networks, 2018. URL <https://arxiv.org/abs/1711.02257>.

526 Zhiyang Chen, Hailong Yao, and Xia Yin. Patlabor: Pareto optimization of timing-driven routing  
527 trees. In *2025 62nd ACM/IEEE Design Automation Conference (DAC)*, pp. 1–7. IEEE, 2025b.

528 Ruoyu Cheng, Xianglong Lyu, Yang Li, Junjie Ye, Jianye Hao, and Junchi Yan. The policy-gradient  
529 placement and generative routing neural networks for chip design. *Advances in Neural Information*  
530 *Processing Systems*, 35:26350–26362, 2022.

531 Minsik Cho, Katrina Lu, Kun Yuan, and David Z Pan. Boxrouter 2.0: Architecture and implementation  
532 of a hybrid and robust global router. In *2007 IEEE/ACM International Conference on Computer-*  
533 *Aided Design*, pp. 503–508. IEEE, 2007.

534 Chris Chu and Yiu-Chung Wong. Fast and accurate rectilinear steiner minimal tree algorithm for vlsi  
535 design. In *Proceedings of the 2005 international symposium on Physical design*, pp. 28–35, 2005.

540 Indraneel Das and John E Dennis. Normal-boundary intersection: A new method for generating the  
 541 pareto surface in nonlinear multicriteria optimization problems. *SIAM journal on optimization*, 8  
 542 (3):631–657, 1998.

543 Sam Daulton, Maximilian Balandat, and Eytan Bakshy. Hypervolume knowledge gradient: a looka-  
 544 head approach for multi-objective bayesian optimization with partial information. In *International  
 545 Conference on Machine Learning*, pp. 7167–7204. PMLR, 2023.

546 Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion schrödinger  
 547 bridge with applications to score-based generative modeling. *Advances in Neural Information  
 548 Processing Systems*, 34:17695–17709, 2021.

549 Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances  
 550 in neural information processing systems*, 34:8780–8794, 2021.

551 Sergei Dolgov, Alexander Volkov, Lutong Wang, and Bangqi Xu. 2019 cad contest: Lef/def based  
 552 global routing. In *2019 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*,  
 553 pp. 1–4. IEEE, 2019.

554 Xingbo Du, Chonghua Wang, Ruizhe Zhong, and Junchi Yan. Hubrouter: Learning global routing  
 555 via hub generation and pin-hub connection. *Advances in Neural Information Processing Systems*,  
 556 36, 2023.

557 Junxi Feng and Lang Feng. A dynamic congestion-aware analytic initial routing flow for vlsi designs.  
 558 In *2025 Conference of Science and Technology of Integrated Circuits (CSTIC)*, pp. 1–3. IEEE,  
 559 2025.

560 Justin Fu and Sergey Levine. Offline model-based optimization via normalized maximum likelihood  
 561 estimation. *arXiv preprint arXiv:2102.07970*, 2021.

562 Joseph L Ganley and James P Cohoon. Routing a multi-terminal critical net: Steiner tree construction  
 563 in the presence of obstacles. In *Proceedings of IEEE International Symposium on Circuits and  
 564 Systems-ISCAS’94*, volume 1, pp. 113–116. IEEE, 1994.

565 Daniel Golovin and Qiuyi Zhang. Random hypervolume scalarizations for provable multi-objective  
 566 black box optimization, 2020. URL <https://arxiv.org/abs/2006.04655>.

567 Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair,  
 568 Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information  
 569 processing systems*, 27, 2014.

570 Nate Gruver, Samuel Stanton, Nathan Frey, Tim GJ Rudner, Isidro Hotzel, Julien Lafrance-Vanassee,  
 571 Arvind Rajpal, Kyunghyun Cho, and Andrew G Wilson. Protein design with guided discrete  
 572 diffusion. *Advances in neural information processing systems*, 36:12489–12517, 2023.

573 Xu Han, Caihua Shan, Yifei Shen, Can Xu, Han Yang, Xiang Li, and Dongsheng Li. Training-  
 574 free multi-objective diffusion model for 3d molecule generation. In *The Twelfth International  
 575 Conference on Learning Representations*, 2023.

576 Maurice Hanan. On steiner’s problem with rectilinear distance. *SIAM Journal on Applied mathematics*,  
 577 14(2):255–265, 1966.

578 Rui Hao, Yici Cai, Qiang Zhou, and Rui Wang. Drplace: A deep learning based routability-driven vlsi  
 579 placement algorithm. *Journal of Computer-Aided Design & Computer Graphics*, 33(4):624–631,  
 580 2021.

581 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in  
 582 neural information processing systems*, 33:6840–6851, 2020.

583 Frank K Hwang. An  $O(n \log n)$  algorithm for rectilinear minimal spanning trees. *Journal of the ACM  
 584 (JACM)*, 26(2):177–182, 1979.

585 Moksh Jain, Sharath Chandra Raparthy, Alex Hernández-García, Jarrid Rector-Brooks, Yoshua  
 586 Bengio, Santiago Miret, and Emmanuel Bengio. Multi-objective gflownets. In *International  
 587 conference on machine learning*, pp. 14631–14653. PMLR, 2023.

594 Yue Jian, Curtis Wu, Danny Reidenbach, and Aditi S Krishnapriyan. General binding affinity  
 595 guidance for diffusion models in structure-based drug design. *arXiv preprint arXiv:2406.16821*,  
 596 2024.

597

598 Jiyan Jiang, Wenpeng Zhang, Shiji Zhou, Lihong Gu, Xiaodong Zeng, and Wenwu Zhu. Multi-  
 599 objective online learning. In *The Eleventh International Conference on Learning Representations*,  
 600 2023. URL <https://openreview.net/forum?id=dKkMnCWfVmm>.

601 Daniel Juhl, David M Warme, Pawel Winter, and Martin Zachariasen. The geosteiner software  
 602 package for computing steiner trees in the plane: an updated computational study. *Mathematical  
 603 Programming Computation*, 10(4):487–532, 2018.

604

605 Ryan Kastner, Elaheh Bozorgzadeh, and Majid Sarrafzadeh. Pattern routing: Use and theory for  
 606 increasing predictability and avoiding coupling. *IEEE Transactions on Computer-Aided Design of  
 607 Integrated Circuits and Systems*, 21(7):777–790, 2002.

608 Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint  
 609 arXiv:1312.6114*, 2013.

610

611 Lingkai Kong, Yuanqi Du, Wenhao Mu, Kirill Neklyudov, Valentin De Bortoli, Dongxia Wu, Haorui  
 612 Wang, Aaron Ferber, Yi-An Ma, Carla P Gomes, et al. Diffusion models as constrained samplers  
 613 for optimization with unknown constraints. *arXiv preprint arXiv:2402.18012*, 2024.

614 MR Kramer and J Van Leeuwen. The complexity of wirerouting and finding minimum area layouts  
 615 for arbitrary vlsicircuits. *Adv. Comput. Res*, 2:129–146, 1984.

616

617 Yann LeCun, Sumit Chopra, Raia Hadsell, M Ranzato, Fujie Huang, et al. A tutorial on energy-based  
 618 learning. *Predicting structured data*, 1(0), 2006.

619 Seul Lee, Jaehyeong Jo, and Sung Ju Hwang. Exploring chemical space with score-based out-of-  
 620 distribution generation. In *International Conference on Machine Learning*, pp. 18872–18892.  
 621 PMLR, 2023.

622

623 Ke Li. A survey of multi-objective evolutionary algorithm based on decomposition: Past and future.  
 624 *IEEE Transactions on Evolutionary Computation*, 2024.

625 Wei Li, Rongjian Liang, Anthony Agnesina, Haoyu Yang, Chia-Tung Ho, Anand Rajaram, and  
 626 Haoxing Ren. Dgr: Differentiable global router. In *Proceedings of the 61st ACM/IEEE Design  
 627 Automation Conference*, pp. 1–6, 2024.

628

629 Ximeng Li, Keyu Peng, Fuxing Huang, and Wenxing Zhu. Pef: Poisson’s equation-based large-scale  
 630 fixed-outline floorplanning. *IEEE Transactions on Computer-Aided Design of Integrated Circuits  
 631 and Systems*, 42(6):2002–2015, 2022a.

632 Yang Li, Liangliang Shi, and Junchi Yan. Iid-gan: an iid sampling perspective for regularizing mode  
 633 collapse. *arXiv preprint arXiv:2106.00563*, 2021.

634

635 Yang Li, Yichuan Mo, Liangliang Shi, and Junchi Yan. Improving generative adversarial networks  
 636 via adversarial learning in latent space. *Advances in neural information processing systems*, 35:  
 637 8868–8881, 2022b.

638

639 Yang Li, Jinpei Guo, Runzhong Wang, and Junchi Yan. From distribution learning in training  
 640 to gradient search in testing for combinatorial optimization. In *Thirty-seventh Conference on  
 641 Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=JtF0ugNMv2>.

642

643 Rongjian Liang, Anthony Agnesina, Wen-Hao Liu, and Haoxing Ren. Gpu/ml-enhanced large scale  
 644 global routing contest. In *Proceedings of the 2024 International Symposium on Physical Design*,  
 645 pp. 269–274, 2024.

646

647 Haiguang Liao, Wentai Zhang, Xuliang Dong, Barnabas Poczos, Kenji Shimada, and Levent Bu-  
 rak Kara. A deep reinforcement learning approach for global routing. *Journal of Mechanical  
 Design*, 142(6):061701, 2020.

648 Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow  
 649 matching for generative modeling. In *The Eleventh International Conference on Learning Repre-*  
 650 *sentations*, 2023.

651 Chih-Hung Liu, Sy-Yen Kuo, DT Lee, Chun-Syun Lin, Jung-Hung Weng, and Shih-Yi Yuan. Obstacle-  
 652 avoiding rectilinear steiner tree construction: A steiner-point-based algorithm. *IEEE Transactions*  
 653 *on Computer-Aided Design of Integrated Circuits and Systems*, 31(7):1050–1060, 2012.

654 Ruizhi Liu, Shizhe Ding, Jingyan Sui, Xingquan Li, Dongbo Bu, et al. Neuralsteiner: Learning  
 655 steiner tree for overflow-avoiding global routing in chip design. In *The Thirty-eighth Annual*  
 656 *Conference on Neural Information Processing Systems*, 2024.

657 Wen-Hao Liu, Wei-Chun Kao, Yih-Lang Li, and Kai-Yuan Chao. Nctu-gr 2.0: Multithreaded collision-  
 658 aware global routing with bounded-length maze routing. *IEEE Transactions on Computer-Aided*  
 659 *Design of Integrated Circuits and Systems*, 32(5):709–722, 2013. doi: 10.1109/TCAD.2012.  
 660 2235124.

661 Shahrzad Mahboubi, Hiroshi Ninomiya, Takeshi Kamio, Hideki Asai, et al. A nesterov’s accelerated  
 662 quasi-newton method for global routing using deep reinforcement learning. *Nonlinear Theory and*  
 663 *Its Applications, IEICE*, 12(3):323–335, 2021.

664 L McMurchie, C Ebeling, and PathFinder. A negotiation-based performance-driven router for fpgas.  
 665 In *Proceedings of the 1995 ACM 3rd International Symposium on Field-Programmable Gate*  
 666 *Arrays*, pp. 111–117, 1995.

667 Gi-Joon Nam, Mehmet Yildiz, David Z Pan, and Patrick H Madden. Ispd placement contest updates  
 668 and ispd 2007 global routing contest. In *Proceedings of the 2007 international symposium on*  
 669 *Physical design*, pp. 167–167, 2007.

670 Walter Lau Neto, Matheus T Moreira, Yingjie Li, Luca Amarù, Cunxi Yu, and Pierre-Emmanuel  
 671 Gaillardon. Slap: A supervised learning approach for priority cuts technology mapping. In *2021*  
 672 *58th ACM/IEEE Design Automation Conference (DAC)*, pp. 859–864. IEEE, 2021.

673 Ji Won Park, Nataša Tagasovska, Michael Maser, Stephen Ra, and Kyunghyun Cho. Botied: Multi-  
 674 objective bayesian optimization with tied multivariate ranks. *arXiv preprint arXiv:2306.00344*,  
 675 2023.

676 Anselm Paulus, Michal Rolínek, Vít Musil, Brandon Amos, and Georg Martius. Comboptnet: Fit the  
 677 right np-hard problem by learning integer programming constraints. In *International Conference*  
 678 *on Machine Learning*, pp. 8443–8453. PMLR, 2021.

679 Aram-Alexandre Pooladian, Heli Ben-Hamu, Carles Domingo-Enrich, Brandon Amos, Yaron Lipman,  
 680 and Ricky TQ Chen. Multisample flow matching: Straightening flows with minibatch couplings.  
 681 *arXiv preprint arXiv:2304.14772*, 2023.

682 Jixiang Qing, Henry B. Moss, Tom Dhaene, and Ivo Couckuyt.  $\{\text{pf}\}^2\text{es}$ : Parallel feasible pareto  
 683 frontier entropy search for multi-objective bayesian optimization, 2023. URL <https://arxiv.org/abs/2204.05411>.

684 Liangliang Shi, Shenhui Zhang, Xingbo Du, Nianzu Yang, and Junchi Yan. Dsbrouter: End-to-end  
 685 global routing via diffusion schrödinger bridge. In *The Forty-second International Conference on*  
 686 *Machine Learning*, 2025.

687 Yunqi Shi, Ke Xue, Song Lei, and Chao Qian. Macro placement by wire-mask-guided black-box  
 688 optimization. *NeurIPS*, 2023.

689 Yang Song, Conor Durkan, Iain Murray, and Stefano Ermon. Maximum likelihood training of  
 690 score-based diffusion models. *Advances in neural information processing systems*, 34:1415–1428,  
 691 2021a.

692 Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben  
 693 Poole. Score-based generative modeling through stochastic differential equations. In *International*  
 694 *Conference on Learning Representations*, 2021b.

702 Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. In *ICML*, pp.  
 703 32211–32252, 2023.

704

705 Nataša Tagasovska, Nathan C Frey, Andreas Loukas, Isidro Hötzl, Julien Lafrance-Vanasse,  
 706 Ryan Lewis Kelly, Yan Wu, Arvind Rajpal, Richard Bonneau, Kyunghyun Cho, et al. A pareto-  
 707 optimal compositional energy-based model for sampling and optimization of protein sequences.  
 708 *arXiv preprint arXiv:2210.10838*, 2022.

709

710 Zhicong Tang, Tiankai Hang, Shuyang Gu, Dong Chen, and Baining Guo. Simplified diffusion  
 711 schrödinger bridge. *CoRR*, abs/2403.14623, 2024. URL <https://doi.org/10.48550/arXiv.2403.14623>.

712

713 Brandon Trabucco, Aviral Kumar, Xinyang Geng, and Sergey Levine. Conservative objective models  
 714 for effective offline model-based optimization. In *International Conference on Machine Learning*,  
 715 pp. 10358–10368. PMLR, 2021.

716

717 Nicolaas G Van Kampen. Stochastic differential equations. *Physics reports*, 24(3):171–228, 1976.

718

719 Jike Wang, Chang-Yu Hsieh, Mingyang Wang, Xiaorui Wang, Zhenxing Wu, Dejun Jiang, Benben  
 720 Liao, Xujun Zhang, Bo Yang, Qiaojun He, et al. Multi-constraint molecular generation based  
 721 on conditional transformer, knowledge distillation and reinforcement learning. *Nature Machine  
 Intelligence*, 3(10):914–922, 2021.

722

723 Shiyu Wang, Xiaojie Guo, Xuanyang Lin, Bo Pan, Yuanqi Du, Yinkai Wang, Yanfang Ye, Ashley  
 724 Petersen, Austin Leitgeb, Saleh AlKhalifa, et al. Multi-objective deep data generation with  
 725 correlated property control. *Advances in neural information processing systems*, 35:28889–28901,  
 2022.

726

727 Yiu-Chung Wong and Chris Chu. A scalable and accurate rectilinear steiner minimal tree algorithm.  
 728 In *2008 IEEE International Symposium on VLSI Design, Automation and Test (VLSI-DAT)*, pp.  
 729 29–34. IEEE, 2008.

730

731 Ke Xue, Rong-Xi Tan, Xiaobin Huang, and Chao Qian. Offline multi-objective optimization. *arXiv  
 preprint arXiv:2406.03722*, 2024.

732

733 Junchi Yan, Xin Liu, Liangliang Shi, Changsheng Li, and Hongyuan Zha. Improving maximum  
 734 likelihood estimation of temporal point process via discriminative and adversarial learning. In  
 735 *IJCAI*, pp. 2948–2954, 2018.

736

737 Yinghua Yao, Yuangang Pan, Jing Li, Ivor Tsang, and Xin Yao. Proud: Pareto-guided diffusion  
 738 model for multi-objective generation. *Machine Learning*, 113(9):6511–6538, 2024.

739

740 Sihyun Yu, Sungsoo Ahn, Le Song, and Jinwoo Shin. Roma: Robust model adaptation for offline  
 741 model-based optimization. *Advances in Neural Information Processing Systems*, 34:4619–4631,  
 2021.

742

743 Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn.  
 744 Gradient surgery for multi-task learning, 2020. URL <https://arxiv.org/abs/2001.06782>.

745

746 Ye Yuan, Can Sam Chen, Zixuan Liu, Willie Neiswanger, and Xue Steve Liu. Importance-aware  
 747 co-teaching for offline model-based optimization. *Advances in Neural Information Processing  
 748 Systems*, 36:55718–55733, 2023.

749

750 Ye Yuan, Youyuan Zhang, Can Chen, Haolun Wu, Zixuan Li, Jianmo Li, James J Clark, and Xue Liu.  
 751 Design editing for offline model-based optimization. *arXiv preprint arXiv:2405.13964*, 2024.

752

753 Ye Yuan, Can Chen, Christopher Pal, and Xue Liu. Paretoflow: Guided flows in multi-objective  
 754 optimization, 2025. URL <https://arxiv.org/abs/2412.03718>.

755 Peng Zhang, Yiyu Qian, and Quan Qian. Multi-objective optimization for materials design with  
 improved nsga-ii. *Materials today communications*, 28:102709, 2021.

756 Qinqing Zheng, Matt Le, Neta Shaul, Yaron Lipman, Aditya Grover, and Ricky TQ Chen. Guided  
757 flows for generative modeling and decision making. *arXiv preprint arXiv:2311.13443*, 2023.  
758

759 Yiheng Zhu, Jialu Wu, Chaowen Hu, Jiahuan Yan, Tingjun Hou, Jian Wu, et al. Sample-efficient  
760 multi-objective molecular optimization with gflownets. *Advances in Neural Information Processing  
761 Systems*, 36:79667–79684, 2023.

762  
763  
764  
765  
766  
767  
768  
769  
770  
771  
772  
773  
774  
775  
776  
777  
778  
779  
780  
781  
782  
783  
784  
785  
786  
787  
788  
789  
790  
791  
792  
793  
794  
795  
796  
797  
798  
799  
800  
801  
802  
803  
804  
805  
806  
807  
808  
809

**Algorithm 2 Training of FM within ParetoRouter**

**Input:** Dataset  $\mathcal{D}$ , epochs  $epos$ , Time scheduler  $T$ , NthuRouter (Chang et al., 2008)  $NTHU$ , NCTU-GR (Liu et al., 2013)  $NCTU$ ;

**Output:** Vector field  $\hat{v}(\mathbf{x}, t; \boldsymbol{\theta})$ ;

- 1: Initialize model parameters  $\boldsymbol{\theta}$ .
- 2: Construct  $q(x_0)$  and  $p_{data}(x_1)$  utilizing  $\mathcal{D}$ .
- 3: **for**  $epo = 1$  to  $epos$  **do**
- 4:     Sample a batch of  $\mathbf{p}_{curr}$  from  $q(x_0)$ .
- 5:     Sample a Gaussian noise  $\mathbf{n}_0$  like  $\mathbf{p}_{curr}$ .
- 6:     Construct  $\mathbf{x}_0$  utilizing  $\mathbf{n}_0$  and  $\mathbf{p}_{curr}$  under  $\mathbf{x}_0 = f_n(\mathbf{n}_0 + \mathbf{p}_{curr})$ .
- 7:     Derive  $\mathbf{x}_1^{NTHU}$  and  $\mathbf{x}_1^{NCTU}$  from  $NTHU(\mathbf{p}_{curr})$  and  $NCTU(\mathbf{p}_{curr})$ , respectively.
- 8:     Sample  $t$  from  $T$ .
- 9:     Derive  $\mathbf{x}_t$  using Eq. 4.
- 10:    Derive output of the vector field  $\hat{v}(\mathbf{x}, t; \boldsymbol{\theta})$ .
- 11:    Compute loss  $\mathcal{L} = \mathbb{E}_{t, p_{data}(\mathbf{x}_0), q(\mathbf{x}_1)} \text{ using Eq. 6.}$
- 12:    Loss.backward().
- 13: **end for**
- 14: **return**  $\hat{v}(\mathbf{x}, t; \boldsymbol{\theta})$ .

## A APPENDIX

## A.1 SUPPLEMENTAL RELATED WORKS

The section reviews works on offline multi-objective optimization and guided generative modeling.

**Offline Multi-Objective Optimization (MOO).** Most MOO research has focused on the online setting, where a black-box function is queried interactively to optimize multiple objectives simultaneously (Jiang et al., 2023; Park et al., 2023; Gruver et al., 2023). By contrast, offline MOO is often more realistic because online queries may be costly or risky (Xue et al., 2024). In the offline regime, a learned predictor serves as the oracle and enables two classical families of methods. Evolutionary algorithms conduct population-based search via iterative parent selection, reproduction, and survivor selection (Zhang et al., 2021; Li, 2024; Yuan et al., 2025). Bayesian optimization instead leverages the predictor within an acquisition function to select promising candidates, iteratively refining the search through sampled evaluations (Daulton et al., 2023; Golovin & Zhang, 2020; Qing et al., 2023). Training the predictor can be further improved by techniques such as COMs (Trabucco et al., 2021), ROMA (Yu et al., 2021), NEMO (Fu & Levine, 2021), ICT (Yuan et al., 2023), Tri-mentoring (Chen et al., 2023), GradNorm (Chen et al., 2018), and PCGrad (Yu et al., 2020), which enhance training efficiency and stability. **Guided Generative Modeling.** A parallel line of work develops generative models that produce samples meeting multiple desired properties. For example, Wang et al. (2021) incorporates structure–property relations into a conditional Transformer to bias generation, and Wang et al. (2022) employs a VAE to recover semantics and property correlations by modeling weights in the latent space. Tagasovska et al. (2022) apply multiple-gradient descent to trained EBMs to synthesize new samples, though training an EBM per property is complex. Han et al. (2023) explores a distinct setting aimed at generating modules that satisfy specified conditions. Zhu et al. (2023) use GFlowNets as acquisition functions, and Jain et al. (2023) integrate multiple objectives into GFlowNets. Yao et al. (2024) induce diversity via hand-crafted penalties rather than uniform weight vectors in a white-box setting. Gruver et al. (2023) investigate online multi-objective optimization within a diffusion framework, using an acquisition function to guide sampling, while Kong et al. (2024) apply multi-objective guidance under diffusion but assume equal weights across properties, which cannot recover the Pareto front. Related work on guided diffusion also targets single-objective optimization (Chen et al., 2024; Yuan et al., 2024). Overall, many of these approaches rely on generators that are either less expressive or difficult to train De Bortoli et al. (2021). In contrast, ParetoRouter pairs a SOTA flow-matching model with classifier-gradient guidance for sampling.

---

864   **Algorithm 3 Classifier Function  $h$**

---

865   **Input:** Vector field  $\hat{v}(\mathbf{x}, t; \boldsymbol{\theta})$ ,  $\mathbf{x}_0$ , scaling factors  $\boldsymbol{\gamma}$ ;

866   **Output:** Weighted score  $\hat{f}_{\boldsymbol{\omega}}(\mathbf{x}_t)$ ;

867   1: Derive weightings  $\boldsymbol{\omega}$  leveraging Das-Dennis approach (Das & Dennis, 1998).

868   2: Derive sampled  $\mathbf{x}'_1$  using the vector field  $\hat{v}(\mathbf{x}, 0; \boldsymbol{\theta})$  and  $\mathbf{x}_0$ .

869   3: Extract predicted routing map  $\mathbf{r}$  from  $\mathbf{x}'_1$ .

870   4: Compute complete routing map  $\mathbf{r}'$  under minimal  $\sum_{i=1}^n \gamma_i \hat{f}_i(\mathbf{x}'_1) \omega_i$ .

871   5: Require  $\mathbf{x}'_1$ .gradient.

872   6: Initialize mask map  $\mathbf{m} = 0$ .

873   7: **for**  $r \in \mathbb{R}' \cup \mathbb{R}$  **do**

874     8:   **if**  $r \in \mathbb{R} \wedge r \in \mathbb{R}'$  **then**

875       9:     pass

876     10:   **else if**  $r \in \mathbb{R} \wedge r \notin \mathbb{R}'$  **then**

877       11:     mask map  $\mathbf{m}(r) = 1$ .

878     12:   **else**  $r \notin \mathbb{R} \wedge r \in \mathbb{R}'$

879       13:     mask map  $\mathbf{m}(r) = -1$ .

880     14:   **end if**

881   15: **end for**

882   16:  $h(\hat{v}(\mathbf{x}, t; \boldsymbol{\theta}) + \mathbf{x}_0) = \mathbf{r} \cdot \mathbf{m}$

883   17: Compute  $\hat{f}_{\boldsymbol{\omega}}(\mathbf{x}_t)$  utilizing Eq. 9.

884   18: **return**  $\hat{f}_{\boldsymbol{\omega}}(\mathbf{x}_t)$ .

---

886

887

888   **A.2 SUPPLEMENTAL ALGORITHMS**

889   This section introduces the training algorithms in Algorithm. 1 (Line 1) and classifier function  $h$  in  
 890   Eq. 7.

891

892   **A.2.1 TRAINING OF FM**

893   Training of the FM within ParetoRouter is summarized in Algorithm 2. We inject scaled noise  $\mathbf{n}_0$   
 894   into the clean data for the following reasons: Through experiments, we find that the intensity of  
 895   the injected noise has a negligible effect on both FM training and Das–Dennis sampling. However,  
 896   without this perturbation the backbone converges poorly (i.e., the trained FM module cannot reliably  
 897   compute, under supervision, the flow that bridges the initial pins  $\mathbf{x}_0$  and the final connected routes  
 898    $\mathbf{x}_1$ .) We therefore conclude that adding appropriately scaled noise to clean data facilitates the model  
 899   training. In practice, we sample Gaussian noise, add it to the clean  $\mathbf{p}_{curr}$ , and then normalize the  
 900   corrupted  $\mathbf{x}_0$  using the normalization function  $f_n$ .

901

902   **A.2.2 CLASSIFIER FUNCTION**

903   ParetoRouter employs a guidance module, first introduced in DSBRouter (Shi et al., 2025), to steer  
 904   route generation, as shown in Algorithm 3. Nevertheless, there are essential differences between the  
 905   guidance used by ParetoRouter and that in DSBRouter.

906   Firstly, DSBRouter applies an SDE-based gradient guidance (Li et al., 2023) to drive DSB generation,  
 907   whereas ParetoRouter operates within the FM framework. DSBRouter adopts SDSB (Tang et al.,  
 908   2024) as its backbone and leverages the series proposed theories in Tang et al. (2024) together with  
 909   the energy-function formalism (LeCun et al., 2006) to justify an SDE-based guidance of the form:

910   
$$p_{\boldsymbol{\theta}}(\mathbf{x}_t \mid \mathbf{x}_{t+1}, \mathbf{g}^*) = Z p_{\boldsymbol{\theta}}(\mathbf{x}_t \mid \mathbf{x}_{t+1}) p(\mathbf{g}^* \mid \mathbf{x}_t) \quad (13)$$

911   where  $\mathbf{g}^*$  denotes the optimal objective score. Since ParetoRouter uses Eq. 7 within the FM framework  
 912   to guide the generation process, the underlying working principles are different, and the DSBRouter  
 913   guidance cannot be directly applied to ParetoRouter.

914   Secondly, ParetoRouter integrates multi-objective optimization (MOO) constraints into the design  
 915   of its classifier-guidance module, whereas DSBRouter considers only the reduction of the objective

918 function (OF). In DSBRouter, the following formula:  
 919

$$\nabla_{\mathbf{x}_{t+1}} \log p(\mathbf{g}^* | \mathbf{x}_{t+1}) = \nabla_{\mathbf{x}_{t+1}} (E_{\mathbf{x}_{t+1} \sim p_r(\mathbf{x}_{t+1} | \mathbf{x}_{t+2})}^o (\eta(\mathbf{x}_{t+1})) - O(\eta(\mathbf{x}_{t+1}))) \quad (14)$$

920 is used to approximate  $p(\mathbf{g}^* | \mathbf{x}_t) \propto \exp([\nabla_{\mathbf{x}_{t+1}} (E_{\mathbf{x}_{t+1} \sim p_r(\mathbf{x}_{t+1} | \mathbf{x}_{t+2})}^o (\eta(\mathbf{x}_{t+1})) - O(\eta(\mathbf{x}_{t+1})))^\top \mathbf{x}_t])$ . Because the gradient of  $O(\eta(\mathbf{x}_{t+1}))$  is zero, Eq. 14 effectively uses the gradient of  $E_{\mathbf{x}_{t+1} \sim p_r(\mathbf{x}_{t+1} | \mathbf{x}_{t+2})}^o (\eta(\mathbf{x}_{t+1}))$  to compute  $\nabla_{\mathbf{x}_{t+1}} \log p(\mathbf{g}^* | \mathbf{x}_{t+1})$ . Here,  $E^o$  in DSBRouter is the expected routing given  $\mathbf{x}_{t+2}$ , and  $E$  is computed under the constraint:  
 921  
 922  
 923  
 924  
 925

$$\arg \min_{\bar{\mathbf{x}}_t} |E_{\bar{\mathbf{x}}_t \sim p_r(\bar{\mathbf{x}}_t | \mathbf{x}_{t+1})}^o (\eta(\bar{\mathbf{x}}_t)) - O(\eta(\bar{\mathbf{x}}_t)) + c(\bar{\mathbf{x}}_t)| \quad (15)$$

926 which indicates that DSBRouter aims solely to minimize the OF. In contrast, ParetoRouter employs a  
 927 neural-network-free classifier that explicitly accounts for MOO constraints:  
 928  
 929

$$\arg \min_{\mathbf{x}'_t} \left| \sum_{i=1}^n \gamma_i \hat{f}_i(\mathbf{x}'_t) \omega_i \right| \quad (16)$$

930 This design aligns with the objective of ParetoRouter’s generation process and enables more diverse  
 931 routing results compared with DSBRouter (Shi et al., 2025).  
 932  
 933

### 934 A.3 DERIVATION OF EQ. 7

935 This section derives Eq. 7. By Lemma 1 in Zheng et al. (2023), the guided vector field takes the form:  
 936

$$\tilde{v}(\mathbf{x}, t; \theta) = a_t \mathbf{x}_t + b_t \nabla_{\mathbf{x}} \log p(\mathbf{x}_t | s) \quad (17)$$

937 where  $a_t = \frac{\dot{a}_t}{a_t}$  and  $b_t = (\dot{a}_t \sigma_t - a_t \dot{\sigma}_t) \frac{\sigma_t}{a_t}$ . Setting  $a_t = t$  and  $\sigma_t = 1 - t$ , Eq. 17 simplifies to:  
 938

$$\tilde{v}(\mathbf{x}, t; \theta) = \frac{1}{t} \mathbf{x}_t + \frac{1-t}{t} \nabla_{\mathbf{x}} \log p(\mathbf{x}_t | s) \quad (18)$$

939 The conditional log-probability function is written as  
 940

$$\log p(\mathbf{x}_t | y) = \log p_{\theta}(\mathbf{x}_t) + \log p(s | h(\mathbf{x}_t, t)) - \log p(s) \quad (19)$$

941 where  $p_{\theta}(\mathbf{x}_t)$  denotes the data distribution learned by the flow matching model and  $p(s | h(\mathbf{x}_t, t))$  is  
 942 the classified property distribution.  
 943

944 Substituting these expressions yields  
 945

$$\tilde{v}(\mathbf{x}, t, y; \theta) = \frac{1}{t} \mathbf{x}_t + \frac{1-t}{t} \nabla_{\mathbf{x}} \log p_{\theta}(\mathbf{x}_t) + \frac{1-t}{t} \nabla_{\mathbf{x}_t} \log p(s | h(\mathbf{x}_t, t)) \quad (20)$$

$$= \tilde{v}(\mathbf{x}, t; \theta) + \frac{1-t}{t} \nabla_{\mathbf{x}_t} \log p(s | h(\mathbf{x}_t, t)) \quad (21)$$

### 950 A.4 PARETOROUTER NETWORK ARCHITECTURE

#### 951 A.4.1 BACKBONE

952 We adopt a symmetric U-Net with time-conditioned residual blocks and attention. Starting from  
 953 64×64 RGB inputs, a 7×7 stem (3→64) feeds four encoder stages with channel widths [64, 64,  
 954 128, 256] and a 512-channel bottleneck. Each encoder stage contains two ResNet blocks with  
 955 adaptive instance normalization (AdaIN) modulated by a learned time embedding, a self-attention  
 956 module (linear attention in the first three stages and full attention at the deepest stage), and a 2×  
 957 downsampling operation, producing the resolution sequence 64→32→16→8→4. Attention uses  
 958 four heads with 32-dimensional subspaces and four learnable memory key-value pairs. At 4×4,  
 959 the bottleneck expands into two hyperconnected residual streams, applies two 512-channel ResNet  
 960 blocks interleaved with full attention, and then reduces back to a single stream. The decoder mirrors  
 961 the encoder: at each resolution it concatenates the corresponding skip features (doubling the channel  
 962 dimensionality), applies two ResNet blocks with 1×1 residual projections, inserts attention (full at  
 963 the first decoder stage, linear thereafter), and upsamples via nearest-neighbor interpolation followed  
 964 by a convolution. RMSNorm is used throughout, linear attention is employed at higher resolutions to  
 965 966 967 968 969 970 971

972 Table 7: **Summary of the test dataset.** We respectively show the scale size, vertical/horizontal  
 973 capacity, number of nets, and average/maximum number of pins for each net.  
 974

CASE	IBM01	IBM02	IBM03	IBM04	IBM05	IBM06	ADA01	ADA02	ADA03	ADA04	ADA05
SIZE	64 × 64	80 × 64	80 × 64	96 × 64	128 × 64	128 × 64	324 × 324	424 × 424	774 × 779	774 × 779	465 × 468
CAP.(V/H)	24/28	44/68	40/60	40/46	84/126	40/66	70/70	80/80	62/62	62/62	110/110
NETS	11507	18429	21621	26163	27777	33354	219794	260159	466295	515304	867441
AVG.PINS	4.31	4.88	4.10	3.86	5.25	4.21	4.29	4.09	4.02	3.71	4.03
MAX.PINS	42	134	55	46	17	35	2271	1935	3713	3974	9863

975 reduce complexity from  $O(n^2)$  to  $O(n)$ , while full attention is retained at the lowest resolution. The  
 976 overall downsampling factor is 16, so input height and width must be divisible by 16.  
 977

978 Temporal conditioning is provided by a 64-D sinusoidal positional encoding passed through a  
 979 two-layer GELU MLP to produce a 256-D time embedding, in each ResNet block, AdaIN applies  
 980 feature-wise scaling and shifting derived from this embedding, i.e.,  $\text{norm}(x) \cdot (\text{scale}+1) + \text{shift}$ . Under  
 981 a conventional instantiation with two 3×3 convolutions per residual block, 1×1 projections where  
 982 required, 128-D attention projections (4×32 heads), and the above time-conditioning MLP, the  
 983 model comprises approximately 21.5 million trainable parameters: 20.9M in convolutional/residual  
 984 pathways, 0.79M in attention projections, and 0.08M in the time-embedding MLP, with normalization  
 985 parameters being negligible.  
 986

#### 987 A.4.2 MOEL PARAMETERS

988 For training, learning rate  $lr$  is fixed to 0.0003 and batch size is set to 256. Training of FM is under  
 989 fp16 precision. For sampling, the sampling steps is set to 1, aligned with the description in 4.2.

#### 990 A.4.3 REASONS FOR CHOOSING NCTU-GR AND NTHURROUTER

991 Previously, single-objective ML-based solvers typically used the solutions produced by a traditional  
 992 solver as supervisory signals. For example, HubRouter Du et al. (2023) is supervised using NCTU-  
 993 GR (Liu et al., 2013), whereas DSBRouter Shi et al. (2025) is supervised using NTHURouter (Chang  
 994 et al., 2008). In our initial experiments, we also used only the outputs of NTHURouter as the  
 995 supervisory signal and observed performance comparable to DSBRouter. This naturally raises the  
 996 question of whether combining the outputs of multiple solvers as a weighted supervisory signal can  
 997 enlarge the effective search space explored by the CFG-guided reverse (denoising) diffusion process,  
 998 thereby improving the results of multi-objective optimization. Motivated by this hypothesis, we  
 999 incorporate supervision from two empirically strong solvers, NTHURouter and NCTU-GR. Our  
 1000 ablation study confirms the effectiveness of the proposed loss design. However, although we do not  
 1001 include additional solver outputs, we argue that the benefit of such supervision is unlikely to grow  
 1002 linearly with the number of solvers. A larger effective search space implies a longer guided generation  
 1003 trajectory, and as the search space grows without bound, the additional exploration becomes almost  
 1004 indistinguishable from the noise injected at time  $t_0$ , ultimately degrading the model’s performance.  
 1005

#### 1006 A.5 EXPERIMENTAL PROTOCOLS

##### 1007 A.5.1 DATASETS AND HARDWARE FOR EXPERIMENTS

1008 We evaluate our approach on the real-world datasets ISPD07 (Nam et al., 2007) and ISPD98 (Alpert,  
 1009 1998). Following Du et al. (2023); Shi et al. (2025); Liu et al. (2024), we build low-overflow expert  
 1010 training datasets by using NthuRouter (Chang et al., 2008) and low-wirelength expert training datasets  
 1011 by using NCTU-GR (Liu et al., 2013) to route a subset of the ISPD07 benchmarks—bigblue1,  
 1012 bigblue2, bigblue3, bigblue4, newblue4, newblue5, newblue6, and newblue7. The dataset pre-  
 1013 processing follows the pipeline outlined in Hubrouter (Du et al., 2023). Interaction of the two solvers’  
 1014 routing results are selected, resulting in nearly 220k samples in total. We initialize the capacities as  
 1015 specified by the benchmarks and route the nets sequentially using the results of Chang et al. (2008).  
 1016 After each capacity update, we generate a condition image comprising the current capacity and the pin  
 1017 locations for the next net, and we simultaneously produce and store the corresponding ground-truth  
 1018 route image. Both images are randomly clipped, when feasible, to a common resolution of 64 × 64.  
 1019 For the evaluation reported in Table. 2, we select newblue1, newblue2 from ISPD07—outside the  
 1020 training set—with a total of 10k samples. A summary of the ISPD98 test cases is provided in Table. 7.  
 1021 We prepare the ISPD98 test cases using the same processing pipeline as in Du et al. (2023).  
 1022

1026 Table 8: **Relative error on ISPD98.** The routing results of Juhl et al. (2018) are treated as the  
 1027 theoretical lower bound. Optimal results are in bold.

MODEL	IBM01	IBM02	IBM03	IBM04	IBM05	IBM06
LOWER BOUND	60142	165863	145678	162734	409709	275868
LABYRINTH	0.262	0.213	0.286	0.203	0.026	0.238
FLUTE+ES	0.023	0.017	<b>0.007</b>	0.027	0.007	0.016
HR-VAE	$0.075 \pm 0.024$	$0.064 \pm 0.04$	$0.098 \pm 0.022$	$0.105 \pm 0.032$	$0.061 \pm 0.007$	$0.091 \pm 0.021$
HR-DPM	$0.105 \pm 0.026$	$0.149 \pm 0.014$	$0.156 \pm 0.017$	$0.128 \pm 0.010$	$0.161 \pm 0.013$	$0.161 \pm 0.010$
HR-GAN	$0.022 \pm 0.002$	<b>0.010 ± 0.001</b>	$0.009 \pm 0.001$	$0.009 \pm 0.002$	<b>0.005 ± 0.001</b>	<b>0.007 ± 0.001</b>
NEURALSTEINER	0.026	0.027	0.016	0.024	0.014	0.028
DSBROUTER	<b>0.021</b>	0.049	0.049	<b>0.007</b>	0.026	0.240
PARETOROUTER	0.053	0.054	0.070	0.056	0.179	0.115

1035 Table 9: **Relative error on ISPD07.** The routing results of Juhl et al. (2018) are treated as the  
 1036 theoretical lower bound. Optimal results are in bold.

MODEL	ADAPTEC01	ADAPTEC02	ADAPTEC03	ADAPTEC04	ADAPTEC05
LOWER BOUND	3389601	3209172	9330748	8865643	9784471
FLUTE+ES	0.008	0.008	0.009	0.003	0.010
HR-GAN	<b>0.005</b>	<b>0.006</b>	<b>0.002</b>	<b>0.002</b>	<b>0.004</b>
NEURALSTEINER	0.014	0.011	0.013	0.015	0.013
DSBROUTER	2.628	2.138	-	2.159	1.738
PARETOROUTER	0.039	0.055	0.018	0.017	0.051

1044 Training of FM and all subsequent experiments are conducted on a machine equipped with an Intel  
 1045 Xeon Platinum 8558 CPU, 8 NVIDIA H200 GPUs (143 GB memory each), and 1600 GB of RAM.

### A.5.2 BASELINES

1048 The baselines referred in Table. 3 are introduced as follows:

1049 1) *GeoSteiner* (Juhl et al., 2018): An optimal RSMT construction solver which get results with SOTA  
 1050 WL.

1052 2) *Labyrinth* (Kastner et al., 2002): A classical routing algorithm that explores how the concept of  
 1053 pattern routing can be utilized to guide the router toward a solution that minimizes interconnect delay  
 1054 while preserving the routability of the circuit.

1055 3) *FLUTE* (Wong & Chu, 2008): A fast and accurate RSMT construction method using a look-up  
 1056 table. It is important to note that this approach can achieve the optimal solution for nets with up to 9  
 1057 degrees.

1059 4) *Edge Shifting* (Chu & Wong, 2005): A fast, practical RSMT-based algorithm that leverages a  
 1060 specialized lookup table for small nets and a refined recursive splitting approach for larger nets.

1061 5) *Hubrouter* (Du et al., 2023): A global router for RST construction based on reinforcement  
 1062 learning. The hub is generated using a diffusion model, followed by reinforcement learning for RST  
 1063 construction.

1064 6) *NeuralSteiner* (Liu et al., 2024): A two-stage global router. The candidate points are predicted by  
 1065 an RCCA-enhanced CNN, and routing is performed by an RST construction algorithm based on a  
 1066 greedy strategy.

1067 7) *DSBRouter* (Shi et al., 2025): An end-to-end global router based on Diffusion Schrödinger Bridge  
 1068 (DSB) which reach SOTA OF reduction, but is behind in generation time.

### A.6 MORE DISCUSSIONS ABOUT TESTED BASELINES.

1072 In this section, we discuss the related error (Du et al., 2023) of selected tested methods and conver-  
 1073 gence of proposed training of FM within ParetoRouter.

#### A.6.1 RELATED ERROR ON ISPD98 AND ISPD07 CASES

1077 To assess improvements relative to the optimal wirelength rather than absolute values, we compare  
 1078 the relative error on ISPD98 and ISPD07 across all tested methods (Table 8 and Table 9). The relative  
 1079 error is defined as  $(WL - LB)/LB$ , where  $LB$  denotes the theoretical lower bound. On the ISPD98  
 benchmarks, ParetoRouter produces slightly more superfluous routes than other SOTA ML-based

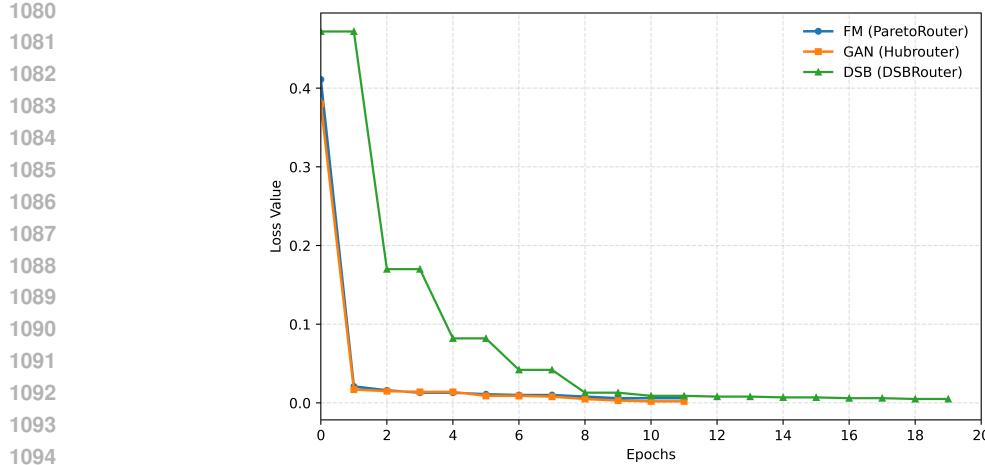


Figure 3: **Convergence of Backbones within different solvers.** Blue line, orange line and green line denote training of FM within ParetoRouter, GAN-based Hubrouter and DSB in DSBRouter.

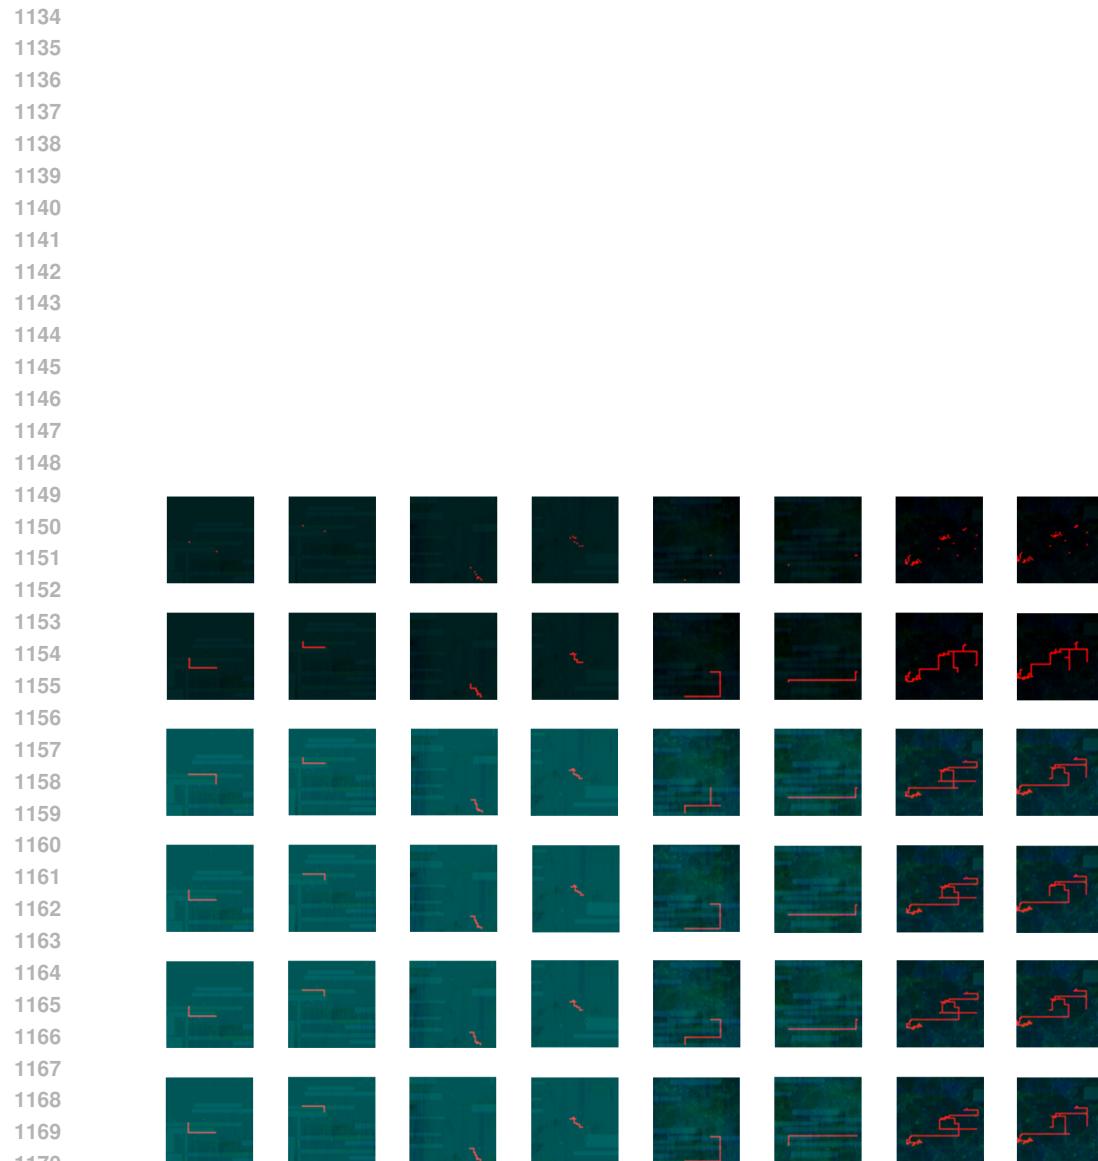
methods (i.e., DSBRouter (Shi et al., 2025), Hubrouter (Du et al., 2023), NeuralSteiner (Liu et al., 2024)). On the ISPD07 benchmarks, GAN-Hubrouter leads, and ParetoRouter eliminates many superfluous routes compared with DSBRouter, but it still lags behind NeuralSteiner, indicating remaining room for improvement.

#### A.6.2 CONVERGENCE OF TRAINING OF FM

To study the convergence of the proposed AF and evaluate the training cost of ParetoRouter, we compare the loss variation during training process across Hubrouter (Du et al., 2023), DSBRouter (Shi et al., 2025) and our proposed ParetoRouter. It needs to be declare that NeuralSteiner (Liu et al., 2024) still disclose the implementation details, so NeuralSteiner is not considered. As shown in Fig. 3, Hubrouter and ParetoRouter can reach convergence in nearly 10 epochs, however, DSBRouter needs 20 epochs to get similar convergence due to alternating training of forward and backward models. This reveals that ParetoRouter can reduce training cost significantly compared to the SOTA ML-based DSBRouter.

#### A.7 ADDITIONAL RESULTS

Generated routing results of ParetoRouter across different net scales, along with initial pins, real routing results (ground truth) are shown in Fig. 4.



1171 **Figure 4: Examples.** Initial pins  $p_{curr}$  (first line), real routing results (second line), generated routing  
 1172 results of ParetoRouter with varying  $\gamma_1 \cdot \omega_1 = 0, 0.2, 0.5, 1$  (third line - sixth line), respectively.  
 1173  
 1174  
 1175  
 1176  
 1177  
 1178  
 1179  
 1180  
 1181  
 1182  
 1183  
 1184  
 1185  
 1186  
 1187