

---

# ORDERING CHAOS: MEMORY-AWARE SCHEDULING OF IRREGULARLY WIRED NEURAL NETWORKS FOR EDGE DEVICES

---

Byung Hoon Ahn<sup>1†</sup> Jinwon Lee<sup>2</sup> Jamie Menjay Lin<sup>2</sup> Hsin-Pai Cheng<sup>3†</sup> Jilei Hou<sup>2</sup> Hadi Esmailzadeh<sup>1</sup>

## ABSTRACT

Recent advance on automating machine learning through Neural Architecture Search and Random Network Generators, has yielded networks that deliver higher accuracy given the same hardware resource constrains, e.g., memory capacity, bandwidth, number of functional units. Many of these emergent networks; however, comprise of irregular wirings (connections) that complicate their execution by deviating from the conventional regular patterns of layer, node connectivity, and computation. The irregularity leads to a new problem space where the schedule and order of nodes significantly affect the activation memory footprint during inference. Concurrently, there is an increasing general demand to deploy neural models onto resource-constrained edge devices due to efficiency, connectivity, and privacy concerns. To enable such a transition from cloud to edge for the irregularly wired neural networks, we set out to devise a compiler optimization that caps and minimizes the footprint to the limitations of the edge device. This optimization is a search for the schedule of the nodes in an intractably large space of possible solutions. We offer and leverage the insight that partial schedules leads to repeated subpaths for search and use the graph properties to generate a signature for these repetition. These signatures enable the use of *Dynamic Programming* as a basis for the optimization algorithm. However, due to the sheer number of neurons and connections, the search space may remain prohibitively large. As such, we devise an *Adaptive Soft Budgeting* technique that during dynamic programming performs a light-weight meta-search to find the appropriate memory budget for pruning suboptimal paths. Nonetheless, schedules from any scheduling algorithm, including ours, is still bound to the topology of the neural graph under compilation. To alleviate this intrinsic restriction, we develop an *Identity Graph Rewriting* scheme that leads to even lower memory footprint without changing the mathematical integrity of the neural network. We evaluate our proposed algorithms and schemes using representative irregularly wired neural networks. Compared to TensorFlow Lite, a widely used framework for edge devices, the proposed framework provides  $1.86\times$  reduction in memory footprint and  $1.76\times$  reduction in off-chip traffic with an average of less than one minute extra compilation time.

## 1 INTRODUCTION

Growing body of work focuses on Automating Machine Learning (AutoML) using Neural Architecture Search (NAS) (Zoph & Le, 2017; Cortes et al., 2017; Zoph et al., 2018; Liu et al., 2019a; Cai et al., 2019; Real et al., 2019; Zhang et al., 2019) and now even, Random Network Generators (Xie et al., 2019; Wortsman et al., 2019) which emit models with *irregular* wirings, and shows that such *irregularly wired neural networks* can significantly enhance classification performance. These networks that deviate from *regular topology* can even adapt to some of the constraints of the hardware (e.g., memory capacity, bandwidth, number of functional units), rendering themselves especially useful in target-

ing edge devices. Therefore, lifting the regularity condition provides significant freedom for NAS and expands the search space (Cortes et al., 2017; Zhang et al., 2019; Xie et al., 2019).

The general objective is to enable deployment of neural intelligence even on stringently constrained devices by trading off regular wiring of neurons for higher resource efficiency. Importantly, pushing neural execution to edge is one way to address the growing concerns about privacy (Miresghallah et al., 2020) and enable their effective use where connectivity to cloud is restricted (Wu et al., 2019). However, the new challenge arises regarding orchestrating execution of these irregularly wired neural networks on the edge devices as working memory footprint during execution frequently surpass the strict cap on the memory capacity of these devices. The lack of multi-level memory hierarchy in these micro devices exacerbates the problem, because the network cannot even be executed if the footprint exceeds the capacity. To that end, despite the significant potential of irregularly wired neural networks, their *complicated execution pattern*,

---

<sup>†</sup>Work done as intern at Qualcomm AI Research. <sup>1</sup>University of California, San Diego <sup>2</sup>Qualcomm AI Research <sup>3</sup>Duke University. Correspondence to: Byung Hoon Ahn <bhahn@eng.ucsd.edu>.

in contrast to previously *streamlined execution of models with regular topology*, renders conventional frameworks futile in taking these networks to edge due to their large *peak memory footprint*. While peak memory footprint is largely dependent on scheduling of neurons, current deep learning compilers (Chen et al., 2018; Vasilache et al., 2018) and frameworks (Abadi et al., 2016; Paszke et al., 2019; Jia et al., 2014) rely on basic topological ordering algorithms that are oblivious to peak memory footprint and instead focus on an orthogonal problem of tiling and kernel level optimization. This paper is an initial step towards embedding peak memory footprint as first-grade constraint in deep learning schedulers to unleash the potential of the emergent irregularly wired neural networks. As such, this paper makes the following contributions:

**(1) Memory-aware scheduling for irregularly wired neural networks.** Scheduling for these networks is a topological ordering problem, which enumerates an intractably large space of possible schedules. We offer and leverage the insight that partial schedules leads to repeated subpaths for search and use the graph properties to generate a signature for these repetition while embedding a notion of the running memory usage. These signatures enable the use of *Dynamic Programming* as a basis for the optimization algorithm.

**(2) Adaptive soft budgeting for tractable compilation time.** Even with the dynamic programming as the base, due to the sheer number of neurons and connections, the search space may remain too large (exponentially large) in practice. As such, we devise an *Adaptive Soft Budgeting* technique that uses a lightweight meta-search mechanism to find the appropriate memory budget for pruning the suboptimal paths. This technique aims to find an inflection point beyond which tighter budgets may lead to no solution and looser budget prolongs the scheduling substantially, putting the optimization in a position of questionable utility.

**(3) Identity graph rewriting for enabling higher potential in memory reduction.** Any scheduling algorithm, including ours, is still bound to the topology of the neural graph under compilation. To relax this intrinsic restriction, we devise an *Identity Graph Rewriting* scheme that exchanges subgraphs leading to a lower memory footprint without altering the mathematical integrity of the neural network.

Results show that our adaptive scheduling algorithm improves peak memory footprint for irregularly wired neural networks by  $1.68\times$  compared to TensorFlow Lite, the de facto framework for edge devices. Our graph rewriting technique provides an opportunity to lower the peak memory footprint by an additional 10.7%. Furthermore, our framework can even bring about  $1.76\times$  reduction in off-chip traffic for devices with multi-level memory hierarchy, and even eliminate the traffic in some cases by confining the memory footprint below the on-chip memory capacity. These gains come at average of less than one minute extra compilation time.

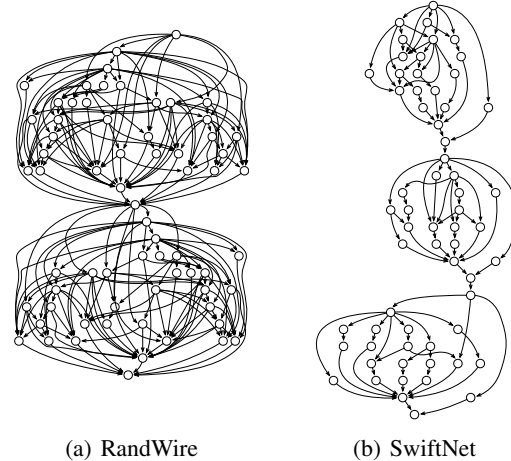


Figure 1. Architecture of network models from NAS and Random Network Generators. Topology of such networks include distinctive *irregular wirings* between the nodes.

## 2 CHALLENGES AND OUR APPROACH

### 2.1 Irregularly Wired Neural Networks

Recent excitement in Automated Machine Learning (AutoML) (Feurer et al., 2015; Dean, 2017; He et al., 2018; Elthakeb et al., 2018; Wang et al., 2019; Laredo et al., 2019) aims to achieve *human out of the loop* in developing machine learning systems. This includes Neural Architecture Search (NAS) (Zoph & Le, 2017; Zoph et al., 2018; Liu et al., 2019a; Cai et al., 2019; Real et al., 2019; Zhang et al., 2019) and Random Network Generators (Xie et al., 2019; Wortsman et al., 2019) that focus on automation of designing neural architectures. Figure 1 demonstrates that networks of this regime are characterized by their distinctive *irregular graph topology* with much more irregular wirings (dataflow) compared to conventional networks with regular graph topology. This paper refers to these networks as *irregularly wired neural networks*.

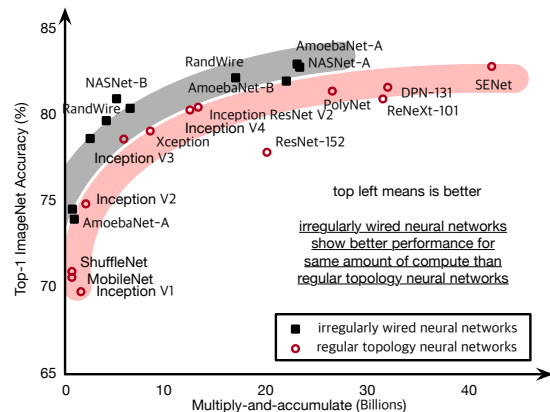


Figure 2. ImageNet accuracy vs number of multiply-and-accumulate, where irregularly wired neural networks show higher performance for same compute than regular topology neural networks. Plot for number of parameters also displays a similar trend.

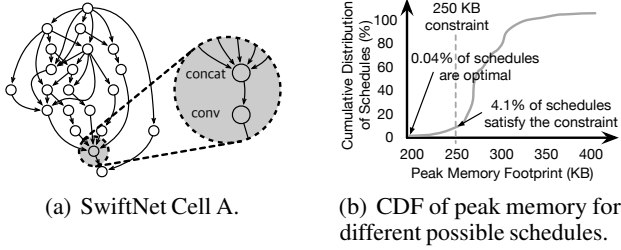


Figure 3. CDF of the peak memory footprint for the different possible schedules of a given *irregularly wired neural network*.

From the performance perspective, these networks have shown to outperform manually designed architectures in terms of accuracy while using less resources. In fact, majority of winning neural architectures in competitions with primary goal of reducing resources (Gauen et al., 2017) rely on NAS, suggesting its effectiveness in that respect. Figure 2 plots the accuracy of different models given their computation. The figure clearly shows that the Pareto frontier of irregularly wired neural networks from NAS and Random Network Generators are better than the hand designed models with regular topology. This indicates that the efficiency in terms of accuracy given fixed resources are better with the irregularly wired neural networks.

## 2.2 Challenges

Many existing compilers (Chen et al., 2018; Vasilache et al., 2018) and frameworks (Paszke et al., 2019; Abadi et al., 2016; Jia et al., 2014) rely on basic topological ordering algorithms to schedule the graph. While the current approach may be sufficient to run conventional networks on server-class machines, such scheme may be unfit for running irregularly wired neural networks on resource-constrained edge devices. This is because, unlike running networks with regular topology, running irregular networks results in varied range of memory footprint depending on the schedule. For instance, given the constraints of a representative edge device (SparkFun Edge: 250KB weight/activation memory and 60M MACs), Figure 3(b) shows that 4.1% of the schedules barely meets the hard memory constraint, while only 0.04% would achieve the optimal peak memory. In reality, such limitation will prevent further exploration regarding the diversity and innovation of network design, and in order to allow edge computing regime to take full advantage of the irregularly wired neural networks, this limitation should be alleviated if not removed.

## 2.3 Design Objectives

**Scheduling algorithm.** To address this issue, our work aims to find a schedule of nodes  $s^*$  from the search space  $\mathcal{S}$  that would minimize peak memory footprint  $\mu_{peak}$ .  $\mathcal{S}$  enumerates all possible orderings of the nodes  $v \in \mathcal{V}$  where  $\mathcal{V}$  is the set of all nodes within a graph  $\mathcal{G}$ .

$$s^* = \underset{s}{\operatorname{argmin}} \mu_{peak}(s, \mathcal{G}), \quad \text{for } s \in \mathcal{S} \quad (1)$$

The most straightforward way to schedule is a *brute force* approach which just enumerates  $\mathcal{S}$  and picks one with the minimum peak memory footprint. While this extreme method may find an optimal solution, it is too costly in terms of time due to its immense complexity:  $\Theta(|V|!)$  where  $|V|$  denotes number of nodes in the graph. One way to improve is to narrow down the search space to just focus on only the *topological orderings*  $\mathcal{S}_T \subset \mathcal{S}$ . However, this will still suffer from a complexity with an upper bound of  $\mathcal{O}(|V|!)$  (takes days to schedule DAG with merely 30 nodes). In fact, previous works (Bruno & Sethi, 1976; Bernstein et al., 1989) already prove optimal scheduling for DAGs is NP-complete. On another extreme are heuristics for topological ordering such as Kahn’s algorithm (Kahn, 1962), with complexity of  $\mathcal{O}(|V| + |E|)$  where  $V$  and  $E$  are number of nodes and edges. However, as demonstrated in Figure 3, such method may yield suboptimal schedule of nodes which will not run on the target hardware. To this end, we explore *dynamic programming* combined with *adaptive soft budgeting* for scheduling to achieve an optimal solution while keeping the graph constant  $s^*$ , without adding too much overhead in terms of time. We explain our algorithms in depth in Section 3.1 and 3.2.

**Graph rewriting.** Any scheduling algorithm including ours is intrinsically bounded by the graph topology. Therefore, we explore to transform the search space through *graph rewriting* (Plump, 1999). Graph rewriting is generally concerned with substituting a certain pattern in the graph with a different pattern to achieve a certain objective. For a computational dataflow graph, leveraging *distributive, associative, and commutative properties* within the computation of the graph, graph rewriting can maintain the semantics while bringing significant improvements regarding some objective. For example, in general programs,  $\sum_i \log x_i$  can be represented as  $\sum_{\text{oddi}} \log x_i + \sum_{\text{eveni}} \log x_i$  or  $\log \prod_i x_i$ , while  $x + x$  can be translated to  $x \times 2$  or  $x \ll 1$ . Likewise, we bring this insight to neural networks to find a set of possible transformations  $\mathcal{X}$  that can rewrite the original graph  $\mathcal{G}$  to a new graph  $\mathcal{G}'$  that would also change our search space  $\mathcal{S}$  to one with a lower peak memory footprint:

$$\mathcal{X}^* = \underset{\mathcal{X}}{\operatorname{argmin}} (\mu_{peak}(s^*, \mathcal{X}(\mathcal{G}))) \quad (2)$$

We identify a set of candidate patterns for transformation  $\chi : g \rightarrow g'$  ( $g \in \mathcal{G}$  and  $g' \in \mathcal{G}'$ ), which constitutes  $\mathcal{X}$ . While transforming the graph, our method keeps the *mathematical integrity* of the graph intact, thus not an approximation method. We embed this systematic way to improve peak memory footprint and the search space as *identity graph rewriting*, and we address this technique in Section 3.3.




Figure 4. Overall workflow of SERENITY, memory-aware scheduling of irregularly wired neural network.

### 3 SERENITY: MEMORY-AWARE SCHEDULING OF IRREGULARLY WIRED NEURAL NETWORKS

As discussed in Section 2, the objective is reducing the peak memory footprint while executing irregularly wired neural networks. We propose SERENITY, memory-aware scheduling that targets devices with restricted resources (e.g., edge devices). Figure 4 summarizes the overall scheduling process, highlighting the major contributions of our approach. Input to SERENITY is a graph of irregularly wired neural network  $\mathcal{G}$ , which in fact acts as an intermediate representation (IR) during the scheduling process. We augment this IR with the metadata of the nodes such as the operation type, input/output edges, input/output shapes, and memory cost. Then the *graph rewriter* transforms the graph  $\mathcal{G} \rightarrow \mathcal{G}'$  to relax the memory costs of memory intensive patterns with the goal of reducing the peak memory footprint  $\mu_{peak}$  of  $\mathcal{G}$ . SERENITY schedules the graph to an optimal schedule  $s^*$  using the *dynamic programming-based scheduler*. However, since the scheduling may be slow due to the complexity, we scale down search space by leveraging *divide-and-conquer* which partitions the graph into multiple subgraphs. Then, we augment the scheduler with an *adaptive soft budgeting* which prunes suboptimal paths by adaptively finding a budget for thresholding through a swift meta-search to speed up the scheduling process. This section focuses on the innovations of SERENITY: dynamic programming-based scheduling, divide-and-conquer, adaptive soft budgeting, and graph rewriting, which are explained in detail in Section 3.1, 3.2, and 3.3, respectively.

#### 3.1 Dynamic Programming-based Scheduling: Achieving Optimal Peak Memory Footprint

Our goal for the scheduling algorithm is to minimize the peak memory footprint  $\mu_{peak}(s, \mathcal{G})$ . As stated in Section 2.3, recursive algorithms that covers the entire search space  $\mathcal{S}$  or the subspace of all topological orderings  $\mathcal{S}_T \subset \mathcal{S}$  takes impractically long time. This is primarily due to the repetitive re-computation of subproblems that upper bounds the algorithm by  $\mathcal{O}(|V|!)$ . Therefore, we leverage *dynamic programming* (Bellman, 1961; 1966; Held & Karp, 1962) which includes a *memoization* scheme that has been shown to be effective in reducing the complexity of time-intensive algorithms by reusing solutions from their subproblems, while still find-

ing optimal solution by sweeping the entire search space.

#### Identifying signature to enable dynamic programming.

The first step to applying dynamic programming to a new problem is characterizing the structure of an optimal solution:  $s^* = s_n^*$  ( $s_n^*$  is an optimal solution for  $n$  number of nodes). Then, it requires identifying a recursive relationship between the optimal solution of a subproblem  $s_i^*$  and the original problem  $s_{i+1}^*$ , and we do this by analyzing the straightforward *recursive topological ordering*, which while inefficient sweeps the entire search space. In essence, topological ordering algorithm is a repeated process of identifying a set of nodes that are available for scheduling and iterating the set for recursion. In graph theory such a set of nodes available for scheduling is called *zero-indegree set*  $z$ , where  $z$  is a set of nodes which all of their incoming edges and the corresponding predecessor nodes (*indegree*) have been scheduled. Figure 5 demonstrates the recursion tree of the different topological ordering algorithms, where the height of the tree is the search step and every path from the root to the leaf is a topological ordering  $s \in \mathcal{S}_T$ . The figure highlights the redundant  $z$  in the recursive topological ordering in the recursion tree, then merges these  $z$  to make them unique, identifying it as the signature for repetition, and prevent the aforementioned re-computation. This makes the scheduling for  $z$  into a unique subproblem, that constitutes

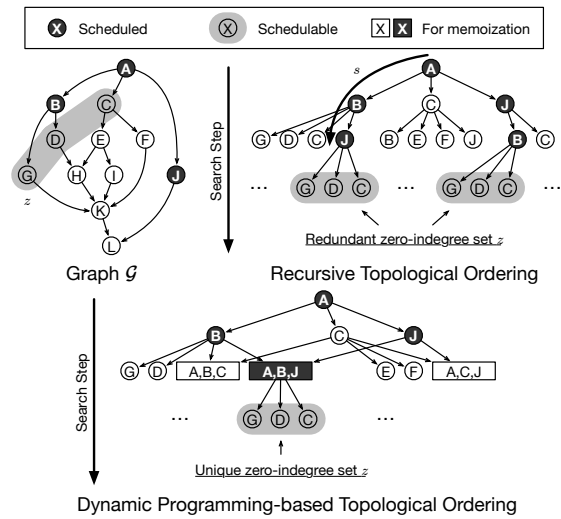


Figure 5. Illustration of identifying redundant zero-indegree set  $z$  and making  $z$  unique (*square*) throughout the topological ordering algorithm to reduce re-computation.

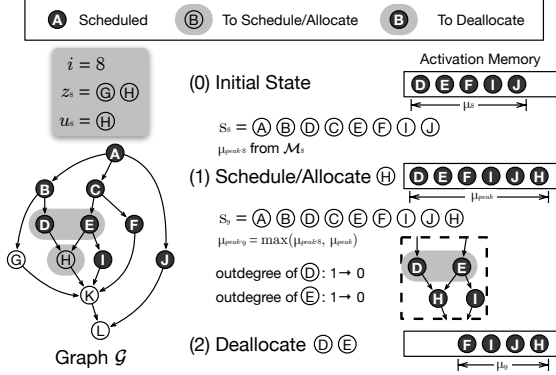


Figure 6. Visualization of scheduling the node  $u_8 = \textcircled{H}$  during the search step  $i = 8$ . Starting from  $s_8$ ,  $\mu_8$ , and  $\mu_{peak,8}$  the figure shows how the algorithm calculates  $s_9$ ,  $\mu_9$ , and  $\mu_{peak,9}$

the *dynamic programming-based topological ordering*.

**Integrating the peak memory footprint constraint.** On top of the dynamic programming formulation that shows potential for optimizing the search space significantly, we overlay the problem specific constraints to achieve the optimal solution. In particular, we calculate the *memory footprint*  $\mu_{i+1}$  and its corresponding *peak*  $\mu_{peak,i+1}$  in each search step  $i$  to select optimal path  $s_{i+1}^*$  for *memoization*. Here, we clarify the process of a search step, explaining the details of calculating  $\mu_{peak,i+1}$  and saving  $s_{i+1}$  for each search step  $i$ . In each search step, we start with number of *unique* zero-indegree sets  $z_i$  (signature), saved in  $i^{th}$  entry of memoization  $\mathcal{M}_i$ . For each  $z_i$ , we append the schedule up to the point  $s_i$ , sum of activations in the memory  $\mu_i$  for the signature  $z_i$ , and the peak memory footprint of the  $s_i$  denoted  $\mu_{peak,i}$ . Therefore, in each search step  $i$ , we start with  $s_i$ ,  $\mu_i$ , and  $\mu_{peak,i}$  for  $s_i$ . Then, when we iterate  $z_i$  to schedule a new node  $u_i$ , its output activation is appended to  $s_i$  to form  $s_{i+1}$ , and is *allocated* in the memory. Size of  $u_i$  is product ( $\prod$ ) of  $u_i.shape$ , where shape is a property of the activation tensor that includes channels, height, width, and the precision (e.g., byte, float), is added to  $\mu_i$ , so  $\mu_{i+1} \leftarrow \mu_i + \prod(u_i.shape)$ . Then we use  $\mu_{i+1}$  as  $\mu_{peak}$  to update  $\mu_{peak,i+1}$  (peak memory footprint for  $s_{i+1}$ ). Since some predecessors of  $u_i$  will not be used anymore after allocating  $u_i$ , we update the *outdegrees* of the node by decrementing them. Having updated the outdegree, we will be left with a *zero-outdegree set* that denotes the nodes that are ready for deallocation. We *deallocate* the nodes in the set and update  $\mu_{i+1}$  accordingly.

To demonstrate scheduling of a node  $u_i$ , Figure 6 simulates scheduling a node  $u_8 = \textcircled{H}$  in  $i = 8$ . In the figure, (1)  $\textcircled{H}$  is appended to  $s_8$  and allocated to memory as it is scheduled, and then the scheduler records maximum of the  $\mu_{peak,8}$  and the sum of all activations in the memory at this point as  $\mu_{peak,9}$ . Then, it recalculates the outdegrees of the predecessor nodes of  $\textcircled{H}$ :  $\textcircled{D}$  and  $\textcircled{E}$ 's outdegree are decremented from one to zero. (2) Then these nodes are deallocated and sum of the

### Algorithm 1 Dynamic Programming-based Scheduling

```

1: Input: graph  $\mathcal{G}$ 
2: Output: optimal schedule  $s^*$ 
3: // initialize memoization
4:  $s_0 \leftarrow []$ ,  $\mu_0, \mu_{peak,0} \leftarrow 0$ ,  $z_0 \leftarrow \text{zero-indegree}(s_0, \mathcal{G})$ 
5:  $\mathcal{M}_0[z_0] \leftarrow (s_0, \mu_0, \mu_{peak,0})$ 
6: // iterate search step
7: for  $i = 0$  to  $n - 1$  do
8:   // iterate (schedule, current memory, peak memory)
9:   for  $z_i, (s_i, \mu_i, \mu_{peak})$  in  $\mathcal{M}_i$  do
10:    for  $u_i$  in  $z_i$  do
11:       $s_{i+1} \leftarrow s_i.append(u_i)$  // allocate
12:       $z_{i+1} \leftarrow \text{zero-indegree}(s_{i+1}, \mathcal{G})$ 
13:       $\mu_{i+1}, \mu_{peak} \leftarrow \mu_i + \prod(u_i.shape)$ 
14:       $\mu_{peak,i+1} \leftarrow \max(\mu_{peak,i}, \mu_{peak})$ 
15:      for  $p_i$  in  $u_i.preds$  do
16:        if  $p_i$  is in  $\text{zero-outdegree}(s_{i+1}, \mathcal{G})$  then
17:           $\mu_{i+1} \leftarrow \mu_{i+1} - \prod(p_i.shape)$  // deallocate
18:        end if
19:      end for
20:      // memoize schedule with least peak memory
21:      if  $\mu_{peak,i+1} \leq \mathcal{M}_{i+1}[z_{i+1}].\mu_{peak,i+1}$  then
22:         $\mathcal{M}_{i+1}[z_{i+1}] \leftarrow (s_{i+1}, \mu_{i+1}, \mu_{peak,i+1})$ 
23:      end if
24:    end for
25:  end for
26: end for
27:  $s^*, \mu_{peak}^* \leftarrow \mathcal{M}[\cdot]_n.s_n, \mathcal{M}[\cdot]_n.\mu_{peak,n}$  // solution
    
```

activation memory here is recorded as  $\mu_9$ .

### Finding schedule with optimal peak memory footprint.

After scheduling  $u_i$ , we save the new signature into the  $\mathcal{M}_{i+1}$  for next search step  $i+1$ . Since the goal of this work is to minimize the overall  $\mu_{peak}$ , we identify the corresponding optimal schedule  $s_{i+1}^*$  for each  $z_{i+1}$  by only saving  $s_{i+1}$  with the minimum  $\mu_{peak,i+1}$ . We integrate the aforementioned step of scheduling  $u_i$  and updating  $\mathcal{M}_{i+1}$  to complete the proposed *dynamic programming-based scheduling algorithm*. Algorithm 1 summarizes the the algorithm. As a first step, the algorithm starts by initializing the memoization table  $\mathcal{M}_0$ , then the algorithm iterates different search steps. In each search step  $i$ , the algorithm performs the above illustrated memory allocation for all  $u_i$  in  $z_i$ , and saving  $s_{i+1}$ ,  $\mu_{i+1}$ , and  $\mu_{peak,i+1}$ . After iterating all search steps to  $n - 1$ ,  $s_*$  is saved in  $\mathcal{M}_n$  with a unique entry, for  $n$  being number of nodes in  $\mathcal{G}$ . We provide the proof for the optimality of the peak memory footprint in the supplementary material.

**Complexity of the algorithm.** The complexity of the proposed dynamic programming-based scheduling is  $\mathcal{O}(|V| \times 2^{|V|})$ , which is significantly faster than the exhaustive search of  $\mathcal{S}_T$  with an upper bound complexity of  $\mathcal{O}(|V|!)$ .




Figure 7. Illustration of *divide-and-conquer*, which divides the graphs into multiple subgraphs (*divide*), schedules each of them using the optimal scheduler (*conquer*), then concatenates the sub-schedules to get the final schedule (*combine*).

Due to the space limitation, we present the derivation of the algorithm complexity in the supplementary material.

### 3.2 Optimizing Scheduling Speed: Speeding up the Dynamic Programming-based Scheduling


While the above scheduling algorithm improves complexity of the search, search space may still be intractable due to the immense irregularity. Therefore, we devise *divide-and-conquer* and *adaptive soft budgeting* to accelerate the search by effectively shrinking and pruning the search space.

**Divide-and-conquer.** We can observe from Figure 1 that the topology of irregularly wired neural networks are *hourglass shaped* ( $\bowtie$ ), because many NAS and Random Network Generators design *cells* with single input and single output then *stack them* to form an hourglass shape topology. (Wilken et al., 2000) shows that, during general purpose code scheduling, graphs can be partitioned (*divide*) into multiple subgraphs and the corresponding solutions (*conquer*) can be concatenated (*combine*) to form an optimal solution for the overall problem. While the complexity of the scheduling algorithm remains the same, this *divide-and-conquer* approach can reduce the number of nodes in each subproblem, speeding up the overall scheduling time. For instance, for a graph that can be partitioned into  $N$  equal subgraphs, the scheduling time will decrease from  $|V| \times 2^{|V|}$  to  $|V| \times 2^{|V|/N}$  that we can speed up scheduling by multiple orders of magnitude compared to the naive approach, depending on the size of the graph and the number of partitions.


As such, Figure 7 shows this insight can be extended to our problem setting, where we can first perform scheduling on each cell and merge those solutions together to form the final solution. First, stage is partitioning the original graph  $G$  into multiple subgraphs  $g$  (*divide*). Then, utilizing the independence among the subgraphs, each subgraph  $g$  can be scheduled separately for their corresponding optimal schedule  $s_g$  (*conquer*). Considering that the number of nodes in the subgraph  $g$  is much smaller than the entire graph  $G$ , the scheduling time will decrease significantly. Finally, the schedules of the subgraphs are concatenated to give optimal schedule  $s^*$  of the entire graph (*combine*).

**Adaptive soft budgeting.** While divide-and-conquer approach scales down the number of nodes, the algorithm may still not be fast enough due to the exponential complexity of the algorithm. Therefore, we explore avoiding suboptimal solutions during the early stage of scheduling without affecting the optimality of the original algorithm. Since our goal is to find a *single* solution that can run within a given *memory budget*  $\tau^* = \mu^*$  while all other solutions can be discarded, setting some budget  $\tau$  that is greater or equal to  $\mu^*$  and pruning suboptimal schedules with which their  $\mu_{peak}$  exceeds  $\tau$  can focus the search to a smaller search space  $S'_T \subset S_T$  while still achieving the optimal schedule  $s^*$ . On top of this, we develop a *meta-search* for  $\tau$ . This is inspired from engineers buying a larger memory (increase  $\tau$ ) if a program fails due to stack overflow (= 'no solution' due to an overly aggressive pruning) and selling out excess memory (decrease  $\tau$ ) if the current budget is prohibitive (= 'timeout' due to lack of pruning). SERENITY takes advantage of this insight to develop an *adaptive soft budgeting* scheme while scheduling to cut down the overall number of explored schedules. Figure 8 illustrates the overall idea by first showing how some schedules are pruned with regard to a given budget  $\tau$  in Figure 8(a) then implication of different  $\tau$  on scheduling time in Figure 8(b).

Figure 8(a) depicts a certain point while scheduling  $G$ , where nodes G, H, F, and J can be scheduled. In particular, the



(a) While both path  $s_1$  and  $s_2$  schedules lead to same  $z'$ , their  $\mu$  and  $\mu_{peak}$  varies and we can prune schedules that yield higher  $\mu_{peak}$  than a given budget  $\tau$ . Numbers next to box or circle are  $\mu$  and numbers next to edges are  $\mu_{peak}$



(b) *Adaptive soft budgeting* starts by setting a hard budget  $\tau_{max}$  as the maximum value for the soft budget  $\tau$ . Then, conducts a binary search for  $\tau$ , higher than  $\tau^*$  that it finds a solution yet not too high that scheduling completes quickly.

Figure 8. Illustration of the *adaptive soft budgeting*. (a) shows how schedules are pruned, and (b) illustrates how the *soft budget*  $\tau$  relates to the number of explored schedules.

figure compares two possible solutions  $s_1$  and  $s_2$  which schedules  $\textcircled{H} \rightarrow \textcircled{F}$  and  $\textcircled{F} \rightarrow \textcircled{H}$ , respectively given  $\tau = 36$ . While  $s_1$  and  $s_2$  both starts from  $z$  with  $\mu = 32$ , scheduling  $\textcircled{H}$  leads to  $\mu_{peak} = 32 + 3 (\text{H}) = 35$ , whereas scheduling  $\textcircled{F}$  or  $\textcircled{J}$  leads to  $\mu_{peak} = 32 + 6 (\text{F or J}) = 38$ . Therefore, since we assume  $\tau = 36$ ,  $s_2$  and  $s_3$  will fail because  $\mu_{peak} = 38$  for  $s_2$  and  $s_3$  exceeds 36. So, as long as we set the budget  $\tau$  higher than  $\mu^*$ , the scheduler still finds a single optimal solution while avoiding many suboptimal paths. On the other hand, too small a  $\tau < \mu^*$  leads to no solution because the optimal path would be pruned away.

Having established the possibility of pruning, our question boils down to discovering  $\tau$  that is greater or equal to  $\mu^*$  which we call an *optimal budget*  $\tau^*$ , yet close enough to shrink the search space effectively. Figure 8(b) and Algorithm 2 summarizes the proposed *adaptive soft budgeting*. Since we start with no information about the approximate range for  $\tau$ , we resort to a commonly used topological ordering algorithm called Kahn’s algorithm (Kahn, 1962) ( $\mathcal{O}(|V| + |E|)$ ) to adaptively gain idea of the range for  $\tau$ . We use the peak memory footprint from this sequence and use it as our *hard budget*  $\tau_{max}$ , and in contrast we call adaptively changing  $\tau$  as a *soft budget*. Since  $\tau_{max} \geq \mu^*$ , we know that any  $\tau \geq \tau_{max}$  do not need to be explored. Having this upper bound for the search, *adaptive soft budgeting* implements a binary search to first run the scheduling algorithm with  $\tau$  and  $T$  as input, where  $T$  is an hyperparameter that limits the scheduling time per search step. The binary search increases  $\tau$  ( $\tau_{new} \leftarrow (\tau_{new} + \tau_{old})/2$ ) if it finds ‘no solution’ and decreases  $\tau$  ( $\tau_{new} \leftarrow \tau_{new}/2$ ) if a search step returns ‘timeout’ (search step duration exceeds  $T$ ). The binary search stops as

---

#### Algorithm 2 Adaptive Soft Budgeting

---

```

1: Input: graph  $\mathcal{G}$ 
2: Output: optimal schedule  $s^*$ 
3:  $\tau_{max} \leftarrow \mu(\text{Kahn'sAlgorithm}(\mathcal{G}), \mathcal{G})$  // hard budget
4:  $\tau_{old}, \tau_{new} \leftarrow \tau_{max}$ 
5:  $flag \leftarrow$  ‘no solution’
6: repeat
7:   // binary search for  $\tau$ : decrease  $\tau$  if ‘timeout’
8:   // and increase  $\tau$  if ‘no solution’
9:   if  $flag$  is ‘timeout’ then
10:    // simultaneous
11:     $\tau_{old} \leftarrow \tau_{new}, \tau_{new} \leftarrow \tau_{new}/2$ 
12:   else if  $flag$  is ‘no solution’ then
13:    // simultaneous
14:     $\tau_{old} \leftarrow \tau_{new}, \tau_{new} \leftarrow (\tau_{new} + \tau_{old})/2$ 
15:   end if
16:   if  $flag$  is ‘solution’ then
17:     $s^* \leftarrow$  schedule // optimal schedule
18:   end if
19: until  $flag$  is ‘solution’
    
```

---

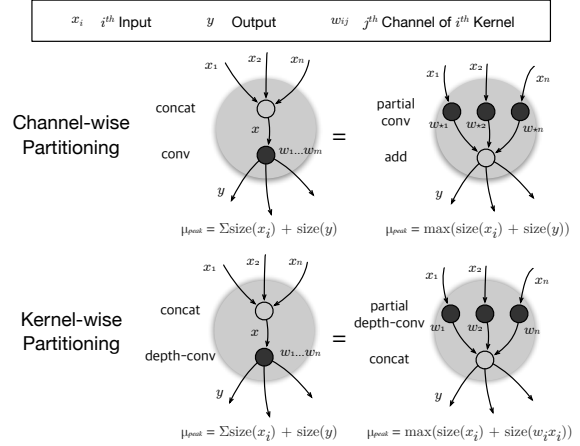


Figure 9. Illustration of the graph rewriting patterns: *channel-wise partitioning* and *kernel-wise partitioning* can reduce the memory cost of convolution and depthwise convolution respectively.

soon as it finds a schedule (‘solution’), and this method using binary search is guaranteed to work due to the monotonically increasing number of explored schedules with  $\tau$ .

### 3.3 Identity Graph Rewriting: Improving the Search Space for Better Peak Memory Footprint

Reorganizing the computational graph of the irregularly wired neural networks may lead to *significant reduction in the peak memory footprint*  $\mu_{peak}$  during computation. For example, it is notable that large stream of NAS-based works (Liu et al., 2019a; Zhang et al., 2019) rely on extensive use of *concatenation* as a natural approach to merge information from multiple branches of the input activations and expand the search space of the neural architectures. However, concatenation with many incoming edges may prolong the liveness of the input activation and increase the memory pressure, which is unfavorable especially for resource constrained scenarios. To address this issue, we propose *identity graph rewriting* to effectively reduce  $\mu_{peak}$  around the concatenation while keeping the arithmetic outputs identical. To this end, we present two main examples of the graph patterns in irregularly wired neural networks that benefits from our technique:

**Channel-wise partitioning (convolution).** One typical pattern in irregularly wired neural networks is *concatenation* ( $concat: [\cdot]$ ) that takes multiple branches of the input prior to a *convolution* ( $conv: *$ ). While executing such pattern, peak memory footprint  $\mu_{peak}$  occurs when the output  $y \in \mathbb{R}^n$  is being computed while concatenated branches of input  $x \in \mathbb{R}^n$  are also mandated to reside in the memory. Our objective is to achieve the same arithmetic results and logical effect as *concat* yet sidestep the corresponding seemingly excessive memory cost. To this end, we *channel-wise partition* the *conv* that follows the *concat* so that the *partitioned conv* can be computed as soon as the input  $x_i$  becomes available.

Equation 3-6 detail the mathematical derivation of this substitution. Specifically, as shown in Equation 3, each kernel iterates and sums up the result of convolving channels in conv. However, using the *distributive property* of  $\sum_i$  and  $*$ , these transform to summation of *channel-wise partitioned convolution*, which we call *partial conv*. This *partial conv* removes *concat* from the graph leading to lower *memory cost*. As illustrated in Figure 9, the memory cost of same computation reduces from  $\sum x_i + y$  to  $\max(w_{*i} * x_i) + y$ , which becomes more effective when there are more incoming edges to concat.

$$y = \left[ \sum_i w_{1i} * x_i, \dots, \sum_i w_{mi} * x_i \right] \text{ (concat+conv)} \quad (3)$$

$$= \sum_i \left[ w_{1i} * x_i, \dots, w_{mi} * x_i \right] \quad (4)$$

$$= \sum_i \left[ w_{1i}, \dots, w_{mi} \right] * x_i \quad (5)$$

$$= \sum_i \left[ w_{*i} * x_i \right] \text{ (partial conv+add)} \quad (6)$$

#### Kernel-wise partitioning (depthwise convolution).

*Depthwise convolution (depthconv)* (Sifre & Mallat, 2014; Howard et al., 2017) has been shown to be effective in reducing computation yet achieve competitive performance, hence its wide use in networks that target extreme efficiency as its primary goal. For *concatenation (concat)* followed by a *depthwise convolution (depthconv)*, similar to above *concat+conv* case, peak memory footprint  $\mu_{peak}$  occurs when the concatenated  $x$  is inside the memory and the result  $y$  additionally gets saved to the memory before  $x$  is deallocated. This time, we leverage the *independence* among different kernels to *kernel-wise partition* the *depthconv* that follows the *concat* so that each input  $x_i$  is computed to smaller feature maps without residing in the memory too long. As such, Equation 7-8 derives this substitution. Equation 7 shows that every component in the  $y$  is *independent* (different subscript index) and is viable for partitioning. In other words, this rewriting simply exposes the *commutative property* between *depthconv* and *concat* plus *kernel-wise partitioning* to reduce  $\mu_{peak}$  significantly.

$$y = \left[ w_1 * x_1, \dots, w_n * x_n \right] \text{ (concat+depthconv)} \quad (7)$$

$$= \left[ [w_1 * x_1], \dots, [w_n * x_n] \right] \text{ (partial depthconv+concat)} \quad (8)$$

**Implementation.** Following the general practice of using pattern matching algorithms in compilers (Lattner & Adve, 2004; Rotem et al., 2018; Jia et al., 2019), we implement *identity graph rewriting* using pattern matching to identify regions of the graph which can be substituted to an operation with *lower computational cost*. Likewise, we make use of this technique to identify regions that leads to *lower memory cost*.

Table 1. Specification of the networks used for evaluation.

NETWORK	TYPE	DATASET	# MAC	# WEIGHT	TOP-1 ACCURACY
DARTS	NAS	IMAGENET	574.0M	4.7M	73.3%
SWIFNET	NAS	HPD	57.4M	249.7K	95.1%
RANDWIRE	RAND	CIFAR10	111.0M	1.2M	93.6%
RANDWIRE	RAND	CIFAR100	160.0M	4.7M	74.5%

## 4 EVALUATION

We evaluate SERENITY with four representative *irregularly wired neural networks* graphs. We first compare the peak memory footprint of SERENITY against TensorFlow Lite (Google) while using the same linear memory allocation scheme<sup>1</sup> for both. Furthermore, we also experiment the impact of such peak memory footprint reduction on off-chip memory communication. We also conduct an in-depth analysis of the gains from the proposed *dynamic programming-based scheduler* and *graph rewriting* using SwiftNet Cell A (Zhang et al., 2019). Lastly, we study the impact of *adaptive soft budgeting* on the scheduling time.

### 4.1 Methodology

**Benchmarks and datasets.** Table 1 lists the details of the networks—representative of the irregularly wired neural networks from Neural Architecture Search (NAS) and Random Network Generators (RAND)—used for evaluation: DARTS (Liu et al., 2019a) for ImageNet, SwiftNet (Zhang et al., 2019) for a dataset comprised of human presence or absence (HPD), and RandWire (Xie et al., 2019) for CIFAR10 and CIFAR100. DARTS (Liu et al., 2019a) is a gradient-based NAS algorithm. In particular we focus on the learned normal cell for image classification on ImageNet: only the first cell because it has the highest peak memory footprint and the reset of the network is just repeated stacking of the same cell following the practice in NASNet (Zoph et al., 2018). SwiftNet (Zhang et al., 2019) is network from NAS by targeting human detection dataset. RandWire (Xie et al., 2019) are from Random Network Generators for image classification on CIFAR10 and CIFAR100. The table also lists their dataset, multiply-accumulate count (# MAC), number of parameters (# WEIGHT), and top-1 accuracy on their respective dataset.

### 4.2 Experimental Results

**Comparison with TensorFlow Lite.** Figure 10 evaluates SERENITY over TensorFlow Lite on different cells of the aforementioned networks in terms of reduction in memory footprint. The figures illustrate that SERENITY’s dynamic programming-based scheduler reduces the memory footprint by a factor of  $1.68\times$  without any changes to the graph. In

<sup>1</sup>TensorFlow Lite implements a linear memory allocator named simple memory arena: [https://github.com/tensorflow/tensorflow/blob/master/tensorflow/lite/s/imple\\_memory\\_arena.cc](https://github.com/tensorflow/tensorflow/blob/master/tensorflow/lite/s/imple_memory_arena.cc)



## Ordering Chaos: Memory-Aware Scheduling of Irregularly Wired Neural Networks for Edge Devices




Figure 10. Reduction in peak memory footprint of SERENITY against TensorFlow Lite (no memory hierarchy).





Figure 11. Reduction in off-chip memory communication of SERENITY against TensorFlow Lite (with memory hierarchy).


addition, the proposed graph rewriting technique yields an average of  $1.86\times$  (extra 10.7%) reduction in terms of peak memory footprint. The results suggest that SERENITY yields significant reduction in terms of the peak memory footprint for irregularly wired neural networks.

**Improvement in off-chip memory communication.** We also show how SERENITY affects the off-chip memory communication, which largely affects both power and inference speed (Chen et al., 2016; Gao et al., 2017; Sharma et al., 2018). To this end, Figure 11 sweeps different on-chip memory configurations to measure the reduction in off-chip communication on systems with multi-level memory hierarchy. Since we know the entire schedule *a priori*, we use Belady’s optimal algorithm (Belady, 1966), also referred to as the clairvoyant algorithm for measuring the off-chip memory communication, to distill the effects of the proposed scheduling. The results show that SERENITY can reduce the off-chip memory communication by  $1.76\times$  for a device with 256KB on-chip memory. In particular, while there were few cases where peak memory footprint was already small enough to fit on-chip (N/A in figure), there were some cases where SERENITY eradicated the off-chip communication by successfully containing the activations in the on-chip memory while TensorFlow Lite failed to do so (marked in figure). This suggests that SERENITY’s effort of reducing memory footprint is also effective in reducing the off-chip memory communication in systems with memory hierarchy, hence the power consumption and inference speed.

**Improvement from dynamic programming-based scheduler and identity graph rewriting.** To demonstrate where the improvement comes from, Figure 12 plots the memory footprint while running Swiftnet Cell A. Figure 12(a) shows the memory footprint of SERENITY with the memory



(a) Memory footprint with the memory allocator (peak memory footprint of TensorFlow Lite = 551.0KB).



(b) Memory footprint without the memory allocator.

Figure 12. Memory footprint while running SwiftNet Cell A with and without the memory allocator (red arrow denotes reduction).

allocation. The figure shows that SERENITY’s dynamic programming-based scheduler brings significant improvement to the peak memory footprint (551.0KB $\rightarrow$ 250.9KB), and the graph rewriting further improves this by 25.1KB (250.9KB $\rightarrow$ 225.8KB) by utilizing patterns that alleviate regions with large memory footprint. In order to focus on the effect of the scheduler and graph rewriting, Figure 12(b) presents the memory footprint of SERENITY without the memory allocation: the sum of the activations while running the network. The figure shows that the proposed scheduler finds a schedule with the optimal (minimum) peak memory footprint without changes to the graph. Then, it shows that the proposed graph rewriting can further reduce the peak memory footprint by 12.5KB (200.7KB $\rightarrow$ 188.2KB). The results suggest that the significant portion of the improvement comes from the proposed dynamic programming-based scheduler and the graph rewriting.

**Scheduling time of SERENITY.** Figure 13 summarizes the (static) scheduling time taken for SERENITY to schedule the networks. Results show that the average scheduling time is 40.6 secs without the graph rewriting and 48.8 secs with graph rewriting, which the difference comes from the increase in the number of nodes from graph rewriting. The results show that all the above gains of SERENITY come at the cost of less than one minute average extra compilation time. While the dynamic programming-based scheduling suffers from an exponential time complexity, SERENITY manages to make the scheduling tractable through the proposed divide-and-conquer and adaptive soft budgeting.

**Speed up from divide-and-conquer and adaptive soft budgeting.** Table 2 summarizes the scheduling time of SwiftNet (Zhang et al., 2019) for different algorithms to demonstrate the speed up from divide-and-conquer and adaptive soft budgeting techniques. As such, the table lists different combination of algorithms, number of




Figure 13. Scheduling time evaluation for SERENITY.

Table 2. Comparison of the scheduling time for different algorithms to schedule SwiftNet. ①, ②, and ③ represent *dynamic programming*, *divide-and-conquer*, and *adaptive soft budgeting* respectively. N/A denotes infeasible within practical time.

GRAPH REWRITING	ALGORITHM	# NODES AND PARTITIONS	SCHEDULING TIME
×	①	62={62}	N/A
×	①+②	62={21,19,22}	56.5 secs
×	①+②+③	62={21,19,22}	<b>37.9 secs</b>
✓	①	92={92}	N/A
✓	①+②	92={33,28,29}	7.2 hours
✓	①+②+③	92={33,28,29}	<b>111.9 secs</b>

nodes, and the corresponding scheduling time. Straightforward implementation of the aforementioned ① *dynamic programming-based scheduling* leads to an immeasurably large scheduling time regardless of the *graph rewriting*. However, additional application of the ② *divide-and-conquer* (①+②) leads to a measurable scheduling time: 56.53 secs and 7.29 hours to schedule without and with the graph rewriting, respectively. Furthermore, we observe that further applying ③ *adaptive soft budgeting* (①+②+③) significantly reduces the scheduling time 37.9 secs and 111.9 secs to schedule without and with the graph rewriting, respectively. Above results indicate that applying the proposed algorithms leads to a scheduling time of practical utility.

## 5 RELATED WORKS

The prevalence of neural networks has led to the development of several compilation frameworks for deep learning (Abadi et al., 2016; Paszke et al., 2019; Rotem et al., 2018; Cyphers et al., 2018). However, even industry grade tools, mostly focus on tiling and fine-grained scheduling of micro-operations on the conventional hardware (NVIDIA, 2017; Google) or accelerators (Chen et al., 2016; 2014; Han et al., 2016a; Judd et al., 2016; Jouppi et al., 2017; Gao et al., 2017; Parashar et al., 2017; Sharma et al., 2018; Fowers et al., 2018). However, these framework are mostly designed for the common regular patterns that have dominated deep learning from almost its conception. As such, these tools inherently had no incentive to deal with the form of irregularities that the emerging NAS (Zoph & Le, 2017; Cortes et al., 2017; Zoph et al., 2018; Liu et al., 2019a; Cai et al., 2019; Real et al., 2019; Zhang et al., 2019) and Random Networks (Xie et al., 2019; Wortsman et al., 2019)

bring about. This paper, in contrast, focuses on this emergent class that breaks the regularity convention and aims to enable their execution on memory constrained edge devices.

**Scheduling and tiling for neural networks.** While prior works on scheduling (Lee et al., 2003; KeBler & Bednarski, 2001; Wilken et al., 2000) focus on classical computing workloads, there have been limited study about the implications of scheduling in the neural networks domain. There is also a significant body of work on scheduling operations on hardware accelerators (Abdelfattah et al., 2018) that also considers tiling (Chen et al., 2018; Vasilache et al., 2018; Liu et al., 2019b; Ahn et al., 2020). However, graph scheduling for irregularly wired neural network, specially with memory constraints, is an emerging problem, which is the focus of this paper.

**Graph rewriting for neural networks.** It has been a common practice to rewrite parts of the graph using rule-based (Abadi et al., 2016; Paszke et al., 2019; Rotem et al., 2018; Cyphers et al., 2018; NVIDIA, 2017) or systematic approaches to expose parallelism and make models more target-aware (Jia et al., 2018; 2019; Schösser & Geiß, 2007). While these approaches may alleviate the complexity of the graph and reduce the peak memory footprint as a side effect, these frameworks do not explore and are not concerned with scheduling. Our work exclusively explores graph rewriting in the context of improving the peak memory footprint.

**Optimizing neural networks.** There are different optimization techniques that aim to simplify the neural network indifferent dimensions. Sparsification/compression (LeCun et al., 1990; Han et al., 2015; Zhu & Gupta, 2018; Anwar et al., 2017), quantization (Han et al., 2016b; Courbariaux et al., 2016; Zhou et al., 2016; Mishra & Marr, 2018; Esser et al., 2020), activation compression (Jain et al., 2018), and kernel modifications reduce the complexity of the individual operations or remove certain computations. However, our focus, the problem of memory-aware graph scheduling still remains orthogonal to these inspiring efforts.

## 6 CONCLUSION

As the new forms of connectivity emerges in neural networks, there is a need for system support to enable their effective use, specially for intelligence at the edge. This paper took an initial step toward orchestrating such network under stringent physical memory capacity constraints. We devised signatures to enable dynamic programming and adaptive soft budgeting to make the optimization tractable. Even more, an identity graph writing was developed to further the potential for gains. The encouraging results for a set of emergent networks suggest that there is significant potential for compiler techniques that enables new forms of intelligent workloads.

## ACKNOWLEDGEMENT

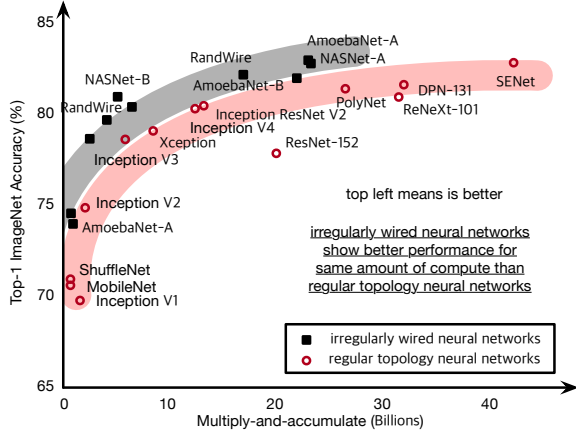
We thank the anonymous reviewers for their insightful comments. We also thank Harris Teague and Jangho Kim for the fruitful discussions and feedbacks on the manuscript, and Parham Noorzad for his help with the mathematical formulations to calculate the complexity of the algorithms.

## REFERENCES

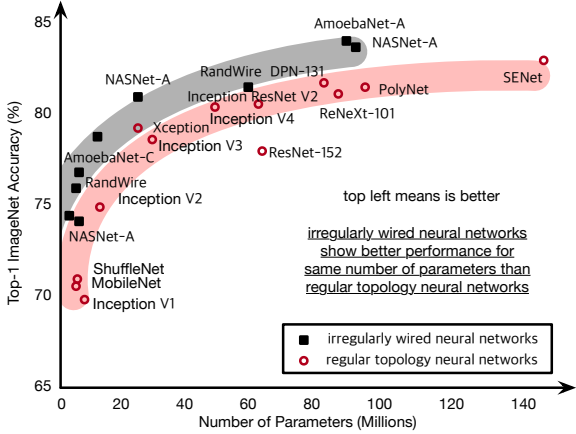
- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., et al. Tensorflow: A system for large-scale machine learning. In *OSDI*, 2016.
- Abdelfattah, M. S., Han, D., Bitar, A., DiCecco, R., O’Connell, S., Shanker, N., Chu, J., Prins, I., Fender, J., Ling, A. C., et al. DLA: Compiler and FPGA overlay for neural network inference acceleration. In *FPL*, 2018.
- Ahn, B. H., Pilligundla, P., and Esmailzadeh, H. Chameleon: Adaptive code optimization for expedited deep neural network compilation. In *ICLR*, 2020. URL <https://openreview.net/forum?id=rygG4AVFvH>.
- Anwar, S., Hwang, K., and Sung, W. Structured pruning of deep convolutional neural networks. *JETC*, 2017.
- Belady, L. A. A study of replacement algorithms for a virtual-storage computer. *IBM Systems Journal*, 1966.
- Bellman, R. Dynamic programming. *Science*, 1966.
- Bellman, R. E. Dynamic programming treatment of the traveling salesman problem. 1961.
- Bernstein, D., Rodeh, M., and Gertner, I. On the complexity of scheduling problems for parallel/pipelined machines. *TC*, 1989.
- Bruno, J. and Sethi, R. Code generation for a one-register machine. *JACM*, 1976.
- Cai, H., Zhu, L., and Han, S. ProxylessNAS: Direct neural architecture search on target task and hardware. In *ICLR*, 2019. URL <https://openreview.net/forum?id=Hy1VB3AqYm>.
- Chen, T., Moreau, T., Jiang, Z., Zheng, L., Yan, E., Shen, H., Cowan, M., Wang, L., Hu, Y., Ceze, L., et al. Tvm: An automated end-to-end optimizing compiler for deep learning. In *OSDI*, 2018.
- Chen, Y., Luo, T., Liu, S., Zhang, S., He, L., Wang, J., Li, L., Chen, T., Xu, Z., Sun, N., et al. Dadiannao: A machine-learning supercomputer. In *MICRO*, 2014.
- Chen, Y.-H., Krishna, T., Emer, J. S., and Sze, V. Eyeriss: An energy-efficient reconfigurable accelerator for deep convolutional neural networks. *JSSC*, 2016.
- Cortes, C., Gonzalvo, X., Kuznetsov, V., Mohri, M., and Yang, S. AdaNet: Adaptive structural learning of artificial neural networks. In *ICML*, 2017.
- Courbariaux, M., Hubara, I., Soudry, D., El-Yaniv, R., and Bengio, Y. Binarized neural networks: Training deep neural networks with weights and activations constrained to +1 or -1. *arXiv*, 2016. URL <https://arxiv.org/pdf/1602.02830.pdf>.
- Cyphers, S., Bansal, A. K., Bhiwandiwalla, A., Bobba, J., Brookhart, M., Chakraborty, A., Constable, W., Convey, C., Cook, L., Kanawi, O., et al. Intel nGraph: An intermediate representation, compiler, and executor for deep learning. *arXiv*, 2018. URL <https://arxiv.org/pdf/1801.08058.pdf>.
- Dean, J. Machine learning for systems and systems for machine learning. In *NIPS Workshop on ML Systems*, 2017.
- Elthakeb, A. T., Pilligundla, P., Yazdanbakhsh, A., Kinzer, S., and Esmailzadeh, H. Releq: A reinforcement learning approach for deep quantization of neural networks. *arXiv*, 2018. URL <https://arxiv.org/pdf/1811.01704.pdf>.
- Esser, S. K., McKinstry, J. L., Bablani, D., Appuswamy, R., and Modha, D. S. Learned step size quantization. In *ICLR*, 2020. URL <https://openreview.net/forum?id=rkg066VKDS>.
- Feurer, M., Klein, A., Eggenberger, K., Springenberg, J., Blum, M., and Hutter, F. Efficient and robust automated machine learning. In *NIPS*, 2015.
- Fowers, J., Ovtcharov, K., Papamichael, M., Massengill, T., Liu, M., Lo, D., Alkalay, S., Haselman, M., Adams, L., Ghandi, M., et al. A configurable cloud-scale dnn processor for real-time ai. In *ISCA*, 2018.
- Gao, M., Pu, J., Yang, X., Horowitz, M., and Kozyrakis, C. TETRIS: Scalable and efficient neural network acceleration with 3d memory. In *ASPLOS*, 2017.
- Gauen, K., Rangan, R., Mohan, A., Lu, Y.-H., Liu, W., and Berg, A. C. Low-power image recognition challenge. In *ASP-DAC*, 2017.
- Google. TensorFlow Lite. URL <https://www.tensorflow.org/mobile/tflite>.
- Han, S., Pool, J., Tran, J., and Dally, W. Learning both weights and connections for efficient neural network. In *NIPS*, 2015.
- Han, S., Liu, X., Mao, H., Pu, J., Pedram, A., Horowitz, M. A., and Dally, W. J. EIE: efficient inference engine on compressed deep neural network. In *ISCA*, 2016a.
- Han, S., Mao, H., and Dally, W. J. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. In *ICLR*, 2016b.
- He, Y., Lin, J., Liu, Z., Wang, H., Li, L.-J., and Han, S. AMC: AutoML for model compression and acceleration on mobile devices. In *ECCV*, 2018.
- Held, M. and Karp, R. M. A dynamic programming approach to sequencing problems. *Journal of the SIAM*, 1962.
- Howard, A. G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., and Adam, H. MobileNets: Efficient convolutional neural networks for mobile vision applications. *arXiv*, 2017. URL <https://arxiv.org/pdf/1704.04861.pdf>.
- Jain, A., Phanishayee, A., Mars, J., Tang, L., and Pekhimenko, G. Gist: Efficient data encoding for deep neural network training. In *ISCA*, 2018.

- Jia, Y., Shelhamer, E., Donahue, J., Karayev, S., Long, J., Girshick, R., Guadarrama, S., and Darrell, T. Caffe: Convolutional architecture for fast feature embedding. In *MM*, 2014.
- Jia, Z., Lin, S., Qi, C. R., and Aiken, A. Exploring hidden dimensions in parallelizing convolutional neural networks. In *ICML*, 2018.
- Jia, Z., Thomas, J., Warszawski, T., Gao, M., Zaharia, M., and Aiken, A. Optimizing dnn computation with relaxed graph substitutions. In *SysML*, 2019.
- Jouppi, N. P., Young, C., Patil, N., Patterson, D., Agrawal, G., Bajwa, R., Bates, S., Bhatia, S., Boden, N., Borchers, A., et al. In-datacenter performance analysis of a tensor processing unit. In *ISCA*, 2017.
- Judd, P., Albericio, J., Hetherington, T., Aamodt, T. M., and Moshovos, A. Stripes: Bit-serial deep neural network computing. In *MICRO*, 2016.
- Kahn, A. B. Topological sorting of large networks. *CACM*, 1962.
- Keßler, C. and Bednarski, A. A dynamic programming approach to optimal integrated code generation. In *LCTES*, 2001.
- Laredo, D., Qin, Y., Schütze, O., and Sun, J.-Q. Automatic model selection for neural networks. *arXiv*, 2019. URL <https://arxiv.org/pdf/1905.06010.pdf>.
- Lattner, C. and Adve, V. LLVM: A compilation framework for lifelong program analysis & transformation. In *CGO*, 2004.
- LeCun, Y., Denker, J. S., and Solla, S. A. Optimal brain damage. In *NIPS*, 1990.
- Lee, C., Lee, J. K., Hwang, T., and Tsai, S.-C. Compiler optimization on vliw instruction scheduling for low power. *TODAES*, 2003.
- Liu, H., Simonyan, K., and Yang, Y. DARTS: Differentiable architecture search. In *ICLR*, 2019a. URL <https://openreview.net/forum?id=S1eYHoC5FX>.
- Liu, Y., Wang, Y., Yu, R., Li, M., Sharma, V., and Wang, Y. Optimizing CNN model inference on CPUs. In *USENIX ATC*, 2019b.
- Mireshghallah, F., Taram, M., Ramrakhiani, P., Jalali, A., Tullsen, D., and Esmailzadeh, H. Shredder: Learning noise distributions to protect inference privacy. In *ASPLOS*, 2020.
- Mishra, A. and Marr, D. Apprentice: Using knowledge distillation techniques to improve low-precision network accuracy. In *ICLR*, 2018. URL <https://openreview.net/forum?id=Blael1ZRb>.
- NVIDIA. TensorRT: Programmable inference accelerator., 2017. URL <https://developer.nvidia.com/tensorrt>.
- Parashar, A., Rhu, M., Mukkara, A., Puglielli, A., Venkatesan, R., Khailany, B., Emer, J., Keckler, S. W., and Dally, W. J. Snn: An accelerator for compressed-sparse convolutional neural networks. In *ISCA*, 2017.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. PyTorch: An imperative style, high-performance deep learning library. In *NeurIPS*, 2019.
- Plump, D. Term graph rewriting. In *Handbook Of Graph Grammars And Computing By Graph Transformation: Volume 2: Applications, Languages and Tools*. World Scientific, 1999.
- Real, E., Aggarwal, A., Huang, Y., and Le, Q. V. Regularized evolution for image classifier architecture search. In *AAAI*, 2019.
- Rotem, N., Fix, J., Abdulrasool, S., Catron, G., Deng, S., Dzhabarov, R., Gibson, N., Hegeman, J., Lele, M., Levenstein, R., et al. Glow: Graph lowering compiler techniques for neural networks. *arXiv*, 2018. URL <https://arxiv.org/pdf/1805.00907.pdf>.
- Schösser, A. and Geiß, R. Graph rewriting for hardware dependent program optimizations. In *AGTIVE*, 2007.
- Sharma, H., Park, J., Suda, N., Lai, L., Chau, B., Chandra, V., and Esmailzadeh, H. Bit Fusion: Bit-level dynamically composable architecture for accelerating deep neural networks. In *ISCA*, 2018.
- Sifre, L. and Mallat, S. Rigid-motion scattering for image classification. *Ph.D. dissertation*, 2014.
- Vasilache, N., Zinenko, O., Theodoridis, T., Goyal, P., DeVito, Z., Moses, W. S., Verdoolaege, S., Adams, A., and Cohen, A. Tensor Comprehensions: Framework-agnostic high-performance machine learning abstractions. *arXiv*, 2018. URL <https://arxiv.org/pdf/1802.04730.pdf>.
- Wang, K., Liu, Z., Lin, Y., Lin, J., and Han, S. HAQ: Hardware-aware automated quantization with mixed precision. In *CVPR*, 2019.
- Wilken, K., Liu, J., and Heffernan, M. Optimal instruction scheduling using integer programming. In *PLDI*, 2000.
- Wortsmann, M., Farhadi, A., and Rastegari, M. Discovering neural wirings. In *NeurIPS*, 2019.
- Wu, C.-J., Brooks, D., Chen, K., Chen, D., Choudhury, S., Dukhan, M., Hazelwood, K., Isaac, E., Jia, Y., Jia, B., et al. Machine learning at facebook: Understanding inference at the edge. In *HPCA*, 2019.
- Xie, S., Kirillov, A., Girshick, R., and He, K. Exploring randomly wired neural networks for image recognition. In *ICCV*, 2019.
- Zhang, T., Yang, Y., Yan, F., Li, S., Teague, H., Chen, Y., et al. Swiftnet: Using graph propagation as meta-knowledge to search highly representative neural architectures. *arXiv*, 2019. URL <https://arxiv.org/pdf/1906.08305.pdf>.
- Zhou, S., Wu, Y., Ni, Z., Zhou, X., Wen, H., and Zou, Y. DoReFa-Net: Training low bitwidth convolutional neural networks with low bitwidth gradients. *arXiv*, 2016. URL <https://arxiv.org/pdf/1606.06160.pdf>.
- Zhu, M. and Gupta, S. To prune, or not to prune: exploring the efficacy of pruning for model compression. In *ICLR Workshop*, 2018. URL <https://openreview.net/forum?id=S11N69AT->.
- Zoph, B. and Le, Q. V. Neural architecture search with reinforcement learning. *ICLR*, 2017. URL <https://openreview.net/forum?id=r1Ue8Hcxg>.
- Zoph, B., Vasudevan, V., Shlens, J., and Le, Q. V. Learning transferable architectures for scalable image recognition. In *CVPR*, 2018.

## A COMPARISON BETWEEN IRREGULARLY WIRED NEURAL NETWORKS AND CONVENTIONAL REGULAR TOPOLOGY NEURAL NETWORKS



(a) ImageNet accuracy vs number of multiply-and-accumulate.



(b) ImageNet accuracy vs number of parameters.

Figure 14. ImageNet accuracy vs number of multiply-and-accumulate or parameters, where irregularly wired neural networks show higher performance for same amount of compute or number of parameters than regular topology neural networks.

## B COMPARISON WITH TENSORFLOW LITE

In addition to the relative reductions provided in Figure 10, Figure 15 provides the raw numbers of the peak memory footprint for the benchmark irregularly wired neural networks.




Figure 15. Peak memory footprint of running irregularly wired neural networks on SERENITY and TensorFlow Lite.

## C PROOF FOR OPTIMAL PEAK MEMORY FOOTPRINT FROM THE DYNAMIC PROGRAMMING-BASED SCHEDULING

Here we prove the optimality of the above dynamic programming-based scheduling algorithm.

**THEOREM 1.** *In order to find a schedule  $s^*$  with an optimal peak memory consumption  $\mu^*$ , it is sufficient to keep just one schedule-peak memory pair  $(s_i, z_i)$  in  $S_{T_i}$  for each zero-indegree set  $z_i$ , and to append subsequent nodes on top of  $s_i$  to get  $s_{i+1}$  in each search step.*

*Proof.* If  $i=0$ , the optimal  $s_0$  is an empty sequence and  $\mu_0$  must be 0. On the other hand, if  $i \geq 1$ , assume that (suboptimal)  $v_i$  constitutes  $s^*$ , substituting  $u_i^* \in z_i$  and achieves  $\mu^*$ . In such case, let  $v_i$  be replaced with (optimal)  $u_i^*$ , which will result in  $\mu_{peak} \leftarrow \min(\mu_i + \prod v_i.shape, \mu_i + \prod u_i^*.shape)$ , and  $\mu_{i+1}$  is calculated by deducting  $\prod p_i.shape, \forall p_i \in (u_i.preds \cap zero-outdegree(s_{i+1}, \mathcal{G}))$ . By recursively applying  $u_k$  for rest of the search steps  $k$ , the algorithm should find an alternative sequence  $s'^*$  with  $\mu'^* \leq \mu^*$  due to the min operator above, contradicting the original assumption on the optimality of  $s^*$ . Therefore, our algorithm finds a schedule with an optimal peak memory consumption. ■

## D COMPLEXITY ANALYSIS OF THE DYNAMIC PROGRAMMING-BASED SCHEDULING AND PROOF

We compare the complexity of exhaustively exploring  $\mathcal{S}_T$  and our dynamic programming-based scheduling. While the algorithm both lists candidate schedules and calculates their peak memory footprint, we consider the peak memory footprint calculation as one operation while deriving the complexity. In order to visualize the analysis, we invent  $\mathcal{G}$  in Figure 16 to demonstrate the upper bound complexity of each algorithm. It has a single entry node and a single exit node  $(A)$  and  $(Z)$ , respectively, and all other nodes constitute independent branches between the entry and the exit node.




Figure 16. Topology of  $\mathcal{G}$  to demonstrate the upper bound complexity of each algorithm.

First, we demonstrate the complexity of the recursive topological sorting that exhaustively explores  $\mathcal{S}_T$ . Since there is a single entry node and a single exit node, there will be  $|V - 2|$  remaining nodes and these nodes can be scheduled independently of one another, thereby the number of candidate schedules become  $(|V - 2|!)$  and the overall

complexity becomes  $\mathcal{O}(|V|!)$ , where  $|V|$  denotes the number of nodes. On the other hand, for the dynamic programming we calculate the number of candidates by utilizing the number of schedules that gets memoized. Our memoization takes advantage of the zero-indegree sets  $z$  for each search step.

For the first and the last search steps, we assume that we have a single entry node and a single exit node. On the other hand, since the number of nodes scheduled in search step  $i$  would be  $i-1$ , the maximum number of entries for memoization is  $\binom{|V|-2}{i-1}$ . On top of this, each step would make an iteration over the set of candidate nodes to discover the next search step's  $z$ . Therefore, search step 1 would explore  $|V|-2$  nodes and the search steps 2 to  $|V|-1$  would iterate over  $|V|-1-i$  nodes. Summarizing this would yield:

$$\begin{aligned}
 & 1 + 1 \times (|V|-2) + \binom{|V|-2}{1} \times (|V|-3) + \\
 & \quad \dots + \binom{|V|-2}{|V|-2} \times 0 + 1 \\
 & = 1 + \binom{|V|-2}{0} \times (|V|-2) + \binom{|V|-2}{1} \times (|V|-3) + \\
 & \quad \dots + \binom{|V|-2}{|V|-2} \times 0 + 1 \\
 & = 2 + \sum_{i=0}^{|V|-2} \binom{|V|-2}{i} \times (|V|-2-i) \\
 & = 2 + (|V|-2) \times 2^{|V|-3} \\
 & \leq (|V|-2) \times 2^{|V|-2} \quad , \text{ for } |V| \geq 4 \\
 & \leq |V| \times 2^{|V|}
 \end{aligned}$$

As a result, we can see that our dynamic programming-based scheduling algorithm is bounded by  $\mathcal{O}(|V| \times 2^{|V|})$ . By using Stirling's approximation on the complexity of the recursive topological sorting, we can prove that the dynamic programming-based scheduling algorithm should be significantly faster than the recursive topological ordering.