FLASH: <u>Fast Neural Architecture Search with Hardware Optimization</u>*

GUIHONG LI, The University of Texas at Austin, USA SUMIT K. MANDAL, University of Wisconsin–Madison, USA UMIT Y. OGRAS, University of Wisconsin–Madison, USA RADU MARCULESCU, The University of Texas at Austin, USA

Neural architecture search (NAS) is a promising technique to design efficient and high-performance deep neural networks (DNNs). As the performance requirements of ML applications grow continuously, the hardware accelerators start playing a central role in DNN design. This trend makes NAS even more complicated and time-consuming for most real applications. This paper proposes FLASH, a very fast NAS methodology that co-optimizes the DNN accuracy and performance on a real hardware platform. As the main theoretical contribution, we first propose the NN-Degree, an analytical metric to quantify the topological characteristics of DNNs with skip connections (e.g., DenseNets, ResNets, Wide-ResNets, and MobileNets). The newly proposed NN-Degree allows us to do training-free NAS within one second and build an accuracy predictor by training as few as 25 samples out of a vast search space with more than 63 billion configurations. Second, by performing inference on the target hardware, we fine-tune and validate our analytical models to estimate the latency, area, and energy consumption of various DNN architectures while executing standard ML datasets. Third, we construct a hierarchical algorithm based on simplicial homology global optimization (SHGO) to optimize the model-architecture co-design process, while considering the area, latency, and energy consumption of the target hardware. We demonstrate that, compared to the state-of-the-art NAS approaches, our proposed hierarchical SHGO-based algorithm enables more than four orders of magnitude speedup (specifically, the execution time of the proposed algorithm is about 0.1 seconds). Finally, our experimental evaluations show that FLASH is easily transferable to different hardware architectures, thus enabling us to do NAS on a Raspberry Pi-3B processor in less than 3 seconds.

 $CCS\ Concepts: \bullet\ Computing\ methodologies \rightarrow Artificial\ intelligence;\ Computer\ vision; \bullet\ Computer\ systems\ organization \rightarrow Embedded\ systems.$

Additional Key Words and Phrases: Neural Networks, Network Science, Hardware Optimization, Neural Architecture Search, Model-Architecture Co-design, Resource-constrained Devices

ACM Reference Format:

Guihong Li, Sumit K. Mandal, Umit Y. Ogras, and Radu Marculescu. 2021. FLASH: Fast Neural Architecture Search with Hardware Optimization. ACM Trans. Embedd. Comput. Syst. 1, 1, Article 1 (January 2021), 25 pages. https://doi.org/10.1145/3476994

*This article appears as part of the ESWEEK-TECS special issue and was presented in the International Conference on Hardware/Software Codesign and System Synthesis (CODES+ISSS), 2021.

Authors' addresses: Guihong Li, lgh@utexas.edu, The University of Texas at Austin, Austin, Texas, USA; Sumit K. Mandal, skmandal@wisc.edu, University of Wisconsin-Madison, Madison, Wisconsin, USA; Umit Y. Ogras, uogras@wisc.edu, University of Wisconsin-Madison, Madison, Wisconsin, USA; Radu Marculescu, radum@utexas.edu, The University of Texas at Austin, Austin, Texas, USA.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Association for Computing Machinery.

1539-9087/2021/1-ART1 \$15.00

https://doi.org/10.1145/3476994

1:2 G. Li, et al.

1 Introduction

During the past decade, deep learning (DL) has led to significant breakthroughs in many areas, such as image classification and natural language processing [6, 21, 25]. However, the existing large models and computation complexity limit the deployment of DL on resource-constrained devices and its large-scale adoption in edge computing. Multiple model compression techniques, such as network pruning [20], quantization [13], and knowledge distillation [22], have been proposed to compress and deploy such complex models on resource-constrained devices without sacrificing the test accuracy. However, these techniques require a significant amount of manual tuning. Hence, neural architecture search (NAS) has been proposed to automatically design neural architectures with reduced model sizes [2, 17, 31, 32, 60].

NAS is an optimization problem with specific targets (e.g., high classification accuracy) over a set of possible candidate architectures. The set of candidate architectures defines the (typically vast) search space, while the optimizer defines the search algorithm. Recent breakthroughs in NAS can simplify the tricky (and error-prone) ad-hoc architecture design process [31, 42]. Moreover, the networks obtained via NAS have higher test accuracy and significantly fewer parameters than the hand-designed networks [32, 44]. These advantages of NAS have attracted significant attention from researchers and engineers alike [55]. However, most of the existing NAS approaches do not explicitly consider the hardware constraints (e.g., latency and energy consumption). Consequently, the resulting neural networks still cannot be deployed on real devices.

To address this drawback, recent studies propose *hardware-aware NAS*, which incorporates the hardware constraints of networks during the search process [27]. Nevertheless, current approaches are time-consuming since they involve training the candidate network, and a tedious search process [56]. To accelerate NAS, recent NAS approaches rely on graph neural networks (GNNs) to estimate the accuracy of a given network [9, 33, 40, 54]. However, training a GNN-based accuracy predictor is still time-consuming (in the order of tens of minutes [12] to hours [36] on GPU clusters). Therefore, adapting existing NAS approaches to different hardware architecture is challenging due to their intensive computation and execution time requirements.

To alleviate the computation cost of current NAS approaches, we propose to analyze the NAS problem from a *network topology* perspective. This idea is motivated by observing that the tediousness and complexity of current NAS approaches stem from the lack of understanding of what actually contributes to a neural network's accuracy. Indeed, the innovations on the topology of neural architecture, especially the introduction of skip connections, have achieved great success in many applications [21, 25]. This is because, in general, the network topology (or structure) strongly influences the phenomena taking place over them [39]. For instance, how closely the social network users are interconnected directly affects how fast the information propagates through the network [3]. Similarly, a DNN architecture can be seen as a network of connected neurons. As discussed in [5], the topology of deep networks has a significant impact on how effectively the gradients can propagate through the network and thus the test performance of neural networks. These observations motivate us to take an approach from network science to quantify the topological property of neural networks to accelerate NAS.

From an application perspective, the performance and energy efficiency of DNN accelerators are other critical metrics besides the test accuracy. In-memory computing (IMC)-based architectures have recently emerged as a promising technique to construct high-performance and energy-efficient hardware accelerators for DNNs. IMC-based architectures can store all the weights on-chip, hence removing the latency occurring from off-chip memory accesses. However, IMC-based architectures face the challenge of a tremendous increase of on-chip communication volume. While most of the state-of-the-art neural networks adopt skip connections in order to improve their

performance [21, 25, 45], the wide usage of skip connections requires large amounts of data transfer across multiple layers, thus causing a significant communication overhead. Prior work on IMC-based DNN accelerators proposed bus-based network-on-chip (NoC) [10] or cmesh-based NoC [46] for communication between multiple layers. However, both bus-based and cmesh-based on-chip communication significantly increase the area, latency, and energy consumption of hardware; hence, they do not offer a promising solution for future accelerators.

Starting from these overarching ideas, this paper proposes FLASH - a fast neural architecture search with hardware optimization - to address the drawbacks of current NAS techniques. FLASH delivers a neural architecture that is co-optimized with respect to accuracy and hardware performance. Specifically, by analyzing the topological property of neural architectures from a network science perspective, we propose a new topology-based metric, namely, the NN-Degree. We show that NN-Degree could indicate the test performance of a given architectures. This makes our proposed NAS training-free during the search process and accelerates NAS by orders of magnitude compared to state-of-the-art approaches. Then, we demonstrate that NN-Degree enables a lightweight accuracy predictor with only three parameters. Moreover, to improve the on-chip communication efficiency, we adopt the mesh-NoC for the IMC-based hardware. Based on the communication-optimized hardware architecture, we measure the hardware performance for a subset of neural networks from the NAS search space. Then, we construct analytical models for the area, latency, and energy consumption of a neural network based on our optimized target hardware platform. Unlike existing neural network-based and black-box style searching algorithms [27], the proposed NAS methodology enable searching across the entire search space via a mathematically rigorous and timeefficient optimization algorithm. Consequently, our experimental evaluations show that FLASH significantly pushes forward the NAS frontier by enabling NAS in less than 0.1 seconds on a 20-core Intel Xeon CPU. Finally, we demonstrate that FLASH could be readily transferred to other hardware platforms (e.g., Raspberry Pi) only by fine-tuning the hardware performance models.

Overall, this paper makes the following contributions:

- We propose a new topology-based analytical metric (*NN-Degree*) to quantify the topological characteristics of DNNs with skip connections. We demonstrate that the NN-Degree enables a *training-free* NAS within seconds. Moreover, we use the NN-Degree metric to build a new lightweight (*three-parameter*) accuracy predictor by training as few as 25 samples out of a vast search space with more than 63 billion configurations. Without any significant loss in accuracy, our proposed accuracy predictor requires 6.88× fewer samples and provides a 65.79× reduction of the fine-tuning time cost compared to existing GNN/GCN based approaches [54].
- We construct analytical models to estimate the latency, area, and energy consumption of various DNN architectures. We show that our proposed analytical models are applicable to multiple hardware architectures and achieve a high accuracy with less than one second fine-tuning time cost.
- We design a hierarchical simplicial homology global optimization (SHGO)-based algorithm, to search for the optimal architecture. Our proposed hierarchical SHGO-based algorithm enables 27729× faster (less than 0.1 seconds) NAS compared to RL-based baseline approach.
- We demonstrate that our methodology enables NAS on a Raspberry Pi 3B with less than 3 seconds computational time. To our best knowledge, this is the first work showing NAS running directly on edge devices with such low computational requirements.

The rest of the paper is organized as follows. In Section 2, we discuss related work and background information. In Section 3, we formulate the optimization problem, then describe the new analytical models and search algorithm. Our experimental results are presented in Section 4. Finally, Section 5 concludes the paper with remarks on our main contributions and future research directions.

1:4 G. Li, et al.

2 Related Work and Background Information

Hardware-aware NAS: Hardware accelerators for DNNs have recently become popular due to high-performance demand for multiple applications [4, 15, 35]; they can reduce the latency and energy associated with DNN inference significantly. The hardware performance (e.g., latency, energy, and area) of accelerators varies with DNN properties (e.g., number of layers, parameters, etc.); therefore, hardware performance also is a crucial factor to consider during NAS.

Several recent studies consider hardware performance for NAS. Authors in [14] introduce a growing and pruning strategy that automatically maximize the test accuracy and minimize the FLOPs of neural architectures during training. A platform-aware NAS targeting mobile devices is proposed in [50]; the objective is to maximize the model accuracy with an upper bound on latency. Authors in [56] create a latency-aware loss function to perform differentiable NAS. The latency of DNNs is estimated through a lookup table which consists of the latency of each operation/layer. However, both of these studies consider latency as the only metric for hardware performance. Authors in [37] propose a hardware-aware NAS framework to design convolutional neural networks. Specifically, by building analytical latency, power, and memory models, they create a hardwareaware optimization methodology to search for the optimal architecture that meets the hardware budgets. Authors in [27] consider latency, energy, and area as metrics for hardware performance while performing NAS. Also, a reinforcement learning (RL)-based controller is adopted to tune the network architecture and device parameters. The resulting network is retrained to evaluate the model accuracy. There are two major drawbacks of this approach. First, RL is a slow-converging process that prohibits fast exploration of the design space. Second, retraining the network further exacerbates the search time leading to hundreds of GPU hours needed for real applications [60]. Furthermore, most existing hardware-aware NAS approaches explicitly optimize the architectures for a specific hardware platform [8, 30, 56]. Hence, if we switch to some new hardware, we need to repeat the entire NAS process, which is very time-consuming under the existing NAS frameworks [8, 30, 56]. The demand for reducing the overhead of adaptation to new hardware motivates us to improve the transferability of hardware-aware NAS methodology.

Accuracy Predictor-based NAS: Several approaches perform NAS by estimating the accuracy of the network [9, 33, 40, 54]. These approaches first train a graph neural network (GNN), or a graph convolution network (GCN), to estimate the network accuracy while exploring the search space. During the searching process, the test accuracy of the sample networks is obtained from the estimator instead of doing regular training. Although by estimating the accuracy, the NAS process is significantly accelerated, the training cost of the accuracy predictor itself remains a bottleneck. GNN requires many training samples to achieve high accuracy, thus involving a significant overhead during training the candidate networks from the search space. Therefore, using accuracy predictors to do NAS still suffers from excessive computation and time requirements.

Time-efficient NAS: To reduce the time cost of training candidate networks, authors in [42, 49] introduced the weight sharing mechanism (WS-NAS). Specifically, candidate networks are generated by randomly sampling part of a large network (supernet). Hence, candidate networks share the weights of the supernet and update these weights during training. By reusing these trained weights instead of training from scratch, WS-NAS significantly improves the time efficiency of NAS. However, the accuracy of these models obtained via WS-NAS is typically far below those obtained from training from scratch. Several optimization techniques have been proposed to fill the accuracy gap between the shared weights and stand-alone training [7, 58]. For example, authors in [7] propose a progressive shrinking algorithm to train the supernet. However, in many cases, the resulting networks still need some fine-tuning epochs to get the final architecture. To further accelerate the NAS process, some works propose the differentiable NAS to accelerate the NAS

process [8, 32]. The differentiable NAS approaches search for the optimal architecture by learning the optimal architecture parameters during the training process. Hence, differentiable NAS only needs to train the supernet once, thus reducing the training time significantly. Nevertheless, due to the significantly large number of parameters of the supernet, differentiable NAS requires a high volume of GPU memory. In order to further improve the time-efficiency of NAS, several approaches have been proposed to do training-free NAS [1, 11]. These approaches leverage some training-free proxy that indicates the test performance of some given architectures; hence, the training time is eliminated from the entire NAS process. However, these methods usually use gradient-based information to build the proxy [1, 11]. Therefore, in order to calculate the gradients, GPUs are still necessary for the backward propagation process. To totally decouple the NAS process from using GPU platforms, our work proposes a GPU-free proxy to do training-free NAS. We provide more details in Section 4.3.

Skip connections and Network Science: Currently, both networks obtained by manual design and NAS have shown that long-range links (i.e., skip connections) are crucial for getting higher accuracy [21, 25, 32, 45]. Overall, there are two commonly used skip connections in neural networks. First, we have the *DenseNet-type* skip connections (DTSC), which concatenate previous layers' outputs as the input for the next layer [25]. To study the topological properties and enlarge the search space, we do *not* use the original DesneNets [25], which contains all-to-all connections. Instead, we consider a generalized version where we vary the number of skip connections by randomly selecting only some channels for concatenation, as shown in Fig. 1(a). The other type of skip connections is the *addition-type* skip connections (ATSC), which consist of links that bypass several layers to be directly added to the output of later layers (see Fig. 1(b)) [21].

In network science, a small-world network is defined as a highly clustered network, thus showing a small distance (typically logarithmic in the number of network nodes) between any two nodes inside the network [53]. Considering the skip connections in neural networks, we propose to use the *small-world network* concept to analyze networks with both short- and long-range (or skip) links. Indeed, small-world networks can be decomposed into: (i) a lattice network $\mathcal G$ accounting for short-range links; (ii) a random network $\mathcal R$ accounting for long-range links (see Fig. 1(c)). The co-existence of a rich set of short- and long-range links leads to both a high degree of clustering

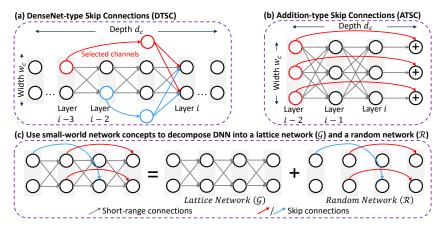


Fig. 1. Modeling a CNN as a network in network science: Each channel is modeled as a node; each convolution kernel/filter is modeled as a link/connection. (a) Illustration of a single cell with DenseNet-type skip connections (DTSC). (b) Illustration of a single cell with Addition-type skip connections (ATSC). (c) Decomposition of a network cell with skip connections into a Lattice Network $\mathcal G$ and a Random Network $\mathcal R$.

1:6 G. Li, et al.

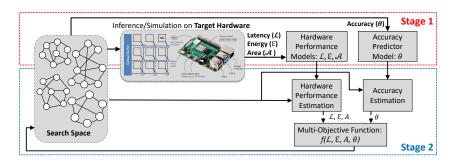


Fig. 2. Overview of the proposed approach. Stage 1 (red box): we build hardware performance model (i.e., latency \mathcal{L} , energy \mathcal{E} , and area \mathcal{A}) and accuracy predictor by randomly sampling candidate networks from the search space to evaluate the hardware characteristics (latency \mathcal{L} , energy \mathcal{E} , and area \mathcal{A}) and test accuracy θ . Stage 2 (blue box): search for the optimal network architecture given the multi-objective function $f(\mathcal{L}, \mathcal{E}, \mathcal{A}, \theta)$.

and short average path length (logarithmic with network size). We use the small-world network to model and analyze the topological property of neural networks in Section 3.

Average Degree: The *average degree* of a network determines the average number of connections a node has, i.e., the total number of edges divided by the total number of nodes. The average degree and degree distribution (i.e., distribution of node degree) are important topological characteristics that directly affect how information flows through a network [3]. Indeed, the small network theory reveals that the average degree of a network has a significant impact on network average path length and clustering behavior [53]. Therefore, we investigate the performance gains due to the topological properties by using network science.

3 Proposed Methodology

3.1 Overview of New NAS Approach

The proposed NAS framework is a two-stage process, as illustrated in Fig. 2: (i) We first quantify the topological characteristics of neural networks by the newly proposed NN-Degree metric. Then, we randomly select a few networks and train them to fine-tune the accuracy predictor based on the network topology. We also build analytical models to estimate the latency, energy, and area of given neural architectures. (ii) Based on the accuracy predictor and analytical performance models in the first stage, we use a simplical homology global optimization (*SHGO*)-based algorithm in a hierarchical fashion to search for the optimal network architecture.

3.2 Problem Formulation of hardware-aware NAS

The overall target of the hardware-aware NAS approach is to find the network architecture that gives the highest test accuracy while achieving small area, low latency, and low energy consumption when deployed on the target hardware. In practice, there are constraints (budgets) on the hardware performance and test accuracy. For example, battery-based devices have very constrained energy capacity [52]. Hence, there is an upper bound for the energy consumption of the neural architecture. To summarize, the NAS problem can be expressed as:

$$\max \quad f_{obj} = \frac{\theta}{\mathcal{A} \times \mathcal{L} \times \mathcal{E}}$$
 subject to: $\theta \ge \theta_M, \ \mathcal{A} \le \mathcal{A}_M, \ \mathcal{L} \le \mathcal{L}_M, \ \mathcal{E} \le \mathcal{E}_M$ (1)

where θ_M , \mathcal{A}_M , \mathcal{L}_M , and \mathcal{E}_M are the constraints on the test accuracy, area, latency, and energy consumption, respectively. We summarize the symbols (and their meaning) in this part in Table 1.

Symbol	Definition
f_{obj}	Objective function of NAS
θ	Test accuracy of a given network
A	Chip area
\mathcal{L}	Inference latency of a given network
3	Inference energy consumption of a given network
θ_M	Constraint of test accuracy for NAS
A_M	Constraint of area for NAS
\mathcal{L}_{M}	Constraint of inference latency for NAS
\mathcal{E}_{M}	Constraint of inference energy consumption for NAS

Table 1. Symbols and their corresponding definition/meaning used in our Problem Formulation.

3.3 NN-Degree and Training-free NAS

This section first introduces our idea of modeling a CNN based on network science [53]. To this end, we define a group of consecutive layers with the same width (i.e., number of output channels, w_c) as a *cell*; then we break the entire network into multiple cells and denote the number of cells as N_c . Similar to MobileNet-v2 [45], we also adopt a width multiplier (w_m) to scale the width of each cell. Moreover, following most of the mainstream CNN architectures, we assume that each cell inside a CNN has the same number of layers (d_c). Furthermore, as shown in Fig. 1, we consider each channel of the feature map as a node in a network and consider each convolution filter/kernel as an undirected link. These notations are summarized in Table 2.

Table 2. Symbols and their corresponding definition/meaning used in our NN-Degree based analytical accuracy predictor.

Symbol	Definition
g	NN-Degree (new metric we propose)
$g_{\mathcal{G}}$	NN-Degree of the lattice network (short-range connections)
$g_{\mathcal{R}}$	NN-Degree of the random network (long-range or skip connections)
N_c	Number of cells
w_c	Number of output channels per layer within cell c (i.e., the width of cell c)
d_c	Number of layers within cell c (i.e., the depth of cell c)
SC_c	Number of skip connections within cell <i>c</i>
$a_{\theta}, b_{\theta}, c_{\theta}$	Learnable parameters for the accuracy predictor

Combining the concept of small-world networks in Section 2 and our modeling of a CNN, we decompose a network cell with skip connections into a lattice network \mathcal{G} and random network \mathcal{R} (see Fig. 1(c)).

Proposed Metrics: Our key objective is two-fold: (i) Quantify which topological characteristics of DNN architectures affect their performance, and (ii) Exploit such properties to accurately predict the test accuracy of a given architecture. To this end, we propose a new analytical metric called NN-Degree, as defined below.

Definition of NN-Degree: Given a DNN with N_c cells, d_c layers per cell, the width of each cell w_c , and the number of skip connections of each cell SC_c , the NN-Degree metric is defined as the sum of the average degree of each cell:

$$g = \sum_{c=1}^{N_c} (w_c + \frac{SC_c}{w_c \times d_c})$$
 (2)

1:8 G. Li, et al.

Intuition: The average degree of a given DNN cell is the sum of the average degrees from lattice network \mathcal{G} and random network \mathcal{R} . Given a cell with d_c convolutional layers and w_c channels per layer, the number of nodes is $w_c \times d_c$. Moreover, each convolutional layer has $w_c \times w_c$ filters (kernels) accounting for the short-range connections; hence, in the lattice network \mathcal{G} , there are $w_c \times w_c \times d_c$ connections (total). Using the above analysis, we can express the NN-Degree as follows:

$$g = g_{\mathcal{G}} + g_{\mathcal{R}}$$

$$= \sum_{c=1}^{N_c} \frac{number\ of\ connections\ in\ \mathcal{G}}{number\ of\ nodes\ in\ cell\ c} + \sum_{c=1}^{N_c} \frac{number\ of\ connections\ in\ \mathcal{R}}{number\ of\ nodes\ in\ cell\ c}$$

$$= \sum_{c=1}^{N_c} \frac{w_c \times d_c \times w_c}{w_c \times d_c} + \sum_{c=1}^{N_c} \frac{number\ of\ skip\ connections}{w_c \times d_c}$$

$$= \sum_{c=1}^{N_c} (w_c + \frac{SC_c}{w_c \times d_c})$$
(3)

Discussion: The first term in Equation 3 (i.e., $g_{\mathcal{G}}$) reflects the width of the network w_c . Many successful DNN architectures, such as DenseNets [25], Wide-ResNets [59], and MobileNets [45], have shown that wider networks can achieve a higher test performance. The second term (i.e., $g_{\mathcal{R}}$) quantifies how densely the nodes are connected through the skip connections. As discussed in [51], networks with more skip connections have more forward/backward propagation paths, thus have a better test performance. Based on the above analysis, we claim that a higher NN-Degree value should indicate networks with higher test performance. We verify this claim empirically in the experimental section. Next, we propose an accuracy predictor based only on the NN-Degree.

Accuracy Predictor: Given the NN-Degree (g) definition, we build the accuracy predictor by using a variant of logistic regression. Specifically, the test accuracy θ of a given architecture is:

$$\theta = \frac{1}{a_{\theta} + \exp(b_{\theta} \times \frac{1}{a} + c_{\theta})} \tag{4}$$

where a_{θ} , b_{θ} , c_{θ} are the parameters that are fine-tuned with the accuracy and NN-Degree of sample networks from the search space. Section 4 shows that by using as few as 25 data samples (NN-Degree and corresponding accuracy values), we can generate an accurate predictor for a huge search space covering more than 63 billion configurations within 1 second on a 20-core Intel Xeon CPU.

Training-free NAS: Section 4 shows that NN-Degree can indicate the test accuracy of a given architecture. Hence, one can use NN-Degree as a proxy of the test accuracy to enable the training-free NAS. Section 4.3 demonstrates that we can do *training-free* NAS within 0.11 seconds on a 20-core CPU.

3.4 Overview of In-memory Computing (IMC)-based Hardware

Fig. 3 shows the IMC architecture considered in this work. We note that the proposed FLASH methodology is not specific to IMC-based hardware. We adopt an IMC architecture since it has been proven to achieve less memory access latency [23]. Due to the high communication volume imposed by deeper and denser networks, communication between multiple tiles is crucial for hardware performance, as shown in [28, 34].

Our architecture consists of multiple tiles connected by network-on-chip (NoC) routers, as shown in Fig. 3(a). We use a mesh-based NoC due to its superior performance compared to bus-based architectures. Each tile consists of a fixed number of compute elements (CE), a rectified linear unit (ReLU), an I/O buffer, and an accumulation unit, as shown in Figure Fig. 3(b).

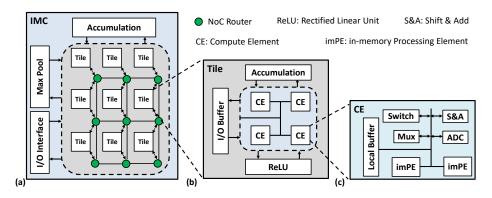


Fig. 3. Details of the IMC hardware. (a) The architecture consists of multiple tiles connected via routers; (b) The structure of a tile. Each tile consists of multiple computing elements (CE), I/O buffer, ReLU unit and accumulation unit; (c) The structure of each CE. Each CE consists of multiple in-memory processing elements (imPE), local buffers, switch, multiplexer, analog to digital converter (ADC), shift and add (S&A) circuit.

Within each CE, there exist a fixed number of im-memory processing elements (imPE), a multiplexer, a switch, an analog-to-digital converter (ADC), a shift and add (S&A) circuit, and a local buffer [10], as shown in Fig. 3(c). The ADC precision is set to four bits to avoid any accuracy degradation. There is no digital-to-analog (DAC) converter used in the architecture. A sequential signaling technique to represent multi-bit inputs is adopted [41]. Each imPE consists of 256×256 IMC crossbars (the memory elements) based on ReRAM (1T1R) technology [10, 28, 34]. This work incorporates a sequential operation between DNN layers since a pipelined operation may cause pipeline bubbles during inference [43, 48].

Table 3. Symbols and their corresponding definition used in our analytical area, latency, and energy models.

Symbol	Definition	Symbol	Definition
N_c	Number of cells	N_i^r	Number of rows of imPE arrays of i^{th} layer
$a_{ heta}, b_{ heta}, c_{ heta}$	Learnable parameters for accuracy predictor	N_i^c	Number of columns of imPE arrays of i th layer
w_m	Width multiplier	Kx_i, Ky_i	Kernel size of <i>i</i> th layer
d_c	Number of layers within cell c	N_i^{if}, N_i^{of}	Number of input and output features of i th layer
w_c	Width of cell <i>c</i>	$(PE_x)_i, (PE_y)_i$	Size of a single imPE of <i>i</i> th layer
SC_c	Number of skip connections within cell c	T_i	Number of tiles of i^{th} layer
$FLOP_c$	Number of FLOPs of cell <i>c</i>	С	Number of CEs in each tile of i^{th} layer
$Comm_c$	The amount of data transferred through NoC inside cell <i>c</i>	p	Number of imPEs in each CE
N_T	Total number of tiles of the chip	A_T	Area of a tile
$F_{\mathcal{E}}$	Features for energy	\mathcal{E}^T	Energy consumption of a tile
$\Lambda_{comp},\ \Lambda_{NoC}$	Weight vectors to estimate computation and NoC latency	F_{Comp}, F_{NoC}	Features to estimate computation and NoC latency

1:10 G. Li, et al.

3.5 Hardware Performance Modeling

This section describes the methodology of modeling hardware performance. We consider three metrics for hardware performance: area, latency, and energy consumption. We use customized versions of NeuroSim [10] for circuit simulation (computing fabric) and BookSim [26] for cycleaccurate NoC simulation (communication fabric). First, we describe the details of the simulator.

Input to the simulator: The inputs to the simulator include the DNN structure, technology node, and frequency of operation. In this work, we consider a layer-by-layer operation. Specifically, we simulate each DNN layer and add its performance at the end to obtain the total performance of the hardware for the DNN.

Simulation of computing fabric: Table 4 shows the parameters considered for the simulation of computing fabric. At the start of the simulation, the number of in-memory computing tiles is computed. Then, the area and energy of one tile are computed through analytical models derived from HSPICE simulation. After that, the area and energy of one tile are multiplied by the total number of tiles to obtain the total area and energy of the computing fabric. The latency of the computing fabric is computed as a function of the workload (the DNN being executed). We note that the original version of NeuroSim considers point-to-point on-chip interconnects, while our proposed work uses mesh-based NoC. Therefore, we skip the interconnect simulation in NeuroSim. Simulation of communication fabric: We consider cycle-accurate simulation for the communication fabric. BookSim is used to perform simulation. First, the number of tiles required for each layer is obtained from the simulation of computing fabric. In this work, we assume that each tile is connected to a dedicated router of the NoC. A trace file is generated corresponding to the particular layer of the DNN. The trace file consists of the information of the source router, destination router, and timestamp when the packet is generated. The trace file is simulated through BookSim to obtain the latency to finish all the transactions between two layers. We also obtain the area and energy of the interconnect through BookSim. Table 4 shows the parameters considered for the interconnect simulator. More details of the simulator can be found in [29].

For hardware performance modeling, first we obtain the performance of the DNN through simulation, then the performance numbers are used to construct the performance models.

Analytical Area Model: An in-memory computing-based DNN accelerator consists of two major components: computation and communication. The computation unit consists of multiple tiles and peripheral circuits; the communication unit includes an NoC with routers and other network components (e.g., buffers, links). To estimate the total area, we first compute the number of rows (N_i^r) and number of columns (N_i^c) of imPEs required for the i^{th} layer of the DNN following Equation 5 and Equation 6.

$$N_i^r = \left\lceil \frac{Kx_i \times Ky_i \times N_i^{if}}{(PE_x)_i} \right\rceil \tag{5}$$

$$N_i^c = \left\lceil \frac{N_i^{of} \times N_{bits}}{(PE_y)_i} \right\rceil \tag{6}$$

where all the symbols are defined in Table 3. Therefore, total number of imPEs required for the i^{th} layer of the DNN is $N_i^r \times N_i^c$. Each tile consists of c CEs, and each CE consists of p number of imPEs. Accordingly, each tile comprises $c \times p$ imPEs. Therefore, the total number of tiles required for the i^{th} layer of the DNN (T_i) is:

$$T_i = \left\lceil \frac{N_i^r \times N_i^c}{c \times p} \right\rceil \tag{7}$$

Hence, the total number of tiles (N_T) required for a given DNN is $N_T = \sum_i T_i$.

Circuit		NoC		
imPE array size	128×128	Bus width	32	
Cell levels	2 bit/cell	Routing algorithm	X-Y	
Flash ADC resolution	4 bits	Number of router ports	5	
Technology used	RRAM	Topology	Mesh	

Table 4. Parameters used for simulation of computation and communication fabric.

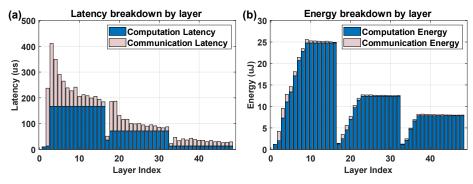


Fig. 4. Layerwise hardware performance breakdown of a DNN with 3 cells ($N_c = 3$), 16 layers per cell ($d_c = 16$), and a total of 48 layers. (a) Latency breakdown layer by layer: the computation latency accounts for 37.9% of the total latency, while communication accounts for 62.1%. (b) Energy consumption breakdown layer by layer: the computation energy accounts for 96.1% of the total latency, while communication accounts for 3.9%.

As shown in Fig. 3(a), each tile is connected to the NoC routers for the on-chip communication. We assume that the total number of required routers is equal to the total number of tiles. Hence, the total chip area is expressed as follows:

$$\mathcal{A} = A_{comp} + A_{NoC}$$

$$= (A_{Tile}^{Tot} + A_{Periphery}) + (A_{Router}^{Tot} + A_{others})$$

$$= N_T \times A_T + N_T \times A_R + (A_{Periphery} + A_{others})$$

$$= N_T \times (A_T + A_R) + A_{rest}$$
(8)

where A_{Tile}^{Tot} is the area accounted for all tiles and A_{Router}^{Tot} is the total area accounted for all routers in the design. The area of a single tile is denoted by A_T ; there are N_T tiles in the design. Therefore $A_{Tile}^{Tot} = N_T \times A_T$. The area of the peripheral circuit $(A_{Periphery})$ consists of I/O interface, max pool unit, accumulation unit, and global buffer. The area of a single router is denoted by A_R ; the number of routers is equal to the number of tiles (N_T) . Therefore $A_{Router}^{Tot} = N_T \times A_R$. The area of other components in the NoC (A_{rest}) comprises links and buffers.

Analytical Latency Model: Similar to area, the total latency consists of computation latency and communication latency, as shown in Fig. 4(a). To construct the analytical model of latency, we use floating-point operations (FLOPs) of the network to represent the computational workload. We observe that the FLOPs of a given network are roughly proportional to the total number of convolution filters (kernels), which is the product of the number of layers and the square of the number of channels per layer (i.e., width value). In the network search space we consider, the width is equivalently represented by the width multiplier w_m and the number of layers is $N_c \times d_c$; hence, we express the number of FLOPs of a given network approximately as the product of the number

1:12 G. Li, et al.

of layers, and the square of width multiplier:

$$FLOPs \sim N_c d_c w_m^2 \tag{9}$$

Moreover, communication volume increases significantly due to the skip connections. To quantify the communication volume due to skip connections, we define $Comm_c$ (the communication volume of a given network cell c) as follows:

$$Comm_c = SC_c \times Feature map size of each SC$$

Combining the above analysis of computation latency and communication latency, we use a linear model to build our analytical latency model as follows:

$$\mathcal{L} = \mathcal{L}_{comp} + \mathcal{L}_{NoC} = \Lambda_{comp}^T F_{comp} + \Lambda_{NoC}^T F_{NoC}$$
 (10)

where Λ_{comp}^T is a weight vector and $F_{comp} = [w_m, d_c, N_c, N_c d_c w_m^2]$ is the vector of features with respect to the computation latency; Λ_{NoC}^T is another weight vector and $F_{comp} = [SC_c, Comm_c]$ is the vector of features corresponding to the NoC latency. We randomly sample some networks from the search space and measure their latency to fine-tune the values of Λ_{comp}^T and Λ_{NoC}^T .

Analytical Energy Model: We divide the total energy consumption into computation energy and communication energy, as shown in Fig. 4(b). Specifically, the entire computation process inside each tile consists of three steps:

- Read the input feature map from the I/O buffer to the CE;
- Perform computations in CE and ReLU unit, then update the results in the accumulator;
- Write the output feature map to the I/O buffer.

Therefore, both the size of feature map and FLOPs contribute to the computation energy of a single cell. Moreover, the communication energy consumption is primarily determined by the communication volume, i.e., $(Comm_c)$. Hence, we use a linear combination of features to estimate the energy consumption of each tile \mathcal{E}^T :

$$\mathcal{E}^T = \Lambda_{\mathcal{E}}^T F_{\mathcal{E}} \tag{11}$$

where $\Lambda_{\mathcal{E}}^T$ is a weight vector and $F_{\mathcal{E}} = [w_m, d_c, N_c, SC_c, Comm_c, FLOP_c, FM_c]$ are the features corresponding to the energy consumption of each tile. We use the measured energy consumption values of several sample networks to fine-tune the values of $\Lambda_{\mathcal{E}}^T$. The total energy consumption (\mathcal{E}) is the product of \mathcal{E}^T and number of tiles:

$$\mathcal{E} = \Lambda_{\mathcal{E}}^T F_{\mathcal{E}} N_T \tag{12}$$

We note that all the features used in both our accuracy predictor and analytical hardware performance model are only related to the architecture of the network through the basic parameters $\{w_m, d_c, N_c, SC_c\}$. Therefore, the analytical hardware models are lightweight. We note that there exist no other lightweight analytical models for IMC platforms. Besides this, FLASH is general and can be applied to different hardware platforms. For a given hardware platform, energy, latency, and area of the DNNs need to be first collected. Then the analytical hardware models need to be trained using the performance data.

3.6 Optimal neural architecture search

Based on the above accuracy predictor and analytical hardware performance models, we perform the second stage of our NAS methodology, i.e., searching for the optimal neural architecture by considering both test accuracy and hardware performance on the target hardware. To this end, we use a modified version of the Simplicial Homology Global Optimization (SHGO [18]) algorithm to search for the optimum architecture. SHGO has mathematically rigorous convergence properties on

Algorithm 1: Our hierarchical SHGO-based search algorithm

```
Input:
        Objective function: f_{obj};
        Global search space:
            SP_{global} = [N_{cmin}, N_{cmax}] \times [w_{m_{min}}, w_{m_{max}}] \times [d_{cmin}, d_{cmax}] \times [SC_{cmin}, SC_{cmax}];
        Search constraints: S_{cons} = \{L_M, E_M, A_M, \theta_M\};
        Coarse-grain search step size: \lambda
Output:
        The optimal architecture \{w_m^*, N_c^*, d_c^*, SC_c^*\};
Search Process:
      Initialize Candidate Architecture Set (CAS) as empty set;
      level 1: Fixed-w_m Search
      for w_m in [w_{m_{min}}, w_{m_{max}}] do | level 2: Coarse-grain Search
           fix w_m, search the optimum N_c^G, d_c^G, SC_c^G with large search step \lambda N_c^G, d_c^G, SC_c^G=SHGO(f_{obj}, SP_{global}, S_{cons}, search step size=\lambda)
                    level 3: Fine-grain Search
                   within the neighbourhood of N_c^G, d_c^G, SC_c^G, search the optimum N_c^L, d_c^L, SC_c^L Local search space: SP_{local} = \{N_c^G \pm 2\lambda, d_c^G \pm 2\lambda, SC_c^G \pm 2\lambda\} N_c^L, d_c^L, SC_c^L=SHGO(f_{obj}, SP_{local}, S_{cons}, search step size=1)
           Add \{w_m, N_c^L, d_c^L, SC_c^L\} to CAS
      Compare the candidate architecture in CAS, find the optimum \{w_m^*, N_c^*, d_c^*, SC_c^*\}.
      Return \{w_m^*, N_c^*, d_c^*, SC_c^*\}
end
```

non-linear objective functions and constraints and can solve derivative-free optimization problems¹. Moreover, the convergence of SHGO requires much fewer samples and less time than reinforcement learning approaches [27]. Hence, we use SHGO for our new *hierarchical* searching algorithm.

Specifically, as shown in Algorithm 1, to further accelerate the searching process, we propose a three-level SHGO-based algorithm instead of using the original SHGO algorithm. At the first level, we enumerate w_m in the search space. Usually, the range of w_m is much more narrow than the other architecture parameters; hence without fixing w_m , we cannot use a large search step size for the second-level coarse-grain search. At the second level, we use SHGO with a large search step size λ to search for a coarse optimum N_c^G , d_c^G , SC_c^G by fixing the w_m . At the third level (fine-grain search), we use SHGO with the smallest search step size (i.e., 1) to search for the optimum N_c^L , d_c^L , SC_c^L values for a specific w_m , within the neighborhood of the coarse optimum N_c^G , d_c^G , SC_c^G , and add it to the candidate set. After completing the three-level search, we compare all neural architectures in the candidate set and determine the (final) optimal architecture $\{w_m^*, N_c^*, d_c^*, SC_c^*\}$. To summarize, given the number of hyper-parameters M and the number of possible values of each hyper-parameter N, the complexity of our hierarchical SHGO-based NAS is roughly proportional to MN, i.e., O(MN).

Experimental results in Section 4 show that our proposed hierarchical search accelerates the overall search process without any decrease in the performance of the obtained neural architecture. Moreover, our proposed hierarchical SHGO-based algorithm involves much less computational workload compared to the original (one-level) SHGO-based algorithm and RL-based approaches [27]; this even enables us to do NAS on a real Raspberry Pi-3B processor.

¹The detailed discussion of SHGO is beyond the scope of this paper. More details are available in [18]

1:14 G. Li, et al.

4 Experimental Results

4.1 Experimental setup

Dataset: Existing NAS approaches show that the test accuracy of CNNs on CIFAR-10 dataset can indicate the test accuracy on other datasets, such as ImageNet [16]. Hence, similar to most of the NAS approaches, we use CIFAR-10 as the primary dataset. Moreover, we also evaluate our framework on CIFAR-100 and Tiny-ImageNet² to demonstrate the generality of our proposed metric NN-Degree and accuracy predictor.

Training Hyper-parameters: We train each of the selected neural networks five times with PyTorch and use the mean test accuracy of these five runs as the final results. All networks are trained for 200 epochs with the SGD optimizer and a momentum of 0.9. We set the initial learning rate as 0.1 and use Cosine Annealing algorithm as the learning rate scheduler.

Search Space: DenseNets are more efficient in terms of model size and computation workload than ResNets while achieving the same test accuracy [25]. Moreover, DenseNets have many more skip connections; this provides us with more flexibility for exploration compared to networks with Addition-type skip connections (ResNets, Wide-ResNets, and MobileNets). Hence, in our experiments, we explore the CNNs with DenseNet-type skip connections.

To enlarge the search space, we generate the generalized version of standard DenseNets by randomly selecting channels for concatenation. Specifically, for a given cell c, we define t_c as the maximum skip connections that each layer can have; thus, we use t_c to control the topological properties of CNNs. Given the definition of t_c , layer i can receive DenseNet-type skip connections (DTSC) from a maximum number of t_c channels from previous layers within the same cell; that is, we randomly select $min\{w_c(i-1), t_c\}$ channels from layers 0, 1, ..., (i-2), and concatenate them at layer i-1. The concatenated channels then pass through a convolutional layer to generate the output of layer i (s_i). Similar to recent NAS research [32], we select links randomly because random architectures are often as competitive as the carefully designed ones. If the skip connections encompass all-to-all connections, this would result in the original DenseNet architecture [25]. An important advantage of the above setup is that we can control the number of DTSC (using t_c) to cover a vast search space with a large number of candidate DNNs.

Like standard DenseNets, we can generalize this setup to contain multiple (N_c) cells of width w_c and depth d_c ; DTSC are present *only within a cell* and not across cells. Furthermore, we increase the width (i.e., the number of output channels per layer) by a factor of 2 and halve the height and width of the feature map cell by cell, following the standard practice [47]. After several cells (groups) of convolutions layers, the final feature map is average-pooled and passed through a fully-connected layer to generate the logits. The width of each cell is controlled using a width multiplier, w_m (like in Wide-ResNets [59]). The base number of channels of each cell is [16,32,64]. For $w_m = 3$, cells will have [48,96,192] channels per layer. To summarize, we control the value $\{w_m, N_c, d_c, t_c\}$ to sample candidate architectures from the entire search space.

Fig. 5 illustrates a sample CNN similar to the candidate architectures in our search space (small values of w_c and d_c are used for clarity). This CNN consists of three cells, each containing $d_c = 4$ convolutional layers. The three cells have a width (i.e., the number of channels per layer) of 2, 3, and 4, respectively. We denote the network width as $w_c = [2, 3, 4]$. Finally, the maximum number of channels that can supply skip connections is given by $t_c = [2, 5, 6]$. That is, the first cell can have a maximum of two skip connection candidates per layer (i.e., previous channels that can supply skip connections), the second cell can have a maximum of five skip connections candidates per layer, and so on. Moreover, as mentioned before, we randomly choose $min\{w_c(i-1), t_c\}$ channels for

 $^{^2}$ Tiny-ImageNet is a downscaled-version ImageNet dataset with 64x64 resolution and 200 classes [15]. For more details, please check: http://cs231n.stanford.edu/tiny-imagenet-200.zip

skip connections at each layer. The inset of Fig. 5 shows for a specific layer, how skip connections are created by concatenating the feature maps from previous layers.

In practice, we use three cells for the CIFAR-10 dataset, i.e., $N_c = 3$. We constrain the $1 \le w_m \le 3$ and $5 \le d_c \le 30$. We also constrain t_c of each cell: $5 \le t_1$, $2t_1 \le t_2$ and $2t_2 \le t_3$ for these three cells, respectively. In this way, we can balance the number of skip connections across each cell. Moreover, the maximum number of skip connections that a layer can have is the product of the width of the cell (w_c) and $d_c - 2$ which happens for the last layer in a cell concatenating all of the output channels except the second last layer. Hence, the upper bound of t_c , for each cell, is $16w_m(d_c - 2)$, $32w_m(d_c - 2)$, $64w_m(d_c - 2)$, respectively. Therefore, the size of the overall search space is:

$$\sum_{w_m=1}^{3} \sum_{d_c=5}^{30} \sum_{t_1=5}^{16w_m(d_c-2)} \sum_{t_2=2t_1}^{32w_m(d_c-2)} (64w_m(d_c-2) - 2t_2 + 1) = 6.39 \times 10^{10}$$

Hardware Platform: The training of the sample neural architectures from the search space is conducting on Nvidia GTX-1080Ti GPU. We use Intel Xeon 6230, a 20-core CPU, to simulate the hardware performance of multiple candidate networks and fine-tune the accuracy predictor and analytical hardware models. Finally, we use the same 20-core CPU to conduct the NAS process.

4.2 Accuracy Predictor

We first derive the NN-Degree (g) for the neural architecture in our search space. Based on Equation 2, we substitute SC_c with the real number of skip connections in a cell as follows:

$$g = \sum_{c=1}^{N_c} (w_c + \frac{SC_c}{w_c \times d_c}) = \sum_{c=1}^{N_c} (w_c + \frac{\sum_{i=2}^{d_c - 1} \min\{(i-1)w_c, t_c\}}{d_c})$$
(13)

In Section 3, we argue that the neural architecture with a higher NN-degree value tends to provide a higher test accuracy. In Fig. 6(a), we plot the test accuracy vs. NN-Degree of 60 randomly sampled neural networks from the search space for CIFAR-10 dataset; our proposed network-topology based metric NN-Degree indicates the test accuracy of neural networks. Furthermore, Fig 6(b)

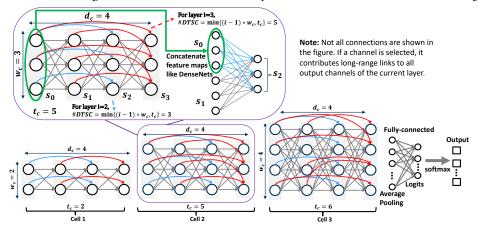


Fig. 5. An example of candidate neural architectures from our search space. (The values of w_c , t_c , and d_c are only for illustration and they do not represent the real search space). Not all skip connections are shown in the figure, for simplicity. The upper inset shows the contribution from all skip and short-range links to layer i=2: The feature maps for the randomly selected channels are concatenated as the input of the current layer i=2 (similar to DenseNets [25]). At each layer in a given cell, the maximum number of channels contributing to skip connections is controlled by t_c .

1:16 G. Li, et al.

Table 5. Our NN-Degree based accuracy predictor for neural architecture search vs. existing predictors implemented by graph-based neural networks. We calculate the improvement ratio for each of the metric by considering the best among all existing approaches in this table. ('-' denotes that the corresponding results are not reported or not applicable.)

Accuracy Estimation	Search S	Search Space (SS) Size		# Training Samples		Training Time (s)	
Technique	Value	% of FLASH SS	Value	Ratio (×)	RMSE (%) Valu		Ratio (×)
	value	% 01 TLA311 33	value	w.r.t FLASH		value	w.r.t FLASH
GNN+MLP [40]	4.2×10^{5}	$6.6 \times 10^{-4} \%$	3.8×10^{5}	15250	-	-	-
GNN [33]	4.2×10^{5}	$6.6 \times 10^{-4} \%$	3.0×10^{5}	11862	0.05	-	-
GCN [9]	1.6×10^{4}	2.5×10^{-5} %	1.0×10^{3}	40	>1.8	-	-
GCN [54]	4.2×10^{5}	$6.6 \times 10^{-4} \%$	1.7×10^{2}	6.88	1.4	25	66
FLASH (NN-Degree +	6.4×10^{10}	100%	2.5×10^{1}	1	0.152	0.38	1
Logistic Regression)	0.4 × 10	100%	2.3 × 10	1	0.132	0.36	1

and Fig 6(c) also show the test accuracy vs. NN-Degree of 20 networks on CIFAR-100 dataset and 27 networks on Tiny-ImageNet randomly sampled from the search space. Clearly, our proposed metric NN-Degree predicts the test accuracy of neural networks on these two datasets as well. Indeed, the results prove that our claim in Section 3 is empirically correct, i.e., networks with higher NN-Degree values have a better test accuracy.

Next, we use our proposed NN-Degree to build the analytical accuracy predictor. We train as few as 25 sample architectures randomly sampled from the entire search space and record their test accuracy and NN-Degree on CIFAR-10, CIFAR-100, and Tiny-ImageNet datasets. Then, we fine-tune our NN-Degree based accuracy predictor described by Equation 7. As shown in Fig. 7(a), Fig 7(b), and Fig 7(c), our accuracy predictor achieves very high performance while using surprisingly few samples with only three parameters on all these datasets.

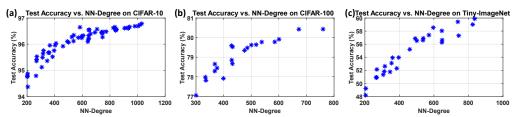


Fig. 6. We randomly select multiple networks from the search space then train and test their accuracy on CIFAR-10, CIFAR-100, Tiny-ImageNet datasets. (a) Real test accuracy vs. NN-Degree: networks with higher NN-Degree values have a higher test accuracy on the CIFAR-10 dataset. (b) Real test accuracy vs. NN-Degree: networks with higher NN-Degree values have a higher test accuracy on the CIFAR-100 dataset. (c) Real test accuracy vs. NN-Degree: networks with higher NN-Degree values have a higher test accuracy on the Tiny-ImageNet dataset.

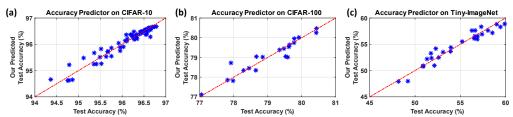


Fig. 7. (a) Predictions of our NN-Degree based accuracy predictor vs. real test accuracy on **CIFAR-10** dataset. (b) Predictions of our NN-Degree based accuracy predictor vs. real test accuracy on **CIFAR-100** dataset. (c) Predictions of our NN-Degree based accuracy predictor vs. real test accuracy on **Tiny-ImageNet** dataset. The red dotted lines in these figures show a very good correlation between the predicted and measured values.

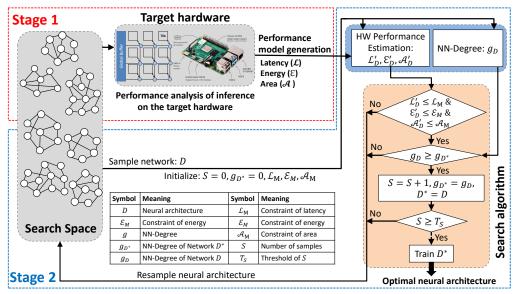


Fig. 8. Overview of the proposed training-free NAS approach. Stage 1 (red box): we build hardware (HW) performance models by randomly sampling candidate networks from the search space to evaluate the hardware characteristics (latency \mathcal{L} , energy \mathcal{E} , and area \mathcal{A}). Stage 2 (blue box): we search for the optimal network architecture with the hardware performance constraints (i.e., \mathcal{L}_M , \mathcal{E}_M , and \mathcal{A}_M); we randomly choose some architectures and use the HW performance models to estimate their hardware performance. Then, we select the neural architecture D^* with the highest NN-Degree which meets the HW performance constraints. Finally, we train the obtained architecture D^* to get the optimal neural architecture.

We also compare our NN-Degree-based accuracy predictor with the current state-of-the-art approaches. As shown in Table 5, most of the existing approaches use Graph-based neural networks to make predictions [9, 33, 40, 54]. However, Graph-based neural networks require much more training data, and they are much more complicated in terms of computation and model structure compared to classical methods like logistic regression. Due to the significant reduction in the model complexity, our predictor requires $6.88\times$ fewer training samples, although a much larger search space $(1.5\times10^5$ larger than the existing work) is covered. Moreover, our NN-Degree based predictor has only three parameters to be updated; hence it consumes $66\times$ less fine-tuning time than the existing approaches. Finally, besides such low model complexity and fast training process, our predictor achieves a very small RMSE (0.152%) as well.

During the search of our NAS methodology, we use the accuracy predictor to directly predict the accuracy of sample architectures as opposed to performing the time-consuming training. The high precision and low complexity of our proposed accuracy predictor also enable us to adopt very fast optimization methods during the search stage. Furthermore, because our proposed metric NN-Degree can predict the test performance of a given architecture, we can use NN-Degree as the proxy of the test accuracy to do NAS without the time-consuming training process. This *training-free* property allows us to quickly compare the accuracy of given architectures and thus accelerate the entire NAS.

4.3 NN-Degree based Training-free NAS

To conduct the training-free NAS, we reformulate the problem described by Equation 1 as follows:

$$\max \theta$$
, subject to: $\mathcal{A} \leq \mathcal{A}_M$, $\mathcal{L} \leq \mathcal{L}_M$, $\mathcal{E} \leq \mathcal{E}_M$ (14)

1:18 G. Li, et al.

Table 6. Our NN-Degree based training-free NAS (**FLASH**) and several representative time-efficient NAS on CIFAR-10 Dataset. We select the optimal architectures with the highest NN-Degree values among 20,000 randomly sampled architectures on a 20-core CPU.

Method	Search Method	#Params	Search Cost	Training needed	Test error (%)
ENAS[42]	RL+weight sharing	4.6M	12 GPU hours	Yes	2.89
SNAS[57]	gradient-based	2.8M	36 GPU hours	Yes	2.85
DARTS-v1[32]	gradient-based	3.3M	1.5 GPU hours	Yes	3.0
DARTS-v2[32]	gradient-based	3.3M	4 GPU hours	Yes	2.76
ProxylessNAS[8]	gradient-based	5.7M	NA	Yes	2.08
Zero-Cost[1]	Proxy-based	NA	NA	Yes	5.78
TE-NAS[11]	Proxy-based	3.8M	1.2 GPU hours	No	2.63
FLASH	NN-Degree based	3.8M	0.11 seconds	No	3.13

To maximize the values of θ , we can search for the network with maximal *NN-Degree* values, which eliminate the training time of candidate architectures. In Fig. 8, we show how we can use the NN-Degree to do training-free NAS. During the first stage, we profile a few networks on the target hardware and fine-tune our hardware performance models. During the second stage, we randomly sample candidate architectures and select those which meet the hardware performance constraints. We use the fine-tuned analytical models to estimate the hardware performance instead of doing real inference, which improves the time efficiency of the entire NAS. After that, we select the optimal architecture with the highest NN-Degree values which meets the hardware performance constraints. We note that the NAS process itself is training-free (hence lightweight), as only the final solution D^* needs to be trained.

To evaluate the performance of our training-free NAS framework, we randomly sample 20,000 candidate architectures from the search space and select the one with the highest NN-Degree values as the obtained/optimal architecture. Specifically, it takes only 0.11 seconds to evaluate these 20,000 samples' NN-Degree on a 20-core CPU to get the optimal architecture (no GPU needed). As shown in Table 6, the optimal architecture among these 20,000 samples achieves a comparable test performance with the representative time-efficient NAS approaches but with much less time cost and computation capacity requirement.

4.4 Analytical hardware performance models

Our experiments show that using 180 samples offers a good balance between the analytical models' accuracy and the number of fine-tuning samples. Hence, we randomly select 180 neural architectures from the search space to build our analytical hardware performance models. Next,

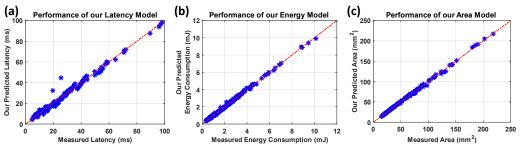


Fig. 9. Performance of our analytical hardware models on **ImageNet** classification networks: (a) Predicted values by our analytical area model vs. measured area. (b) Predicted values by our analytical latency model vs. measured latency. (c) Predicted values by our analytical energy model vs. measured energy consumption. The red lines demonstrate that our proposed models generalize well for networks evaluated on ImageNet-scale datasets.

Table 7. Summary of the performance of our proposed analytical models for Area, Latency, and Energy.

Model	#Features	Mean Error (%)	Max Error (%)	Fine-tuning Time (s)
Area	2	0.1	0.2	0.49
Latency	9	3.0	20.8	0.52
Energy	16	3.7	24.4	0.56

Table 8. Estimation error with different ML models for ImageNet with IMC as target hardware platform.

	SVM	Random Forest (Max. Depth = 16)	Analytical Models (Proposed)
Latency Est. Error (%)	58.98	8.23	6.7
Energy Est. Error (%)	78.49	11.01	3.5
Area Est. Error (%)	36.99	13.37	1.7

we perform the inference of these selected 180 networks on our simulator [29] to obtain their area, latency, and energy consumption. After obtaining the hardware performance of 180 sample networks, we fine-tune the parameters of our proposed analytical area, latency, and energy models discussed in Section 3. To evaluate the performance of these fine-tuned models, we randomly select another 540 sample architectures from the search space then conduct inference and obtain their hardware performance.

Table 7 summarizes the performance of our analytical models. The mean estimation error is always less than 4%. Fig. 9 shows the estimated hardware performance obtained by our analytical model for the ImageNet dataset. We observe that the estimation coincides with the measured values from simulation. Our analytical models enable us to obtain very accurate predictions of hardware performance with the time cost of less than 1 second on a 20-core CPU. The high performance and low computation workload enable us to directly adopt these analytical models to accelerate our searching stage instead of conducting real inference.

Comparison with other machine learning models: Table 8 compares the estimation error for SVM, random forest with a maximum tree depth of 16 and the proposed analytical hardware models for ImageNet dataset. A maximum tree depth of 16 is chosen for random forest since it provides the best accuracy among random forest models. We observe that our proposed analytical hardware models achieve the smallest error among all three modeling techniques. SVM performs poorly since it tries to classify the data with a hyper-plane, and no such plane may exist given the complex relationship between the features and performance of the hardware platform.

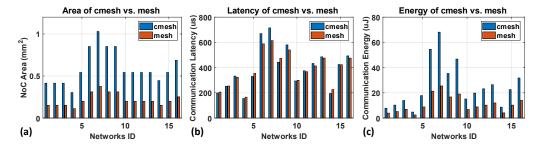


Fig. 10. Performance comparison between our mesh-NoC and cmesh-NoC [46] on **CIFAR-10** classification networks for 16 different networks: (a) Our mesh-NoC needs much less area than the cmesh-NoC; (b) Our mesh-NoC has almost the same latency as the cmesh-NoC; (c) Our mesh-NoC consumes much less energy consumption than the cmesh-NoC.

1:20 G. Li, et al.

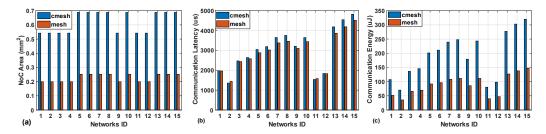


Fig. 11. Performance comparison between our mesh-NoC and cmesh-NoC [46] on **ImageNet** classification networks for 15 different networks: (a) Our mesh-NoC needs much less area than the cmesh-NoC; (b) Our mesh-NoC has almost the same latency as the cmesh-NoC; (c) Our mesh-NoC consumes much less energy consumption than the cmesh-NoC.

4.5 On-chip communication optimization

As shown in Fig. 10 and Fig. 11, we compare the NoC performance (area, energy, and latency) of our FLASH with respect to the cmesh-NoC [46] for 16 randomly selected networks from the search space for CIFAR-10 dataset and ImageNet dataset, respectively. We observe that the mesh-NoC occupies on average only 37% area and consumes only 41% energy with respect to the cmesh-NoC. Since the cmesh-NoC uses extra links and repeaters to connect diagonal routers, the area and energy with the cmesh-NoC are significantly higher than the mesh-NoC. Additional links and routers in the cmesh-NoC result in lower hop counts than the mesh-NoC. However, the lower hop count reduces the latency at low congestion. As the congestion in the NoC increases, the latency of the cmesh-NoC becomes higher than the mesh-NoC due to increased utilization of additional links. This phenomenon is also demonstrated in [19]. Therefore, the communication latency with the cmesh-NoC is higher than the mesh-NoC for most of the DNNs. The communication latency with the mesh-NoC is on average within 3% different from the communication latency with the cmesh-NoC. Moreover, we observe that the average utilization of the queues in the mesh-NoC varies between 20%-40% for the ImageNet dataset. Furthermore, the maximum utilization of the queues ranges from 60% to 80%. Therefore, the mesh-NoC is heavily congested. Thus, our proposed communication optimization strategy outperforms the state-of-the-art approaches.

4.6 Hierarchical SHGO-based neural architecture search

After we fine-tune the NN-Degree based accuracy predictor and analytical hardware performance models, we use our proposed hierarchical SHGO-based searching algorithm to do the neural architecture search.

Baseline approach: Reinforcement Learning (RL) is widely used in NAS [24, 27, 61]; hence we have implemented a RL-based NAS framework as a baseline. For the baseline, we consider the objective function in Equation 1. Specifically, we incorporate a deep-Q network approach for the baseline-RL [38]. We construct four different controllers for the number of cell (N_c), cell depth (d_c), width multiplier (w_m) and number of long skip connections (SC_c). The training hyper-parameters for the baseline-RL are shown in Table 9. The baseline-RL approach estimates the optimal parameters (N_c , d_c , w_m , SC_c). We tune the baseline-RL approach to obtain the best possible results. We also implement a one-level SHGO algorithm (i.e., original SHGO) as another baseline to show the efficiency of our hierarchical algorithm.

We compare the baseline-RL approach with our proposed SHGO-based optimization approach. As shown in Table 10, when there is no constraint in terms of accuracy and hardware performance, our hierarchical SHGO-based algorithm brings negligible overhead compared to the one-level SHGO algorithm. Moreover, our hierarchical SHGO-based algorithm needs much fewer samples

Performance of our Latency I	Model **	(b) Performance of our E	nergy Mod
Optimizer	ADAM	Loss	MSE
Number of neurons in each layer	20	Activation	softmax
Number of layers	3	Learning rate	0.001
Metric	Value	Metric	Value

Table 9. Parameters chosen for the baseline-RL approach.

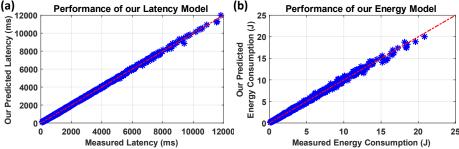


Fig. 12. (a) Predictions of our analytical latency models vs. measured values for RPi-3B. (b) Predictions of our analytical energy consumption models vs. measured values for RPi-3B. The red dotted lines in these two figures show a high correlation between predicted and measured values.

(144.93×) during the search process than RL-based methods. Our proposed search algorithm is as fast as 0.07 seconds and 27929× faster than the RL-based methods, while achieving the same quality of the solution! As for the searching with specific constraints, the training of RL-based methods cannot even converge after training with 10000 samples. Furthermore, our hierarchical SHGO-based algorithm obtains a better-quality model with 7.03× fewer samples and 14.7× less search time compared to the one-level SHGO algorithm. The results show that our proposed hierarchical strategy further improves the efficiency of the original SHGO algorithm.

4.7 Case study: Raspberry Pi and Odroid MC1

As discussed in previous sections, each component and stages of FLASH are very efficient in terms of both computation and time costs. To further demonstrate the efficiency of our FLASH methodology, we implement FLASH on two typical edge devices, namely, the Raspberry Pi-3 Model-B (RPi-3B) and Odroid MC1 (MC1).

Table 10. Comparison between RL-based search, one-level SHGO-based search, and our proposed hierarchical SHGO-based search. No constraint means that we don't set any bounds for the accuracy, area, latency, and energy consumption of the networks; we compare FLASH with RL when there are no constraints. For searching with constraints, we set the minimal accuracy being 95.8% ($\theta \ge \theta_M = 95.8\%$) as an example; we compare FLASH with one-level SHGO because RL does not converge. The quality of the model is calculated by the objective function in Equation 1 (higher is better).

Constraints	Method	Search cost	Search	Quality of obtained	Converge?
involved?		(#Samples)	Time (s)	model (Eq. 1)	
	RL	10000	1955	20984	Yes
No	one-level SHGO	23	0.03	20984	Yes
NO	hierarchical SHGO (FLASH)	69	0.07	20984	Yes
	Improvement	144.93×	27929×	1×	-
	RL	>10000	-	-	No
Yes. $\theta \geq \theta_M$	one-level SHGO	1195	3.82	10550	Yes
	hierarchical SHGO (FLASH)	170	0.26	11969	Yes
	Improvement	7.03×	14.7×	1.13×	-

1:22 G. Li, et al.

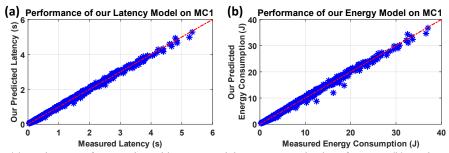


Fig. 13. (a) Predictions of our analytical latency models vs. measured values for MC1. (b) Predictions of our analytical energy consumption models vs. measured values for MC1. The red dotted lines in these two figures show a very good correlation between the predicted and measured values.

Setup: RPi-3B has an Arm Cortex-A53 quad-core processor with a nominal frequency of 1.2GHz and 1GB of RAM. Furthermore, we use the Odroid Smart Power 2 to measure voltage, current, and power. We use TensorFlow-Lite (TF-Lite) as the run-time framework on RPi-3B. To achieve this, we first define the architecture of the models by TensorFlow (TF). Then we convert the TF model into the TF-Lite format and generate the binary file deployed on the RPi-3B.

Odroid MC1 is powered by Exynos 5422, a heterogeneous system-on-a-chip (MPSoC). This SoC consists of two clusters of ARM cores and a small GPU core. Besides the hardware platform itself, we use the same setup as for the RPi-3B.

Accuracy predictor and analytical hardware performance models: We adopt the same accuracy predictor used in Section 4.6. We only consider latency and energy consumption as the hardware performance metrics because the chip area is fixed. Hence, the objective function of searching on RPi-3B and MC1 is:

$$f_{obj} = \frac{Accuracy}{Latency \times Energy} \tag{15}$$

To fine-tune the analytical latency and energy models, we randomly select 180 sample networks from the search space. Then we convert them into the TF-Lite format and record their latency and energy consumption on the RPi-3B. Based on the recorded data, we update the parameters of the analytical latency and energy models. Fig. 12 and 13 show that our analytical hardware performance models almost coincide with the real performance of both the RPi-3B and MC1.

Search Process on RPi-3B and MC1: We do not show the results of RL-based methods because the training of RL models requires intensive computation resources; thus, they cannot be deployed on RPi-3B and MC1. As shown in Table 11, for searching without any constraint, our hierarchical SHGO-based algorithm has only a minimal overhead compared with the basic (one-level) SHGO algorithm. Moreover, our hierarchical SHGO-based algorithm is faster than the one-level SHGO algorithm on MC1.

Table 11. Comparison between one-level and hierarchical SHGO-based search on RPi-3B and Odroid MC1. For searching with constraints, we set the minimal accuracy being 96% ($\theta \ge \theta_M = 96\%$) as an example. The quality of the model is calculated by Equation 15 (higher is better).

Constraints involved?		Search Cost		Search time		Model Quality	
	Method	(# Samples)		(s)		(Equation 15)	
		RPi-3B	MC1	RPi-3B	MC1	RPi-3B	MC1
No	one-level SHGO	112	113	1.68	0.71	4.74	4.13
NO	hierarchical SHGO (FLASH)	180	135	2.21	0.45	4.74	4.13
	one-level SHGO	1309	1272	45.98	9.65	0.35	0.38
Yes, $\theta \geq \theta_M$	hierarchical SHGO (FLASH)	261	414	2.33	1.32	0.48	0.57
	Improvement	5.01 ×	3.07 ×	19.73 ×	20.5 ×	1.37 ×	1.51 ×

For searching with constraints, the hierarchical SHGO-based algorithm obtains a better-quality model with $5.01\times$ fewer samples and $19.73\times$ less search time on the RPi-3B; we achieve similar improvements on MC1 as well. These results prove the effectiveness of our hierarchical strategy again. Overall, the total searching time on RPi-3B and MC1 are as short as 2.33 seconds and 1.32 seconds, respectively on such resource-constrained edge devices. To our best knowledge, this is the first time when a neural architecture search is reported on edge devices.

5 Conclusions and Future Work

This paper presented a very fast methodology, called FLASH, to improve the time efficiency of NAS. To this end, we have proposed a new topology-based metric, namely the *NN-Degree*. Using the NN-Degree, we have proposed an analytical accuracy predictor by training as few as 25 samples out of a vast search space with more than 63 billion configurations. Our proposed accuracy predictor achieves the same performance with 6.88× fewer samples and 65.79× reduction in fine-tuning time cost compared to state-of-the-art approaches. We have also optimized the on-chip communication by designing a mesh-NoC for communication across multiple layers; based on the optimized hardware, we have built new analytical models to predict area, latency, and energy consumption.

Combining the accuracy predictor and the analytical hardware performance models, we have developed a hierarchical simplicial homology global optimization (SHGO)-based algorithm to optimize the co-design process while considering both test accuracy and the area, latency, and energy figures of the target hardware. Finally, we have demonstrated that our newly proposed hierarchical SHGO-based algorithm enables 27729× faster (less than 0.1 seconds) NAS compared to the state-of-the-art RL-based approaches. We have also shown that FLASH can be readily transferred to other hardware platforms by doing NAS on a Raspberry Pi-3B and Odroid MC1 in less than 3 seconds. To our best knowledge, our work is the first to report NAS performed directly and efficiently on edge devices.

We note that there is no fundamental limitation to apply FLASH to other machine learning tasks. However, no IMC-based architectures are widely adopted yet for other machine learning tasks like speech recognition or object segmentation. Therefore, the current work focuses on DNN inference and leaves the extension to other machine learning tasks as future work. Finally, we plan to incorporate more types of networks such as ResNet and MobileNet-v2 as part of our future work.

6 Acknowledgments

This work was supported in part by the US National Science Foundation (NSF) grant CNS-2007284, and in part by Semiconductor Research Corporation (SRC) grants GRC 2939.001 and 3012.001.

References

- [1] Mohamed S Abdelfattah, Abhinav Mehrotra, Łukasz Dudziak, and Nicholas Donald Lane. 2021. Zero-Cost Proxies for Lightweight NAS. In *International Conference on Learning Representations*.
- [2] Bowen Baker, Otkrist Gupta, Nikhil Naik, and Ramesh Raskar. 2016. Designing Neural Network Architectures using Reinforcement Learning. arXiv preprint arXiv:1611.02167 (2016).
- [3] Albert-László Barabási and Eric Bonabeau. 2003. Scale-free Networks. Scientific american 288, 5 (2003), 60–69.
- [4] Hadjer Benmeziane et al. 2021. A Comprehensive Survey on Hardware-Aware Neural Architecture Search. arXiv preprint arXiv:2101.09336 (2021).
- [5] Kartikeya Bhardwaj, Guihong Li, and Radu Marculescu. 2021. How Does Topology Influence Gradient Propagation and Model Performance of Deep Networks With DenseNet-Type Skip Connections? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).*
- [6] Tom B Brown et al. 2020. Language Models are Few-Shot Learners. arXiv preprint arXiv:2005.14165 (2020).
- [7] Han Cai, Chuang Gan, Tianzhe Wang, Zhekai Zhang, and Song Han. 2020. Once-for-All: Train One Network and Specialize it for Efficient Deployment. In *International Conference on Learning Representations*.

1:24 G. Li, et al.

[8] Han Cai, Ligeng Zhu, and Song Han. 2019. ProxylessNAS: Direct Neural Architecture Search on Target Task and Hardware. In *International Conference on Learning Representations*.

- [9] Thomas Chau, Łukasz Dudziak, Mohamed S Abdelfattah, Royson Lee, Hyeji Kim, and Nicholas D Lane. 2020. BRP-NAS: Prediction-based NAS using GCNs. arXiv preprint arXiv:2007.08668 (2020).
- [10] Pai-Yu Chen, Xiaochen Peng, and Shimeng Yu. 2018. NeuroSim: A Circuit-level Macro Model for Benchmarking Neuro-Inspired Architectures in Online Learning. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems 37, 12 (2018), 3067–3080.
- [11] Wuyang Chen, Xinyu Gong, and Zhangyang Wang. 2021. Neural Architecture Search on ImageNet in Four GPU Hours: A Theoretically Inspired Perspective. In International Conference on Learning Representations.
- [12] Wei-Lin Chiang et al. 2019. Cluster-gcn: An Efficient Algorithm for Training Deep and Large Graph Convolutional Networks. In Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. 257–266.
- [13] Matthieu Courbariaux, Itay Hubara, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. 2016. Binarized Neural Networks: Training Deep Neural Networks with Weights and Activations Constrained to+ 1 or-1. arXiv preprint arXiv:1602.02830 (2016).
- [14] Xiaoliang Dai, Hongxu Yin, and Niraj K. Jha. 2020. Grow and Prune Compact, Fast, and Accurate LSTMs. IEEE Trans. Comput. 69, 3 (2020), 441–452.
- [15] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A Large-Scale Hierarchical Image Database. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). 248–255.
- [16] Xuanyi Dong and Yi Yang. 2020. NAS-Bench-201: Extending the Scope of Reproducible Neural Architecture Search. arXiv preprint arXiv:2001.00326 (2020).
- [17] Thomas Elsken, Jan Hendrik Metzen, and Frank Hutter. 2019. Neural architecture search: A survey. The Journal of Machine Learning Research 20, 1 (2019), 1997–2017.
- [18] Stefan C Endres, Carl Sandrock, and Walter W Focke. 2018. A Simplicial Homology Algorithm for Lipschitz Optimisation. Journal of Global Optimization 72, 2 (2018), 181–217.
- [19] Boris Grot and Stephen W Keckler. 2008. Scalable On-Chip Interconnect Topologies. In 2nd Workshop on Chip Multiprocessor Memory Systems and Interconnects.
- [20] Song Han, Huizi Mao, and William J Dally. 2015. Deep Compression: Compressing Deep Neural Networks with Pruning, Trained Quantization and Huffman Coding. arXiv preprint arXiv:1510.00149 (2015).
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition.* 770–778.
- [22] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the Knowledge in a Neural Network. arXiv preprint arXiv:1503.02531 (2015).
- [23] Mark Horowitz. 2014. 1.1 Computing's Energy Problem (and what we can do about it). In *IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*. 10–14.
- [24] Chi-Hung Hsu et al. 2018. Monas: Multi-objective Neural Architecture Search using Reinforcement Learning. arXiv preprint arXiv:1806.10332 (2018).
- [25] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. 2017. Densely Connected Convolutional Networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 4700–4708.
- [26] Nan Jiang et al. 2013. A Detailed and Flexible Cycle-Accurate Network-on-Chip Simulator. In IEEE ISPASS. 86–96.
- [27] Weiwen Jiang et al. 2020. Device-circuit-architecture Co-exploration for Computing-in-memory Neural Accelerators. *IEEE Trans. Comput.* (2020).
- [28] Gokul Krishnan et al. 2020. Interconnect-aware Area and Energy Optimization for In-memory Acceleration of DNNs. *IEEE Design & Test* 37, 6 (2020), 79–87.
- [29] Gokul Krishnan et al. 2021. Interconnect-Centric Benchmarking of In-Memory Acceleration for DNNS. In 2021 China Semiconductor Technology International Conference (CSTIC). IEEE, 1–4.
- [30] Yuhong Li et al. 2020. EDD: Efficient Differentiable DNN Architecture and Implementation Co-Search for Embedded AI Solutions. In *Proceedings of the 57th ACM/EDAC/IEEE Design Automation Conference*. IEEE Press, Article 130, 6 pages.
- [31] Chenxi Liu et al. 2018. Progressive Neural Architecture Search. In Proceedings of the European Conference on Computer Vision (ECCV). 19–34.
- [32] Hanxiao Liu, Karen Simonyan, and Yiming Yang. 2018. Darts: Differentiable Architecture Search. arXiv preprint arXiv:1806.09055 (2018).
- [33] Jovita Lukasik, David Friede, Heiner Stuckenschmidt, and Margret Keuper. 2020. Neural Architecture Performance Prediction Using Graph Neural Networks. arXiv preprint arXiv:2010.10024 (2020).
- [34] Sumit K Mandal et al. 2020. A Latency-Optimized Reconfigurable NoC for In-Memory Acceleration of DNNs. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems* 10, 3 (2020), 362–375.

- [35] Christopher D Manning, Christopher D Manning, and Hinrich Schütze. 1999. Foundations of Statistical Natural Language Processing. MIT press.
- [36] Hongzi Mao, Malte Schwarzkopf, Shaileshh Bojja Venkatakrishnan, Zili Meng, and Mohammad Alizadeh. 2019. Learning Scheduling Algorithms for Data Processing Clusters. In ACM Special Interest Group on Data Communication. 270–288.
- [37] Diana Marculescu, Dimitrios Stamoulis, and Ermao Cai. 2018. Hardware-Aware Machine Learning: Modeling and Optimization. In *Proceedings of the International Conference on Computer-Aided Design (ICCAD '18)*.
- [38] Volodymyr Mnih et al. 2013. Playing Atari with Deep Reinforcement Learning. arXiv preprint arXiv:1312.5602 (2013).
- [39] Mark Newman, Albert-László Barabási, and Duncan J Watts. 2006. The Structure and Dynamics of Networks. Princeton University Press.
- [40] Xuefei Ning, Yin Zheng, Tianchen Zhao, Yu Wang, and Huazhong Yang. 2020. A Generic Graph-based Neural Architecture Encoding Scheme for Predictor-based NAS. (2020).
- [41] Xiaochen Peng et al. 2019. Inference Engine Benchmarking Across Technological Platforms from CMOS to RRAM. In *Proceedings of the International Symposium on Memory Systems*. 471–479.
- [42] Hieu Pham, Melody Guan, Barret Zoph, Quoc Le, and Jeff Dean. 2018. Efficient Neural Architecture Search via Parameters Sharing. In *International Conference on Machine Learning*. PMLR, 4095–4104.
- [43] Ximing Qiao, Xiong Cao, Huanrui Yang, Linghao Song, and Hai Li. 2018. Atomlayer: A Universal Reram-based CNN Accelerator with Atomic Layer Computation. In IEEE/ACM DAC.
- [44] Esteban Real et al. 2017. Large-scale Evolution of Image Classifiers. In International Conference on Machine Learning. PMLR, 2902–2911.
- [45] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. 2018. Mobilenetv2: Inverted Residuals and Linear Bottlenecks. In Proceedings of the IEEE conference on computer vision and pattern recognition. 4510–4520.
- [46] Ali Shafiee et al. 2016. ISAAC: A Convolutional Neural Network Accelerator with in-situ Analog Arithmetic in Crossbars. ACM/IEEE ISCA (2016).
- [47] Karen Simonyan and Andrew Zisserman. 2014. Very Deep Convolutional Networks for Large-scale Image Recognition. arXiv preprint arXiv:1409.1556 (2014).
- [48] Linghao Song, Xuehai Qian, Hai Li, and Yiran Chen. 2017. Pipelayer: A Pipelined Reram-based Accelerator for Deep Learning. In *IEEE HPCA*. 541–552.
- [49] Dimitrios Stamoulis et al. 2019. Single-Path NAS: Designing Hardware-Efficient ConvNets in less than 4 Hours. arXiv preprint arXiv:1904.02877 (2019).
- [50] Mingxing Tan et al. 2019. Mnasnet: Platform-aware neural architecture search for mobile. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2820–2828.
- [51] Andreas Veit, Michael Wilber, and Serge Belongie. 2016. Residual Networks Behave Like Ensembles of Relatively Shallow Networks. arXiv preprint arXiv:1605.06431 (2016).
- [52] Siqi Wang, Anuj Pathania, and Tulika Mitra. 2020. Neural Network Inference on Mobile SoCs. *IEEE Design & Test* 37, 5 (2020), 50–57.
- [53] Duncan J Watts and Steven H Strogatz. 1998. Collective Dynamics of 'Small-World' Networks. Nature 393, 6684 (1998), 440–442.
- [54] Wei Wen, Hanxiao Liu, Yiran Chen, Hai Li, Gabriel Bender, and Pieter-Jan Kindermans. 2020. Neural Predictor for Neural Architecture Search. In *European Conference on Computer Vision*. Springer, 660–676.
- [55] Martin Wistuba, Ambrish Rawat, and Tejaswini Pedapati. 2019. A Survey on Neural Architecture Search. arXiv preprint arXiv:1905.01392 (2019).
- [56] Bichen Wu et al. 2019. Fbnet: Hardware-aware Efficient Convnet Design via Differentiable Neural Architecture Search. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 10734–10742.
- [57] Sirui Xie, Hehui Zheng, Chunxiao Liu, and Liang Lin. 2019. SNAS: stochastic neural architecture search. In *International Conference on Learning Representations*.
- [58] Jiahui Yu et al. 2020. BigNAS: Scaling up Neural Architecture Search with Big Single-Stage Models. In *Computer Vision ECCV 2020*. 702–717.
- [59] Sergey Zagoruyko and Nikos Komodakis. 2016. Wide Residual Networks. arXiv preprint arXiv:1605.07146 (2016).
- [60] Barret Zoph and Quoc V Le. 2016. Neural Architecture Search with Reinforcement Learning. arXiv preprint arXiv:1611.01578 (2016).
- [61] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V Le. 2018. Learning Transferable Architectures for Scalable Image Recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). 8697–8710.