From Specification to Topology: Automatic Power Converter Design via Reinforcement Learning

Shaoze Fan†*, Ningyuan Cao†*, Shun Zhang†*, Jing Li†**, Xiaoxiao Guo‡, Xin Zhang‡**

†*New Jersey Institute of Technology;*   ‡*IBM T. J. Watson Research Center;*

Abstract—The tidal waves of modern electronic/electrical devices have led to increasing demands for ubiquitous application-specific power converters. A conventional manual design procedure of such power converters is computation- and labor-intensive, which involves selecting and connecting component devices, tuning component-wise parameters and control schemes, and iteratively evaluating and optimizing the design. To automate and speed up this design process, we propose an automatic framework that designs custom power converters from design specifications using reinforcement learning. Specifically, the framework embraces upper-confidence-bound-tree-based (UCT-based) reinforcement learning to automate topology space exploration with circuit design specification-encoded reward signals. Moreover, our UCT-based approach can exploit small offline data via the specially designed default policy to accelerate topology space exploration. Further, it utilizes a hybrid circuit evaluation strategy to substantially reduces design evaluation costs. Empirically, we demonstrated that our framework could generate energy-efficient circuit topologies for various target voltage conversion ratios. Compared to existing automatic topology optimization strategies, the proposed method is much more computationally efficient — it can generate topologies with the same quality while being up to 67% faster. Additionally, we discussed some interesting circuits discovered by our framework.

Index Terms—design automation, power converter topology design, upper-confidence-bound tree (UCT), reinforcement learning

I. INTRODUCTION

Power converters are ubiquitous in electronic/electrical devices. With the proliferation of customized electrical systems [1], such as electric vehicles, self-powered IoT, wearable/implantable biosensors, the need for custom power converters is rapidly increasing to provide diverse supply power standards. The design specifications, such as voltage conversion ratio, power efficiency, and output ripple, differ significantly from application to application. To meet these specifications, conventional manual circuit optimization is time-consuming and relies heavily on existing circuit topologies, as illustrated in Figure 1. The expensive design process has dramatically hindered the development of novel power converters for fast-paced and innovative custom designs. Hence, there is a pressing need for an automatic circuit design framework
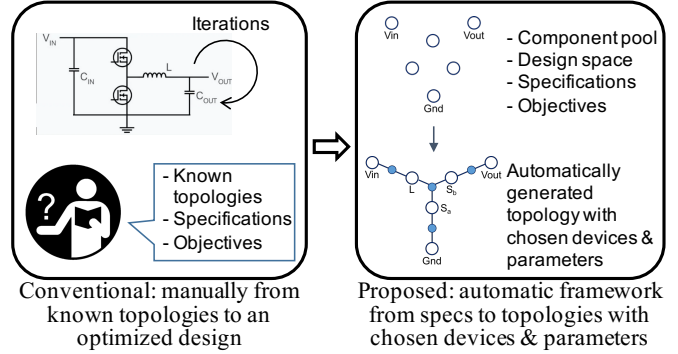
Fig. 1: Given a custom power converter design task, the conventional manual approach relies heavily on known topologies and is labor-intensive, computationally expensive, and time-consuming. In contrast, our automatic power converter design framework can explore the design space more effectively, thereby immensely decreasing the development time and cost without compromising the performance.

that can efficiently search and generate high-quality power converter topologies from the design specifications.

However, how to automate the power converter topology design remains a challenging task. Firstly, topology generation for electronic circuits or integrated circuits (IC) lacks thorough investigation. The state-of-art analog/mixed-signal (AMS) IC design automation mainly addresses device sizing or parameter optimization for a fixed circuit structure [2]–[9]. People have also investigated automation methods to accelerate the physical implementation of AMS ICs when schematic/topology design is already done [10]–[14]. More recently, researchers have started looking into circuit synthesis [15]–[17]. But some of them require substantial domain knowledge, which greatly hindered their generality. Others explore the enormous topology space via exhaustive search, metaheuristic search, or gradient descent, which may not be as efficacious in the non-continuous topology space as for other design tasks.

Fundamentally, automated topology generation is inherently difficult, as it faces challenges due to the immense search space and severe data discontinuity. In fact, the search space increases exponentially with circuit complexity, and the enormous number of possible topologies prohibits exhaustive or random search. Further, metaheuristic search strategies may get stuck in a restricted set of topologies and thus output suboptimal results unless the number of random samples becomes large enough. Moreover, unlike device parameters such as transistor width and length, a small shift in the component connections of one topology will very likely lead to significant changes in the circuit performance. As such, search algorithms

or optimization methods that rely on continuity between "similar topologies", such as genetic search or gradient descent, may become less effective in reducing search efforts.

Finally, it is time-consuming to evaluate the performance of generated topologies properly. For power converters, this is often a more severe problem because they are usually nonlinear and dynamically controlled switching circuits, which require long simulations to reach their steady states. The conventional Spice simulation [18] is able to provide high-fidelity evaluation results, but this comes with the cost of a long simulation time to achieve the desired precision and control scheme exploration. The evaluation cost per topology can be as high as minutes, making the topology exploration process prohibitively time-consuming.

Contribution. We present a design automation framework for power converter circuit design and optimization to address the above challenges. The main contributions are as follows:

- We propose the first automatic power converter design framework that intelligently explores the power converter topology space and generates high-quality candidate circuits based on custom design specifications. As shown in Figure 1, our framework can efficiently locate well-performing topologies with appropriate control schemes and also has the potential to generate novel topologies under specific design constraints.
- For the first time, our framework applies reinforcement learning (RL) using the upper-confidence-bound-tree (UCT) to circuit topology generation. We construct the UCT structure to sufficiently capture the semantic of topologies to explore the topology space more efficiently. Moreover, this UCT structure is able to exploit offline knowledge, which is obtained from a few suitable topologies with smaller sizes and encapsulated into our specially designed default rollout policy, and further accelerate the topology space exploration.
- As the long-running circuit evaluation is the bottleneck of fast topology exploration, we detect isomorphic topologies and adopt a hybrid circuit evaluation approach. A State-Space Averaging method is used during the exploration, which reduces the time cost of circuit evaluation by orders of magnitude, and a high-fidelity Spice transient simulation is used to validate circuit candidates generated by the exploration and filter out the over-optimistic ones.
- We conduct extensive experiments on 5-component (13-port) power converter design tasks. Evaluation results demonstrate that our proposed automatic framework can produce energy-efficient circuits for varying voltage conversion ratios. Furthermore, compared to baseline strategies (i.e., genetic search and random search algorithms) adapted from other circuit design tasks, our framework can generate constraint-satisfied and highly efficient topologies while needing fewer queries to circuit evaluation. Hence, it is up to 70%, 56%, and 50% faster than the baseline strategies for the experimented buck, boost, and buck-boost converter design tasks, respectively.

II. BACKGROUND

With the demands of custom electronics, application-specific design automaton of analog/mixed-signal (AMS) and radio-frequency (RF) power management circuits starts to play vital roles in accelerating high-quality electronic circuit designs. However, the traditional manual design routines are inherently time-consuming and rely heavily on domain expertise. To reduce the cost and improve the design quality, mainstream research about circuit design automation are three folded: (1) automating the device parameter optimization for known circuit topologies; (2) automating the physical implementation for known circuit topologies and parameters; (3) automating the circuit synthesis that directly generates topologies.

Parameter Optimization. Great efforts have been devoted to automating the parameter optimization for predetermined topologies. For example, [2] proposed geometric-programming-based optimization, [3], [4] used regression and convex/polynomial optimization, and [5] applied a Bayesian optimizer, [6], [7] both adopted model-based reinforcement learning to find the optimal device parameter combinations for analog circuits. Additionally, [8] encoded circuits using graph convolutional neural networks to transfer the parameter optimization knowledge learned between two topologies or between technology nodes of the same circuit.

Physical Implementation. The physical implementation automation for integrated circuits (IC), such as device placement and routing, also plays a vital role for many high-performance AMS/RF circuits. For example, analog generators were proposed in [10], [11], which directly build analog circuit layouts. [12] applied a data-driven approach to check layout symmetry, which is crucial in high-quality AMS physical layout. A feed-forward equalization transmitter layout generator was introduced in [13], which significantly reduced layout time. [14] presented a novel detailed routing framework for AMS layout synthesis to address the sensitive net coupling issues.

Topology Optimization. In contrast to parameter optimization and physical implementation that target fixed topologies, recent works have started investigating circuit topology optimization. Specifically, [15] proposed a graph neural network (GNN) model that learns to simulate the electromagnetic properties of distributed circuits. Via back-propagating the gradient, this GNN model can also be used to optimize the circuit parameters and topology. However, due to the special electromagnetic property of coupling decays in distributed circuits, an "edge" in a circuit topology is determined by the physical distance between two nodes and its impact on the circuit is continuous and decomposable. In comparison, the edge in a power converter topology is determined by whether the two components are connected in the circuit, and removing one edge may utterly change the performance of the circuit (e.g., from valid to invalid). Hence, the gradient back-propagation approach with the GNN model in [15] cannot be directly applied to the power converter design task. For topology synthesis for large analog integrated circuits, [16] presented a graph-grammar-based circuit topology representation, which

2

hierarchically decomposes a circuit until reaching the basic predefined building block. To reduce the search space, this work focused on designing meaningful decomposition rules and circuit formation rules that domain experts manually add while the generation is performed using an *exhaustive search* within the representation space. For searching in the circuit topology space, [17] developed a *genetic search* algorithm, where the device types of components in the topology is essentially fixed. We extend it to allow changing component types and compare it with our proposed framework.

AI-assisted Power Electronic Design. The advancements of artificial intelligence (AI) and IC design automation have introduced marvelous opportunities for power electronic circuit design automation [19]. Existing works have been looking into methods that greatly reduce circuit evaluation time by advancing surrogate models that approximate the system dynamics (e.g., electromagnetic properties, thermal characterization, and wire costs) with lower computational efforts [15], [20]–[22]. Researchers have also applied AI techniques and shown some successes in modeling and optimizing other aspects of power electronic systems, such as component model, system parameters, and post-layout performance [8], [9], [23], [24].

III. PROBLEM STATEMENT

We investigate the automatic power converter design problem with topology generation, device type selection, and control parameter tuning. This section first describes the custom design specifications of power converters considered in this work, followed by their topological representation and parameters. Next, we formulate the problem as an optimization problem with encoded custom design specifications.

Custom Design Specifications. We consider two primary design metrics of power converters, namely, *voltage conversion ratio* γ and *power conversion efficiency* η, as shown in Figure 2. The voltage conversion ratio is the ratio of the output voltage to the input voltage, i.e., $\gamma = V_{\text{out}}/V_{\text{in}}$, and is the main constraint for the generated power converter. Generating power converter designs with higher power conversion efficiencies is the optimization goal of the custom design task. Other constraints include the number of components in the converter topology and the types of available devices. The design task evaluated in this work only considers devices including capacitors C, inductors L, phase-I switches S_a, and phase-II switches S_b, but it can be easily extended to other device types.

Topological Representation and Parameters. A candidate power converter design consists of a *topological representation s* and a *switching control parameter d*. Specifically, the topological representation contains a set of components with ports and edges connecting the ports. Each component has a device type (from the set of available devices) with device parameters (e.g., inductance, capacitance, and transistor dimensions) and two ports (i.e., left and right ports). Note that each component is nondirectional — switching all the connections of its two ports results in the same circuit, despite needing different indexes to distinguish the two ports in
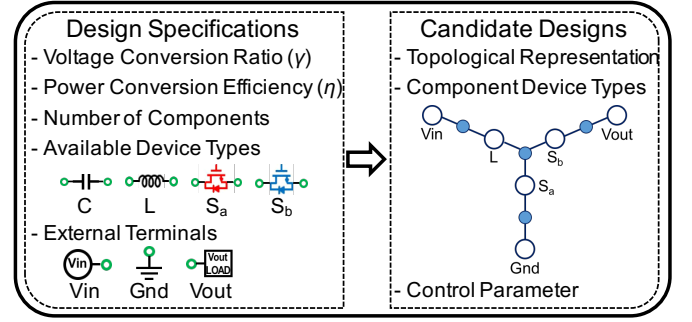

Fig. 2: Custom power converter design task from design specifications to candidate designs.

the topological representation. Additionally, there are three *external terminal ports*: the input voltage port Vin, the output voltage port Vout, and the ground port Gnd. The edges in the topological representation specify the connections between the components' ports and the terminal ports. The switching control parameter specifies the duty cycle of a candidate design, which often affects the output voltage. In this work, the design task involves designing the device types of components, the edges connecting ports, and the control parameter, while the device parameters for each device type are predefined. Thus, this work focuses mainly on the challenging topology design problem and leaves the integration with existing device parameter optimization methods as future work.

Circuit Evaluation. Given a generated power converter circuit with topological representation s and control parameter d, the Spice-based transient simulation is conducted to obtain the voltage conversion ratio $\gamma_{s,d}$ and power efficiency $\eta_{s,d}$.

Circuit Generation Objective. Given the custom design task with a target voltage conversion ratio γ_0, the goal of our framework is to automatically generate the topological representation s (with chosen component types and edges) and configure the control parameter d of the power converter circuit. Specifically, the objective can be described as:

$$s^*, d^* = \arg \max_{s \in S, d \in D} [U_{\gamma_0}(\gamma_{s,d}, \eta_{s,d})], \quad (1)$$

where S and D denote the set of topological representations and the set of control parameter configurations. U_{γ_0} is a utility function that evaluates a circuit design's conversion ratio and efficiency for the custom design task with target conversion ratio γ_0. Specifically, the utility function is formulated as:

$$U_{\gamma_0}(\gamma_{s,d}, \eta_{s,d}) = \eta_{s,d} \cdot \delta(\gamma_{s,d}, \gamma_0). \quad (2)$$

In our formulation, δ measures how close the obtained conversion ratio $\gamma_{s,d}$ and the target conversion ratio γ_0 are. In our experiments, we use $\delta(\gamma, \gamma_0) = 1.1^{-\left(\frac{15(\gamma - \gamma_0)}{|\gamma_0|}\right)^2}$. The utility function equals 0 when the topology is invalid or incomplete.

IV. FRAMEWORK

In this work, we propose an automatic power converter design framework that utilizes a reinforcement learning mechanism to efficiently explore the topological representation
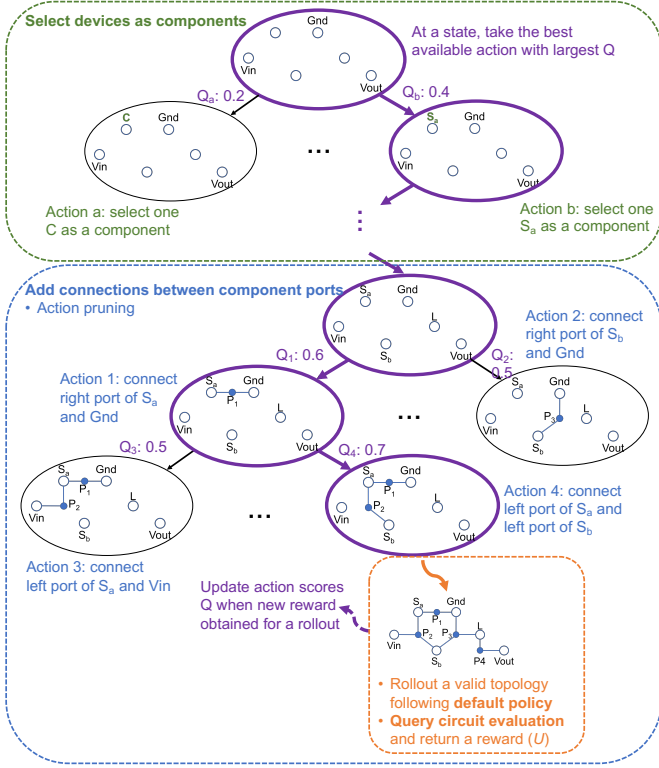
Fig. 3: UCT-based Topology Generation Overview: The generation starts with a topology having empty components and 3 terminals. It first makes device selections (shown in green) and then connects ports of components and terminals (shown in blue), where the actions are specially designed to reduce isomorphism and invalid circuits. The look-ahead tree construction and action score calculation (shown in purple) follow the UCT algorithm, which utilizes upper confidence bound in exploitation and exploration trade-offs. To improve the sample efficiency and rollout speed, we designed a default policy and adopted a fast evaluation technique (shown in orange).

space and locate candidate designs with high utility scores. We formulate the circuit generation task as a sequential decision-making problem. Instead of synthesizing the entire topology all at once, the component device types and connections between ports are added step by step, as illustrated in Figure 3. This multi-step formulation allows physics-aware connection pruning and removes many isomorphic topologies by construction. We further improve the topology generation by incorporating offline knowledge from the pre-collected dataset through a default policy. Finally, to address the issue of costly circuit evaluation, we use a State-Space Averaging technique to evaluate different parameter configurations of the same circuit and validate the output candidates via Spice simulations.

A. Sequential Circuit Topology Generation

We formulate topology generation as a sequential decision task. In particular, we model the topology generation as a Markov Decision Process, namely a 4-tuple of $\langle S, A, T, R \rangle$, representing the state set S, action set A, state transition function T, and reward function R. As shown in Figure 3, the topology generation always starts with an empty topology

of the power converter and has two phases with multiple steps: the device type selection phase and connection selection phase.

In the t-th step, the state $s_t \in S$ is a partial or complete topology of the power converter circuit. Inspired by the simple fact that circuit topology is a graph, each state s_t maintains a component set, a port set, and an adjacency matrix specifying the connections between each pair of ports. The action set A_t depends on the current state s_t. For a state in the device type selection phase, an action $a_t \in A_t$ decides whether a device type is selected for a component. For the connection selection phase, an action either decides to skip adding more connections or decides which port is connected to the port under consideration. Given our state and action formulation, the state transition is a deterministic function $s_{t+1} = T(s_t, a_t)$, which maps the current partial topology and the action choice to the next topology. The reward function R_{γ_0} encodes the custom design objective U_{γ_0} in Equation 2 such that $R_{\gamma_0}(s_t, a_t) = \max_d U_{\gamma_0}(\gamma_{s_{t+1},d}, \eta_{s_{t+1},d}) - \max_{d'} U_{\gamma_0}(\gamma_{s_t,d'}, \eta_{s_t,d'})$. The optimization objective is to find an action sequence that maximizes the final topology's power efficiency:

$$a^*_{0:T} = \arg\max_{a_{0:T}} \sum_{t=0}^{T} R_{\gamma_0}(s_t, a_t) \quad (3)$$

Note that the above objective requires computing the maximum utility over the different control parameters. The effective way to compute this utility will be presented in Section IV-E.

B. Upper-Confidence-Bound-Tree-based Topology Generation

A UCT-based Reinforcement Learning method is used to optimize the action sequence. UCT is a popular Monte Carlo Tree Search algorithm for large state-space sequential decision-making optimization problems, which suits the large topological representation space in our problem formulation. At a high level, UCT builds a look-ahead tree and greedily selects an action based on the actions' estimated scores at each step. In our proposed framework shown in figure 3, each tree-node of UCT corresponds to a partial or complete topology (i.e., a state) and each tree-edge corresponds to a device type or connection selection (i.e., an action).

To accurately estimate the action scores Q, UCT performs multiple *look-ahead rollouts* (also called simulations in RL literature), which sample the remaining action sequence following a *default policy* and evaluate the quality of the samples in terms of reward. UCT is an anytime optimization algorithm with three parameters, the number of look-ahead rollouts, the maximum depth (uniform for each rollout), and an exploration parameter. In general, the larger the number of rollouts and the depth parameter are, the slower UCT is, but the better it is. Compared to other Monte Carlo Tree Search algorithms, UCT utilizes the Upper Confidence Bound in the tree building process to provide improved estimations of action scores. UCT computes a score for each action a_t at a state s_t as the sum of the exploitation and exploration terms, as follows:

$$Q^{\text{UCT}}(s_t, a_t) = Q^{\text{MC}}(s_t, a_t) + \sqrt{\frac{\log(n(s_t))}{n(s_t, a_t)}} \quad (4)$$

4

Here, Q^{MC} is the Monte Carlo average of the sum of rewards obtained from the look-ahead rollouts, i.e., $Q^{\text{MC}}(s_t, a_t) = \frac{1}{M} \sum_{m=1}^{M} \sum_{t'=t}^{T} R(s_{t'}^{m}, a_{t'}^{m})$, where m identifies a look-ahead rollout in the total M rollouts. The exploration term $\sqrt{\log(n(s_t))/n(s_t, a_t)}$ is the Upper Confidence Bound, where $n(s_t)$ is the number of visits for the state node s_t and $n(s_t, a_t)$ the number of visits of the action a_t at state s_t. UCT selects the action to rollout greedily with respect to this summed score using the look-ahead tree. Once the input-parameter number of rollouts are produced each to the maximum depth, UCT returns the exploitation term for each action at the root node.

C. Combining Knowledge into UCT via Action Pruning

As discussed above, UCT estimates a state-action pair's score $Q^{\text{MC}}(s_t, a_t)$ more accurately when more rollouts are performed and the total number of rollouts is fixed according to the affordable computation cost. Hence, pruning the action space can improve UCT's search efficiency. Therefore, we design several physics-informed action pruning strategies to reduce the actions that lead to invalid or isomorphic topologies.

Reduce Invalid Topologies. Based on the knowledge of electronic circuits, we pose a set of constraints when adding connections to avoid generating invalid power converter circuits. Specifically, we do not include an action a of connecting two ports into the candidate action set A if adding this connection leads to one of the following situations: (1) a shortcut; (2) a direct connection between terminal ports VIN, VOUT, or GND; (3) a prohibited path; (4) a disconnected circuit. Here, the prohibited paths are the paths that violate basic circuit principles. For example, if we connect VIN and GND only with an inductor, it is equivalent to a power shortcut. A circuit is considered disconnected if there is at least one port not connected to any of the terminal ports via any paths of connections after the generation process. Together with the action pruning rules described below for connection selection, we can identify such disconnected topologies early in the generation process. Specifically, if there is no more allowed connection that can be added to a port and there is no path connecting this port to any of the terminal ports, then this partial topology will eventually become a disconnected circuit even if additional connections are selected between other ports. Thus, we mark the corresponding state as a terminal state with a reward of 0, so no more actions can be taken from this state.

Reduce Isomorphic Topologies. Combining the states representing isomorphic topologies can also improve UCT's search efficiency, since the rewards of rollouts from all these states can be collectively used to estimate the score more accurately. Hence, we propose the following methods to reduce the generated isomorphic topologies by construction.

First, we split the device type selection phase into multiple rounds, where each round has an ordered set of available device types, e.g., $\{S_a, C, L\}$ in the second round. The set in the first round includes all device types. The ordering of device types is set to S_a, S_b, C, and L. In each step of a round, each available device type is considered for selection, e.g., selecting

one S_a for a component or skipping adding S_a. If a device type has been skipped in the current round, it will be removed from the set of available device types for the next round. For instance, if we have added one S_a and one L but skipped C in the second round, the set for the third round becomes $\{S_a, L\}$. The device type selection phase ends when the number of selected devices is equal to the number of components. In this way, every state in this phase is unique (i.e., representing different device selections), while all combinations of device selections can be generated.

Next, for the connection selection phase, we number all the ports where the terminal ports have the smallest indexes. We consider each port for adding connections with other ports one by one. Although a connection in a converter topology is not directed, we only allow a port with a smaller index to be connected to a port with a larger index. Thus, this connection can only be added once. In addition, since a device is nondirectional, we always have its left port be connected before its right port when both ports have no connection. Following the above rules, many actions that lead to states representing isomorphic topologies are pruned.

D. Combining Knowledge into UCT via Default Policy

Due to the large space of topologies, the computation cost for finding a good topology can be high even after action pruning. Hence, we further improve the effectiveness of UCT by replacing the random default policy with a data-driven one. In particular, we randomly generate some topologies with fewer components and evaluate their efficiencies and conversion ratios. For example, we collect a small data set with 3-component topologies for the design task of 5-component topologies. Among them, we find the *good topologies* with higher rewards and collect their device selections and all paths of connections between the terminal ports. Exploiting the information obtained from good topologies with smaller sizes, we develop the following two default policies for the two topology generation phases.

Node Selection. We collect the number of times a device selection combination C (e.g., $\{S_a, S_a, L\}$) occurring in the good topologies, denoted by $n_d(C)$. This information is encoded into the default policy, such that its probability of a device selection combination C' is approximately proportional to the collected device selection distribution. Specifically, we calculate the weight of a device selection combination C' as $w(C') = \sum_{C \subseteq C'} n_d(C)$. The total weight is defined as $W_d = \sum_{C''} w(C'')$. Then the probability of the device selection combination C' generated by the default policy is $\mathbb{P}_d(C') = \frac{w(C')/W_d + \epsilon_d}{\sum_{C''}(w(C'')/W_d + \epsilon_d)}$, where ϵ_d is a small constant that decides by how much the default policy follows the collected distribution. In our experiments, ϵ_d is set to 0.01.

Edge Selection. We also collect the number of occurrences of a *good path* of connections between any two of the terminal ports (VIN, VOUT, and GND) in the good topology data set, denoted as $n_d(p)$. When considering a port e for adding a connection onto the current partial topology s_t, we first construct
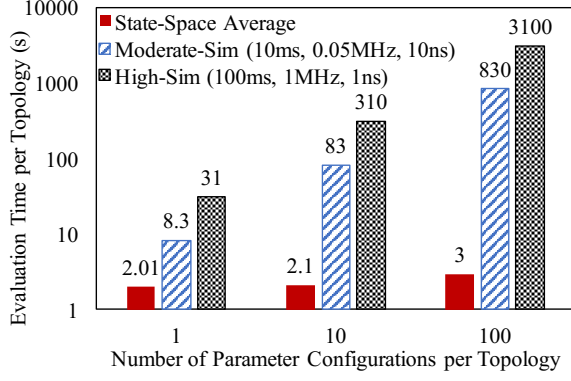
Fig. 4: Evaluation time per topology for different numbers of parameter configurations per topology under State-Space Averaging method, moderate-fidelity simulation (with 10ms transient time, 0.05MHz switching frequency, and 10ns minimum time-step), and high-fidelity simulation (with 100ms transient time, 1MHz switching frequency, and 1ns minimum time-step).

a set of all possible good paths $P(e, s_t)$. A path in $P(e, s_t)$ must contains port e and adding this path to s_t cannot lead to invalid topologies (e.g., shortcuts). Next, for any connection (e, e') allowed to be added to s_t, we calculate the weight of taking this action as $w(e, e') = \sum_{p \in P(e, s_t) \& (e, e') \subseteq p} n_d(p)$. Similar to node selection, we can calculate the total weight of all allowed actions as $W_p = \sum_{(e, e'')} w(e, e'')$. Then the probability of the connection selection action (e, e') generated by the default policy is $\mathbb{P}_p(e, e') = \frac{w(e, e')/W_p + \epsilon_p}{\sum_{(e, e'')}(w(e, e'')/W_p + \epsilon_p)}$, where ϵ_p is a small constant and is set to 0.4 in the evaluation. Note that although the weights are calculated using the good paths in the data set with fewer components, longer paths with more components can be generated with the help of ϵ_p.

With the default policies, we can bias UCT towards searching the topology space that may contain high-reward ones.

E. Efficient Evaluation via State-Space Averaging

The reward function used by UCT is a computational bottleneck as it requires evaluating different control parameters for different topologies sampled by all the rollouts. Simulation software, such as NGSpice [18], often takes a long time to evaluate one topology with fixed parameters. To speed up the circuit evaluation, we instead adopt the *State-Space Averaging* approach [25]. This approach shares the major computation among the circuits with the same topology but different parameters, which makes it a faster surrogate model in estimating a large number of parameters.

Specifically, State-Space Averaging characterizes the transfer properties of switching stages of a power converter circuit. By dividing the circuit into phases-I and phase-II sub-circuits, deriving state-space equations, and averaging state changes with the corresponding duty-cycles (i.e., d and $1-d$), it derives the estimated output voltage and power efficiency.

We compared the computation times of State-Space Averaging versus NGSpice. Figure 4 shows that the evaluation time per topology under State-Space Averaging is significantly lower compared to moderate-fidelity and high-fidelity simula-

tions. The advantage of State-Space Averaging is even higher when more parameter configurations of one topology (e.g., the same topology with different duty cycles) need to be evaluated. This is because deriving state-space equations is the most expensive step of State-Space Averaging, but these equations remain the same as long as the topology remains the same. Hence, increasing the number of parameter configurations per topology does not increase the evaluation time much under State-Space Averaging, while simulations must be performed for each parameter configuration of the same topology.

Despite its speed advantage, State-Space Averaging has a disadvantage — it is only accurate if the circuit is linear during operations. Nonlinear effects, such as discontinuity in conduction or enforced switch dead-zone, will introduce errors. As such, the state-space method usually overestimates the efficiencies of topologies. To address this issue, we take a hybrid circuit evaluation approach. During topology search, a topology is evaluated by State-Space Averaging. With the optimistic nature of State-Space Averaging, it covers all qualified topologies. Once the top topology candidates are generated, NGSpice simulation will be invoked to offer ground-truth evaluation for the final topology and control scheme selection.

V. EVALUATION

We evaluate our framework across a spectrum of custom design tasks with different conversion ratios and compare its performance with baseline algorithms. We also conduct an ablative study on the effectiveness of the proposed data-driven default policy. Finally, we discuss the nonconventional power converter topologies discovered by our framework.

A. Experiment Setup

We conduct the evaluation on the power converter design task with five components, each with two ports. Together with the three external ports (Vin, Vout, and Gnd), the design space contains topologies with a total of thirteen ports. The component device types have fixed device parameters and include capacitors C ($10\mu F$), inductors L ($100\mu H$), phase-I switches S_a, and phase-II switches S_b For external ports, we consider an input resistor of 0.1Ω for Vin and an output resistor of 100Ω and an output capacitor of $10\mu F$ for Vout. The candidate control parameter (i.e., duty cycle) ranges from 0.1 to 0.9, with a step size of 0.1. We configure the frequency as $1MHz$ for both analytic evaluation and simulation. We set the input voltage to $100V$, and the target voltage outputs are chosen from the range of $-300V$ to $300V$. Additionally, a topology with a conversion ratio smaller than -5 or larger than 5 will be considered invalid. In NGSpice simulation, the transient simulation time is set as $60s$.

B. Implementation Details and Competitive Algorithms

In this section, we describe the implementation details of our UCT-based framework and the baseline algorithms.

Hash Table. Both the UCT-based and baseline algorithms may query the circuit evaluation about a topology more than once. To avoid unnecessary circuit evaluation costs,

we implemented a hash table to cache the efficiencies and conversion ratios of the topologies that have been queried during each experiment. Therefore, in our experiments for both the proposed algorithm and baselines, all the queries to the circuit evaluation are about unique topologies.

Implementation Details of Proposed Algorithm. There are two important implementation details of our UCT-based topology generation algorithm. First, we use the number of rollouts to decide the number of explorations each run of the UCT-based topology generation. In each rollout, we expand a UCT tree-node and execute a rollout to get a complete topology. However, the number of rollouts can be larger than the number of queries, since we only query the circuit simulation when a rollout leads to a valid topology that is not in the hash table. Second, different from the common UCT design that completely reconstructs a new tree for each step, we inherit the sub-tree corresponding to the selected action from the tree of the previous step to make full use of the collected information. This operation makes our UCT-DP algorithm converge faster.

Random Search (RS). Random Search is a strategy that starts with an empty topology, randomly selects device types, and then randomly connects ports until reaching a complete topology. Note that, if a port already has connections to other ports, then it is not required to connect to yet another port. Thus, in this case, RS may skip adding more connections to this port with a probability of 0.8. RS uses the same reward function as UCT, along with the same prohibitive paths that prevent generating invalid topologies. After searching a prefixed number of complete circuit topologies, RS outputs the one that has the highest reward.

Genetic Search (GS). We implemented another popular heuristic-based search algorithm, namely Genetic Search, for the power converter design task [26]. GS starts with 15 random topologies. In each round, GS selects 3 topologies with the highest rewards from the current generation and uses them as parents to generate offspring by mutation. We implemented the following mutation types: (1) A crossover randomly selects a component that exists in two parents and exchanges the connections of this component [17]; (2) A component change randomly changes the device type of a random component in a parent; (3) A connection insertion randomly selects and connects two unconnected ports in a parent; (4) A connection removal removes a random connection of a parent; (5) A connection switch selects a random connection of a parent and changes one of the connection's ports to another random port; In each round, the above mutation types are chosen with probabilities of {20%, 10%, 30%, 20%, 20%}, respectively.

C. Experiment Results

Custom Design Tasks. We first evaluated our UCT-based framework on different custom design tasks for power converters. The design tasks can be classified into three settings, including buck-boost, buck, and boost converters. For each target voltage conversion ratio, we run our framework five


Fig. 5: Average power efficiencies of power converters generated by the proposed framework for custom design tasks with increasing voltage conversion ratios. From left to right in different colors and shades, the target converter types are buck-boost, buck, and boost.

times and compute the average efficiency of the generated topologies. Figure 5 shows the average power efficiencies under our framework for different voltage conversion ratios. Results show that our framework can successfully find high-performing topologies for most settings. The average efficiencies for smaller conversion ratios tend to be lower, mainly because these buck-boost converters are more sensitive to their duty cycles. For example, if we change the step size of the duty cycle from 0.1 to 0.05, our framework can discover a better topology with an efficiency of 0.89 for the design task with a conversion ratio of -2.5.

Comparison with Baseline Algorithms. We compared our UCT-based approach with RS and GS on buck-boost, buck, and boost converters. Since all these algorithms are anytime algorithms (i.e., the performance monotonically increases as more computation is used), we compared their performances conditioned on their computation costs, measured by the number of queries to circuit evaluations. Figure 6 reports the average reward of topologies generated by each algorithm. For all types of converters, our UCT-based approach, named **UCT-DP**, outperforms both RS and GS, while GS is comparable with or slightly outperforms RS. When the number of queries is around 110, UCT-DP achieves 83%, 35%, and 37% higher rewards compared with GS for buck-boost, buck, and boost converters, respectively. In terms of computation efficiency, results show that, compared to GS, UCT-DP needs up to 63%, 52%, and 67% fewer queries to obtain the same average rewards for buck-boost, buck, and boost converters, respectively. We also observe that all the algorithms (especially RS and GS) perform slightly worse for buck-boost converters, which may be caused by the step size of duty cycles as discussed above. Overall, the results demonstrate the efficacy of our UCT-based approach to discover high-quality topologies using only a small number of queries.

Ablative studies on default policies. In Section IV, we described how we incorporate offline knowledge using default policies. In this experiment, we examine the effectiveness of the default policies by performing an ablation study. Figure 7 presents the average rewards with and without the two default

(a) Buck-Boost Converter   (b) Buck Converter   (c) Boost Converter

Fig. 6: Average rewards under UCT-DP vs. GS vs. RS with increasing numbers of unique queries to circuit evaluations. The x-axis shows increasing numbers of queries to circuit evaluations for unique topologies. The y-axis is the obtained reward calculated using the power efficiency and voltage conversion ratio of the best candidate circuit generated by an algorithm, averaging from 200 runs.


Fig. 7: Ablative studies of UCT-based approaches without default policy (UCT), with node selection default policy (UCT-DP-Node), with edge selection default policy (UCT-DP-Edge), and with both node and edge selection default policies (UCT-DP) for boost converters.

(a) Generated Circuit   (b) Topological Representation

(c) Circuit Power Efficiency   (d) Circuit Output Voltage

Fig. 8: An example of unconventional power converter circuit discovered by our framework.

policies, namely node selection and edge selection, for the boost converter design task. The results for other power converter types are similar. Not surprisingly, UCT-DP with both default policies performs the best. We also observe that UCT with edge selection only (UCT-DP-Edge) performs significantly better than with node selection only (UCT-DP-Node). This is partly because choosing good connections among the exponentially many connections is more challenging. Moreover, since device type selections are performed in the first phase before connection selections, the node selection default policy is only used by the rollouts during the first phase. In contrast, the edge selection default policy is used by all the rollouts, so it has a higher impact on the performance.

D. Discussion on New Topologies

Our automatic power converter design framework generates topologies that meet design targets. Besides classical power converter topologies, it is also able to find interesting unconventional topologies t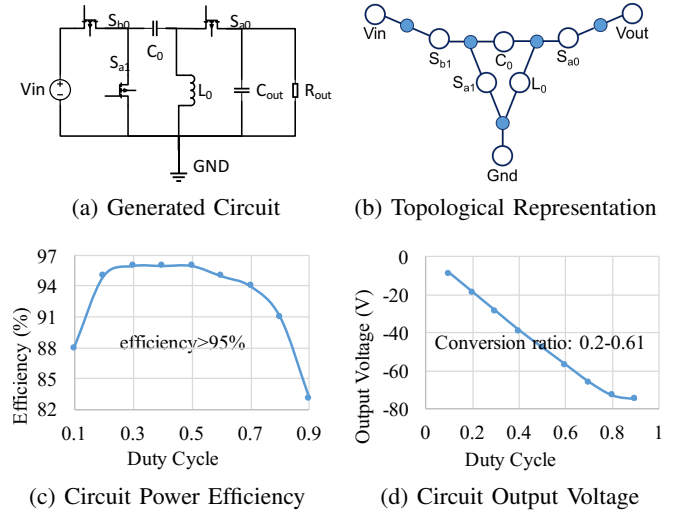hat both satisfy specifications and have high power efficiencies, such as the one in Figure 8. These automatically generated topologies have the potential to shed light on fundamental circuit innovations. With close collaboration with human experts, our framework can help to discover innovative circuits that have not been studied.

VI. CONCLUSION

In this work, we proposed a UCT-based power converter topology generation framework, which explores the design space automatically. We incorporated physics-informed constraints and data-driven default policies to reduce the design space and improve the efficiency of our framework. Additionally, we adopted a hybrid circuit evaluation with both the fast State-Space Averaging method and the accurate high-fidelity simulation. Finally, evaluations showed that our framework can generate near-optimal circuit topology for buck-boost, buck, and boost converters. Compared to the alternative approaches, our framework can discover better circuit topologies with reduced computational costs.

REFERENCES

[1] B. Nikolic, E. Alon, and K. Asanovic, "Generating the Next Wave of Custom Silicon," in *ESSCIRC 2018 - IEEE 44th European Solid State Circuits Conference (ESSCIRC)*, Sep. 2018, pp. 6–11, iSSN: 1930-8833.

[2] J. Kim, J. Lee, L. Vandenberghe, and C.-K. K. Yang, "Techniques for improving the accuracy of geometric-programming based analog circuit design optimization," in *IEEE/ACM International Conference on Computer Aided Design, 2004. ICCAD-2004.* IEEE, 2004, pp. 863–870.

[3] A. K. Singh, K. Ragab, M. Lok, C. Caramanis, and M. Orshansky, "Predictable equation-based analog optimization based on explicit capture of modeling error statistics," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 31, no. 10, pp. 1485–1498, 2012.

[4] Y. Wang, M. Orshansky, and C. Caramanis, "Enabling efficient analog synthesis by coupling sparse regression and polynomial optimization," in *Proceedings of the 51st Annual Design Automation Conference*, 2014, pp. 1–6.

[5] 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*. PMLR, 2018, pp. 3306–3314.

[6] K.-E. Yang, C.-Y. Tsai, H.-H. Shen, C.-F. Chiang, F.-M. Tsai, C.-A. Wang, Y. Ting, C.-S. Yeh, and C.-T. Lai, "Fast Design Space Adaptation with Deep Reinforcement Learning for Analog Circuit Sizing," *arXiv preprint arXiv:2009.13772*, 2020.

[7] W. Lee and F. A. Oliehoek, "Analog Circuit Design with Dyna-Style Reinforcement Learning," in *NeurIPS 2020 Workshop: Machine Learning for Engineering Modeling, Simulation, and Design*, 2020.

[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 *2020 57th ACM/IEEE Design Automation Conference (DAC)*, 2020, pp. 1–6.

[9] Z. Zhou, S. Belakaria, A. Deshwal, W. Hong, J. R. Doppa, P. P. Pande, and D. Heo, "Design of Multi-Output Switched-Capacitor Voltage Regulator via Machine Learning," in *2020 Design, Automation Test in Europe Conference Exhibition (DATE)*, 2020, pp. 502–507.

[10] K. Hakhamaneshi, N. Werblun, P. Abbeel, and V. Stojanović, "Late Breaking Results: Analog Circuit Generator based on Deep Neural Network enhanced Combinatorial optimization," in *2019 56th ACM/IEEE Design Automation Conference (DAC)*, 2019, pp. 1–2.

[11] E. Chang, J. Han, W. Bae, Z. Wang, N. Narevsky, B. NikoliC, and E. Alon, "BAG2: A process-portable framework for generator-based AMS circuit design," in *2018 IEEE Custom Integrated Circuits Conference (CICC)*, 2018, pp. 1–8.

[12] M. Liu, W. Li, K. Zhu, B. Xu, Y. Lin, L. Shen, X. Tang, N. Sun, and D. Z. Pan, "S3DET: Detecting System Symmetry Constraints for Analog Circuits with Graph Similarity," in *2020 25th Asia and South Pacific Design Automation Conference (ASP-DAC)*, 2020.

[13] S. Han, S. Jeong, C. Kim, H.-J. Park, and B. Kim, "GUI-Enhanced Layout Generation of FFE SST TXs for Fast High-Speed Serial Link Design," in *2020 57th ACM/IEEE Design Automation Conference (DAC)*, 2020, pp. 1–6.

[14] H. Chen, K. Zhu, M. Liu, X. Tang, N. Sun, and D. Z. Pan, "Toward Silicon-Proven Detailed Routing for Analog and Mixed-Signal Circuits," in *2020 IEEE/ACM International Conference On Computer Aided Design (ICCAD)*, 2020, pp. 1–8.

[15] G. Zhang, H. He, and D. Katabi, "Circuit-GNN: Graph Neural Networks for Distributed Circuit Design," in *International Conference on Machine Learning*. PMLR, 2019, pp. 7364–7373.

[16] Z. Zhao and L. Zhang, "An Automated Topology Synthesis Framework for Analog Integrated Circuits," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)*, vol. 39, no. 12, pp. 4325–4337, 2020.

[17] z. Rojec, A. Burmen, and I. Fajfar, "Analog circuit topology synthesis by means of evolutionary computation," *Engineering Applications of Artificial Intelligence*, vol. 80, pp. 48–65, 2019.

[18] P. Nenzi and H. Vogt, "Ngspice users manual version 23," 2011.

[19] S. Zhao, F. Blaabjerg, and H. Wang, "An Overview of Artificial Intelligence Applications for Power Electronics," *IEEE Transactions on Power Electronics*, vol. 36, no. 4, pp. 4633–4658, 2021.

[20] Y. Zhang, Z. Wang, H. Wang, and F. Blaabjerg, "Artificial Intelligence-Aided Thermal Model Considering Cross-Coupling Effects," *IEEE Transactions on Power Electronics*, vol. 35, no. 10, pp. 9998–10 002, 2020.

[21] T. Wu, Z. Wang, B. Ozpineci, M. Chinthavali, and S. Campbell, "Automated Heatsink Optimization for Air-Cooled Power Semiconductor Modules," *IEEE Transactions on Power Electronics*, vol. 34, no. 6, pp. 5027–5031, 2019.

[22] Y.-C. Lu, S. S. K. Pentapati, L. Zhu, K. Samadi, and S. K. Lim, "Tp-gnn: a graph neural network framework for tier partitioning in monolithic 3d ics," in *2020 57th ACM/IEEE Design Automation Conference (DAC)*. IEEE, 2020, pp. 1–6.

[23] K. Hakhamaneshi, N. Werblun, P. Abbeel, and V. Stojanović, "Bagnet: Berkeley analog generator with layout optimizer boosted with deep neural networks," in *2019 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*. IEEE, 2019, pp. 1–8.

[24] K. Settaluri, A. Haj-Ali, Q. Huang, K. Hakhamaneshi, and B. Nikolic, "Autockt: Deep reinforcement learning of analog circuit designs," in *2020 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 2020, pp. 490–495.

[25] W. M. Polivka, P. R. Chetty, and R. D. Middlebrook, "State-Space Average modelling of converters with parasitics and storage-time modulation," in *1980 IEEE Power Electronics Specialists Conference*, 1980, pp. 119–143.

[26] T. McConaghy, P. Palmers, M. Steyaert, and G. G. E. Gielen, "Trustworthy Genetic Programming-Based Synthesis of Analog Circuit Topologies Using Hierarchical Domain-Specific Building Blocks," *IEEE Transactions on Evolutionary Computation*, vol. 15, no. 4, pp. 557–570, 2011.