MigSizer: Physical-aware Sizing Improving Analog/RF Circuit Migration

Yanfang Liu[1], Peng Xu[1], Yapeng Li[2], Mingjun Wang[1], Leilei Jin[1], Rongliang Fu[1*], Tinghuan Chen[2], Wei Xing[3], Bei Yu[1], Tsung-Yi Ho[1]
[1] The Chinese University of Hong Kong, [2] The Chinese University of Hong Kong (Shenzhen), [3] The University of Sheffield

***Abstract*—Migrating analog circuits with high-quality, silicon-proven layouts from legacy nodes to newer processes typically requires non-linear device resizing to recover electrical performance. However, most existing migration frameworks assume that the legacy placement can be reused as-is while only sizing, routing, or design-rule spacings are updated. In practice, the sizing adjustments needed to meet pre-layout specifications in the target technology perturb device geometries and break this fixed-placement assumption, which can degrade placement quality and even render the migrated layout infeasible. This exposes a largely unexplored perspective: sizing should be co-optimized with placement feasibility during migration, rather than treated as an independent electrical-only task. We present MigSizer, the first physical-aware framework that formulates and addresses this sizing and placement co-optimization problem. MigSizer proposes the Effective Area Utilization Rate (EAUR), a deterministic metric that estimates how a sizing solution affects the feasibility of the migrated placement without full layout generation. Built on EAUR, MigSizer employs a multi-objective Bayesian optimization scheme with customized hypervolume-based acquisition functions that exploit EAUR's deterministic nature to navigate the trade-off between electrical and placement quality. Experiments on diverse analog and RF circuits show that MigSizer attains favorable EAUR and FoM across benchmarks, with up to 34.2% FoM improvement and 8.16× speedup over state-of-the-art baselines, while also yielding better post-migration placement.**

I. Introduction

With the continual advancement of technology nodes, many analog circuits that already possess high-quality, silicon-proven layouts on legacy nodes must be rapidly migrated to newer processes while recovering, or even improving, their performance. During this migration, the non-linear device scaling introduced by the process change forces designers to re-tune transistor widths, lengths, and passive values to restore electrical performance in the target technology. Traditionally, these sizing adjustments are handled independently from the downstream physical design stage. In reality, as illustrated in Fig. 1, this decoupling entirely overlooks that a sizing solution dictated by the electrical specifications may be wholly unsuitable for the placement migrated from the original design: the resized devices no longer fit the legacy placement topology, manifesting as poor area utilization, fragmented whitespace, and, in the worst case, infeasible placements that demand weeks of manual redesign per circuit block. This not only inflates design cost but also squanders the valuable design experience embodied in the legacy layout.

* Corresponding author: Rongliang Fu (rlfu@cse.cuhk.edu.hk).

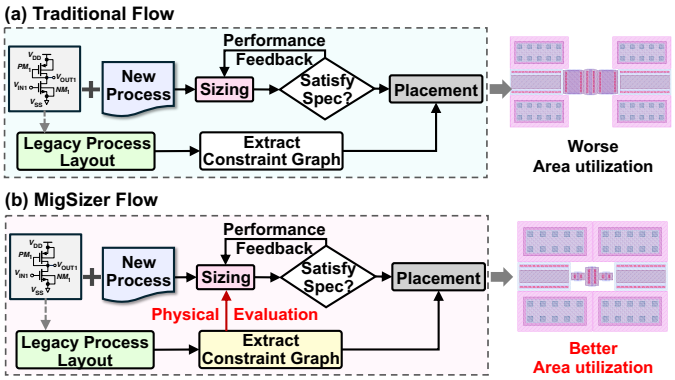

Fig. 1 Comparison between the traditional sizing flow and MigSizer for analog migration. (a) The traditional flow tunes sizing solely toward electrical specifications and then feeds the fixed result to placement; the resulting geometry mismatch against the legacy topology leads to worse area utilization. (b) MigSizer injects physical evaluation into the sizing loop, co-optimizing electrical performance and placement feasibility so that each sizing candidate is admissible by the migrated placement, yielding better area utilization.

Existing analog migration works [1]–[4], which build upon the layout design experience of the legacy process, almost uniformly rely on a strong assumption: the legacy placement is reused as-is, and migration only updates spacing rules, routing, or process-dependent parameters. Constraint-graph-based methods [1], [2] extract horizontal/vertical constraint graphs (HCG/VCG) from the legacy layout and solve a linear program to re-pack the same modules under new design rules, while prototyping-oriented variants explore alternative placement configurations but still treat device dimensions as predetermined inputs. Regardless of the specific method, however, all of them take the sizing result as a fixed input to placement, so their achievable quality is tightly bound by the quality of that sizing result. Among all stages of analog physical design, placement is the most sensitive to such geometric perturbations: it dictates the aspect ratio, whitespace distribution, parasitic length, and the feasibility of the routing and symmetry constraints that follow. An electrically attractive sizing solution can therefore yield an unusable placement, invalidating the fixed-layout assumption on which every downstream migration step depends.

Consequently, the sizing that adapts a design to the new

process is critical to obtaining a satisfactory migration outcome. Meeting the pre-layout specifications in a new node routinely requires non-trivial width/length re-tuning, varactor re-sizing, and passive re-dimensioning. On the sizing side, the literature has progressed along a separate track. Heuristic methods such as evolutionary algorithms [5]–[7], reinforcement learning (RL)-based sizers [8]–[11], and recent sample-efficient Bayesian optimization frameworks [12]–[15] have made significant progress in electrical performance optimization, batch evaluation, and transfer across technologies. Yet all of them target electrical objectives alone; none incorporates any notion of the resulting placement feasibility. When coupled with a migration flow, these sizers readily produce electrically compliant but physically unrealizable solutions, forcing the expensive redesign iterations noted above. This limitation is especially severe in cross-technology migration, where design rules and device aspect ratios differ substantially between source and target nodes.

Taken together, two parallel bodies of work have grown in isolation: migration assumes sizing is fixed, while sizing assumes placement is someone else's concern. Their intersection, namely co-optimizing sizing and placement feasibility during migration so that a sizing solution is admissible by the migrated placement in the first place, has, to our knowledge, never been formulated. This exposes an orthogonal and previously unexplored problem: sizing must be co-optimized with placement feasibility during migration, rather than treated as an independent electrical-only task.

This problem cannot be solved by naively appending placement to the sizing loop. Existing sizers offer no notion of placement feasibility, and, more fundamentally, no lightweight early-stage metric exists to quantify how a sizing candidate will affect the migrated placement before a full layout is generated. To close this gap and ultimately obtain a high-quality migrated layout, the sizing stage must be made aware of whether each candidate is sufficiently feasible and performant for the downstream physical design. Achieving this requires answering two open questions: (i) how to assess placement feasibility during sizing, long before a full layout is generated; and (ii) how to build a unified optimization framework that jointly navigates electrical performance and placement feasibility under a single budget.

To answer both questions, we propose MigSizer, the first framework to bring placement feasibility directly into the analog sizing loop for cross-technology migration. MigSizer proposes the Effective Area Utilization Rate (EAUR), a theoretically-grounded metric computed via constraint-graph-based linear programming (LP), which deterministically quantifies how a candidate sizing solution stresses the migrated placement without full layout synthesis. Built on EAUR, MigSizer develops a multi-objective Bayesian optimization (MOBO) algorithm with a weighted hypervolume (HV)-based acquisition function that exploits the deterministic nature of EAUR to attain a principled exploration–exploitation balance with linear-time complexity.

Overall, our contributions are summarized as follows:

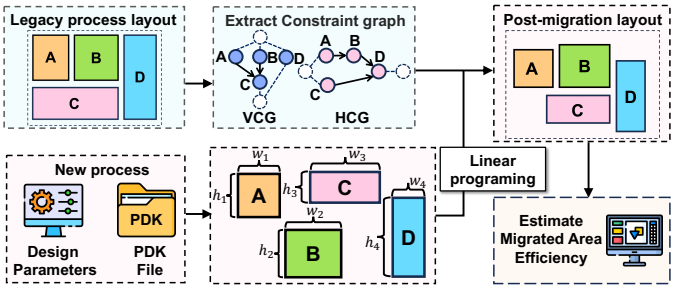

Fig. 2 Constraint graph-based analog migration workflow.

- We propose **MigSizer**, the first physical-aware sizing framework for analog migration, which formulates and solves the previously unexplored sizing and placement co-optimization problem.
- We introduce EAUR, a theoretically grounded metric that quantifies how a sizing candidate affects the feasibility of the migrated placement, at low cost and without full layout generation.
- We develop a customized MOBO algorithm that exploits the evaluation-cost asymmetry between the two objectives via weighted HV-based acquisition functions with linear-time complexity.
- We conduct extensive experiments on diverse circuit benchmarks, showing that MigSizer consistently achieves the highest EAUR and FoM, with up to 34.2% FoM improvement and 8.16× speedup over state-of-the-art (SOTA) baselines.

II. Preliminaries

A. Analog Circuit Sizing

Analog circuit sizing is usually formulated as a constrained optimization problem that maximizes target performance while satisfying design constraints:

$$\begin{aligned} &\max f(\mathbf{x}), \quad \mathbf{x} \in \mathcal{X}, \\ &\text{s.t. } G_i(\mathbf{x}) \le 0, i \in \{1, \cdots, K\}, \end{aligned} \tag{1}$$

where $\mathcal{X} = [0,1]^d$ denotes the normalized design space, $\mathbf{x}$ is the d-dimensional vector of normalized device sizing variables, $f(\mathbf{x})$ is the objective function to be maximized, and $G_i(\mathbf{x})$ is the i-th constraint. The constrained problem is transformed into an unconstrained formulation through a Lagrangian Figure of Merit (FoM) [14]:

$$F(\mathbf{x}) = f(\mathbf{x}) - \sum_{i=1}^{K} \lambda_i \cdot \max\{G_i(\mathbf{x}), 0\}, \tag{2}$$

where $\lambda_i > 0$ denotes the penalty weight for the i-th constraint. Existing analog sizing methods focus predominantly on pre-layout electrical performance, but fail to address physical constraints across technology nodes.

B. Constraint Graph-based Analog Migration

Analog layout migration reuses silicon-proven layouts in new technologies by preserving topological relationships while adapting to new design rules [4]. Beyond fixed-topology

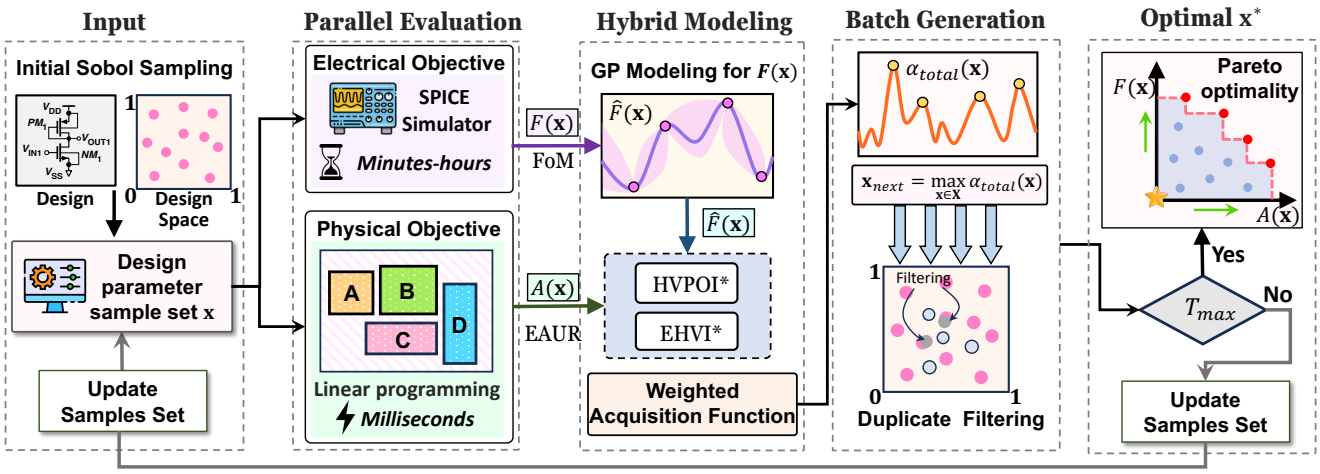

Fig. 3 The overall flow of MigSizer: joint optimization of electrical performance and layout migration feasibility through MOBO with hybrid surrogate modeling. The framework exploits evaluation-cost asymmetry between expensive SPICE simulations for the Lagrangian FoM $F(\mathbf{x})$ (combining circuit performance with electrical constraints) and cheap linear programming for EAUR computation $A(\mathbf{x})$, enabling efficient exploration of the electrical and physical design space with parallel batch evaluation.

migration, layout prototyping methods [1], [2] extend this framework by exploring alternative placement and routing configurations through augmentations of the extracted constraint graphs. As illustrated in Fig. 2, the migration process extracts constraint graphs (HCG and VCG) from the legacy layout, updates them with target technology parameters, and synthesizes the new layout via linear programming or graph optimization methods. While these approaches ensure design rule compliance and can explore layout variations, they operate purely at the layout level and device sizes are assumed to be predetermined. However, in cross-technology migration, the sizing choices critically impact placement feasibility, as poor sizing can lead to excessive area overhead or even layout failure. This motivates the EAUR metric (introduced in the method section) to guide sizing decisions during the optimization process.

C. Multi-objective Bayesian Optimization

MOBO efficiently handles problems with multiple competing objectives by employing surrogate models to minimize expensive evaluations [16]–[18]. Consider K objectives $f_1(\mathbf{x}), f_2(\mathbf{x}), \ldots, f_K(\mathbf{x})$ to be maximized. MOBO builds probabilistic surrogate models, typically Gaussian processes (GPs), for each objective to approximate the expensive-to-evaluate functions. These surrogates provide predictive distributions ϕ_k with mean $\mu_k(\mathbf{x})$ and variance $\sigma_k^2(\mathbf{x})$ for each objective $k \in \{1, \ldots, K\}$.

In MOBO, acquisition functions guide the sequential sampling process by quantifying the potential benefit of evaluating new points. HV-based acquisition functions provide a principled approach to multi-objective optimization by measuring the volume of objective space dominated by the current solution set. The HV indicator $\mathcal{H}(\mathcal{P})$ measures the dominated volume relative to a reference point $\mathbf{r}$, where $\mathcal{P}$ represents the set of non-dominated solutions observed so far. For a candidate solution $\mathbf{x}$ with predicted objective vector $\mathbf{y} = [f_1(\mathbf{x}), \ldots, f_K(\mathbf{x})]^\top$, the exclusive HV contribution is:

$$\mathcal{H}_{exc}(\mathbf{y}, \mathcal{P}) = \mathcal{H}(\mathcal{P} \cup \{\mathbf{y}\}) - \mathcal{H}(\mathcal{P}). \tag{3}$$

The improvement function quantifies the HV gain:

$$I(\mathbf{y}, \mathcal{P}) = \begin{cases} \mathcal{H}_{exc}(\mathbf{y}, \mathcal{P}) & \text{if } \mathbf{y} \text{ is non-dominated} \\ 0 & \text{otherwise} \end{cases} \tag{4}$$

Two commonly used HV-based acquisition functions are the Expected Hypervolume Improvement (EHVI) and Hypervolume Probability of Improvement (HVPOI) [19]:

$$\alpha_{EHVI}(\mathbf{x}) = \int_{\mathbf{y} \in \mathcal{A}} I(\mathbf{y}, \mathcal{P}) \prod_{k=1}^{K} \phi_k(y_k) d\mathbf{y}, \tag{5}$$

$$\alpha_{HVPOI}(\mathbf{x}) = I(\boldsymbol{\mu}(\mathbf{x}), \mathcal{P}) \int_{\mathbf{y} \in \mathcal{A}} \prod_{k=1}^{K} \phi_k(y_k) d\mathbf{y}, \tag{6}$$

where $\boldsymbol{\mu}(\mathbf{x}) = [\mu_1(\mathbf{x}), \ldots, \mu_K(\mathbf{x})]^\top$ is the mean predictions and $\mathcal{A}$ denotes the non-dominated region in objective space.

III. MigSizer Framework

To address the fundamental gap between electrical optimization and physical feasibility in cross-technology migration, we propose MigSizer, a framework that co-optimizes circuit electrical performance and placement physical feasibility from the outset. The overall flow of MigSizer is illustrated in Fig. 3. In this section, we first introduce EAUR, a deterministic metric quantifying placement migration feasibility (Section III-A). We then present our physical-aware MOBO algorithm that simultaneously optimizes EAUR and electrical performance objectives (Section III-B). Finally, we summarize the overall MigSizer framework (Section III-C).

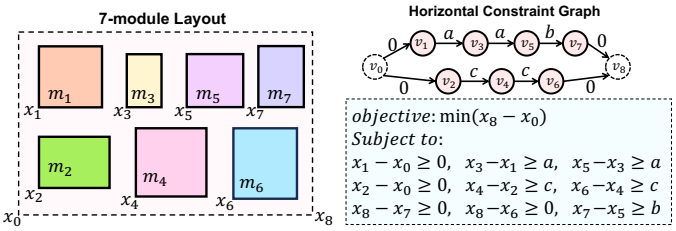

Fig. 4 Migrated layout placement solved by linear programming for a 7-module circuit. The horizontal constraint graph (HCG) encodes relative positioning constraints preserved during migration while respecting target technology design rules.

A. EAUR for Placement Feasibility Assessment

The effectiveness of MigSizer depends on a predictive metric that rapidly assesses placement feasibility during sizing optimization without full layout generation. We propose EAUR, which quantifies area utilization efficiency under target technology geometric constraints.

Metric Definition and Physical Interpretation. EAUR is defined as the ratio of the effective area of devices to the minimum achievable outline area under topological and spacing constraints:

$$\text{EAUR} = \frac{A_{\text{active}}}{A_{\text{total}}} \times 100\%, \tag{7}$$

where $A_{\text{active}} = \sum_{i=1}^{N} w_i \times h_i$ represents total active device area determined by sizing variables $\mathbf{x}$, and A_{total} denotes the minimum outline area required to accommodate all devices while satisfying topological relationships and design-rule spacings.

EAUR provides a proxy for placement quality. Higher values indicate efficient packing with minimal whitespace, while lower values signal excessive area overhead from oversized devices or poor technology compatibility. Unlike electrical metrics requiring expensive SPICE simulations with process corners and Monte Carlo variations, EAUR is computed deterministically via constraint-graph analysis and linear programming.

Constraint-Graph Based Computation. Computing A_{total} without full layout synthesis leverages constraint-graph representations. The high-quality, silicon-proven layout topology of the legacy process is extracted as horizontal and vertical constraint graphs (HCG G_H and VCG G_V). Each directed edge $(i \to j)$ in G_H enforces that module j must be positioned to the right of module i with minimum horizontal separation determined by device widths and target PDK spacing rules. VCG edges similarly enforce vertical ordering.

Fig. 4 illustrates a 7-module example. The HCG encodes constraints such as module m_3 positioned to the right of m_1 with spacing a. By preserving these topological relationships while updating device dimensions according to new sizing $\mathbf{x}$ and enforcing target technology spacing rules, we formulate a linear program for determining the minimum feasible outline. Given candidate sizing $\mathbf{x}$, we solve:

$$\begin{aligned} \min \quad & W + H \\ \text{s.t.} \quad & x_i + w_i \le W, && \forall i \in V, \\ & y_i + h_i \le H, && \forall i \in V, \\ & x_i + w_i + s_{ij}^x \le x_j, && \forall (i \to j) \in E_H, \\ & y_i + h_i + s_{ij}^y \le y_j, && \forall (i \to j) \in E_V, \end{aligned} \tag{8}$$

where decision variables include module positions (x_i, y_i) and outline dimensions (W, H). The first two constraints ensure modules fit within the outline; the latter enforce pairwise separations from constraint graphs. Minimizing $W + H$ maintains linearity while serving as an effective proxy for area minimization.[1] The solution yields $A_{\text{total}} = W \times H$, providing a lower bound on achievable outline area under given topology and design rules.

Computational Efficiency. EAUR evaluation complexity is dominated by solving Equation (8). Constraint-graph extraction occurs once during preprocessing with $\mathcal{O}(|E|)$ complexity. For each sizing candidate, the LP contains $\mathcal{O}(N)$ variables and $\mathcal{O}(|E|)$ constraints. For practically relevant analog circuit designs, the resulting LPs remain small to medium scale and can be solved in well under 10 milliseconds on a standard CPU. This is three to four orders of magnitude faster than SPICE-level simulation, so that EAUR can be evaluated exactly for every candidate at low cost and used directly as one of the two optimization objectives, without any surrogate approximation.

B. MOBO for Physical-aware Analog Sizing

Having established EAUR as a deterministic metric for placement feasibility, we now address the joint optimization of electrical performance and placement quality. We develop a MOBO framework that exploits the pronounced *evaluation-cost asymmetry* between the two objectives to efficiently explore the electrical and placement trade-off.

Problem Formulation with Hybrid Surrogate Modeling. We formulate physical-aware sizing as a bi-objective optimization problem:

$$\mathbf{x}^* = \arg\max_{\mathbf{x} \in \mathcal{X}} [F(\mathbf{x}), A(\mathbf{x})]^\top, \tag{9}$$

where $\mathbf{x} \in \mathcal{X}$ represents device sizing parameters (e.g., transistor widths and lengths and passive values), $F(\mathbf{x})$ denotes the Lagrangian FoM that combines the target performance objective with constraint violations as defined in Equation (2) and $A(\mathbf{x})$ represents the EAUR value computed via Equation (7).

The key observation is the evaluation-cost asymmetry between the two objectives: evaluating $F(\mathbf{x})$ requires expensive SPICE simulation with process corners and Monte Carlo analysis (minutes to hours), whereas $A(\mathbf{x})$ requires only linear programming (milliseconds). Since placement feasibility has never been incorporated into the sizing loop, existing MOBO methods are not tailored to this setting and cannot be applied as-is: they either fit uniform GP surrogates to both objectives,

[1]For analog layouts, minimizing perimeter yields near-optimal area solutions as rectangular outlines naturally tend toward balanced aspect ratios.

introducing needless approximation error on the deterministic EAUR, or fold placement quality into a penalty term, discarding the trade-off information between electrical performance and placement quality.

We propose a hybrid surrogate strategy that exploits this asymmetry. For the expensive performance metric $F(\mathbf{x})$, we construct a GP surrogate:

$$\hat{F}(\mathbf{x}) \sim \mathcal{GP}(m_F(\mathbf{x}), k_F(\mathbf{x}, \mathbf{x}')), \tag{10}$$

where the mean function m_F captures the expected performance and the covariance kernel k_F models spatial correlation. The GP provides predictive distributions $\mathcal{N}(\mu_F(\mathbf{x}), \sigma_F^2(\mathbf{x}))$ that quantify both expected values and uncertainties.

For the cheap EAUR objective $A(\mathbf{x})$, we bypass surrogate modeling entirely and compute it deterministically during optimization. This eliminates approximation errors in feasibility assessment while enabling rapid evaluation of thousands of candidate points without simulation overhead. The hybrid approach fundamentally changes the optimization dynamics: exact placement feasibility guides the search while uncertain performance predictions drive exploration.

Efficient Acquisition Functions with Parallel Batch Evaluation. MOBO requires acquisition functions that balance exploration of unknown regions with exploitation of promising areas. Standard EHVI and HVPOI formulations require expensive multi-dimensional integration over all objective uncertainties. Since EAUR is deterministic, we derive customized variants of EHVI and HVPOI, denoted as EHVI* and HVPOI*, respectively, to exploit this deterministic objective.

For EHVI, the standard double integration over both objectives reduces to a single integral over performance uncertainty. Since EAUR is deterministic with $\phi_A(a) = \delta(a - A(\mathbf{x}))$, applying the sifting property of the Dirac delta yields the simplified EHVI*:

$$\alpha_{EH}^*(\mathbf{x}) = \int_{-\infty}^{\infty} I([f, A(\mathbf{x})], \mathcal{P}) \cdot \phi_F(f)\, df, \tag{11}$$

where $I(\cdot, \mathcal{P})$ measures the hypervolume improvement over the current Pareto set $\mathcal{P}$. This reduces the per-candidate evaluation cost from $\mathcal{O}(n^2)$ (for double Monte Carlo integration) to $\mathcal{O}(n)$ (single integration) for n samples.

Similarly, HVPOI* simplifies to direct evaluation without integration:

$$\alpha_{HP}^*(\mathbf{x}) = I([\mu_F(\mathbf{x}), A(\mathbf{x})], \mathcal{P}), \tag{12}$$

achieving $\mathcal{O}(1)$ complexity per candidate. These simplified formulations preserve the theoretical convergence guarantees of hypervolume-based optimization while dramatically reducing computational cost, making large-scale batch construction tractable.

To balance exploration and exploitation adaptively, we employ a weighted combination:

$$\alpha_{total}(\mathbf{x}) = \omega(t) \cdot \alpha_{EH}^*(\mathbf{x}) + (1 - \omega(t)) \cdot \alpha_{HP}^*(\mathbf{x}) + \beta \cdot \sigma_F(\mathbf{x}), \tag{13}$$

where $\omega(t) = \omega_{min} + (\omega_{max} - \omega_{min}) \cdot \exp(-\gamma \cdot t/T)$ provides smooth transition based on optimization progress, and the uncertainty term $\sigma_F(\mathbf{x})$ ensures continued exploration of high-uncertainty regions.

Algorithm 1 MigSizer Framework

Require: Design space $\mathcal{X}$, evaluation budget T, batch size q
Ensure: Pareto-optimal solutions $\mathcal{P}^*$
Initialize $\mathcal{D}_0 = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^{n_0}$ via Sobol sampling;
for $t = 1$ to T **do**
  Fit GP model $\mathcal{GP}(\mu_F, \sigma_F^2)$ for $F(\mathbf{x})$, compute $A(\mathbf{x})$ deterministically;
  Extract Pareto front $\mathcal{P}_t$ and build representative set;
  Compute weighted acquisition $\alpha_{total}(\mathbf{x})$ via Equation (13)
  $\mathbf{x}_1 = \arg\max_{\mathbf{x}} \alpha_{total}(\mathbf{x})$;
  for $k = 2$ to q **do**
    $\mathbf{x}_k = \arg\max_{\mathbf{x}} [\alpha_{total}(\mathbf{x}) \prod_{i=1}^{k-1} \rho(\mathbf{x}, \mathbf{x}_i)]$;
  end for
  Apply duplicate filtering with Gaussian perturbation if needed;
  Evaluate batch $\mathbf{y}_{new} = [F(\mathbf{x}_{new}), A(\mathbf{x}_{new})]$;
  $\mathcal{D}_t = \mathcal{D}_{t-1} \cup \{(\mathbf{x}_{new}, \mathbf{y}_{new})\}$;
end for
return Final Pareto front $\mathcal{P}^*$;

To fully exploit parallel simulation capabilities in modern analog design environments, we extend our framework with batch evaluation. The key challenge is selecting diverse candidates that avoid redundant exploration. We address this through diversity-weighted acquisition maximization:

$$\mathbf{x}_{k+1} = \arg\max_{\mathbf{x} \in \mathcal{X}} \left[\alpha_{total}(\mathbf{x}) \cdot \prod_{i=1}^{k} \rho(\mathbf{x}, \mathbf{x}_i) \right], \quad k = 1, ..., q-1, \tag{14}$$

where $\rho(\mathbf{x}, \mathbf{x}_i) = \max(\|\mathbf{x} - \mathbf{x}_i\|_2, \epsilon)$ is a repulsion term with $\epsilon = 10^{-4}$ for numerical stability.

Critically, our computational optimizations make batch construction practical: the $\mathcal{O}(1)$ complexity of HVPOI* and linear-time EHVI* allow evaluation of thousands of candidates per batch member in seconds. Because $A(\mathbf{x})$ is computed exactly at low cost, it enters the acquisition functions on equal footing with the surrogate-predicted $F(\mathbf{x})$, so that the two objectives are jointly optimized within a single batch-construction step rather than evaluated in separate stages.

C. The Overall MigSizer Framework

MigSizer employs an iterative framework that jointly optimizes electrical performance and placement feasibility, as summarized in Algorithm 1. The optimization initializes with Sobol sampling [20], evaluating the FoM through SPICE simulation while computing exact EAUR via linear programming. A GP surrogate models the expensive FoM, whereas EAUR is obtained deterministically at low cost, so that thousands of candidates can be screened per iteration. Each iteration constructs a diverse candidate batch using our weighted acquisition functions (Equations (11) to (14)), in which the exact EAUR value and the surrogate-predicted FoM jointly determine the HV improvement, and then evaluates the selected batch in parallel for both objectives before updating the GP and the Pareto set. This hybrid scheme attains linear-time complexity while preserving exact feasibility assessment.

TABLE I Benchmark Circuit Specifications

VCOs			
Circuit	#Modu.	Design Parameters	Optimization Objective
ColVCO	28	$W_N, L, C, W_{\text{var}}, I_{\text{tail}}, V_b$	min DCP (s.t. TR, OutP)
IFVCO	25	$W_N, L_1, L_2, C_1, C_2, W_{\text{var}}$	min DCP (s.t. TR, OutP)
Mixers			
DBAMixer	36	W_{N1}, W_{N2}, R, C	max CGain (s.t. NF, DCP)
SBAMixer	32	$W_{N1}, W_{N2}, R, C, I_{\text{tail}}$	max CGain (s.t. NF, DCP)

The resulting Pareto front characterizes the electrical and placement trade-off, enabling designers to select solutions according to project requirements. By embedding placement feasibility directly into sizing, MigSizer substantially reduces the iterative redesign cycles of conventional migration flows.

IV. Experimental Results

A. Experimental Setting

Benchmark. We evaluate MigSizer on four analog/RF circuits from [21], migrated to the TSMC 45nm PDK. TABLE I summarizes the benchmark specifications. The two VCOs (ColVCO and IFVCO) adopt LC-tank architectures with varactor tuning and minimize DC power (DCP) under tuning-range (TR) and output-power (OutP) constraints. The two mixers (DBAMixer and SBAMixer) maximize conversion gain (CGain) under noise-figure (NF) and DCP constraints. All circuits share the FoM defined in Equation (2). Electrical metrics are obtained from SPICE simulation, while EAUR is computed via linear programming.

Baselines. We evaluate MigSizer against six SOTA methods spanning analog-specific and general multi-objective optimization baselines:

- **Analog sizing methods:** MACE [12] (batch Bayesian optimization), TLMBO [13] (transfer learning-based MOBO), cVTS [14] (constrained Voronoi tree search), and KATO [15] (knowledge alignment transfer)
- **General multi-objective optimizers:** ParEGO [16] (scalarization-based) and EHVI [22] (HV-based)

As none of them natively models placement feasibility, we extend each with EAUR for a fair comparison. For the weighted-sum methods (MACE, cVTS, KATO), we assign equal weights (0.5 for the electrical FoM and 0.5 for EAUR) following their original multi-objective formulation; equal weighting avoids method-specific tuning and keeps the comparison consistent. For the Pareto-based methods (TLMBO, ParEGO, EHVI, MigSizer), we treat EAUR as an additional objective and report the solution with the maximum hypervolume contribution value (HCV) [23].

Experimental Protocol. All methods operate under an identical budget of 20 Sobol initial samples followed by 100 iterations with batch size 4, totaling 420 SPICE evaluations. We report the mean over 10 independent runs with different random seeds, using the published default hyperparameters for every baseline. Experiments run on Ubuntu 20.04 with an NVIDIA RTX 3090 (24GB).

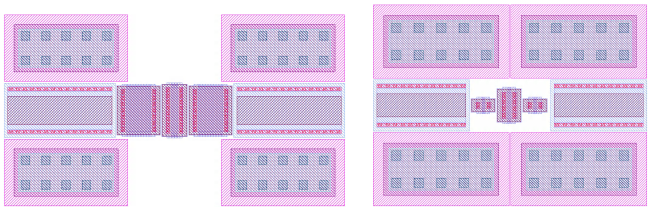

(a) (b)

Fig. 5 SBAMixer's GDSII layouts: (a) without MigSizer (EAUR=78.11%) and (b) with MigSizer (EAUR=93.57%).

B. Physical-aware Sizing Results

TABLE II reports the sizing results across all benchmarks. MigSizer attains favorable EAUR and FoM on every circuit while meeting all electrical specifications.

VCO Circuits. For ColVCO, MigSizer achieves the highest FoM of 2.01, surpassing TLMBO by 3.6% and MACE by 128%. The optimization reduces DC power consumption to 9.86mW, a 9.5% improvement over cVTS, while satisfying all tuning range and output power specifications. MigSizer also achieves the best EAUR at 94.52%, marginally exceeding TLMBO's 94.51%. For IFVCO, MigSizer attains a FoM of 2.54, exceeding MACE by 3.3% and TLMBO by 13.9%, achieving the lowest power consumption at 77.6mW. This represents a 16.6% power reduction compared to TLMBO while achieving the highest EAUR of 88.79% and meeting all constraints.

Mixer Circuits. For DBAMixer, MigSizer achieves a FoM of 3.20, outperforming TLMBO by 3.9% and KATO by 10.0%. The optimization yields 6.85dB conversion gain, 2.7% higher than TLMBO's 6.67dB, while achieving the best EAUR at 80.01%, surpassing cVTS by 0.2 percentage points. For SBAMixer, MigSizer delivers the highest FoM of 4.94, surpassing KATO by 2.3% and cVTS by 34.2%. The conversion gain of 6.52dB represents a 12.6% improvement over TLMBO. Most significantly, MigSizer achieves the best EAUR at 93.57%, exceeding cVTS by 1.75 percentage points while simultaneously delivering superior electrical metrics.

Computational Efficiency. We evaluate computational efficiency by comparing the GPU time under an identical simulation budget (420 simulations). MigSizer demonstrates superior efficiency from two key perspectives. Firstly, compared to the Pareto-based baseline TLMBO, it achieves a dramatic 7.37× to 8.16× speedup, validating that our hybrid modeling strategy is computationally lightweight and highly effective at reducing runtime. Secondly, its runtime (2.77-3.27 minutes) is comparable to the fastest weighted-sum methods, such as MACE and KATO. This combination of a massive speedup over its direct Pareto-based competitor and strong competitiveness against fast, simpler methods confirms the excellent computational efficiency of MigSizer.

TABLE II Physical-aware Sizing Results on All Benchmark Circuits

Method	ColVCO EAUR(%)	ColVCO FoM	ColVCO DCP(mW) min	ColVCO TR(GHz) > 1.17	ColVCO OutP(dBm) > 14.59	ColVCO GPU time (min)	IFVCO EAUR(%)	IFVCO FoM	IFVCO DCP(mW) min	IFVCO TR(GHz) > 1.97	IFVCO OutP(dBm) > 9	IFVCO GPU time (min)
MACE [12]	94.23	0.88	11.0	1.35	15.59	3.07	86.42	2.46	79.0	3.79	9.35	3.30
TLMBO [13]	94.51	1.94	14.1	1.18	**16.07**	24.01	82.04	2.23	93.1	3.31	**10.34**	24.11
cVTS [14]	93.67	0.61	10.9	1.34	15.41	12.72	86.53	2.05	86.9	3.29	9.75	10.33
KATO [15]	93.98	1.29	13.3	1.20	15.79	**3.05**	88.72	2.42	79.1	3.68	9.28	3.33
MigSizer	**94.52**	**2.01**	**9.86**	**1.39**	16.04	3.10	**88.79**	**2.54**	**77.6**	**3.81**	9.47	**3.27**

Method	DBAMixer EAUR(%)	DBAMixer FoM	DBAMixer CGain(dB) max	DBAMixer DCP(mW) < 120	DBAMixer NF(dB) < 12	DBAMixer GPU time (min)	SBAMixer EAUR(%)	SBAMixer FoM	SBAMixer CGain(dB) max	SBAMixer DCP(mW) < 5.5	SBAMixer NF(dB) < 20	SBAMixer GPU time (min)
MACE [12]	73.33	2.74	6.41	102	**11.04**	2.89	81.86	4.58	6.20	3.85	18.99	2.93
TLMBO [13]	76.22	3.08	6.67	**98.4**	11.18	22.61	91.51	4.33	5.79	3.64	**18.74**	22.99
cVTS [14]	79.81	3.18	6.83	98.7	11.31	4.36	91.82	3.68	5.43	4.44	19.15	4.85
KATO [15]	77.52	2.91	6.54	101	11.09	**2.43**	86.36	4.83	6.45	3.75	18.99	2.91
MigSizer	**80.01**	**3.20**	**6.85**	98.8	11.23	2.77	**93.57**	**4.94**	**6.52**	**3.63**	18.96	**2.89**

Note: "min"/"max" indicates the optimization direction of the FoM, and ">"/"<" denotes the corresponding constraint bound that each performance metric must satisfy. For each circuit, the best value in every column is shown in **bold**.

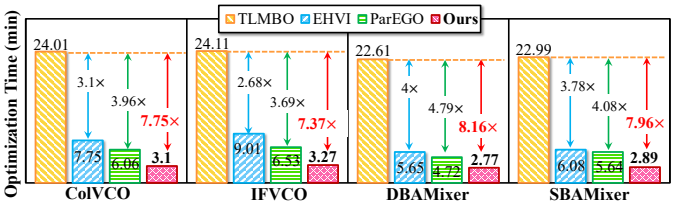

Fig. 6 The comparison of GPU time (min).

TABLE III Ablation Study on Acquisition Functions

Method	DBAMixer	SBAMixer	ColVCO	IFVCO
Full (Ours)	**4.31**	**4.33**	**0.86**	**0.77**
w/o EHVI*	3.95	4.03	0.52	0.60
w/o HVPOI*	3.83	4.23	0.79	0.66
w/o Both	3.08	4.11	0.57	0.64

C. Visualization on Placement Results

Fig. 5 provides direct visual evidence of physical-aware sizing benefits for SBAMixer. Without MigSizer (Fig. 5(a)), the final layout achieves only 78.11% EAUR, showing visible placement conflicts. With MigSizer's physical-aware sizing (Fig. 5(b)), the EAUR increases to 93.57%, a gain of 15.46 percentage points. The resulting layout exhibits significantly better device alignment.

D. Multi-objective Performance Analysis

Computational Efficiency. In terms of computational efficiency, Fig. 6 shows the comparison of GPU time under identical simulation budgets. MigSizer demonstrates significant efficiency improvements over TLMBO, achieving speedups of 7.75×, 7.37×, 8.16×, and 7.96× for ColVCO, IFVCO, DBAMixer and SBAMixer respectively. This acceleration stems from two sources: treating EAUR as an exactly evaluated objective rather than fitting it with a GP, and the simplified acquisition functions reduce the per-candidate cost.

HV Convergence Analysis. To comprehensively evaluate MigSizer's multi-objective optimization effectiveness, we compare it against three representative methods: TLMBO [13], ParEGO [16], and EHVI [22]. Fig. 7 illustrates HV convergence over iterations, averaged across ten random seeds. MigSizer achieves the highest HV values with significantly faster convergence, reaching superior performance within 40-60 iterations compared to 60-80 iterations required by baselines. This shows MigSizer steers the search toward the electrical and placement trade-off front more efficiently.

E. Ablation Study

To validate our composite acquisition function design, we conduct ablation studies by removing different components. TABLE III presents final HV values across all benchmarks. Our complete formulation achieves the highest HV consistently. Removing EHVI* degrades HV by 6.9%-39.5%, while removing HVPOI* causes 2.3%-14.3% drops. Disabling both components reverts to standard MOBO with 5.1%-33.7% degradation, confirming the synergy of our weighted acquisition strategy. Fig. 7 further shows the superior convergence of our complete formulation throughout optimization.

V. Conclusion

We presented MigSizer, the first physical-aware sizing framework for cross-technology analog circuit migration, which formulates and addresses the previously unexplored sizing and placement co-optimization problem. MigSizer proposes EAUR, a deterministic metric that assesses placement

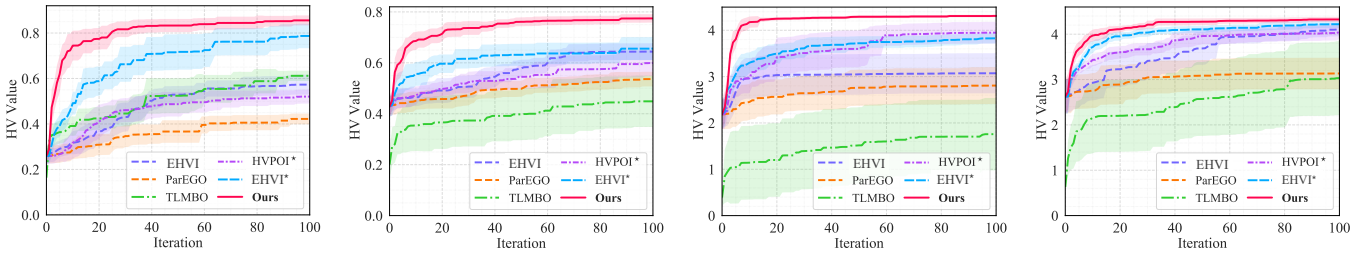

(a) ColVCO (b) IFVCO (c) DBAMixer (d) SBAMixer

Fig. 7 Comparison of HV convergence over optimization iterations on the four benchmark circuits. Each curve reports the mean HV value across 10 independent runs with different random seeds, and the shaded band denotes the standard deviation.

feasibility at the sizing stage without full layout generation, and embeds it into a customized HV-based MOBO algorithm that jointly navigates the electrical and placement trade-off. By coupling a GP surrogate for the expensive FoM with exact EAUR evaluation, the hybrid scheme attains accurate yet efficient optimization under a single budget. Experiments on four analog/RF circuits show that MigSizer consistently attains the highest EAUR and FoM, with up to 34.2% FoM improvement and 8.16× speedup over SOTA baselines, while yielding better post-migration placement. Future work will extend EAUR toward a multi-factor physical score that incorporates symmetry and routing congestion, integrate yield-aware objectives, and couple sizing with downstream placement refinement.

Acknowledgments

The research work described in this paper was conducted in the JC STEM Lab of Intelligent Design Automation funded by The Hong Kong Jockey Club Charities Trust.

References

[1] Y.-P. Weng, H.-M. Chen, T.-C. Chen, P.-C. Pan, C.-H. Chen, and W.-Z. Chen, "Fast analog layout prototyping for nanometer design migration," in *IEEE/ACM International Conference on Computer-Aided Design*, 2011, pp. 517–522.

[2] P.-C. Pan, C.-Y. Chin, H.-M. Chen, T.-C. Chen, C.-C. Lee, and J.-C. Lin, "A fast prototyping framework for analog layout migration with planar preservation," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 34, no. 9, pp. 1373–1386, 2015.

[3] S.-H. Chang, S.-Y. Chen, C.-W. Yang, H.-W. Huang, Y.-C. Yang, W.-L. Chen, C.-N. Liu, and H.-M. Chen, "Mitigating power and process variation for analog CIM design migration," in *International Conference on Synthesis, Modeling, Analysis and Simulation Methods and Applications to Circuit Design*, 2024, pp. 1–4.

[4] X. Dong and L. Zhang, "Lithography-aware analog layout retargeting," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 35, no. 2, pp. 232–245, 2016.

[5] B. Liu, Y. Wang, Z. Yu, L. Liu, M. Li, Z. Wang, J. Lu, and F. V. Fernández, "Analog circuit optimization system based on hybrid evolutionary algorithms," *Integration*, vol. 42, no. 2, pp. 137–148, 2009.

[6] R. A. Vural and T. Yildirim, "Analog circuit sizing via swarm intelligence," *AEU-International Journal of Electronics and Communications*, vol. 66, no. 9, pp. 732–740, 2012.

[7] R. Zhou, P. Poechmueller, and Y. Wang, "An analog circuit design and optimization system with rule-guided genetic algorithm," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 41, no. 12, pp. 5182–5192, 2022.

[8] H. Wang, K. Wang, J. Yang, L. Shen, N. Sun, H.-S. Lee, and S. Han, "GCN-RL circuit designer: Transferable transistor sizing with graph neural networks and reinforcement learning," in *ACM/IEEE Design Automation Conference*, 2020, pp. 1–6.

[9] A. F. Budak, P. Bhansali, B. Liu, N. Sun, D. Z. Pan, and C. V. Kashyap, "DNN-Opt: An RL inspired optimization for analog circuit sizing using deep neural networks," in *ACM/IEEE Design Automation Conference*, 2021, pp. 1219–1224.

[10] K. Settaluri, Z. Liu, R. Khurana, A. Mirhaj, R. Jain, and B. Nikolic, "Automated design of analog circuits using reinforcement learning," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 41, no. 9, pp. 2794–2807, 2022.

[11] J. Gao, W. Cao, and X. Zhang, "RoSE: Robust analog circuit parameter optimization with sampling-efficient reinforcement learning," in *ACM/IEEE Design Automation Conference*, 2023, pp. 1–6.

[12] W. Lyu, F. Yang, C. Yan, D. Zhou, and X. Zeng, "Batch Bayesian optimization via multi-objective acquisition ensemble for automated analog circuit design," in *International Conference on Machine Learning*, 2018, pp. 3306–3314.

[13] Z. Zhang, T. Chen, J. Huang, and M. Zhang, "A fast parameter tuning framework via transfer learning and multi-objective Bayesian optimization," in *ACM/IEEE Design Automation Conference*, 2022, pp. 133–138.

[14] A. Zhao, X. Wang, Z. Lin, Z. Bi, X. Li, C. Yan, F. Yang, L. Shang, D. Zhou, and X. Zeng, "cVTS: A constrained Voronoi tree search method for high dimensional analog circuit synthesis," in *ACM/IEEE Design Automation Conference*, 2023, pp. 1–6.

[15] W. W. Xing, W. Fan, Z. Liu, Y. Yao, and Y. Hu, "KATO: Knowledge alignment and transfer for transistor sizing of different design and technology," in *ACM/IEEE Design Automation Conference*, 2024, pp. 1–6.

[16] J. Knowles, "ParEGO: A hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems," *IEEE Transactions on Evolutionary Computation*, vol. 10, no. 1, pp. 50–66, 2006.

[17] D. Hernández-Lobato, J. Hernandez-Lobato, A. Shah, and R. Adams, "Predictive entropy search for multi-objective Bayesian optimization," in *International Conference on Machine Learning*, 2016, pp. 1492–1501.

[18] S. Belakaria, A. Deshwal, N. K. Jayakodi, and J. R. Doppa, "Uncertainty-aware search framework for multi-objective Bayesian optimization," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 06, 2020, pp. 10 044–10 052.

[19] I. Couckuyt, D. Deschrijver, and T. Dhaene, "Fast calculation of multiobjective probability of improvement and expected improvement criteria for Pareto optimization," *Journal of Global Optimization*, vol. 60, no. 3, pp. 575–594, 2014.

[20] L. Paulin, D. Coeurjolly, J.-C. Iehl, N. Bonneel, A. Keller, and V. Ostromoukhov, "Cascaded Sobol' sampling," *ACM Transactions on Graphics*, vol. 40, no. 6, Dec. 2021.

[21] A. Mehradfar, X. Zhao, Y. Huang, E. Ceyani, Y. Yang, S. Han, H. Aghasi, and S. Avestimehr, "FALCON: An ML framework for fully automated layout-constrained analog circuit design," in *Advances in Neural Information Processing Systems*, vol. 38, 2025, pp. 60 631–60 664.

[22] M. Balandat, B. Karrer, D. 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*, vol. 33, 2020, pp. 21 524–21 538.

[23] K. Bringmann and T. Friedrich, "An efficient algorithm for computing hypervolume contributions," *Evolutionary Computation*, vol. 18, no. 3, pp. 383–402, 2010.