ALARM: Attention and Line Search Based Analog Circuit Optimization on a Riemannian Manifold

Undergraduate Capstone Technical Report

Soohyun Choi[1], Minseok Oh[2], and Ickhyun Song[1,*]

[1]Department of Electronic Engineering, Hanyang University, Seoul, Republic of Korea
[2]AI Team, Deeply Inc., Seoul, Republic of Korea

*Corresponding author.
Contact: petersun0221@hanyang.ac.kr

Document status	Undergraduate capstone technical report documenting a project completed at Hanyang University in 2024.
Authors	Soohyun Choi, Minseok Oh, and Ickhyun Song.
Associated honors	Outstanding Graduation Project Award, Hanyang University; Excellence Award, College of Engineering Capstone Design.
Revision	August 2026. Language, layout, and the quotient-space interpretation of the original design-manifold construction were clarified; the method and experimental claims follow the archived 2024 manuscript.

Abstract

Analog-circuit sizing requires repeated high-cost simulator evaluations and can therefore be slow even when the number of design variables is modest. This report documents ALARM, an undergraduate capstone project that combines a learned SPICE surrogate, Riemannian line search, an attention network, and a buffer of previously trained critics. Its geometric premise is designer-relative: two circuits are regarded as equivalent when they have the same values for all performance attributes selected by the designer, including specifications, figures of merit, and physical cost when relevant. The resulting equivalence classes form a quotient design space on which distance is measured through the selected performance representation. The critic supplies local performance estimates and gradients, while the attention network adjusts the coordinate-wise metric used by the line-search agent. Archived 2024 experiments optimized an 11-parameter voltage-controlled oscillator in a 130-nm CMOS setting. Across ten runs, ALARM satisfied the project criteria in 10/10 runs with 33.2 SPICE evaluations on average, compared with 9/10 runs and 145.9 evaluations for DNN-OPT. The Euclidean-metric ablation succeeded in 7/10 runs. These results demonstrate the intended capstone proof of concept: performance-aware geometry can guide simulator-efficient circuit optimization.

1 Project Overview

The schematic optimization of analog and radio-frequency circuits is typically a black-box problem. A vector of transistor dimensions and bias values is proposed, a SPICE simulator evaluates circuit performance, and the design is adjusted until all specifications are met. The simulator is far more expensive than the arithmetic needed to propose a new design, so the number of SPICE calls is a practical measure of sample efficiency.

Prior automated sizing methods include evolutionary search, Gaussian-process optimization, and reinforcement-learning-inspired optimizers. Differential Evolution (DE) is straightforward but may require many evaluations [2]. AutoCkt and GCN-RL learn policies for circuit design [3, 4]. DNN-OPT instead trains a neural surrogate and searches through the resulting learned objective surface [1]. ALARM was developed as a capstone extension of this surrogate-based direction.

The archived project makes four engineering contributions:

(1) a critic network that approximates the scalarized SPICE objective;
(2) an attention module that produces coordinate-dependent weights;
(3) a performance-induced Riemannian metric coupled with backtracking line search; and
(4) a critic buffer that reuses locally relevant surrogate parameters.

The design manifold is intentionally defined from the viewpoint of a circuit designer rather than from raw component parameters alone. If two physical designs agree on every target the designer has chosen to consider, then the optimizer has no reason to distinguish them. Such designs belong to the same equivalence class. A performance-induced distance is therefore a pseudometric on the raw design set and a genuine metric on the quotient by these equivalence classes. This quotient interpretation is the geometric premise of the original manuscript.

2 Optimization Setting

2.1 Bounded design variables

Let $\xi \in \mathcal{X}$ denote a circuit design with fixed topology, and let $z = \psi(\xi) \in \mathbb{R}^d$ be its design chart. Its coordinates are bounded physical parameters such as bias voltages, transistor widths, and transistor lengths. Because these coordinates may carry different units, ALARM does not define circuit similarity through an ordinary Euclidean distance between raw parameter vectors. A simulator instead maps a design to measured specifications,

$$s(z) = \big[P,\, V_{\text{peak}},\, f_c,\, L,\, \text{TR}, \ldots\big], \tag{1}$$

where P is power, V_{peak} is peak output voltage, f_c is center frequency, L is the archived noise measure, and TR is tuning range.

2.2 Performance equivalence and quotient design space

Let $\mathcal{X}$ be the set of circuit designs with a fixed topology, and let

$$F : \mathcal{X} \to \mathcal{Y} \tag{2}$$

be the designer-selected performance representation. Depending on the task, F may contain gain, bandwidth, oscillation frequency, noise, power, area, tuning range, constraint margins, or any other attribute that the designer regards as relevant. Define

$$x_1 \sim x_2 \quad \Longleftrightarrow \quad F(x_1) = F(x_2). \tag{3}$$

The design manifold used by ALARM is the quotient

$$\mathcal{M} = \mathcal{X} / \sim . \tag{4}$$

If $(\mathcal{Y}, d_{\mathcal{Y}})$ is a metric performance space, then

$$D\big([x_1],[x_2]\big) := d_{\mathcal{Y}}\big(F(x_1), F(x_2)\big) \tag{5}$$

is well defined on $\mathcal{M}$: changing the representative inside either equivalence class does not change F. Moreover, $D([x_1],[x_2]) = 0$ implies $[x_1] = [x_2]$ by Eq. (3). Thus the zero-distance directions in the raw component space are precisely the directions that the designer has declared irrelevant.

For a scalar performance map f, Eq. (5) becomes

$$D\big([x_1],[x_2]\big) = |f(x_1) - f(x_2)|. \tag{6}$$

For an infinitesimal displacement dx,

$$ds^2 = df^2 = dx^{\mathsf{T}} \nabla f(x) \nabla f(x)^{\mathsf{T}} dx. \tag{7}$$

The rank-one form in Eq. (7) is expected in the scalar case: locally, the quotient retains the direction that changes the selected performance and collapses directions tangent to a performance-equivalence class. A vector-valued F analogously yields the pullback $J_F^{\mathsf{T}} W J_F$ for a chosen metric W on performance space.

2.3 Scalar objective and success criteria

The project combines a primary cost $c_0(s)$ and specification violations $c_j(s)$ into

$$J(z) = w_0 c_0\big(s(z)\big) + \sum_{j=1}^{m} w_j \min\Big\{\bar{c}_j,\, \big[c_j\big(s(z)\big)\big]_+\Big\}, \tag{8}$$

where $[a]_+ = \max(a, 0)$ and $\bar{c}_j$ prevents one violation from dominating the score. The optimization target recorded in the project was an active-inductor-based voltage-controlled oscillator (VCO). A run was considered successful when the design met all of the following archived criteria:

$$\begin{aligned} \text{FoM}_{\text{proj}} &> 150 \text{ dB}, \\ 2.35 \text{ GHz} < f_c &< 2.45 \text{ GHz}, \\ V_{\text{peak}} &> 0.5 \text{ V}, \\ \text{TR} &> 10\%. \end{aligned} \tag{9}$$

Here FoM_{proj} denotes the figure of merit used by the capstone optimizer. More generally, the performance representation and scalarization are chosen to retain the circuit properties that matter for the design target; properties outside that choice are intentionally quotiented out.

3 ALARM Framework

3.1 System overview

ALARM alternates between inexpensive model-based proposal steps and expensive simulator evaluations. An initial population is evaluated in SPICE. The critic is trained on all accumulated pairs $(z, J(z))$. The line-search agent then uses the critic and attention module to propose a candidate, which is evaluated in SPICE. The new observation is appended to the population, and the cycle repeats.

3.2 Critic model

The critic $Q_\theta(z)$ approximates the scalar SPICE objective. With population $\mathcal{D} = \{(z_i, J_i)\}_{i=1}^N$, its mean-squared loss is

$$\mathcal{L}_{\text{critic}}(\theta) = \frac{1}{N_b} \sum_{i \in \mathcal{B}} \big(Q_\theta(z_i) - J_i\big)^2, \tag{10}$$

where $\mathcal{B}$ is a minibatch. Because a surrogate can be inaccurate away from observed samples, ALARM uses it to rank local candidates rather than replacing SPICE as the final evaluator.

3.3 Riemannian metric and attention network

Let $g = \nabla_z Q_\theta(z)$. Following the quotient-space construction in Eq. (7), the critic induces the local performance term

$$G_M = g g^{\mathsf{T}}. \tag{11}$$

This term measures motion across performance-equivalence classes. To realize a search direction in

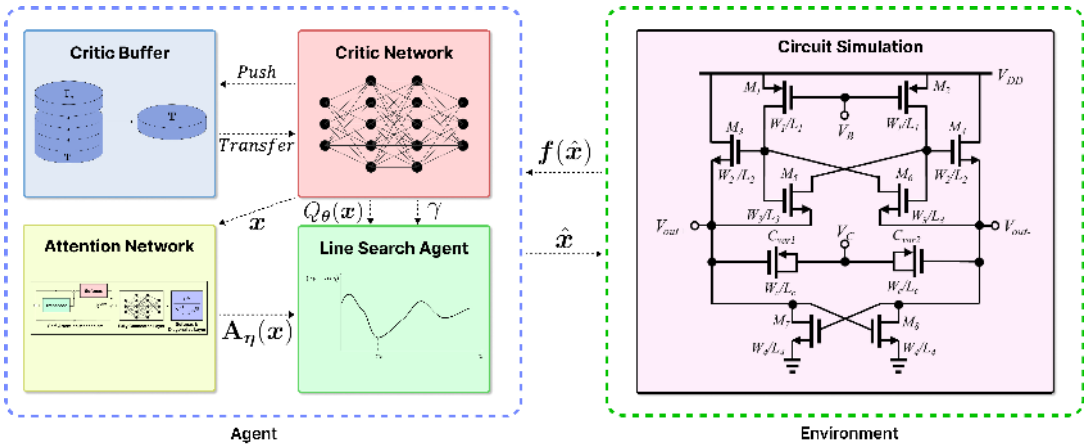

Figure 1: Archived ALARM architecture. The agent contains the critic buffer, critic network, attention module, and line-search component; the circuit simulator supplies the authoritative evaluation.

the full design chart, ALARM augments it with a coordinate term

$$G_D = \mathrm{Diag}(\boldsymbol{g} \odot \boldsymbol{g}), \qquad G = G_M + \mu G_D, \quad \mu > 0. \tag{12}$$

The diagonal term chooses how quotient-space descent is lifted back to component coordinates. It does not change the definition of circuit equivalence; it determines which representative motion is taken in the ambient design chart.

The attention module maps a design to positive coordinate weights

$$\boldsymbol{a}_\eta(z) = \mathrm{softmax}\big(h_\eta(z)\big), \qquad A_\eta(z) = \mathrm{Diag}\big(\boldsymbol{a}_\eta(z)\big). \tag{13}$$

The metric used by the line-search agent is

$$\widetilde{G}_\eta(z) = \boldsymbol{g}\boldsymbol{g}^\mathsf{T} + \mu A_\eta(z)\mathrm{Diag}(\boldsymbol{g} \odot \boldsymbol{g}), \tag{14}$$

and the proposal direction is

$$\boldsymbol{d} = -\widetilde{G}_\eta(z)^{-1}\boldsymbol{g}. \tag{15}$$

Because the attention weights are positive, the coordinate term supplies positive curvature in each locally active design coordinate. On the active subspace, $\widetilde{G}_\eta$ is positive definite and

$$\boldsymbol{g}^\mathsf{T}\boldsymbol{d} = -\boldsymbol{g}^\mathsf{T}\widetilde{G}_\eta^{-1}\boldsymbol{g} < 0 \quad \text{whenever } \boldsymbol{g} \neq \mathbf{0}. \tag{16}$$

Thus, $\boldsymbol{d}$ is the Riemannian steepest-descent direction in the chosen chart metric. If a coordinate is locally inactive, the same construction is interpreted on the active subspace, equivalently through the restricted inverse. The quotient metric and its ambient-coordinate lift therefore play different but compatible roles.

3.4 Backtracking line search

Starting from $\alpha_0 > 0$, the agent repeatedly sets $\alpha \leftarrow \rho\alpha$ with $0 < \rho < 1$ until the surrogate Armijo condition is satisfied [5, 6]:

$$Q_\theta(z + \alpha\boldsymbol{d}) \leq Q_\theta(z) + c_1\alpha\boldsymbol{g}^\mathsf{T}\boldsymbol{d}, \quad 0 < c_1 < 1. \tag{17}$$

Multiple surrogate steps may be taken before a SPICE query, but the final candidate is always checked by the simulator. The Armijo condition controls the critic-based line-search step, while SPICE supplies the authoritative circuit evaluation.

3.5 Critic buffer

Repeatedly training a critic from random initialization can waste information. ALARM stores earlier critic parameters together with the design neighborhood and the observed candidate quality. For a stored critic record T_i, the critic score is

$$\begin{aligned}\mathrm{CS}(z \mid T_i) = \kappa_1 &\left[Q_{\theta_i}(\widehat{z}_i) - J(\widehat{z}_i)\right]^2 \\ &+ \kappa_2\left[J(\widehat{z}_i) - \gamma_i\right] \\ &+ \kappa_3 D(z, z_i)^2,\end{aligned} \tag{18}$$

where the three terms represent prediction error, improvement relative to the previous objective, and proximity to the current search point. Records with favorable scores are preferred when selecting a critic for transfer. The selected network supplies a warm start for retraining, as in the original critic-buffer mechanism.

3.6 Procedure

ALARM optimization procedure

1. Sample an initial population of bounded designs and evaluate each design in SPICE.
2. Train the critic on all observed design-objective pairs; optionally warm-start it from a locally relevant buffered critic.
3. Train or update the attention module while holding the critic fixed.
4. Select the best observed design, compute the preconditioned direction in Eq. (15), and backtrack until Eq. (17) holds.
5. Generate the candidate along the line-search direction and evaluate it in SPICE.
6. Add the observation to the population, update the critic buffer, and retain the best measured design.
7. Stop when all specifications in Eq. (9) are met or the SPICE budget is exhausted.

4 Archived VCO Experiment

4.1 Target circuit and setup

The experiment optimized the 11 design variables listed in Table 1 for the VCO in Figure 2. The archived setup used LTspice and a 130-nm CMOS process design kit. Every method began with the same randomly generated population of 60 circuits. DE, DNN-OPT, an Euclidean-gradient ablation (ALARM-EM), and ALARM were each run ten times, with a maximum budget of 300 SPICE evaluations.

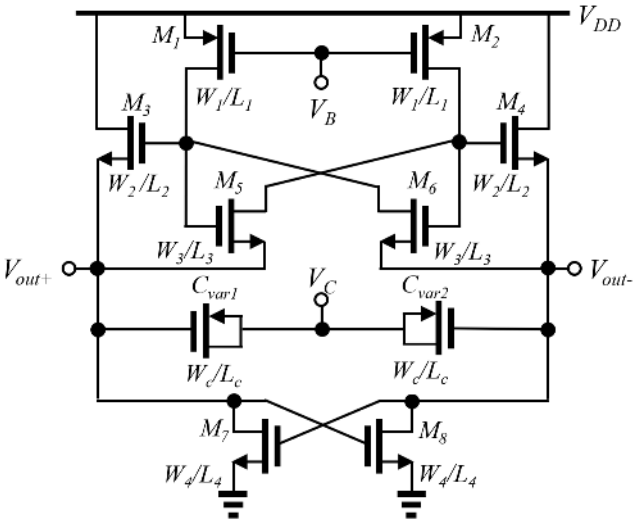

Figure 2: Active-inductor-based VCO used in the archived experiment. Eleven variables control the bias and transistor dimensions.

Table 1: Archived design-variable bounds.

Variable	Lower bound	Upper bound
V_b (V)	0	1.8
L_1, L_2, L_3, L_4, L_c (μm)	0.15	1000
W_1, W_2, W_3, W_4, W_c (μm)	0.42	1000

For the learned methods, the archived study used the same critic architecture to reduce a major source of implementation variance. Reported simulation counts for DNN-OPT, ALARM-EM, and ALARM are computed over successful runs. Mean wall-clock times are retained exactly as recorded in the 2024 manuscript.

4.2 Optimization results

Table 2 is the principal quantitative record. ALARM satisfied the project criteria in all ten trials and required 33.2 SPICE evaluations on average. DNN-OPT succeeded in nine trials with 145.9 evaluations on average. On these archived values, ALARM used approximately $145.9/33.2 = 4.39$ times fewer simulator calls among successful runs. The ALARM-EM ablation reached a competitive mean when successful but failed in three of ten trials, whereas ALARM succeeded in every recorded trial.

Table 2: Archived optimization performance over ten runs. Simulation statistics exclude failed runs; the DE time is retained as reported despite the absence of a successful run.

Method	Success	Min. sim.	Max. sim.	Mean sim.	Mean time (min)
DE	0/10	–	–	–	22.0
DNN-OPT	9/10	24	283	145.9	14.0
ALARM-EM	7/10	34	112	78.7	7.8
ALARM	**10/10**	**12**	**59**	**33.2**	**3.3**

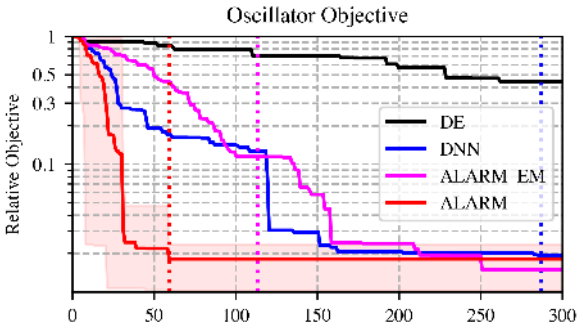

Figure 3: Archived mean optimization trajectories. The vertical dotted lines indicate the mean successful evaluation count recorded for each learned method; the horizontal dashed line indicates the worst successful terminal objective among the ten runs.

The trajectory plot in Figure 3 is descriptive rather than inferential. With only ten runs, it does not support fine-grained statistical ranking, but it is consistent with the success-rate and simulator-count summary.

4.3 Representative optimized circuit

The archived representative design moved the center frequency from 0.52 GHz to 2.40 GHz while retaining the required peak voltage and tuning range. Its project-specific FoM increased from 144.4 dB to 151.5 dB, and the scalar objective decreased from 4.16 to 0.15. Table 3 reproduces these values without reinterpretation.

Table 3: SPICE evaluation of one archived initial and optimized VCO.

Quantity	Initial	Optimized
Power (mW)	0.32	0.72
Peak output voltage (V)	0.68	0.56
Center frequency (GHz)	0.52	2.40
Archived noise measure (dBc)	-145.2	-142.5
$\mathrm{FoM}_{\mathrm{proj}}$ (dB)	144.4	151.5
Tuning range (%)	25.8	31.9
Scalar objective	4.16	0.15

5 Discussion

The VCO experiment supports the design motivation behind ALARM. DNN-OPT uses a learned critic but searches without the same performance-aware geometry, while ALARM-EM replaces the design metric

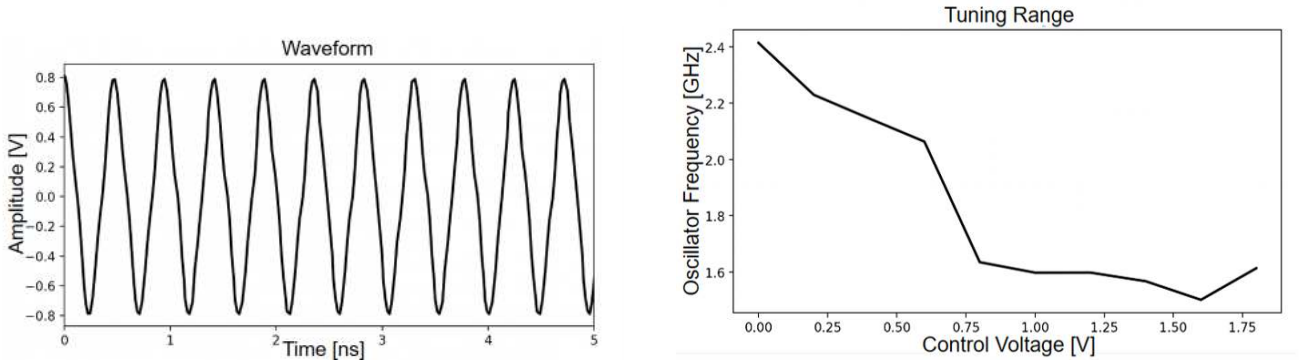

(a) Output waveform.

(b) Oscillation frequency versus control voltage.

Figure 4: Archived SPICE outputs for the representative optimized VCO.

with the Euclidean metric. ALARM reached the target in all ten trials and required fewer SPICE evaluations on average. The comparison between ALARM and ALARM-EM is consistent with the claim that the quotient-induced metric and attention-based coordinate lift can improve search stability.

This was an undergraduate capstone study, so its scope was deliberately focused: one active-inductor VCO, one process setting, and ten trials per method. Natural extensions include evaluation across additional circuit topologies, process nodes, and specification sets, together with separate ablations of the attention network, multi-step search, and critic buffer. These are follow-up directions rather than prerequisites for the capstone result reported here.

6 Conclusion

ALARM is an undergraduate capstone method for reducing simulator use in analog-circuit sizing. Its central idea is to identify circuits through designer-selected performance, optimize on the resulting quotient design manifold, lift the Riemannian descent direction to component coordinates with an attention-weighted metric, and reuse critics across nearby search regions. In the archived VCO study, ALARM achieved 10/10 success with 33.2 SPICE evaluations on average, compared with 9/10 and 145.9 evaluations for DNN-OPT. The experiment demonstrates the intended proof of concept and motivates broader evaluation of performance-aware circuit geometry.

Author and Project Record

This technical-report version lists Soohyun Choi, Minseok Oh, and Ickhyun Song, in that order, with Ickhyun Song as the corresponding author. The project received the Outstanding Graduation Project Award at Hanyang University and the Excellence Award in the College of Engineering Capstone Design program.

Acknowledgements

The original project acknowledged support in part from Samsung Electronics Co., Ltd. (IO231020-07441-01), the Institute of Information & Communications Technology Planning & Evaluation under the artificial-intelligence-semiconductor talent program (IITP-(2024)-RS-2023-00253914), and the National Research Foundation of Korea (RS-2024-00409492). EDA tools were supported by the IC Design Education Center (IDEC), Korea. This technical report preserves that acknowledgement from the archived manuscript; it does not independently reinterpret the scope of each sponsor's support.

References

[1] 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 *Proceedings of the 58th ACM/IEEE Design Automation Conference*, 2021, pp. 1219–1224.

[2] R. Storn and K. Price, "Differential evolution - a simple and efficient heuristic for global optimization over continuous spaces," *Journal of Global Optimization*, vol. 11, pp. 341–359, 1997.

[3] K. Settaluri, A. Haj-Ali, Q. Huang, K. Hakhamaneshi, and B. Nikolic, "AutoCkt: Deep reinforcement learning of analog circuit designs," in *Proceedings of DATE*, 2020, pp. 490–495.

[4] 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 *Proceedings of the 57th ACM/IEEE Design Automation Conference*, 2020, pp. 1–6.

[5] L. Armijo, "Minimization of functions having Lipschitz continuous first partial derivatives," *Pacific Journal of Mathematics*, vol. 16, no. 1, pp. 1–3, 1966.

[6] J. Nocedal and S. J. Wright, *Numerical Optimization*, 2nd ed. New York: Springer, 2006.

[7] A. Vaswani et al., "Attention is all you need," in *Advances in Neural Information Processing Systems*, vol. 30, 2017.

[8] S.-I. Amari, "Natural gradient works efficiently in learning," *Neural Computation*, vol. 10, no. 2, pp. 251–276, 1998.

[9] P.-A. Absil, R. Mahony, and R. Sepulchre, *Optimization Algorithms on Matrix Manifolds*. Princeton, NJ: Princeton University Press, 2008.

[10] L.-H. Lu, H.-H. Hsieh, and Y.-T. Liao, "A wide tuning-range CMOS VCO with a differential tunable active inductor," *IEEE Transactions on Microwave Theory and Techniques*, vol. 54, no. 9, pp. 3462–3468, 2006.

[11] Analog Devices, "LTspice simulator," 2024. [Online]. Available: https://www.analog.com/en/resources/design-tools-and-calculators/ltspice-simulator.html

[12] SkyWater Technology Foundry and Google, "SkyWater open-source PDK," 2023. [Online]. Available: https://github.com/google/skywater-pdk