

Practical Layout-Aware Analog/Mixed-Signal Design Automation with Bayesian Neural Networks

Ahmet F. Budak
The University of Texas at Austin
Austin, TX, USA
ahmetfarukbudak@utexas.edu

Keren Zhu
The University of Texas at Austin
Austin, TX, USA
keren.zhu@utexas.edu

David Z. Pan
The University of Texas at Austin
Austin, TX, USA
dpan@ece.utexas.edu

Abstract—The high simulation cost has been a bottleneck for practical analog/mixed-signal design automation. Many learning-based algorithms require thousands of simulated data points, which is impractical for expensive-to-simulate circuits. We propose a learning-based algorithm that can be trained on a small amount of data and, therefore, scalable to task-expensive simulations. Our efficient algorithm solves the layout performance optimization problem where simulation is known to be expensive. Our comprehensive study also addresses the schematic-level sizing problem. For efficient optimization, we utilize Bayesian Neural Networks as a regression model to approximate circuit performance. For layout-aware optimization, we handle the problem as a multi-fidelity optimization problem and improve efficiency by exploiting the correlation between cheaper evaluations. We present three test cases to demonstrate the efficiency of our algorithms. Our tests prove that the proposed approach is more efficient than conventional baselines and state-of-the-art algorithms.

Index Terms—electronic design automation, analog/signal optimization, analog sizing automation, analog layout automation

I. INTRODUCTION

Analog/Mixed-signal (AMS) integrated circuit (IC) design typically follows a process flow visualized in Figure 1. A combination of designer experience and computer simulation feedback is iterated to determine the design that meets the performance requirements. A large portion of design time is spent on the sizing and layout phases, where multiple iterations are possible due to potential loop-backs in the design flow. This is a labor-intensive process in industry practice with little to no automation. To address this costly exercise, a considerable effort in academia is focused on introducing automated solutions.

Analog sizing automation is the task of optimizing AMS design variables, e.g., transistor widths, lengths, resistor, and capacitor values. The aim is to satisfy the performance constraints and optimize the design objective. In general, sizing automation is run through schematic-level simulations. However, AMS IC performance is also sensitive to layout implementation [1]. Especially in the advanced process nodes, layout-induced parasitics may greatly affect the final design performance. Therefore, sizing the AMS design variables considering the layout effects is also crucial.

The majority of the recent sizing and post-layout performance optimization algorithms have simulation feedback in




Fig. 1: AMS Design Flow

the loop. Due to advanced scaling, simulations are required to obtain accurate performance evaluations. Simulation-based AMS automation algorithms adapted various methods from the optimization and Machine Learning (ML) communities. The earlier approaches include population-based methods such as particle swarm optimization [2] and evolutionary algorithms [3]. Although these algorithms have good convergence behavior, they are inefficient in sampling since they explore the design space randomly. To mitigate sample inefficiency, model-based methods gained popularity [4]–[6]. These methods employ surrogate-models between the solution space and performance space and provide efficiency in exploring the solution space. A typical surrogate model is Gaussian Process Regression (GPR) [7], which is a well-studied model in Bayesian Optimization (BO) field [8] and is adapted by several analog sizing algorithms. The main drawback of GPR modeling is its computational complexity.

Recent research trend in analog sizing introduces ML to simulation-based methodology [9]. However, the literature review reveals that most of these methods require thousands of simulation data to train Deep Neural Network (DNN) models that approximate the relations between the design variables and the performance metrics. Therefore, the practicality of these algorithms is severely reduced when the optimization task has a high simulation cost. For example, drawing/generating the layout, extracting the parasitics, and running post-layout simulations is typically an expensive procedure. Therefore, optimization algorithms designed for schematic-level sizing can not be adapted by simply changing how data is generated.

This paper presents a Machine Learning-based simulation-in-the-loop automation method for the AMS design problem. Overall, we formalize two stand-alone recipes for schematic-level sizing and post-layout performance optimization, i.e., layout-aware sizing. We integrate the state-of-the-art analog layout generator, MAGICAL [10], into our flow to handle layout-aware sizing. Our algorithms do not assume the pre-existence of any dataset, and we generate all training data during the optimization. We employ Bayesian Neural Networks (BNN) for modeling design performances. Bayesian Neural Networks allow error quantification, and compared to Deep Neural Networks, BNN are shown to be effective in handling scarce datasets and preventing overfitting [11]. Therefore, BNN can be trained on smaller datasets, significantly improving the practicality and scalability. We also introduce a batch-optimization framework and design space sampling strategy that is compatible with BNN modeling. Further, when optimizing the design variables based on post-layout performance, we exploit the correlation between schematic-level simulations and post-layout simulations. Our algorithm introduces a co-learning scheme that reduces the need for costly post-layout simulations and boosts efficiency even further. We compile our contributions as follows:

- We use Bayesian Neural Network-based modeling to obtain performance approximations. Different learning strategies are adapted for schematic-level sizing and post-layout performance optimization.
- We adopt a scalable sampling strategy and query the optimization batches by utilizing a trust region and Thompson sampling.
- The post-layout sizing is handled as a multi-fidelity optimization problem, and an efficient co-learning strategy is developed.
- The efficiency of the proposed methods is demonstrated on three circuits by providing comparisons to previous state-of-the-art.

The rest of the paper is organized as follows. Section II introduces the backgrounds and previous work. Section III describes our algorithms for handling schematic-level sizing and post-layout performance-based sizing problems. Section IV provides the experiments on circuit examples to demonstrate the efficiency of our algorithms. Finally, Section V concludes the paper.

II. BACKGROUND & RELATED WORK

In this section, we first formally define the AMS design automation problem. Then we review the recent approaches to schematic-level sizing and layout-aware sizing. We summarize the state-of-the-art algorithms' advantages and shortcomings.

A. Problem Formulation

In this paper, we assume that the existence of post-layout performance implies the existence of schematic-level performance values. However, the reverse implication does not hold.

We formulate the AMS schematic-level sizing and layout-aware sizing task as a constrained optimization problem succinctly as below.

$$\begin{aligned} & \text{minimize } f_0(\mathbf{x}) \\ & \text{subject to } f_i(\mathbf{x}) \leq 0 \quad \text{for } i = 1, \dots, m \end{aligned} \quad (1)$$

where, $\mathbf{x} \in \mathbb{R}^d$ is the parameter vector and d is the number of design variables of sizing task. Thus, $\mathbb{R}^d$ is the design space. $f_0(\mathbf{x})$ is the objective performance metric we aim to minimize. Without loss of generality, we denote i^{th} constraint by $f_i(\mathbf{x})$. Notice that if the problem is schematic-level optimization, the f_i values are obtained from schematic simulations. If the problem is post-layout optimization, the f_i values are determined by post-layout simulations.

Through this paper, we will evaluate the quality of a design by defining a Figure of Merit (FoM) in the following form:

$$\text{FoM}(\mathbf{x}) = w_0 \times f_0(\mathbf{x}) + \sum_{i=1}^m \min(1, \max(0, w_i \times f_i(\mathbf{x}))) \quad (2)$$

where w_i is the weighting factor. Note, a $\max(\cdot)$ clipping is used for equating designs after constraints are met, and $\min(\cdot)$ is used to prevent single constraint violation from dominating FoM value.

B. Schematic-Level Sizing

The recent methods for AMS sizing can be collected under two algorithm classes: Bayesian Optimization methods and Deep Learning methods.

Bayesian Optimization methods are tested on AMS problems and are proven to be sample efficient. For example, GASPAD [4] is a hybrid algorithm using a combination of evolutionary space exploration and GPR surrogate-based selection. WEIBO [5] method also employs GPR as a surrogate and introduces a Bayesian Optimization framework where a weighted acquisition function is tailored to comply with the performance-constrained nature of sizing problem. In [6], the authors introduced a multi-fidelity GPR algorithm where the fidelity of the performance is varied with the simulation accuracy. However, this work did not address the layout effects. The disadvantage shared by all GPR models is their cubic complexity to the number of samples, $\mathcal{O}(N^3)$.

Deep Learning based sizing methods includes supervised learning and reinforcement learning (RL) methods [12]–[16]. GCN-RL [13] is a Graph Neural Network algorithm where state representation is built via device index, type, and selected electrical properties. They also propose methods to transfer the optimization experience between different topologies and processes. However, their training graphs show that they use up to 10^4 simulations for sizing academic circuits. AutoCkt [14] is a discrete action space policy gradient method. The RL agent is trained on different optimization tasks where the task is randomly sampled from a predefined set. The trained agent is then tested for the particular tests during deployment. We also observe from the training graphs that AutoCkt requires up to 10^5 simulated samples for training. In [17], the authors successfully applied BNN on multi-objective analog sizing.




Fig. 2: Proposed AMS Automation Framework

However, they did not consider handling constraints, and layout effects are ignored. DNN-Opt [15] is introduced as an RL-inspired supervised learning optimization method that shows high sample efficiency and can be trained during optimization. It uses less than a thousand iterations to optimize academic benchmarks. DNN-Opt does not quantify the variance on approximated values and has no methodic way to balance exploration/exploitation during design space exploration.

C. Post-Layout Based Sizing

Several works in AMS sizing proposed solutions to include layout-induced parasitics. The studies proposed in [18] and [19] embedded a layout generator in the automation loop, and performance metrics to be optimized are obtained through post-layout simulations. However, they did not consider the correlations between the schematic-level and post-layout simulations; therefore, their efficiency is limited. The work in [20] employs a less accurate parasitic prediction during sizing, so the finalized post-layout performance is not guaranteed. In [16], the authors propose a Transfer Learning strategy where a DNN is first trained on schematic-level simulations. Then this knowledge is transferred to improve the learning of a relatively small number of post-layout data. Although this work provides a suggestion to improve the efficiency of post-layout optimization, it requires up to 5×10^3 schematic-level data for initial DNN training, which suffers from the scalability concerns mentioned before. In summary, a scalable solution to optimize AMS design parameters under layout parasitics is yet to be studied.

III. ANALOG/MIXED-SIGNAL IC AUTOMATION FLOW

In this work, we provide solutions to two problems in Analog/Mixed-Signal design automation: schematic-level sizing automation (Task 1) and layout-aware sizing automation (Task 2). The high-level frameworks of proposed solutions to both tasks are summarized together in Figure 2. Section III-A introduces and elaborates on the core principles that complete our proposed automation flow for schematic-level sizing tasks. Then, in section III-B, we explain how to solve the post-layout performance optimization problem efficiently by transforming the BNN learning scheme.

A. Schematic-Level Sizing Automation

We propose a BNN-based sizing algorithm to optimize AMS design on schematic-level simulations. The complete

flow of the proposed approach is summarized in Algorithm 1. The algorithm starts with sampling random points in the design space and simulating them via the SPICE simulator. An initial dataset for training the BNN performance model is built from these samples. Then a trust-region state is initialized before algorithm iterations start. The trust region determines the bounds of the exploration space. The following subsections will provide more details regarding the BNN modeling and trust-region search.

Our algorithm models each performance metric at each optimization iteration with an individual BNN model. Then a batch of samples is collected based on the posterior realization of points lying inside the trust region. Candidate design performance realizations are obtained using the Thompson sampling method, and the candidates are ranked based on the utility values (FoM). A batch of q points is collected, and their real performances are obtained through simulation. The new data is added to the database, and the trust region is updated based on the real simulation outputs of the last batch.

Algorithm 1 BNN-Based Sizing Algorithm

Require: An initial sample set $\mathbf{X}^{\text{init}}$ of N_{init} designs and their evaluations $f(\mathbf{X}^{\text{init}})$

- 1: Assign the solution with maximum utility
- 2: Initialize trust-region state
- 3: **for** $t = 1, 2, \dots, t_{\text{max}}$ **do**
- 4: Train BNN for each performance metric
- 5: Generate r candidate points $x_1, \dots, x_r \in \Omega$ in the trust region.
- 6: **for** $b = 1, 2, \dots, q$ **do**
- 7: For each of the r points of the next batch, sample a realization $\{(\hat{f}(x_i), \hat{f}_1(x_i), \dots, \hat{f}_m(x_i))^T \mid 1 \leq i \leq r\}$ from the posterior over each candidate and add a point of maximum utility to the batch.
- 8: **end for**
- 9: Simulate the q new query points and obtain specs $f(\mathbf{X}_t)$ via SPICE sims
- 10: Update database with new designs and evaluations
- 11: Update trust region state by comparing the current best and $f(\mathbf{X}_t)$
- 12: **end for**
- 13: **return** The design with the highest utility

1) *Performance Modeling with Bayesian Neural Networks:* We base our Bayesian Neural Network regression method on the assumption that the observed function values follow a Gaussian distribution and the probabilistic model on the observations are in the following form:

$$p(f(x) | x, \theta) = \mathcal{N}(\phi(x; \theta), \tau^{-1}) \quad (3)$$

where θ is the BNN parameters, i.e., weights and biases, $\phi(x; \theta)$ is the output of the BNN with parameters θ and τ is the noise parameter. We assign a standard Gaussian prior distribution over each element of the NN parameters, θ , and a Gamma prior over each noise precision, $p(\tau) = \text{Gam}(\tau | a_0, b_0)$. Let define $y_n = f(x_n)$. Given the dataset $\mathcal{D} = \{(x_n, y_n)\}_{n=1}^N$, the joint probability of our model is given by

$$p(\theta, \mathcal{Y}, \tau, | \mathcal{X}) = \mathcal{N}(\text{vec}(\theta) | 0, \mathbf{I}) p(\tau) \prod_{n=1}^N \mathcal{N}(y_n | \phi(x_n), \tau^{-1}) \quad (4)$$

where $\mathcal{X} = \{x_n\}$, $\mathcal{Y} = \{y_n\}$, and $\text{vec}(\cdot)$ is vectorization.

Due to its unbiased, high-quality uncertainty quantification, we use Hamiltonian Monte Carlo (HMC) [21] sampling to perform posterior inference and generate samples of $\theta^i \sim p(\theta | \mathcal{D})$ from the posterior of BNN parameters. Then, using the samples of θ , we make a Gaussian approximation to the function value as follows:

$$\begin{aligned} \mu(f(x) | \mathcal{D}) &= \frac{1}{M} \sum_{i=1}^M \phi(x; \theta^i) \\ \sigma^2(f(x) | \mathcal{D}) &= \frac{1}{M} \sum_{i=1}^M (\phi(x; \theta^i) - \mu(f(x) | \mathcal{D}))^2 + \frac{1}{\tau} \end{aligned} \quad (5)$$

where μ is the mean and σ^2 is the variance approximation.

2) *Trust-Region Search Engine:* We follow the trust region approach introduced in [22] and confine the candidate points locally. The trust-region assigns a localized subset of the search space and proceeds in rounds. We denote the trust region by Ω . In each round, a batch of q designs in Ω are selected by the BNN algorithm and then simulated in parallel. Note that this procedure is easily extended to asynchronous batch evaluations, and we adapt asynchronous evaluation for the multi-fidelity BNN algorithm (will be discussed), where evaluation times show significant differences. The trust-region is centered around the best design explored, i.e., the design with minimum FoM where the ties are handled according to the design objective. This approach mitigates common issues of Bayesian optimization in high-dimensional settings, where popular acquisition functions fail to focus on promising regions and spread out samples due to large prediction uncertainty.

Thompson Sampling-based Exploration: We employ Thompson sampling to obtain performance approximations for untested design candidates. Thompson sampling scales to large batches at low computational cost and has shown to be as effective as the expected improvement acquisition function [22]. Further, the Thompson sampling naturally extends to constrained settings which is usually the case for AMS automation. To select a point for the next batch, we sample r

candidate points in Ω . Let $x_1, \dots, x_r$ be the sampled candidate points. Then we use the predictive model given in Equation 5, and sample a realization $(\hat{f}_0(x_i), \hat{f}_1(x_i), \dots, \hat{f}_m(x_i))^T$ for all x_i with $1 \leq i \leq r$ from the respective posterior distributions on the functions $f_0, f_1, \dots, f_m$. Let $\hat{F} = \{x_i | \hat{f}_j(x_i) \leq 0 \text{ for } 1 \leq j \leq m\}$ be the set of points whose realizations are feasible. If $\hat{F} \neq \emptyset$ holds, we select an $\text{argmin}_{x \in \hat{F}} \hat{f}(x)$, i.e., the design with minimum objective. Otherwise, we select a point of minimum total violation based on the FoM definition given in Equation 2.

Maintaining the trust-region: We initialize a trust region as a hypercube with side length L around the maximum utility point. As the optimization progresses, we track the number of successes n_s and failures n_f since the last time the trust-region is updated. A success is when the algorithm improves the solution quality, and by construction, this point must be inside the trust region. We call it a failure when the last batch of simulated designs is worse than the current best solution. The center C of the trust region is updated as follows. If there exist feasible designs, the one with the minimum objective is assigned as the center. Otherwise, the design with minimum FoM, i.e., minimum scaled constraint violation, is chosen as the center. Therefore, the center of the trust-region is updated every time the design performance is improved. The side length of the trust region is updated as follows: if $n_s = \tau_s$ then the side length is set to $L = \min\{2L, L_{\max}\}$ and we reset $n_s = 0$. If $n_f = \tau_f$, then we set $L = L/2$ and $n_f = 0$. If the side length drops below a set threshold $L_{\min}$, we initialize a new trust region.

B. Post-Layout Performance Optimization

We start our discussion by defining the modifications necessary to automate post-layout performance-based AMS sizing. We tailor the classical sizing flow to include the post-layout effects on the performance during sizing. Instead of optimizing the design variables based on the schematic level simulations, we utilize the layout automation tool MAGICAL to modify performance evaluation steps. The suggested flow is shown in Fig 3. First, an automated layout is generated via MAGICAL to obtain the post-layout performance of each new design. This step is followed by parasitic extraction, and circuit simulations are run on the updated netlist with parasitic elements.




Fig. 3: Post-Layout Performance Based Optimization

However, this new flow is much more expensive than the schematic-level sizing task since the additional steps (layout generation, parasitic extraction, and post-layout simulations)

are typically computationally expensive. Therefore, methods are sought to further increase the efficiency of the (BNN-based) optimization algorithm. As a solution, we treat this problem as a multi-fidelity problem where we have access to two different information sources for calculating circuit performance metrics. Considering that the schematic-level simulations are less accurate approximations of post-layout level simulations, we define these information sources as schematic-level simulations having the lower fidelity and post-layout simulations are the highest fidelity.

We modify the BNN architecture to capture two levels of fidelities (Figure 2) at the output and propose a co-learning scheme similar to multi-task BNN learning [23]. The multi-fidelity BNN model has two output nodes where $\phi(x)[1]$ models the lower fidelity prediction, i.e., schematic-level performance prediction, and $\phi(x)[2]$ models the high fidelity prediction, i.e., post-layout performance prediction. Under the assumption that we have access to two levels of information sources, we denote the new dataset by $\mathcal{D} = \mathcal{D}_1 \cup \mathcal{D}_2$, where $\mathcal{D}_i = \{(\mathbf{x}_n^k, y_n^k)\}_{n=1}^{N_k}$ and the joint probability of the updated BNN model is given by:

$$p(\theta, \mathcal{Y}, \tau, | \mathcal{X}) = p(\theta)p(\tau) \prod_{k=1}^2 \prod_{n=1}^{N_k} \mathcal{N}(y_n^k | \phi(\mathbf{x}_n^k)[k], \tau^{-1}) \quad (6)$$

where $p(\theta) = \mathcal{N}(\text{vec}(\theta) | \mathbf{0}, \mathbf{I})$ and N_k is the number training points in given fidelity level k . The joint probability expression is a combination of the data sourced from both types of simulations; therefore, we utilize the full dataset to train multi-fidelity BNN. In this way, both fidelities are learned together, and the correlations between them are captured due to shared BNN parameters.

To handle the multi-fidelity problem, we adopt the following modifications to Algorithm 1:

- 1) We train multi-fidelity BNN models using the whole history of simulations, $\mathcal{D}$.
- 2) The trust-region centering and length updates are based on the post-layout simulation results, i.e., highest-level fidelity results.
- 3) We determine the candidate selection by modifying the work of [24] where they propose an upper-confidence-bound selection criteria for a single objective BO. We obtain Thompson sampling-based realizations for each fidelity, i.e., $\{\hat{f}_i^{(1)}(\mathbf{x}), \hat{f}_i^{(2)}(\mathbf{x}), \text{ for } i = 0, 1, \dots, m\}$ where $\{1, 2\}$ indicate the fidelity level (schematic-level simulations and post-layout level simulations) and then calculate the low fidelity and high fidelity FoM approximations, $FoM(\hat{f}^L(x))$ and $FoM(\hat{f}^H(x))$ using the corresponding realizations. The candidate selection is queried according to the following utility expression:

$$\mathcal{U}(x) = \max(FoM(\hat{f}^L(x)) - \Delta, FoM(\hat{f}^H(x)))$$

where Δ is the FoM difference between the samples with the best utility at each fidelity. In this step, we take a practical approach to convert two fidelities to each other by defining a reduction term and assign the conservative prediction as the utility value. Finally, the argmin selection is conducted on the

candidate utility values to determine the next batch.

4) The current literature on multi-fidelity Bayesian optimization lacks in handling large number of constraints. Therefore, we randomly assign the fidelity (simulation type) for selected candidates and leave the fidelity selection as future work. Note that this action does not prevent us from studying the benefits of multi-fidelity handling of layout-aware sizing. However, we sacrifice potential cost-aware improvements through intelligent fidelity selection.

IV. EXPERIMENTS

Experiment Setup and Algorithm Settings: We run our tests using 3 different AMS circuits designed with different technologies. A Two-Stage Folded Cascode Operational Transconductance Amplifier (OTA), and a Strong-Arm Latch Comparator are designed with TSMC 180nm process and used to test schematic-level sizing algorithms. Then, we demonstrate the results for layout-aware algorithms on a Two-Stage Miller OTA. This circuit is designed in TSMC 40nm technology since the layout generator used in this work, MAGICAL, is crafted for TSMC 40nm. The schematic designs for these circuits are included in Figure 4.

We run experiments to study the effectiveness of both of the proposed algorithms. First, we test for the schematic-level sizing algorithm, which is given by Algorithm 1, and we refer to our Bayesian Neural Network Based Bayesian Optimization algorithm as "BNN-BO". Then, we run tests for our post-layout performance-based sizing algorithm. Since we utilize a multi-fidelity BNN for this task, we will refer to this algorithm as "MF-BNN-BO".

We implemented several state-of-the-art baseline algorithms to compare and quantify the quality of our proposed algorithms. We selected the baseline algorithms to cover the different categories of approaches. We list the compared baseline algorithms as follows: 1) A differential evolution global optimization algorithm (DE), 2) Bayesian Optimization with weighted expected improvement (BO) [5], and, 3) RL-based sizing algorithm, DNN-Opt [15]. All algorithms are implemented using Python. We implemented DNN-Opt via PyTorch [25], Bayesian Optimization algorithm is implemented using BoTorch [26] package and BNN-BO and MF-BNN-BO are implemented using PyTorch and Hamiltorch [27] packages.

We configured BNN-BO and MF-BNN-BO to evaluate a batch of $q = 8$ designs in parallel. For fairness, DNN-Opt and BO are also configured to do parallel evaluations. Both our algorithms use 200 HMC samples to train BNN models. All BNN models are feedforward neural networks with 2 hidden layers and 100 nodes at each hidden layer. Trust-region is initiated with $L = 0.8$ and L_{min} and L_{max} are chosen to be 0.5^4 and 1.6, respectively. Failure and success tolerances as chosen as $n_f = 2$ and $n_s = 3$. All experiments are run on the same machine using CPU for training learning (DNN and BNN) models. During experiments, the model-based algorithms BO, DNN-Opt, and BNN-BO are run until exploring 500 designs, and DE is run for 5000 new samples.




Fig. 4: Schematics of Tested AMS Circuits




Fig. 5: Folded Cascode OTA Power Modeling




Fig. 6: Folded Cascode Optimization

Schematic-Level Sizing Automation: We tested our algorithm, BNN-BO, and other baseline algorithms on two circuits: two-stage folded cascode ota and strong-arm latch comparator. All transistors in both designs are parameterized for optimization. The Folded Cascode OTA has 20 independent design variables, and the Strong-Arm Latch Comparator has 13 independent design variables after respecting the symmetry constraints. The parameterized device sizes include: transistor lengths & widths, capacitor values, and multipliers (integer valued).

The schematic-level constrained sizing problem for Folded Cascode OTA is defined as follows:

$$\begin{aligned}
 &\text{minimize Power} \\
 &\text{s.t. DC Gain} > 60 \text{ dB} & \text{Settling Time} < 30 \text{ ns} \\
 &\text{CMRR} > 80 \text{ dB} & \text{Saturation Margin} > 50 \text{ mV} \\
 &\text{PSRR} > 80 \text{ dB} & \text{Unity Gain Freq.} > 30 \text{ MHz} \\
 &\text{Out. Swing} > 2.4 \text{ V} & \text{Out. Noise} < 30 \text{ mV}_{\text{rms}} \\
 &\text{Static error} < 0.1 & \text{Phase Margin} > 60 \text{ deg.}
 \end{aligned} \quad (7)$$

The schematic-level constrained sizing problem for Strong-

Arm Latch Comparator is defined as follows:

$$\begin{aligned}
 &\text{minimize Power} \\
 &\text{s.t. Set Delay} < 10 \text{ ns} \\
 &\text{Reset Delay} < 6.5 \text{ ns} \\
 &\text{Input-referred Noise} < 50 \mu\text{Vrms} \\
 &\text{Differential Reset Voltage} < 1 \mu\text{V} \\
 &\text{Differential Set Voltage} > 1.195 \text{ V} \\
 &\text{Positive-Integration Node Reset Voltage} < 60 \mu\text{V} \\
 &\text{Negative-Integration Node Reset Voltage} < 60 \mu\text{V} \\
 &\text{Positive-Output Node Reset Voltage} < 0.35 \mu\text{V} \\
 &\text{Negative-Output Node Reset Voltage} < 0.35 \mu\text{V}.
 \end{aligned} \quad (8)$$

We show the accuracy of the BNN modeling by demonstrating the training metrics. Training Mean Squared Error (MSE) and the logarithmic likelihood of the fitted model are given in Figure 5. Collecting new HMC samples from the posterior increases the likelihood and reduces the training error. We observed very similar training schemes for all other circuits and performance metrics.

We repeat all experiments 10 times to account for the randomization involved in tested algorithms. The statistical results of our tests are shown in Table I. Testing on both

TABLE I: Schematic-Level Sizing Optimization Statistics

Circuit Name	Folded Cascode OTA				Strong-Arm Latch Comparator			
Algorithm	DE	BO	DNN-Opt	BNN-BO	DE	BO	DNN-Opt	BNN-BO
success rate	10/10	7/10	10/10	10/10	7/10	1/10	9/10	10/10
# of simulations	3200	340	151	82	2800	>500	154	68
Min power (mW)	0.75	0.88	0.64	0.60	3.02	3.67	2.45	2.5
Max power (mW)	1.53	1.43	0.8	0.75	4.1	3.67	2.66	2.55
Mean pow. (mW)	1.14	1.19	0.72	0.69	3.44	3.67	2.54	2.52
Modeling time (h)	NA	30	0.6	1.5	NA	17	0.3	0.7
Simulation time (h)	54	2.7	2.7	2.7	72	3.6	3.6	3.6
Total runtime (h)	54	32.7	3.3	4.2	72	20.6	3.9	4.3




Fig. 7: Strong-Arm Latch Comparator Optimization




Fig. 8: Miller OTA Post-Layout Performance Optimization

circuits suggests that BNN-BO can achieve feasible solutions in all runs, and it uses the smallest number of simulations to achieve this. Compared to Differential Evolution (DE), BNN-BO can find feasible solutions using up to 40x less number of simulations. Compared to the closest baseline algorithm, DNN-Opt, BNN-BO reduces the simulation time for finding similar results by up to 55%, proving its high efficiency. It is also demonstrated in Table I that, on average, the final design proposed by BNN-BO draws up to 40% less power. The only disadvantage of BNN-BO to DNN-Opt is the modeling time as DNN-Opt maintains a single DNN model to approximate all performance metrics. Note that all reported times consider the full simulation budget (500 new samples). Therefore, although it takes longer time for BNN-BO to do a single iteration, the required real time for BNN-BO to find a feasible solution is still smaller than other approaches.

In addition to experiment statistics, we further include the FoM convergence curves of both tests in Figure 6 and Figure 7. The y-axis in the graphs represents the total constraint violation; therefore, FoM=0 represents a feasible solution. We observe that, compared to DNN-Opt, BNN-BO has 65% and 33% smaller area under the curve for Folded Cascode OTA and SA Latch Comparator, respectively.

Layout-Aware Design Automation: In order to demonstrate the importance of layout effects on the final performance, we perform experiments on a Miller OTA circuit designed in 40nm technology (Fig. 4). The optimization problem has 17 independent design variables and the optimization problem is defined as follows:

$$\begin{aligned}
 &\text{minimize Power} \\
 &\text{s.t. } \begin{array}{ll} \text{DC Gain} > 45 \text{ dB} & \text{Settling Time} < 100 \text{ ns} \\ \text{CMRR} > 55 \text{ dB} & \text{Saturation Margins} > 50 \text{ mV} \\ \text{PSRR} > 55 \text{ dB} & \text{Unity Gain BW.} > 40 \text{ MHz} \\ \text{Out. Swing} > 1 \text{ V} & \text{RMS Noise} < 400 \text{ uV}_{\text{rms}} \\ \text{Static error} < \%2 & \text{Phase Margin} > 60 \text{ deg.} \end{array}
 \end{aligned} \tag{9}$$

Obtaining the post-layout performance of the Miller OTA is around 9 times more expensive than obtaining the schematic-level performance. Therefore this experiment is to prove the efficiency by utilizing multiple information sources. We initialize all algorithms with 50 high-fidelity random samples, and MF-BNN-BO has additional 50 samples from low-fidelity source (schematic-level simulations). We demonstrate the FoM evolution for the rest of the optimization steps in Figure 8. We observe that our Multi-Fidelity BNN algorithm provides even

more efficiency compared to already efficient BNN-BO. Our analysis shows that the area under the curve is 45% smaller for MF-BNN-BO compared to BNN-BO. Further, we observe that BNN-BO's average best solution after 150 high-fidelity iterations is surpassed by MF-BNN-BO only using 84 simulations. This also implies close to 45% improved efficiency due to utilizing correlations between the schematic-level evaluations and post-layout level evaluations. Note that there is an equal number of schematic-level simulations while running MF-BNN-BO that are not reflected in Figure 8. Considering these simulations, the time efficiency is slightly reduced to around 38%. This efficiency figure serves as a lower-bound since we leave the improvements on fidelity selection as a future work.

V. CONCLUSION

In this work, we presented Bayesian Neural Network-based solutions for schematic-level analog sizing automation and post-layout performance optimization. We targeted the scalability issue of the learning-based automation methods and provided a sample efficient optimization flow. We demonstrated the efficiency of the proposed approaches on academic benchmarks. Compared to the state-of-the-art, we improved the sizing automation efficiency by up to 45%. The Multi Fidelity BNN algorithm analysis proved that utilizing cheaper (schematic-level) simulations reduces the need for expensive (post-layout) simulations considerably, further boosting the efficiency.

REFERENCES

- [1] A. F. Budak, K. Zhu, H. Chen, S. Poddar, L. Zhao, Y. Jia, and D. Z. Pan, "Joint optimization of sizing and layout for ams designs: Challenges and opportunities," in *Proceedings of the 2023 International Symposium on Physical Design*, ser. ISPD '23. New York, NY, USA: Association for Computing Machinery, 2023, p. 84–92.
- [2] A. Vural *et al.*, "Analog circuit sizing via swarm intelligence," *AEU - International Journal of Electronics and Communications*, 2012.
- [3] B. Liu, G. Gielen, and F. V. Fernandez, *Automated Design of Analog and High-frequency Circuits: A Computational Intelligence Approach*. Springer, 2013.
- [4] B. Liu, D. Zhao, P. Reynaert, and G. G. E. Gielen, "Gaspad: A general and efficient mm-wave integrated circuit synthesis method based on surrogate model assisted evolutionary algorithm," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)*, Feb 2014.
- [5] W. Lyu, F. Yang, C. Yan, D. Zhou, and X. Zeng, "Multi-objective bayesian optimization for analog/rf circuit synthesis," in *ACM/IEEE Design Automation Conference (DAC)*, 2018.
- [6] B. He, S. Zhang, Y. Wang, T. Gao, F. Yang, C. Yan, D. Zhou, Z. Bi, and X. Zeng, "A batched bayesian optimization approach for analog circuit synthesis via multi-fidelity modeling," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 42, no. 2, pp. 347–359, 2023.
- [7] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. The MIT Press, 2005.
- [8] J. Snoek, H. Larochelle, and R. P. Adams, "Practical bayesian optimization of machine learning algorithms," in *Conference on Neural Information Processing Systems (NIPS)*, vol. 25, 2012.
- [9] A. F. Budak, S. Zhang, M. Liu, W. Shi, K. Zhu, and D. Z. Pan, "Machine learning for analog circuit sizing," in *Machine Learning for Analog Circuit Sizing*, H. Ren and J. Hu, Eds. Springer, 2022, pp. 307–335.
- [10] B. Xu, K. Zhu, M. Liu, Y. Lin, S. Li, X. Tang, N. Sun, and D. Z. Pan, "MAGICAL: Toward fully automated analog IC layout leveraging human and machine intelligence," in *IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, 2019.
- [11] E. Goan and C. Fookes, "Bayesian neural networks: An introduction and survey," in *Case Studies in Applied Bayesian Data Science*. Springer International Publishing, 2020, pp. 45–87.
- [12] A. Budak, M. Gandara, W. Shi, D. Pan, N. Sun, and B. Liu, "An efficient analog circuit sizing method based on machine learning assisted global optimization," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)*, 2021.
- [13] H. Wang, K. Wang, J. Yang, L. Shen, N. Sun, H. Lee, and S. Han, "GCN-RL circuit designer: Transferable transistor sizing with graph neural networks and reinforcement learning," *ACM/IEEE Design Automation Conference (DAC)*, 2020.
- [14] K. Settaluri, A. Haj-Ali, Q. Huang, K. Hakhamaneshi, and B. Nikolić, "AutoCkt: deep reinforcement learning of analog circuit designs," *IEEE/ACM Proceedings Design, Automation and Test in Europe (DATE)*, 2020.
- [15] A. F. Budak, P. Bhansali, B. Liu, N. Sun, D. Z. Pan, and C. V. Kashyap, "FNN-Opt: an RL inspired optimization for analog circuit sizing using deep neural networks," in *ACM/IEEE Design Automation Conference (DAC)*, 2021.
- [16] J. Liu, S. Su, M. Madhusudan, M. Hassanpourghadi, S. Saunders, Q. Zhang, R. Rasul, Y. Li, J. Hu, A. K. Sharma, S. S. Sapatnekar, R. Harjani, A. Levi, S. Gupta, and M. S.-W. Chen, "From specification to silicon: Towards analog/mixed-signal design automation using surrogate nn models with transfer learning," in *2021 IEEE/ACM International Conference On Computer Aided Design (ICCAD)*, 2021, pp. 1–9.
- [17] Z. Gao, J. Tao, F. Yang, Y. Su, D. Zhou, and X. Zeng, "Efficient performance trade-off modeling for analog circuit based on bayesian neural network," in *2019 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, 2019, pp. 1–8.
- [18] N. Lourenço, R. Martins, A. Canelas, R. Póvoa, and N. Horta, "Aida: Layout-aware analog circuit-level sizing with in-loop layout generation," *Integration*, vol. 55, pp. 316–329, 2016.
- [19] K. Hakhamaneshi, N. Werblun, P. Abbeel, and V. Stojanović, "BagNet: Berkeley analog generator with layout optimizer boosted with deep neural networks," in *IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, 2019.
- [20] M. Liu, W. J. Turner, G. F. Kokai, B. Khailany, D. Z. Pan, and H. Ren, "Parasitic-aware analog circuit sizing with graph neural networks and bayesian optimization," in *IEEE/ACM Proceedings Design, Automation and Test in Europe (DATE)*, 2021.
- [21] R. M. Neal, "MCMC using Hamiltonian dynamics," *Handbook of Markov Chain Monte Carlo*, vol. 54, pp. 113–162, 2010.
- [22] D. Eriksson and M. Poloczec, "Scalable constrained bayesian optimization," in *The 24th International Conference on Artificial Intelligence and Statistics, AISTATS 2021, April 13-15, 2021, Virtual Event*, ser. Proceedings of Machine Learning Research, A. B. 0001 and K. Fukumizu, Eds., vol. 130. PMLR, 2021.
- [23] J. T. Springenberg, A. Klein, S. Falkner, and F. Hutter, "Bayesian optimization with robust bayesian neural networks," in *Advances in Neural Information Processing Systems*, D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, Eds., vol. 29. Curran Associates, Inc., 2016.
- [24] K. Kandasamy, G. Dasarthy, J. B. Oliva, J. Schneider, and B. Poczos, "Gaussian process bandit optimisation with multi-fidelity evaluations," in *Advances in Neural Information Processing Systems*, D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, Eds., vol. 29. Curran Associates, Inc., 2016.
- [25] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "Pytorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems 32*. Curran Associates, Inc., 2019, pp. 8024–8035.
- [26] M. Balandat, B. Karrer, D. R. Jiang, S. Daulton, B. Letham, A. G. Wilson, and E. Bakshy, "BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization," in *Advances in Neural Information Processing Systems 33*, 2020.
- [27] A. D. Cobb and B. Jalaian, "Scaling hamiltonian monte carlo inference for bayesian neural networks with symmetric splitting," *Uncertainty in Artificial Intelligence*, 2021.