APOLLO: AUTOMATIC PARTITION-BASED OPERATOR FUSION THROUGH LAYER BY LAYER OPTIMIZATION

Jie Zhao [1] **Xiong Gao** [2] **Ruijie Xia** [2] **Zhaochuang Zhang** [2] **Deshi Chen** [2] **Lei Chen** [3] **Renwei Zhang** [2]
Zhen Geng [2] **Bin Cheng** [2] **Xuefeng Jin** [2]

ABSTRACT

We study fusion for deep neural networks (DNNs) in a just-in-time (JIT) compilation framework APOLLO. It considers both memory- and compute-bound tensor operators for fusion, and integrates graph-level node grouping and operator-level loop fusion closely, widening the fusion search space. APOLLO enables the upward feedback from the downstream loop optimizer, enforcing the graph engine to regenerate partition patterns amenable to the downstream pass and thus resolving the scalability issue. Besides data locality, APOLLO also exploits the parallelism between independent tensor operators, further improving the performance of DNN workloads. Experimental results on training workloads show that APOLLO outperforms TensorFlow and XLA by 1.86× and 1.37× on a single GPU, and 1.96× and 1.18× on multiple GPUs. APOLLO also improves the performance of a vendor-provided DNN framework by 19.7% on a domain-specific accelerator. In addition, the results of inference workloads demonstrate the general applicability of our fusion framework.

1 INTRODUCTION

DNN frameworks (Abadi et al., 2016; Paszke et al., 2019) offer the ease-of-use interfaces by obscuring hardware information from users, but the ever-increasing depth of network layers and amount of data demand for the extensive exploration of architectural features to harness the computing power provided by the target platform. Such a widening gap is calling for an effective end-to-end compilation infrastructure (Lattner et al., 2021), of which the *fusion* transformation has fascinated massive attentions.

Fusion, or *node grouping* (Jia et al., 2019a; Jangda & Bondhugula, 2020), is used to optimize a *computational graph*, where each node expresses an operator (*op*) and is associated with others through edges of producer-consumer relations. Many graph compilers (Google, 2017; Wei et al., 2018; Rotem et al., 2019) only perform node grouping between memory-bound *op*s, since they assume the kernel launch overhead of compute-bound *op*s like convolution (*conv*) and matrix multiplication (*matmul*) is insignificant compared to the kernel execution. They dispatch each compute-bound *op* to a vendor library kernel, *missing the*

opportunities to fuse with compute-bound ops and exploiting fusion within an incomplete space (**challenge 1**). An example that demonstrates the profit of such fusion strategies was described by Zhao & Di (2020).

Fusion is also referred to as *loop fusion* (McKinley et al., 1996) when a computational graph is lowered to *multi-dimensional tensor expressions*, where an *op* is instantiated using arithmetic operations encompassed by nested loops. Existing tensor compilers (Chen et al., 2018a; Vasilache et al., 2019) compensate the weakness of graph compilers by exploiting fusion patterns involving compute-bound *op*s. A routine transformation orchestration of node grouping followed by loop fusion is adopted by these compilers, exerting overwhelming force from the graph engine to the downstream loop optimizer, *challenging the scalability of the loop fusion heuristics* (**challenge 2**). For instance, a large loop shifting factor (§4.1) that may increase compilation time is mandatory in some cases, which is enforced by the sub-graph generated by the upstream graph compiler.

Fusion, named as *stitching* (Zheng et al., 2020), also contributes to the exploitation of parallelism between independent *op*s. Unlike traditional compilers that only optimize the producer-consumer relations, TASO (Jia et al., 2019a) and Rammer (Ma et al., 2020) also pack independent *op*s into a single kernel and execute them concurrently, maximizing the utilization of hardware parallelism when these *op*s cannot saturate the target device. These approaches compile *op*s ahead of time due to the reliance on manual schedule templates, *falling short in supporting custom ops in training scenarios* (**challenge 3**; see Listing 1 and 2).

---

[1] State Key Laboratory of Mathematical Engineering and Advanced Computing, Zhengzhou, China [2] Huawei Technologies Co., Ltd., Hangzhou, Beijing and Shenzhen, China [3] Hong Kong University of Science and Technology, Hong Kong, China. Correspondence to: Jie Zhao <yaozhujiajie@gmail.com>, Lei Chen <leichen@cse.ust.uk>, Xuefeng Jin <jinxuefeng@huawei.com>.

We present APOLLO, an **A**utomatic **P**artition-based **O**perator fusion framework through **L**ayer by **L**ayer **O**ptimization, to address the above challenges. It splits a sub-graph extracted from a DNN model into a batch of micro-graphs using a rule-based algorithm, which considers both memory- and compute-bound *op*s and results in a much broader search space of fusion. The rules are designed with the requirements from a downstream loop optimizer fully considered, removing the constraints from the upstream graph compiler. The *op*s within a generated micro-graph can thus be always fused smoothly by the polyhedral model (Feautrier & Lengauer, 2011) and the scalability issue is fully addressed.

APOLLO obtains optimized sub-graphs through a series of bottom-up, multi-layer fusion transformations. First, the bottom layer exploits fusion within each micro-graph using a *polyhedral loop fusion* heuristic, optimizing memory hierarchy by allowing the tensors carried by producer-consumer dependences on faster local memory. Second, the middle layer studies fusion between micro-graphs using a *memory stitching* approach, compensating the drawback of polyhedral compilation. Finally, the top layer combines disjoint micro-/sub-graphs generated by the first/second layer, maximizing the utilization of hardware parallelism. The coupled implementation of multi-layer fusion goes beyond prior work (Ragan-Kelley et al., 2013; Chen et al., 2018a;b; Vasilache et al., 2019) by simultaneously exploiting data locality and parallelism between *op*s.

Instead of generating code for each micro-/sub-graph individually, APOLLO delivers the intermediate representations (IRs) generated by each previous layer to the next, enabling faster compilation of multiple micro-/sub-graphs using a *piecewise compilation strategy*. The ultimate integrated IR is passed to the code generator, producing one or multiple optimized kernel implementations. The fully automated, high-efficiency compilation of APOLLO makes itself competent to be used as a JIT compiler and applicable to both training and inference workloads.

In summary, this work makes the following contributions.

- APOLLO extends the search space of fusion by considering more *op* types, generating more profitable across-layer schedules originally hindered by *op* boundaries;

- APOLLO addresses the scalability issue of fusion by allowing reverse feedback from the operator-level optimizer, achieving a fully automatic fusion framework;

- APOLLO enhances the performance of DNN workloads by modeling both data locality and parallelism, producing more efficient code than the state of the art;

- APOLLO exhibits reasonable JIT compilation overhead, demonstrating its effectiveness using rather difficult real-life training workloads.

We conduct experiments on training workloads and compare the performance with TensorFlow and XLA, over which APOLLO achieves $1.86\times$ and $1.37\times$ speedups on a single GPU, and $1.96\times$ and $1.18\times$ on multiple GPUs. APOLLO outperforms a vendor DNN framework by 19.7% on a domain-specific accelerator. We also provide some preliminary results of inference workloads, demonstrating the general applicability of APOLLO.

The paper is organized as follows. §2 overviews the architecture of APOLLO. §3 and §4 introduce the partition and fusion phases, respectively. §5 puts everything together, followed by the experimental results reported in §6 and related work discussed in §7. §8 concludes the work.

2 ARCHITECTURE OF APOLLO

The colored *op*s shown in Fig.1 constitute two disjoint sub-graphs, with each including a compute-bound *op* (op_3 or op_5). op_2, op_4 and op_7 are *compound ops* and op_1 and op_6 are *primitive ops*. A compound *op* is a function, *e.g.*, SoftMax, composed of multiple primitive *op*s, each of which is an arithmetic operation like addition, multiplication, division, *etc.* op_2 is composed of two (blue) primitive *op*s (op_{21} and op_{22}), op_4 consists of two (gray) primitive *op*s (op_{41} and op_{42}), and op_7 is made up of three (yellow) primitive *op*s, op_{71}, op_{72} and op_{73}. A (green/orange/red) primitive *op* is not affiliated to any compound *op*s.
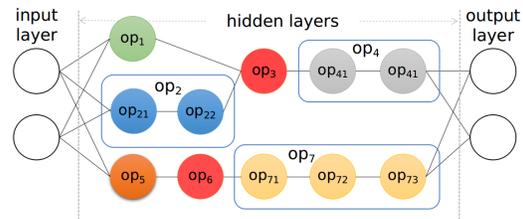

Figure 1: An illustrative DNN computational graph.

Many graph compilers (Google, 2017; Wei et al., 2018) did not consider op_3 or op_5 for fusion. They resort to vendor tuned libraries for such *op*s and naturally isolate each sub-graph into multiple components. A dynamic programming strategy (Ding et al., 2021) is used to evaluate each fusion possibility with the growing of each sub-graph's complexity, but its compilation overhead may become an issue.

Tensor compilers (Chen et al., 2018a;b; Vasilache et al., 2019; Baghdadi et al., 2019) perform fusion together with tiling (Irigoin & Triolet, 1988), but their fusion heuristics are subject to the constraints imposed by upstream graph compilers and thus suffer from the scalability issue (Mehta et al., 2014; Zhao & Di, 2020). For example, an upstream graph compiler expects for a single kernel implementation for op_7, but the downstream loop optimizer may not be able to meet this requirement when op_{71} is a reduction *op*, as will be explained in §3.2.

Another major weakness of existing compilers is their inability to effectively use the available hardware parallelism when given smaller batch sizes. The recent work (Jia et al., 2019a; Ma et al., 2020; Zheng et al., 2020) follows this direction to pack multiple independent branches (*e.g.*, op_1 and op_2) of a sub-graph into a single kernel, but they rarely consider training scenarios or dedicated chips.

We design APOLLO, the architecture of which is depicted in Fig.2. Its *partition phase* (§3) first extracts the maximum set of sub-graphs $\mathcal{P}$ and next splits $\mathcal{P}$ into m individual sub-graphs $\mathcal{F}_x$ ($1 \leq x \leq m$). Unlike XLA, compute-bound *op*s are considered when performing fusion. The m sub-graphs are split into n micro-graphs $\mathcal{G}_y$ ($1 \leq y \leq n$) by a rule-based algorithm, with the challenges for the scalability of the fusion phase fully considered, different from the graph engine of TVM (Chen et al., 2018a) that partitions a sub-graph without the awareness of its loop optimizer's requirements, which in turn may lead to the scalability issue.
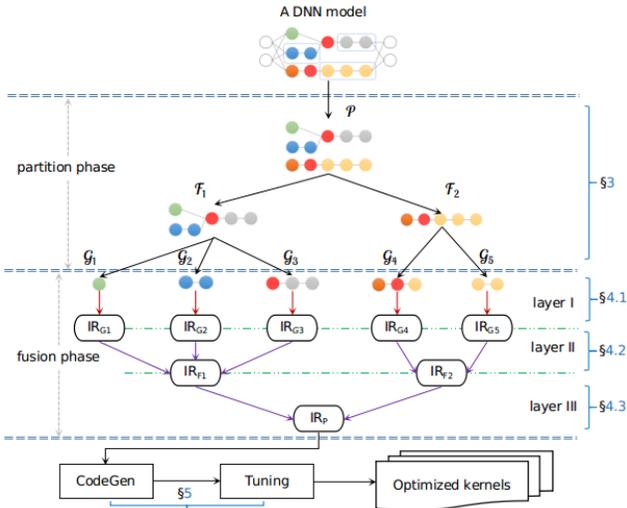

Figure 2: Architecture of APOLLO.

The fusion phase (§4) includes three layers. Layer I (§4.1) carries out loop fusion for each $\mathcal{G}_y$. It addresses the scalability by always producing a single group for a $\mathcal{G}_y$ using a polyhedral loop fusion heuristic in the polyhedral model (Verdoolaege & Janssens, 2017). Layer II (§4.2) implements node grouping by aggregating Layer I's outputs to implement fusion between a reduction *op* with its following *op*s, which was not studied by prior approaches (Zhao et al., 2021; Vasilache et al., 2019).

The outputs of Layer II are passed to Layer III (§4.3) for exploiting the parallelism between independent *op*s. In contrast to TASO (Jia et al., 2019a), we consider parallelism stitching between *op*s of different types and support code generation for both GPU and a dedicated accelerator (§5). APOLLO finally generates one or multiple kernels that are optimized by our auto-tuner, with the execution of Layer I and II parallelized for JIT compilation.

3 PARTITION PHASE

Before a DNN computational graph is lowered to the partition phase, we also perform some pre-processing optimizations to simplify the graph. They are as follows.

- ♣ *Algebraic simplification* acts as function inlining by making use of associativity, commutativity and distributivity, as also adopted by Google (2017).
- ♣ *Data-flow optimization* performs common subexpression elimination and constant folding, which is also considered by nGraph (Cyphers et al., 2018).
- ♣ *Control-flow optimization* attempts to simplify a graph by eliminating dead producer *op*s and branches, which may be caused by algebraic simplification.
- ♣ *Data-layout transformation* changes the way a tensor is stored as demand to keep dimension alignment between the producer and consumer data spaces.

The preparation simplifies a DNN computational graph, reduces the compilation overhead, and meets the "static affine control" requirement (Verdoolaege, 2010) of the polyhedral model, though it might be incomplete.

3.1 Extracting Sub-graph Cluster

APOLLO first extracts the set of eligible nodes from a computational graph. Two kinds of *op*s are not considered by APOLLO. First, the user-defined and/or extraordinary *op*s with complex computational logic. A typical example is `all-reduce` used for training speech recognition (Amodei et al., 2016). The most appropriate solution to deploy such *op*s is wrapping a highly-crafted library like that of Cho et al. (2019). Second, control flow *op*s like TensorFlow's `RefSwitch` should be excluded. The remaining node types may constitute a group of sub-graphs $\mathcal{P}$ that are disconnected from others. We use fusion to minimize the producer-consumer distances between *op*s within each sub-graph.

3.2 Opening Compound Operators

A dependence-based fusion algorithm can be applied within each sub-graph $\mathcal{F}_x$, but the dependence patterns between *op*s are usually complex, disabling the generation of a single kernel for an $\mathcal{F}_x$ or even for a single compound *op*.

The use of activation functions is one of the major reasons. An activation function usually involves many arithmetic operations and can be split into multiple primitive *op*s. We use a stable variant of `SoftMax`, `LogSoftMax`, to explain this issue. Its formula can be expressed as

$$S(t_i) = t_i - \ln(\sum_{j=1}^{N} e^{t_j}) \qquad (1)$$

where t_i is i-th element of an input vector of length N and e the exponential constant. Formula (1) requires two operations, one computing the logarithm of the reduction over all

vector elements and the other performing the subtraction, that have to be decomposed into multiple parallel execution units through loop tiling. The tiled subtraction must wait for the completion of all simultaneously executed tiles of the reduction, preventing the fusion between the two tiled operations. A loop fusion heuristic coupled with loop tiling thus fails to generate a single kernel for this compound *op*. As such, `LogSoftMax` is not a fusible candidate.

A compound *op* is expressed using a (blue) rounded box in Fig. 1. The fusion patterns that can be explored by a loop fusion heuristic are constrained by such a box, which we refer to as an *op boundary*. An *op* boundary results in suboptimal fusion patterns by isolating the internal primitive *op*s of a compound *op* from external ones. We break up each compound *op* to get rid of every *op* boundary. For instance, they have been removed in each $\mathcal{F}_x$ of Fig. 2.

Example 1. *Suppose that* op_3 *composed of a reduction* op_{31} *and a subtraction* op_{32} *in Fig.3(a) is a* `LogSoftMax` *function preceded by two primitive ops. The fusion pattern when the op boundary exists is shown in Fig.3(b), with each fusion group denoted using a red block. Another fusion pattern, Fig.3(c), is possible if the boundary is removed.*


(a)   (b)   (c)
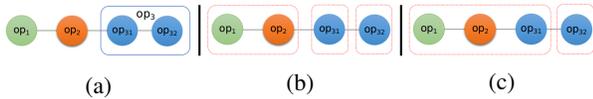
Figure 3: Effect of *op* boundaries. (a) `LogSoftMax`; and the fusion patterns when *op* boundaries (b) exist or (c) not.

3.3 Aggregating Primitive Operators

The final step is to partition each $\mathcal{F}_x$ into micro-graphs. Instead of producing individual micro-graphs by splitting an $\mathcal{F}_x$, we instantiate an individual micro-graph $\mathcal{G}$ with each primitive *op* of a given $\mathcal{F}_x$, and progressively merge these disjoint micro-graphs using *aggregation rules* that imitate the fusion behaviors tolerated by a loop fusion heuristic. This approach allows us to spell out each fusion pattern that can be solved by a polyhedral heuristic, evading unseen scenarios that may cause the scalability issue at the expense of losing fusion opportunities in Layer I (§4.1). Such a side effect will either be overcome at Layer II (§4.2), or be learned offline to update the rules of the partition phase.

Our experience reveals that the scalability issue is closely related to the types of primitive *op*s to be fused. A primitive *op* takes as input p multi-dimensional tensors and outputs the result to another tensor. In other words, each primitive *op* maintains p input data spaces $I_k \in V_{\mathbb{Z}}^d$ $(1 \leq k \leq p)$ and one output data space $O \in V_{\mathbb{Z}}^d$, with $V_{\mathbb{Z}}^d$ $(d \geq 1)$ representing the positive integer vector space. One can build a *dataflow relation* $I_k \to O$ between the k-th input and the output tensor; or the k-th input tensor can be discarded.

It is safe to assume each I_k is with the same dimension-

ality d as O, since we can perform *tensor broadcasting* to transform two tensors into compatible shapes. A dataflow relation can thus be written as the conjunction of d 1D functions:

$$\bigwedge_{1 \leq l \leq d} f_l := I_k^l \to O^l \quad \left(I_k^l, O^l \in V_{\mathbb{Z}} \right) \qquad (2)$$

Integers $e_i, e_o \geq 1$ denote the loop trips of I_k^l and O^l. f_l is supposed as either *bijective* $(e_i = e_o)$, *injective* $(e_i = 1 \wedge e_o \neq 1)$, *surjective* $(e_i \neq 1 \wedge e_o = 1)$ or a *general function*. A dataflow relation $I_k \to O$ is thus considered as

1. a bijective relation if each f_l is bijective, or
2. an injective relation if u $(u \geq 1)$ out of d f_l's are injective and each of the remaining is bijective, or
3. a surjective relation if u $(u \geq 1)$ out of its d 1D functions are subjective and each of the remaining $d - u$ f_l's is either a bijective or injective function, or
4. a general relation if at least one f_l is a general function.

Given a primitive *op*, we can determine its type using

Definition 1. *op is an* element-wise *operator iff each* $I_k \to O$ $(1 \leq k \leq p)$ *is a bijective dataflow relation.*

Definition 2. *op is considered as a* broadcast *operator iff there exist* u $(u \geq 1)$ *out of the* p *dataflow relations are injective and each of the remaining is a bijective function.*

Definition 3. *The type of op is* reduction *iff at least one of its* p *dataflow relations is a surjective dataflow relation and each of the remaining is either bijective or injective.*

Definition 4. *op is referred to as an* opaque *operator iff it is with at least one general dafaflow relation.*

Our definition classifies reshaping operations, (batched) *matmul* and *conv* as opaque *op*s. Once the type of a primitive *op* is determined, one can define the type of the aggregation result of two micro-graphs. We summarize the rules to merge two micro-graphs ($\mathcal{G}_p$, $\mathcal{G}_c$) in Table 1.

Table 1: Aggregation rules. $\mathcal{G}_p$ and $\mathcal{G}_c$ hold a producer-consumer relation; $\mathcal{G}_a$ is the merged micro-graph.

Rules	$\mathcal{G}_p$	$\mathcal{G}_c$	$\mathcal{G}_a$
❶	element-wise	element-wise	element-wise
❷	broadcast	element-wise	broadcast
❸	broadcast	broadcast	broadcast
❹	element-wise	reduction	reduction
❺	broadcast	reduction	reduction
❻-transpose	element-wise/broadcast	transpose	transpose
❻-matmul	matmul	element-wise	matmul
❻-matmul	element-wise	matmul	matmul
❻-conv	conv	element-wise	conv
❻-conv	element-wise	conv	conv

These rules do not need to cover all composition patterns of *op*s, since some pair of *op*s should not be fused. For example, we originally defined an aggregation rule for a reshaping *op* followed by an element-wise or broadcast *op*, but we found that the polyhedral loop optimizer cannot finish the scheduling process within a reasonable time. This rule was thus removed from Table 1, which was suggested

by the feedback from the downstream optimizer, as will be introduced in §4.1. This removal results in an incomplete set of rules covered by Table 1, which may generate more sub-graphs in practice but avoids the scalability issue of fusion. This is important to a JIT compilation tool.

The rightmost column indicates the type of the aggregation result, while the preceding two enumerate each possible type combination of the micro-graph pair to be aggregated. Such a definition enables the recursive adoption of these rules and guarantees the termination of the algorithm when undefined combination patterns are encountered. A micrograph has only a single output tensor when initialized using a primitive *op*, but multiple output tensors within a micrograph may be created during the merging process. Algo.1 in Appendix A describes how to aggregate micro-graphs.

Example 2. *Algo.1 obtains Fig.3(c)'s first fusion group by applying* ❶ *and* ❹*. The second one made up of a single op is separated from the remaining since no rules aggregating a reduction with a follow-up op are defined in Table 1.*

4 FUSION PHASE

The fusion phase performs fusion in a bottom-up manner. Layer I also performs loop tiling. The selected tile sizes determine how much memory is occupied by each micrograph, which exposes the amount of used faster memory to the high-level fusion heuristics of Layer II and III. The later two layers would otherwise not be aware of such information. Such a bottom-up fusion strategy is more flexible with respect to the support of custom *op*s for training scenarios: the fusion phase can feedback to the partition phase when any $\mathcal{G}_y$ produced by the later is not acceptable or a single fusion group cannot be generated.

4.1 Layer I: Polyhedral Loop Fusion

A $\mathcal{G}_y$ has been converted into a sequence of loop nests when lowered to this layer, and we use AKG (Automatic Kernel Generator) (Zhao et al., 2021), a polyhedral optimizer, to perform loop fusion and tiling.

The polyhedral model determines the parallelism and tilability of each loop nest by solving an integer linear programming (ILP) problem (Bondhugula et al., 2008) and introduces auxiliary transformations *e.g.*, loop interchange, shifting and scaling, to ensure the alignment between the loop nests when necessary. Expressed using affine relations (Verdoolaege & Janssens, 2017) in the polyhedral model, some of these loop transformations are not accessible in manual scheduling approaches, but they are sometimes essential to minimize the producer-consumer distances. Writing manual schedule templates may not be able to minimize the producer-consumer distances between loop nests even though more schedule primitives can be complemented, since auxiliary loop transformations have to be

used together with loop tiling and fusion (Zhao & Cohen, 2019; Jangda & Bondhugula, 2020), requiring a systematical composition of the modeled transformations. Listing 1 shows an example that calls for the systematical composition of loop scaling and fusion; the polyhedral model can transform it into the form show in Listing 2.

Unfortunately, AKG still follows the routine transformation orchestration formulated by TVM since it inherits the graph engine of the later. TVM's rules used to perform graph-level node grouping did not allow the feedback from the downstream AKG, whose loop fusion heuristic may sometimes generate very inefficient fusion patterns or cannot terminate within a reasonable time. Similar scalability issue is also faced by other polyhedral compilers (Vasilache et al., 2019; Verdoolaege et al., 2013).

```
for i in [0,M)                    for i in [0,M)
 for j in [0,N)                    for j in [0,N){
  a(i,j)=a(i,j)+bias; //S1           a(i,j)=a(i,j)+bias;
for i in [0,M/2)                     if(i+1) mod 2 = 0 and
 for j in [0,N/2)                       (j+1) mod 2 = 0
  pool(i,j)=max(a(2i,2j),            pool((i-1)/2,(j-1)/2)=
    a(2i,2j+1),                        max(a(i-1,j-1),a(i,j-1),
    a(2i+1,2j),                          a(i-1,j),a(i,j)); //S2
    a(2i+1,2j+1)); //S2              }
```

Listing 1: Original loop nests.   Listing 2: After fusion.

Example 3. *Still consider the example in Fig.3(a). Without our aggregation rules, all of the four ops will be delivered by the graph engine of TVM to AKG, which constructs a single fusion group for these four ops. A very large loop shifting factor is required to guarantee the execution order of tiled ops as explained in §3.2. We observe that not only does the compilation time increase but the single fusion group obtained with the help of a large loop shifting factor also degrades the execution performance in practice.*

Our aggregation rules overcome this weakness by always producing a micro-graph whose composition of loop nests is predictable. The fusion heuristic of the polyhedral model will thus never be challenged by the scalability issue. Once the polyhedral scheduler cannot finish within a reasonable time, its *op* composition pattern will be feedback to a developer, who can then use such information to update the aggregation rules in Table 1.

We optimize the fusion of reduction *op*s using PANAMERA (Zhao et al., 2022), which can effectively fuse a reduction *op* with its preceding element-wise/broadcast *op*s, with the help of a dimension flattening optimization that always coalesces the loop nest of a reduction *op* into three canonical forms, simplifying the scheduling process and mitigating the polyhedral complication complexity. This optimization was not considered by manual scheduling approaches or vendor libraries (Chetlur et al., 2014), but it allows APOLLO to find better fusion patterns in practice.

Example 4. *Listing 3 shows a reduction op preceded by an element-wise op. APOLLO first converts it into one of the three canonical forms (referred to as y-reduce) through*

loop interchange, as shown in Listing 4, allowing APOLLO *to obtain fusion result in Listing 5. The tensor subscripts are updated automatically by the polyhedral model.*

```
for i in [0,M] and j in [0,N] and k in [0,P] and l in [0,Q]
  a(i,j,k,l) = a(i,j,k,l) + bias
for i in [0,M] and j in [0,N] and k in [0,P] and l in [0,Q]
  b(i,k) += a(i,j,k,l)
```

Listing 3: Original loop nests.

```
for x in [0,M*P] and y in [0,N*Q]
  a(x/P,y/Q,x%P,y%Q) =
    a(x/P,y/Q,x%P,y%Q) + bias
for x in [0,M*P] and y in [0,N*Q]
  b(x/P,x%P) += a(x/P,y/Q,x%P,y%Q)
```

Listing 4: The canonical form.

```
for x in [0,M*P]
  and y in [0,N*Q]{
  a(x/P,y/Q,x%P,y%Q) = ...
  b(x/P,x%P) += ...
}
```

Listing 5: After fusion.

Once each micro-graph is fused by the polyhedral model, we can generate the optimized IR and deliver it to Layer II for exploring fusion possibilities between micro-graphs.

4.2 Layer II: Memory Stitching

We define *complementary aggregation rules* to exploit the fusion possibilities between micro-graphs. The micro-graph type combinations considered at this layer are summarized in Table 2. It allows the recursive application of its each rule, which greedily searches the fusion possibilities between micro-graphs. These rules only consider micro-graph type combinations starting with a reduction *op*, thereby guaranteeing the termination of the greedy fusion approach when other cases are encountered.

Table 2: The complementary aggregation rules.

Rules	$\mathcal{G}_p$	$\mathcal{G}_c$	$\mathcal{G}_a$
❼	reduction	element-wise/broadcast	reduction
❽	reduction	reduction	reduction

Nonetheless, it is still not straightforward to design a fusion algorithm for Layer II due to the diverse reduction scenarios of DNN models. A reduction can take place along any one or multiple dimensions of the enclosing loop nest of an *op*. For example, reduction is performed along the second and fourth dimensions of the loop nest shown in Listing 3. In addition, loop tiling at Layer I binds individual loop dimensions of a reduction micro-graph to different hardware dimensions. Layer II has to investigate and ensure both loop dimensions and hardware parameters are matched.

The matching between the loop dimensions of two micro-graphs is guaranteed by the three canonical forms generated by PANAMERA, as introduced in §4.1. APOLLO always transforms a reduction micro-graph into an *all*-reduce form, where the enclosing loop nest a reduction micro-graph is flattened into a single reduced loop, or an *x-/y*-reduce form, where the multi-dimensional loop nest of a fused micro-graph has been coalesced into a 2D loop nest and reduction is performed only along the outer/inner loop dimension. We are allowed to fuse two micro-graphs if they have been flattened into the same canonical form.

The hardware setting is specified at Layer I and heavily

related to the parameters of specific loop transformations. For example, the tile sizes found by the loop optimizer have a significant impact on the configuration of grid/block dimensions when targeting GPU. We mainly concern about the tuned tile sizes found by the previous layer, which determine where the tensors relating two micro-graphs will be allocated. $(\mathcal{G}_p, \mathcal{G}_c)$ can be fused when their tile sizes and hardware parameter setting are both identical. Algo.2 in Appendix A delineates the memory stitching method.

Example 5. *Denoted by a (red) dotted box, each micro-graph in Fig.5 is extracted from a real-world DNN workload. Primitive ops in the same color belongs to the same micro-graph. Algo. 2 fuses these micro-graphs using the rule* ❼*, with the intermediate tensors carried by each blue edge allocatable on local memory.*

4.3 Layer III: Parallelism Stitching

Layer I and II did not consider the intrinsic parallelism between independent *op*s/micro-graphs. This is also the weakness of many existing tensor compilers (Chen et al., 2018a; Vasilache et al., 2019), which are thus usually used to generate code for a single device.

Determining the independence between each pair of *op*s is non-trivial, which may result in the combinatorial explosion issue. We observe that such parallelism mainly exists between the branches of a multi-head/-tail *op*. Another scenario that should be considered is the independence between $\mathcal{F}_1$ and $\mathcal{F}_2$ in Fig.2, where no common producer/-consumer exists between the two sub-graphs. A virtual common consumer/producer can be introduced to covert such scenario into a multi-head/-tail case without violating the semantic. We thus only discuss the handling of multi-head/-tail scenarios in the following context.

To pick up each suitable *op* along one branch that can be executed in parallel with those of another branches, we traverse backward/forward along a branch and terminate until another multi-head/-tail *op* is reached. The common head-/tail is not considered during the traverse. The direction is respect to that of the producer-consumer dependence between *op*s. A multi-head *op* is either considered as can be executed in parallel with those *op*s on other branches when it is a terminator of a backward traverse, or excluded when it appears at the end of a forward path. Conversely, a multi-tail terminator is collected in a forward path or excluded in a backward traverse. Besides, a compute-bound *op* should also be excluded, since its huge amount of data usually consumes up the hardware resources.

Example 6. *The two branches of the right bottom* mul *op of Fig.6(b) are first traversed. They are considered as parallelizable, represented using blue boxes. The two independent branches are then evaluated by introducing a virtual tail, with the parallelizable parts denoted using red boxes.*

APOLLO can now try to compose the collected paralleliz-able candidates to exploit the inter-*op* parallelism. The *op*s along the same branch cannot be executed in parallel, since they depend on each other. We repeatedly choose one *op* from each branch and consider their composition possibilities. Ideally, all of these extracted *op*s can be ex-ecuted simultaneously, but the parallelization is restricted by the available hardware resources. Dispatching more numbers of parallelizable *op*s onto the limited hardware resources, *e.g.*, GPU streaming multiprocessors, does not imply higher performance. We evaluate the potential per-formance gain of parallelization using a cost model

$$gain = \sum_{op=m}^{k} cost_{op} - \max_{m \le op \le k}(cost_{op}) \qquad (3)$$

where k is the number of *op*s evaluated and m the starting number of these k sorted *op*s. $cost_{op}$ can be inspected as the execution time of an *op* and can be estimated accord-ing to the shapes of its tensors and the allocated hardware resources. The summation part computes the sequential ex-ecution time of all evaluated *op*s. The maximum part repre-sents the maximal execution time among these *op*s, which determines the execution time after parallelization. Their subtraction is the performance gain. Parallelism stitching is formally described by Algo.3 in Appendix A.

Example 7. *Suppose that we have five branches and the descending order of the extracted ops from each branch is* $\{op_3, op_1, op_5, op_2, op_4\}$. *Algo.3 attempts to compose all ops (m=1, k=5) but the performance gain is not possitive. It updates k (at line 8) and continues. We assume that it is still not able to compose* $\{op_3, op_1, op_5, op_2\}$ *and retries using* $\{op_3, op_1, op_5\}$. *ops* $\{op_2, op_4\}$ *will be evaluated if Algo.3 succeeds, and the resulting fusion groups may be composed of* $\{op_3, op_1, op_5\}$ *and* $\{op_2, op_4\}$.

Note that the optimization performed by Layer III is platform-neutral, since the polyhedral model can manage both traditional and emerging multi-directional memory systems. On the contrary, existing frameworks (Jia et al., 2019a; Zheng et al., 2020) are only evaluated on typical platforms, since they rely heavily on TVM that can only manage the traditional memory hierarchy pyramid of GPU or similar targets.

5 PUTTING IT ALL TOGETHER

To make APOLLO applicable to both training and inference scenarios, we have to provide promising performance while guaranteeing lightweight compilation overhead. To achieve this, we complement APOLLO through the following steps.

Auto-Tuning. APOLLO's versatility of loop transforma-tions resides in the polyhedral engine. Long compilation time is still inevitable in AKG, but APOLLO addresses its scalability issue. and thus always produces an effective so-lution within seconds (or minutes in the worst case), which is further optimized by the piecewise compilation strategy as will be introduced later. We also offer an auto-tuning strategy. During the tuning process, APOLLO may cap-ture the composition patterns of *op*s that are prevented from parallelization by Cost Model (3). The developers can use such information to design new (complementary) rules.

Piecewise Compilation. Piecewise compilation achieves the faster compilation of multiple micro-/sub-graphs. It is devised to accelerate the compilation overhead of APOLLO when its complexity becomes an issue. Piecewise compila-tion along each red/violet arrow shown in Fig.2. It further reduces the compilation overhead by allowing the polyhe-dral model to solve the ILP problems in a much shorter period. Each compilation process can make its own de-cision on target-specific optimization trade-offs. The syn-chronization of piecewise compilation happens along the green dash-dot lines in Fig.2.

Code Generation. APOLLO generates CUDA code on NVIDIA V100 GPU and so-called CCE code executable on Huawei Ascend 910 (Liao et al., 2021). CUDA code generation under the polyhedral model has a history of near a decade: the backends of PPCG (Verdoolaege et al., 2013) and TC (Vasilache et al., 2019) can be borrowed to sup-port GPU devices. The automatic memory management for both traditional memory hierarchy pyramid on GPU and emerging multi-level, multi-directional memory hier-archy on Ascend 910 chips is guaranteed by *isl* (integer set library) (Verdoolaege, 2010). Some complementary code optimization strategies including making use of SIMD hardware intrinsics and low-level optimization of synchro-nization between emitted instructions that are beyond the ability of the polyhedral model are also integrated, further improving the performance of the generated code.

Example 8. *We finally recall Fig.1 to illustrate the end-to-end compilation of APOLLO. The partition phase first extracts $\mathcal{P}$ from the computational graph and obtains two sub-graphs, $\mathcal{F}_1$ and $\mathcal{F}_2$. $\mathcal{F}_1$ is split into three micro-graphs, the third one of which, $\mathcal{G}_3$, contains a compute-bound op considered for fusion. This was not considered by graph compilers. Similarly, $\mathcal{F}_2$ is also set apart by the partition phase into two micro-graphs. As op_{71} is a primitive re-duction op, it has be separated from its original compound op op_7. Our rules in Table 1 allow it to be considered as fusible with its preceding element-wise op (op_6). As a re-sult, we obtain $\mathcal{G}_4$ and $\mathcal{G}_5$. Layer I of the fusion phase is then used to search fusion plans within each $\mathcal{G}_y$ in a short period. The output is then delivered to Layer II, which per-forms memory stitching between $\mathcal{G}_1/\mathcal{G}_2$ and $\mathcal{G}_3$, and $\mathcal{G}_4$ and $\mathcal{G}_5$, respectively. Finally, Layer III exploits the parallelism between $\mathcal{G}_1$ and $\mathcal{G}_3$, with the parallelism between $\mathcal{F}_1$ and $\mathcal{F}_2$ also studied. The auto-tuning process happening dur-ing the compilation will trigger the feedback if any.*

6 RESULTS

APOLLO is implemented within MindSpore (Huawei, 2020). It is written in $14.4k$ lines of C++ and $2k$ lines of Python. The code is accessible at https://gitee.com/mindspore/mindspore, with APOLLO enabled when the parameter enable_graph_kernel is set true. We evaluate the performance using training workloads on V100 GPUs and Ascend 910 chips. Preliminary results of inference workloads are collected to demonstrate the general applicability. The code executed on GPU is compiled using CUDA Toolkit 10.1 with *-O3* enabled, and the CCE code is compiled using a native compiler of Ascend chips. The geometric mean of 10 executions is reported.

6.1 Sub-graph Case Study

We use some sub-graphs extracted from BERT (Devlin et al., 2019) to illustrate the effect of each layer on GPU. The result on Ascend chips is similar. Note that the scalability issue is first challenged by tiled reductions; we resolve it by decomposing a sub-graph into micro-graphs and use Eg. 3 to demonstrate the effectiveness of APOLLO. We now use another experiment to validate APOLLO can also address the cases caused by opaque *op*s.

Fig. 4 is a sub-graph obtained by AKG. The reshaping *op* transforms a 1D vector into a 2D matrix that is used by the follow-up multiplication *op*. Due to the mismatching between loop dimensions, the fusion heuristic of the polyhedral model fails to find a fusion plan for this sub-graph within 30 minutes. Such a reshaping *op* is considered as an opaque *op*. We find no rules in Table 1 to construct a sub-graph for it; APOLLO thus never delivers such a subgraph to its fusion phase. Instead, it returns a fusion plan as the red boxes within one second. Note that the removal of a rule that aggregates a reshaping *op* with its follow-up element-wise *op*s is triggered by the feedback, which results in the current set of rules in Table 1.

Figure 4: A case study for the scalability issue.

Fig. 5 is used to study the effect of Layer II. Each micro-graph is represented using a red dotted box. Like Eg. 5, each pair of micro-graphs is separated by a reduction *op*. One can apply Algo.2 to each micro-graph pair. The tensors passed by each blue edge can be promoted to GPU shared memory/registers. The sub-graphs shown in Fig.6 are used to illustrate the effect of Layer III, with the branches that can be executed in parallel encompassed by boxes with the same color. Fig.6(a) is a multi-head case, for which Algo.3 considers all branches as parallelizable. Fig.6(b) has been explained in Eg.6.

Fig.6(c) is a mixture of multi-head and multi-tail scenarios,
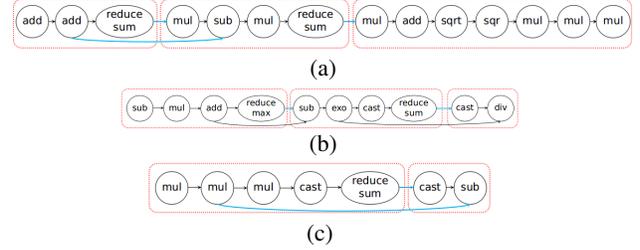
Figure 5: Examples for effect of memory stitching.
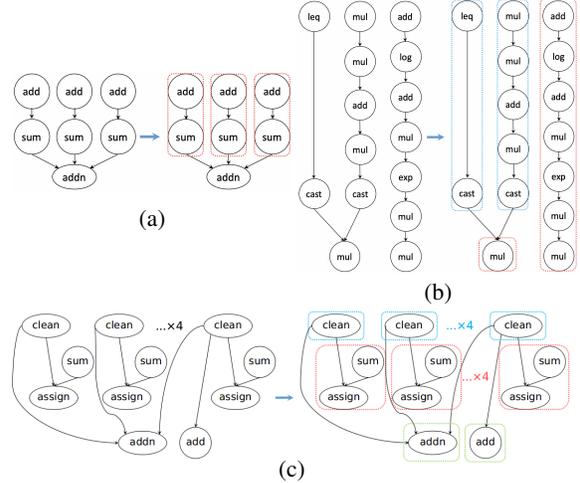
Figure 6: Examples for effect of parallelism stitching.

composed of 7 branches. A group of parallelizable candidates composed of *clean op*s is obtained, as they have a common tail–*addn*. *addn* and *add* are then stitched by introducing a virtual common tail, followed by a final round of invocation of Algo.3 on 7 independent sub-graphs, each of which composed of *assign* and *sum*. Algo. 3 then performs parallelism stitching as explained in Eg.7.

The performance of these sub-graphs is shown in Table 3. The bold percentage numbers are improvements (*imp.*) over MindSpore (MS), and the red ones are improvements of APOLLO (full) over the (partial) version with Layer II or III disabled. Layer II achieves an average improvement of 69% for the sub-graphs in Fig. 5 by allocating the results generated by micro-graphs on GPU shared memory and registers, and Layer III provides a mean speedup of 92% for the sub-graphs in Fig. 6. The benefit brought by parallelism stitching is marginal when the number of parallelizable candidates is relative smaller (Fig.6(a)); it grows with the increase of the later (Fig.6(b) and 6(c)).

Table 3: Execution time (in μs) of sub-graphs.

cases	MS	disabling	partial (*imp.*)	full (*imp.*)
Fig.5(a)	527	Layer II	224 (**135%**)	98 (**438%**, **129%**)
Fig.5(b)	2018	Layer II	342 (**490%**)	221 (**813%**, **56%**)
Fig.5(c)	1532	Layer II	233 (**558%**)	190 (**706%**, **23%**)
Fig.6(a)	249.8	Layer III	87.6 (**185%**)	85.6 (**192%**, **2%**)
Fig.6(b)	86.1	Layer III	4.46 (**1830%**)	3.34 (**2478%**, **34%**)
Fig.6(c)	97.6	Layer III	26.7 (**266%**)	7.85 (**1143%**, **240%**)

6.2 Results on Single GPU

We evaluate APOLLO using BERT, Transformer (Vaswani et al., 2017), Wide&Deep (Cheng et al., 2016), the Dark-Net of Yolo-v3 (Redmon et al., 2016) and DeepFM (Guo et al., 2017). Except BERT that uses mixed precision (FP16&FP32), the remaining four workloads are trained using single precision (FP32). We compare the throughput with TensorFlow version 1.15 and XLA for all workloads except Wide&Deep, the performance of which optimized by XLA suffers from significant degradation when experimenting using TensorFlow version 1.15. We thus implement this model using TensorFlow version 2.0. The data is collected in Table 4. We report sentences/s for BERT with 12 encoder layers (BT-base), tokens/s for Transformer (TR), samples/s for Wide&Deep (WD) and DeepFM (FM) and images/s for Yolo-v3 (YO), with two batch sizes (*b.s.*) considered for each model. The percentage numbers are improvements of XLA over TensorFlow (TF) and APOLLO over MS. The (red) improvement (*imp.*) of APOLLO over XLA is listed in the rightmost column.

Table 4: Throughput on single GPU.

models	*b.s.*	TF	XLA/TF	MS	APOLLO/MS	*imp.*
BT-base	32	167	**105%**	135	**252%**	**39%**
	64	200.8	**129%**	183.6	**212%**	**23%**
TR	8	6750	**16%**	5122	**84%**	**20%**
	16	9500	**11%**	10868	**59%**	**64%**
WD	16000	1133696	**15%**	762086	**123%**	**48%**
	32000	1470221	**5%**	836820	**121%**	**20%**
YO	4	33.11	**15%**	39.48	**46%**	**51%**
	8	56.00	**12%**	75.01	**10%**	**31%**
FM	8192	26117	**-1%**	479744	**151%**	**-**
	16384	30279	**-2%**	543024	**167%**	**-**

APOLLO outperforms MindSpore by 2.23× on average. Layer I contributes most by 1.91× to the overall performance. The remaining improvements are from Layer II and III. For instance, the execution time of BERT is reduced from 215 ms/step to 93.5 ms/step when only Layer I is enabled, and further reduced to 88.1 ms/step when Layer II is also turned on. It eventually falls down to 82.4 ms/step when Layer III is also enabled. However, Layer I cannot finish its compilation within a reasonable time for training scenarios without the cooperation between Layer II/III and the partition phase. Addressing its scalability issue using our work maximizes the power of the polyhedral engine.

XLA outperforms TensorFlow by 1.31×. The network architectures of each workload expressed using TensorFlow and MindSpore can be considered as identical and have the same convergence properties, except that MindSpore introduces a specific *inplace_assign op* somewhere in a network that results in the performance difference between the two frameworks. They both are backed by cuBLAS (Nvidia, 2013) and cuDNN (Chetlur et al., 2014). The implementations of DeepFM greatly differ from each other, lead-

ing to a much better throughput of MindSpore than TensorFlow. The optimized code of XLA suffers from slight performance degradation. We have no ideas of the reasons. APOLLO improves the MindSpore implementation of DeepFM by 2.51× and 2.67× for the two batch sizes. This example is excluded when computing the average number, and APOLLO still helps MindSpore outperform TensorFlow and XLA by 1.86× and 1.37×. The superiority to XLA is due to two reasons. First, APOLLO performs each fusion strategy considered by XLA. Second, APOLLO also perfectly models the fusion of compute-bound *op*s not considered by XLA without resulting in the scalability issue.

We also use APOLLO to optimize the workloads in the model zoo of MindSpore, and the performance comparison is collected in Fig.7. We report the execution time for a single training epoch. These workloads cover many application domains, including computer vision, speech recognition, natural language processing (NLP), recommendation system and neural search architecture. Some of them and the BERT version that will be used in §6.3 have the same architectures as those in the MLPerf benchmarks (Reddi et al., 2020). We did not compare the performance with any frameworks or compilers due to the missing of official implementations using TensorFlow or Pytorch. However, we will complement the comparison with other tools in the future. We report the results of these workloads to illustrate the general applicability of APOLLO. On average, APOLLO produces an improvement of 29.6% over MindSpore.
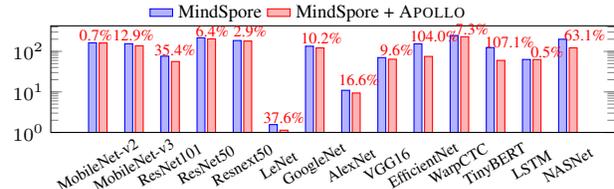

Figure 7: Execution times of MindSpore's model zoo (*y* axis: log scaled time in ms; lower is better).

Unlike MobileNet-v3 (Howard et al., 2017), There exist many consecutive *conv op*s in MobileNet-v2 which APOLLO does not fuse. The number of consecutive primitive *op*s that can be fused by Layer I is also much smaller. Most of the execution time of the LSTM model (Hochreiter & Schmidhuber, 1997) is consumed by a so-called compound LSTM *op*, which has not yet been decomposed into primitive *op*s and thus cannot be fused. These two models thus observe less performance improvement than others.

6.3 Results on Multiple GPUs

We also conduct experiments on multiple GPUs. MindSpore does not provide the implementation of the DarkNet of Yolo-v3 on multiple GPUs. We experiment using the remaining four workloads evaluated in Table 4, the result

of which is collected in Table 5. Besides, we also consider another version of BERT with 24 encoder layers (BERT-large). The batch sizes are put in parenthesis, and the number of used GPUs is listed in the second column.

Table 5: Throughput on multiple GPUs.

models	GPUs	TF	XLA/TF	MS	APOLLO/MS	imp.
BT-base(32)	8	1244.9	**96%**	944.4	**247%**	**34%**
BT-base(64)	8	1555.4	**117%**	1333.1	**222%**	**27%**
BT-large(4)	4	66.94	**33%**	37.62	**133%**	**-2%**
WD(16000)	8	8086178	**1%**	4964319	**87%**	**13%**
FM(16384)	4	31767	**-7%**	2117685	**130%**	**-**

The results of multiple GPUs follow the single GPU case, and the reasons of the performance improvement are the same. The two differences are listed below. First, the performance degradation (7%) of DeepFM is worsened on multiple GPUs. Second, the throughput of our approach falls behind that of XLA by 2% for BERT-large, due to the introduced *inplace_assign op*s by MindSpore. MindSpore thus suffers from a much more serve degradation by 44% than TensorFlow. We obtain a mean speedup of 2.64× over MindSpore, the performance of which falls behind those of TensorFlow and XLA. With APOLLO, MindSpore outperforms TensorFlow by 1.96× and XLA by 1.18×.

We report the data for each single epoch, and a complete training pass behaviors similarly. For instance, the overall training time of BERT-base on 8 GPUs is decreased from 342 hours to 106 hours when given batch size 64 and epoch number 40. Other workloads observe similar results.

6.4 Results on Ascend 910 Chips

The throughputs and improvements of BERT and PanGu-α (Zeng et al., 2021) on Ascend 910 chips are shown in Fig. 8, with both single and multiple chips considered. PanGu-α is the Chinese-language equivalent of GPT-3 (Radford et al., 2018), which solves many NLP problems without fine tuning. Their exist three different model configurations of PanGu-α, of which we use the 2.6B version with 32 layers. It is trained using mixed precision (FP16&FP32). We did not consider TensorFlow or XLA here since they are not tailored to this accelerator.
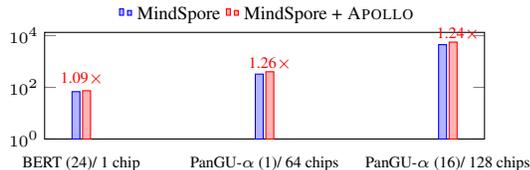
Figure 8: Throughput of BERT and PanGu-α (examples/s) on Ascend. Batch sizes are in parentheses. Higher is better.

The off-chip data movement on Ascend chips is more expensive than the GPU case, which can be optimized through the fusion optimization. The vendor libraries for Ascend chips, however, cannot support the fused custom *op*s, the performance of which thus falls behind that of APOLLO. Our framework provides a mean improvement of 19.7% over MindSpore that is backed by vendor libraries.

The difference between Ascend and GPU lies in the memory system: the multi-level, multi-directional memory hierarchy of Ascend chips requires a complicated data-flow management, which has been facilitated by AKG. APOLLO goes beyond AKG from two aspects. First, APOLLO exploits the parallelism between independent *op*s that executed by the *vector* unit of the Ascend chip. Second, the decoupled handling of reduction *op*s in Layer I and II avoids the need for the computation of large loop shifting factors, which AKG leverages to achieve fusion in some cases.

6.5 Compilation Overhead

We now report the compilation overhead of APOLLO in Table 6. We only list the data of four models due to the limited space, each of which represents the compilation time when generating code for GPU except BERT. The compilation time of MindSpore is reported as a baseline. APOLLO can always generate code within seconds/minutes for training scenarios, validating its JIT compilation efficiency. MindSpore always consumes less compilation time since it does not exploit fusion between *op*s.

Table 6: Compilation overhead in seconds.

Workloads	MS	APOLLO	Workloads	MS	APOLLO
BT(24)/Ascend	186.206	237.691	WD(16000)	2.3	10.01
TR(16)	82.29	188.83	YO(8)	10.06	31.93

6.6 Results on Inference Workloads

We use Wide&Deep and Yolo-v3 to demonstrate the effectiveness of APOLLO for inference scenarios on GPU. We also consider an internal inference workload called EPP-MVSNet used for the 3D reconstruction of real objects, the architecture of which can be retrieved from the project repository of MindSpore (Huawei, 2020). Table 7 shows the experimental results on inference workloads. As a DL framework under development, MindSpore supports the training of a DNN model but does not excel at expressing inference workloads. The front-ends of existing optimizing compilers (Chen et al., 2018a; Ma et al., 2020) currently cannot import an inference workload expressed using MindSpore. These compilers are thus not consider for comparison in this experiment. One can observe that APOLLO can still obtain superior performance on inference workloads.

Table 7: Throughput of inference workloads on GPU.

Workloads	b.s.	Throughput	MS	APOLLO	imp.
Wide&Deep	16000	sample/s	629349.5	695229.4	**10.5%**
Yolo-v3	32	images/s	19.2	19.7	**2.6%**
EPPMVSNet	1	pictures/s	1.42	2.32	**63.4%**

7 RELATED WORK

The first challenge faced by polyhedral compilers (Vasilache et al., 2019; Zhao et al., 2021) is the scalability issue of loop fusion, which has been proved as NP-complete (Darte, 2000). Similar recent work (Acharya et al., 2020; Mehta et al., 2014) tried to address this problem but did not study the impact of tiled reductions. Our framework resolves this issue.

The recent work DNNFusion (Niu et al., 2021) also classifies *op*s by employing similar rules to those in Table 1 for graph rewriting, but it did not distinguish primitive *op*s and compound *op*s. In other words, DNNFusion did not consider the complementary rules in Table 2 that model the fusion of reductions with its follow-up *op*s. DNNFusion may still result in the scalability issue when coupled with a polyhedral loop optimizer for tensors, though it is scalable when working with manual scheduling approaches.

Coarser grained kernel fusion was also studied by Ashari et al. (2015) and Sivathanu et al. (2019). The former exploited graph-level fusion for a specific computation pattern, while the latter only considered the fusion between *matmul* and element-wise *op*s. APOLLO covers a broader set of fusion patterns than these approaches.

Combining independent *op*s to saturate the target device is also exploited by APOLLO. Prior work referred to this technique as stitching (Zheng et al., 2020) or horizontal fusion (Li et al., 2020); TensorRT (Nvidia, 2016) also falls into this category. However, these approaches only target inference scenarios on GPU, and some of them (Jia et al., 2019b;a) only consider the parallelism between *op*s with the same type. Our work does not have such limitations and is with a stronger ability to support *op*s in training scenarios than Rammer (Ma et al., 2020).

The recent fusion framework Deepcuts (Jung et al., 2021) also considers training scenarios, but it only targets a single GPU. APOLLO goes beyond Deepcuts by supporting multiple GPU devices and a domain-specific accelerator. Nimble (Kwon et al., 2020) considers both training and inference workloads, but it uses the ahead-of-time (AOT) compilation approach. As a summary, we finally compare the closest related graph compilers in Table 8.

Table 8: Comparison with closest related graph compilers

graph compilers	memory stitching	parallelism stitching	training support	compilation approach
TVM's	×	×	×	AOT
XLA	×	✓	✓	JIT/AOT
TASO	×	✓	×	AOT
FusionStitching	✓	×	×	AOT
Rammer	×	✓	×	AOT
Nimble	×	×	✓	AOT
DNNFusion	×	✓	×	AOT
Deepcuts	×	✓	✓	AOT
APOLLO	✓	✓	✓	JIT

8 CONCLUSION

In this work, we unified loop fusion, memory stitching and parallelism stitching and conducted extensive experiments. The proposed framework outperforms prior work by exploiting both locality and parallelism, and it made progress by breaking the *op* boundaries, leading to a much wider fusion space. The picewise compilation strategy makes APOLLO suitable for both training and inference scenarios.

APOLLO leverages AKG as its code generator in Layer I. However, APOLLO made significant contributions over AKG from many aspects. First, APOLLO addresses the scalability issue of AKG. As mentioned in §4.1, long compilation time still exists in AKG due to the presence of tiled reduction *op*s in a sub-graph. Second, AKG only exploits data locality using polyhedral loop fusion heuristics, but APOLLO also enables parallelism stitching in Layer III, further improving the performance of the generated code. Finally, the upward feedback enabled by APOLLO may trigger the updates the rules in Table 1 and 2. This makes it possible to better manage the black-box loop transformations performed by the polyhedral model, allowing users to better interact with the framework.

APOLLO suffers from two limitations. First, the aggregation scenarios covered by Table 1 and 2 are incomplete, but they are sufficient to handle the DNN workloads we have ever seen. One can easily complement the rule set. Second, Cost Model (3) is still rather simple, which roughly models the trade-off between parallelism and communication/synchronization. The multi-layer fusion phase, however, still made progress over prior work. We leave these tasks as future work. Nonetheless, our work still offers insight into DNN compiler design and implementation.

REFERENCES

Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., Kudlur, M., Levenberg, J., Monga, R., Moore, S., Murray, D. G., Steiner, B., Tucker, P., Vasudevan, V., Warden, P., Wicke, M., Yu, Y., and Zheng, X. Tensorflow: A system for large-scale machine learning. In *Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation*, OSDI'16, pp. 265–283, Berkeley, CA, USA, 2016. USENIX Association. ISBN 978-1-931971-33-1. URL http://dl.acm.org/citation.cfm?id=3026877.3026899.

Acharya, A., Bondhugula, U., and Cohen, A. Effective loop fusion in polyhedral compilation using fusion conflict graphs. *ACM Trans. Archit. Code Optim.*, 17(4), September 2020. ISSN 1544-3566. doi: 10.1145/3416510. URL https://doi.org/10.1145/3416510.

Amodei, D., Ananthanarayanan, S., Anubhai, R., Bai, J., Battenberg, E., Case, C., Casper, J., Catanzaro, B., Cheng, Q., Chen, G., Chen, J., Chen, J., Chen, Z., Chrzanowski, M., Coates, A., Diamos, G., Ding, K., Du, N., Elsen, E., Engel, J., Fang, W., Fan, L., Fougner, C., Gao, L., Gong, C., Hannun, A., Han, T., Johannes, L., Jiang, B., Ju, C., Jun, B., LeGresley, P., Lin, L., Liu, J., Liu, Y., Li, W., Li, X., Ma, D., Narang, S., Ng, A., Ozair, S., Peng, Y., Prenger, R., Qian, S., Quan, Z., Raiman, J., Rao, V., Satheesh, S., Seetapun, D., Sengupta, S., Srinet, K., Sriram, A., Tang, H., Tang, L., Wang, C., Wang, J., Wang, K., Wang, Y., Wang, Z., Wang, Z., Wu, S., Wei, L., Xiao, B., Xie, W., Xie, Y., Yogatama, D., Yuan, B., Zhan, J., and Zhu, Z. Deep speech 2 : End-to-end speech recognition in english and mandarin. In Balcan, M. F. and Weinberger, K. Q. (eds.), *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 173–182, New York, New York, USA, 20–22 Jun 2016. PMLR. URL http://proceedings.mlr.press/v48/amodei16.html.

Ashari, A., Tatikonda, S., Boehm, M., Reinwald, B., Campbell, K., Keenleyside, J., and Sadayappan, P. On optimizing machine learning workloads via kernel fusion. In *Proceedings of the 20th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, PPoPP 2015, pp. 173–182, New York, NY, USA, 2015. Association for Computing Machinery. ISBN 9781450332057. doi: 10.1145/2688500.2688521. URL https://doi.org/10.1145/2688500.2688521.

Baghdadi, R., Ray, J., Romdhane, M. B., Del Sozzo, E., Akkas, A., Zhang, Y., Suriana, P., Kamil, S., and Amarasinghe, S. Tiramisu: A polyhedral compiler for expressing fast and portable code. In *Proceedings of the 2019 IEEE/ACM International Symposium on Code Generation and Optimization*, CGO 2019, pp. 193–205, Piscataway, NJ, USA, 2019. IEEE Press. ISBN 978-1-7281-1436-1. URL http://dl_acm.gg363.site/citation.cfm?id=3314872.3314896.

Bondhugula, U., Hartono, A., Ramanujam, J., and Sadayappan, P. A practical automatic polyhedral parallelizer and locality optimizer. In *Proceedings of the 29th ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI'08, pp. 101–113, New York, NY, USA, 2008. ACM. ISBN 978-1-59593-860-2. doi: 10.1145/1375581.1375595. URL http://doi.acm.org/10.1145/1375581.1375595.

Chen, T., Moreau, T., Jiang, Z., Zheng, L., Yan, E., Cowan, M., Shen, H., Wang, L., Hu, Y., Ceze, L., Guestrin, C., and Krishnamurthy, A. Tvm: An automated end-to-end optimizing compiler for deep learning. In *Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation*, OSDI'18, pp. 579–594, Berkeley, CA, USA, 2018a. USENIX Association. ISBN 978-1-931971-47-8. URL http://dl.acm.org/citation.cfm?id=3291168.3291211.

Chen, T., Zheng, L., Yan, E., Jiang, Z., Moreau, T., Ceze, L., Guestrin, C., and Krishnamurthy, A. Learning to optimize tensor programs. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NIPS'18, pp. 3393–3404, Red Hook, NY, USA, 2018b. Curran Associates Inc.

Cheng, H.-T., Koc, L., Harmsen, J., Shaked, T., Chandra, T., Aradhye, H., Anderson, G., Corrado, G., Chai, W., Ispir, M., Anil, R., Haque, Z., Hong, L., Jain, V., Liu, X., and Shah, H. Wide & deep learning for recommender systems. In *Proceedings of the 1st Workshop on Deep Learning for Recommender Systems*, DLRS 2016, pp. 7–10, New York, NY, USA, 2016. Association for Computing Machinery. ISBN 9781450347952. doi: 10.1145/2988450.2988454. URL https://doi.org/10.1145/2988450.2988454.

Chetlur, S., Woolley, C., Vandermersch, P., Cohen, J., Tran, J., Catanzaro, B., and Shelhamer, E. cudnn: Efficient primitives for deep learning, 2014.

Cho, M., Finkler, U., Kung, D., and Hunter, H. Blueconnect: Decomposing all-reduce for deep learning on heterogeneous network hierarchy. In Talwalkar, A., Smith, V., and Zaharia, M. (eds.), *Proceedings of Machine Learning and Systems*, volume 1, pp. 241–251, 2019. URL https://proceedings.mlsys.org/paper/2019/file/9b8619251a19057cff70779273e95aa6-Paper.pdf.

Cyphers, S., Bansal, A. K., Bhiwandiwalla, A., Bobba, J., Brookhart, M., Chakraborty, A., Constable, W., Convey, C., Cook, L., Kanawi, O., Kimball, R., Knight, J., Korovaiko, N., Kumar, V., Lao, Y., Lishka, C. R., Menon, J., Myers, J., Narayana, S. A., Procter, A., and Webb, T. J. Intel ngraph: An intermediate representation, compiler, and executor for deep learning, 2018.

Darte, A. On the complexity of loop fusion. *Parallel Computing*, 26(9):1175–1193, 2000. ISSN 0167-8191. doi: https://doi.org/10.1016/S0167-8191(00)00034-X. URL https://www.sciencedirect.com/science/article/pii/S016781910000034X.

Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1423. URL https://www.aclweb.org/anthology/N19-1423.

Ding, Y., Zhu, L., Jia, Z., Pekhimenko, G., and Han, S. Ios: Inter-operator scheduler for cnn acceleration. In Smola, A., Dimakis, A., and Stoica, I. (eds.), *Proceedings of Machine Learning and Systems*, volume 3, pp. 1–14, 2021. URL

https://proceedings.mlsys.org/paper/2021/file/38b3eff8baf56627478ec76a704e9b52-Paper.pdf.

Feautrier, P. and Lengauer, C. *Polyhedron Model*, pp. 1581–1592. Springer US, Boston, MA, 2011. ISBN 978-0-387-09766-4. doi: 10.1007/978-0-387-09766-4_502. URL https://doi.org/10.1007/978-0-387-09766-4_502.

Google. Xla: Optimizing compiler for machine learning, 2017. URL https://www.tensorflow.org/xla.

Guo, H., TANG, R., Ye, Y., Li, Z., and He, X. Deepfm: A factorization-machine based neural network for ctr prediction. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pp. 1725–1731, 2017. doi: 10.24963/ijcai.2017/239. URL https://doi.org/10.24963/ijcai.2017/239.

Hochreiter, S. and Schmidhuber, J. Long Short-Term Memory. *Neural Computation*, 9(8):1735–1780, 11 1997. ISSN 0899-7667. doi: 10.1162/neco.1997.9.8.1735. URL https://doi.org/10.1162/neco.1997.9.8.1735.

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, 2017.

Huawei. Mindspore, 2020. URL https://www.mindspore.cn/en.

Irigoin, F. and Triolet, R. Supernode partitioning. In *Proc. of the 15th ACM SIGPLAN-SIGACT Symp. on Principles of Programming Languages*, POPL'88, pp. 319–329, New York, NY, USA, 1988. ACM. ISBN 0-89791-252-7. doi: 10.1145/73560.73588. URL http://doi.acm.org/10.1145/73560.73588.

Jangda, A. and Bondhugula, U. An effective fusion and tile size model for polymage. *ACM Trans. Program. Lang. Syst.*, 42(3), November 2020. ISSN 0164-0925. doi: 10.1145/3404846. URL https://doi.org/10.1145/3404846.

Jia, Z., Padon, O., Thomas, J., Warszawski, T., Zaharia, M., and Aiken, A. Taso: Optimizing deep learning computation with automatic generation of graph substitutions. In *Proceedings of the 27th ACM Symposium on Operating Systems Principles*, SOSP'19, pp. 47–62, New York, NY, USA, 2019a. ACM. ISBN 9781450368735. doi: 10.1145/3341301.3359630. URL https://doi.org/10.1145/3341301.3359630.

Jia, Z., Thomas, J., Warszawski, T., Gao, M., Zaharia, M., and Aiken, A. Optimizing dnn computation with relaxed graph substitutions. In Talwalkar, A., Smith, V., and Zaharia, M. (eds.), *Proceedings of Machine Learning and Systems*, volume 1, pp. 27–39, 2019b. URL https://proceedings.mlsys.org/paper/2019/file/b6d767d2f8ed5d21a44b0e5886680cb9-Paper.pdf.

Jung, W., Dao, T. T., and Lee, J. Deepcuts: A deep learning optimization framework for versatile gpu workloads. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation*, PLDI 2021, pp. 190–205, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450383912. doi: 10.1145/3453483.3454038. URL https://doi.org/10.1145/3453483.3454038.

Kwon, W., Yu, G.-I., Jeong, E., and Chun, B.-G. Nimble: Lightweight and parallel gpu task scheduling for deep learning. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M. F., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 8343–8354. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper/2020/file/5f0ad4db43d8723d18169b2e4817a160-Paper.pdf.

Lattner, C., Amini, M., Bondhugula, U., Cohen, A., Davis, A., Pienaar, J., Riddle, R., Shpeisman, T., Vasilache, N., and Zinenko, O. Mlir: Scaling compiler infrastructure for domain specific computation. In *2021 IEEE/ACM International Symposium on Code Generation and Optimization (CGO)*, pp. 2–14, 2021. doi: 10.1109/CGO51591.2021.9370308.

Li, A., Zheng, B., Pekhimenko, G., and Long, F. Automatic horizontal fusion for gpu kernels, 2020.

Liao, H., Tu, J., Xia, J., Liu, H., Zhou, X., Yuan, H., and Hu, Y. Ascend: a scalable and unified architecture for ubiquitous deep neural network computing : Industry track paper. In *2021 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, pp. 789–801, 2021. doi: 10.1109/HPCA51647.2021.00071.

Ma, L., Xie, Z., Yang, Z., Xue, J., Miao, Y., Cui, W., Hu, W., Yang, F., Zhang, L., and Zhou, L. Rammer: Enabling holistic deep learning compiler optimizations with rtasks. In *14th USENIX Symposium on Operating Systems Design and Implementation (OSDI 20)*, pp. 881–897. USENIX Association, November 2020. ISBN 978-1-939133-19-9. URL https://www.usenix.org/conference/osdi20/presentation/ma.

McKinley, K. S., Carr, S., and Tseng, C.-W. Improving data locality with loop transformations. *ACM Trans. Program. Lang. Syst.*, 18(4):424–453, July 1996. ISSN 0164-0925. doi: 10.1145/233561.233564. URL https://doi.org/10.1145/233561.233564.

Mehta, S., Lin, P.-H., and Yew, P.-C. Revisiting loop fusion in the polyhedral framework. In *Proceedings of the 19th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, PPoPP'14, pp. 233–246, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-2656-8. doi: 10.1145/2555243.2555250. URL http://doi.acm.org/10.1145/2555243.2555250.

Niu, W., Guan, J., Wang, Y., Agrawal, G., and Ren, B. Dnnfusion: Accelerating deep neural networks execution with advanced operator fusion. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation*, PLDI 2021, pp. 883–898, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450383912. doi: 10.1145/3453483.3454083. URL https://doi.org/10.1145/3453483.3454083.

Nvidia. cublas, 2013. URL https://developer.nvidia.com/cublas.

Nvidia. Nvidia tensorrt, 2016. URL https://developer.nvidia.com/tensorrt.

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. Pytorch: An imperative style, high-performance deep learning library. In *Advances in neural information processing systems*, pp. 8026–8037, 2019.

Radford, A., Narasimhan, K., Salimans, T., and Sutskever, I. Improving language understanding by generative pre-training. 2018.

Ragan-Kelley, J., Barnes, C., Adams, A., Paris, S., Durand, F., and Amarasinghe, S. Halide: A language and compiler for optimizing parallelism, locality, and recomputation in image processing pipelines. In *Proceedings of the 34th ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI'13, pp. 519–530, New York, NY, USA, 2013. ACM. ISBN 978-1-4503-2014-6. doi: 10.1145/2491956.2462176. URL http://doi.acm.org/10.1145/2491956.2462176.

Reddi, V. J., Cheng, C., Kanter, D., Mattson, P., Schmuelling, G., Wu, C.-J., Anderson, B., Breughe, M., Charlebois, M., Chou, W., Chukka, R., Coleman, C., Davis, S., Deng, P., Diamos, G., Duke, J., Fick, D., Gardner, J. S., Hubara, I., Idgunji, S., Jablin, T. B., Jiao, J., John, T. S., Kanwar, P., Lee, D., Liao, J., Lokhmotov, A., Massa, F., Meng, P., Micikevicius, P., Osborne, C., Pekhimenko, G., Rajan, A. T. R., Sequeira, D., Sirasao, A., Sun, F., Tang, H., Thomson, M., Wei, F., Wu, E., Xu, L., Yamada, K., Yu, B., Yuan, G., Zhong, A., Zhang, P., and Zhou, Y. Mlperf inference benchmark, 2020.

Redmon, J., Divvala, S., Girshick, R., and Farhadi, A. You only look once: Unified, real-time object detection. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 779–788, June 2016. doi: 10.1109/CVPR.2016. 91. URL https://www.computer.org/10.1109/CVPR.2016.91.

Rotem, N., Fix, J., Abdulrasool, S., Catron, G., Deng, S., Dzhabarov, R., Gibson, N., Hegeman, J., Lele, M., Levenstein, R., Montgomery, J., Maher, B., Nadathur, S., Olesen, J., Park, J., Rakhov, A., Smelyanskiy, M., and Wang, M. Glow: Graph lowering compiler techniques for neural networks, 2019.

Sivathanu, M., Chugh, T., Singapuram, S. S., and Zhou, L. Astra: Exploiting predictability to optimize deep learning. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, ASPLOS'19, pp. 909–923, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450362405. doi: 10.1145/3297858.3304072. URL https://doi.org/10.1145/3297858.3304072.

Vasilache, N., Zinenko, O., Theodoridis, T., Goyal, P., Devito, Z., Moses, W. S., Verdoolaege, S., Adams, A., and Cohen, A. The next 700 accelerated layers: From mathematical expressions of network computation graphs to accelerated gpu kernels, automatically. *ACM Trans. Archit. Code Optim.*, 16(4), October 2019. ISSN 1544-3566. doi: 10.1145/3355606. URL https://doi.org/10.1145/3355606.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, u., and Polosukhin, I. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17,

pp. 6000–6010, Red Hook, NY, USA, 2017. Curran Associates Inc. ISBN 9781510860964.

Verdoolaege, S. Isl: An integer set library for the polyhedral model. In *Proceedings of the Third International Congress Conference on Mathematical Software*, ICMS'10, pp. 299–302, Berlin, Heidelberg, 2010. Springer-Verlag. ISBN 3-642-15581-2, 978-3-642-15581-9. URL https://doi.org/10.1007/978-3-642-15582-6_49.

Verdoolaege, S. and Janssens, G. Scheduling for ppcg. *Report CW*, 706, 2017.

Verdoolaege, S., Carlos Juega, J., Cohen, A., Ignacio Gómez, J., Tenllado, C., and Catthoor, F. Polyhedral parallel code generation for cuda. *ACM Trans. Archit. Code Optim.*, 9 (4):54:1–54:23, January 2013. ISSN 1544-3566. doi: 10. 1145/2400682.2400713. URL http://doi.acm.org/10.1145/2400682.2400713.

Wei, R., Schwartz, L., and Adve, V. Dlvm: A modern compiler infrastructure for deep learning systems, 2018.

Zeng, W., Ren, X., Su, T., Wang, H., Liao, Y., Wang, Z., Jiang, X., Yang, Z., Wang, K., Zhang, X., Li, C., Gong, Z., Yao, Y., Huang, X., Wang, J., Yu, J., Guo, Q., Yu, Y., Zhang, Y., Wang, J., Tao, H., Yan, D., Yi, Z., Peng, F., Jiang, F., Zhang, H., Deng, L., Zhang, Y., Lin, Z., Zhang, C., Zhang, S., Guo, M., Gu, S., Fan, G., Wang, Y., Jin, X., Liu, Q., and Tian, Y. Pangu-α: Large-scale autoregressive pretrained chinese language models with auto-parallel computation, 2021.

Zhao, J. and Cohen, A. Flextended tiles: A flexible extension of overlapped tiles for polyhedral compilation. *ACM Trans. Archit. Code Optim.*, 16(4), December 2019. ISSN 1544-3566. doi: 10.1145/3369382. URL https://doi.org/10.1145/3369382.

Zhao, J. and Di, P. Optimizing the memory hierarchy by compositing automatic transformations on computations and data. In *Proceedings of the 53rd IEEE/ACM International Symposium on Microarchitecture*, MICRO-53, pp. 427–441, Piscataway, NJ, USA, 2020. IEEE Press. doi: 10.1109/MICRO50266.2020.00044. URL https://www.microarch.org/micro53/papers/738300a427.pdf.

Zhao, J., Li, B., Nie, W., Geng, Z., Zhang, R., Gao, X., Cheng, B., Wu, C., Cheng, Y., Li, Z., Di, P., Zhang, K., and Jin, X. Akg: Automatic kernel generation for neural processing units using polyhedral transformations. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation*, PLDI 2021, pp. 1233–1248, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450383912. doi: 10.1145/3453483.3454106. URL https://doi.org/10.1145/3453483.3454106.

Zhao, J., Bastoul, C., Yi, Y., Hu, J., Nie, W., Zhang, R., Geng, Z., Li, C., Tachon, T., and Gan, Z. Parallelizing neural network models effectively on gpu by implementing reductions atomically. In *Proceedings of the 31st International Conference on Parallel Architectures and Compilation Techniques (Submitted)*, PACT'22. ACM, 2022.

Zheng, Z., Zhao, P., Long, G., Zhu, F., Zhu, K., Zhao, W., Diao, L., Yang, J., and Lin, W. Fusionstitching: Boosting memory intensive computations for deep learning workloads, 2020.

A ALGORITHMS

This section summarizes the algorithms mentioned in the paper. Algo. 1 formally describes the approach to aggregate micro-graphs in §3.3. The algorithm takes as input an $\mathcal{F}_x$ and updates its initialized micro-graphs. Each $\mathcal{G}$ is instanced using a primitive op, with the type of $\mathcal{G}$ set by the counterpart of op. These micro-graphs are used to initialize the intermediate set M (line 2), which is then used to update $\mathcal{F}_x$ (line 3). The algorithm recursively aggregates (with the **while** loop) each pair of $(\mathcal{G}_p, \mathcal{G}_c)$, which is related with each other by a producer-consumer edge e, in $\mathcal{F}_x$ using a rule. The result is output to $\mathcal{G}_a$, with its type specified by the rule r (line 7) and $\mathcal{F}_x$ updated at line 8. Aggregating micro-graphs is also subject to the absence of introduced cycles after the aggregation, which has been enforced by the **if** conditional (The *acyclic* function). The order in which the **for** loop between lines 5 and 9 is iterated has been defined at line 4, which specifies the priority of each rule in Table 1. All rules involving a *matmul op* {⑥-matmul} are allocated with the equal priority, which also applies to {⑦-matmul}.

Algorithm 1 Aggregation Algorithm

Input: $\mathcal{F}_x$
1   $M = \varnothing$
  foreach $op \in \mathcal{F}_x$ **do**
2 | $\mathcal{G} = \{op\}$; $\mathcal{G}$.type=op.type; $M = M \cup \mathcal{G}$
3   $\mathcal{F}_x = M$
4   $Rules$={①,②,③,④,⑤,⑥-transpose, {⑥-matmul}, {⑥-conv}}
  foreach $r \in Rules$ **do**
5 | $changed$ = **true**
  | **while** $changed$ **do**
6 | | $changed$ = **false**
  | | **foreach** e between $(\mathcal{G}_p, \mathcal{G}_c) \wedge \mathcal{G}_p, \mathcal{G}_c \in \mathcal{F}_x$ **do**
  | | | **if** $\mathcal{G}_p$.type=$r.p \wedge \mathcal{G}_c$.type=$r.c \wedge$ acyclic$(\mathcal{G}_p, \mathcal{G}_c)$ **then**
7 | | | | $\mathcal{G}_a$ = aggregate$(\mathcal{G}_p, \mathcal{G}_c)$; $\mathcal{G}_a$.type = $r(\mathcal{G}_p, \mathcal{G}_c)$
8 | | | | $\mathcal{F}_x = \mathcal{F}_x - \{\mathcal{G}_p, \mathcal{G}_c\}$; $\mathcal{F}_x = \mathcal{F}_x \cup \mathcal{G}_a$
9 | | | | $changed$ = **true**
Output: $\mathcal{F}_x$

Algo.2 delineates the memory stitching approach (§4.2). It takes as input IR_{F_x}, a collection of the IR_{G_y} generated by Layer I for each $\mathcal{G}_y$, and tries to update it by aggregating $(\mathcal{G}_p, \mathcal{G}_c)$ that is connected by an edge e. The outermost **if** conditional is used to ensure that $\mathcal{F}_x$ and the IRs will only be updated when $(\mathcal{G}_p, \mathcal{G}_c)$ conforms the rules defined in Table 2 of the paper and no cycles will be introduced if they are merged. Line 2 is used to guarantee that the two micrographs have consistent numbers of (parallel and sequential) loop dimensions and identical hardware parameters, and line 3 is used to ensure the fused two reduction micrographs have the same canonical form. The fusion happens between lines 5-7, with both $\mathcal{F}_x$ and IR_{F_x} updated accordingly. Finally, the tensors of the producer-consumer edge e that have been transformed into intermediate variables

are allocated either on (faster) local memory or (slower) global memory depending on their sizes (after loop tiling of Layer I). Unlike the fusion heuristics of the polyhedral model, Algo.2 is used as a complementary strategy by trying to stitch the tensors between micro-graphs obtained by Layer I on faster memory, further optimizing the memory hierarchy.

Algorithm 2 Memory Stitching Algorithm

Input: $IR_{F_x} = \{IR_{G_y}\}$
1   $Rules$={⑦,⑧}
  foreach $r \in Rules$ **do**
  | **foreach** e between $(\mathcal{G}_p, \mathcal{G}_c) \wedge \mathcal{G}_p, \mathcal{G}_c \in F_x$ **do**
  | | **if** $\mathcal{G}_p$.type=$r.p \wedge \mathcal{G}_c$.type=$r.c \wedge$ acyclic$(\mathcal{G}_p, \mathcal{G}_c)$ **then**
  | | | **if** !consist_dim_and_param$(\mathcal{G}_p, \mathcal{G}_c)$ **then**
2 | | | | **continue**
  | | | **if** r = ⑧ $\wedge$!have_identical_canonical_form$(\mathcal{G}_p, \mathcal{G}_c)$ **then**
3 | | | | **continue**
4 | | | $\mathcal{G}_a$ = aggregate$(\mathcal{G}_p, \mathcal{G}_c)$; $\mathcal{G}_a$.type = $r(\mathcal{G}_p, \mathcal{G}_c)$
5 | | | IR_{G_a} = concat(IR_{G_p}, IR_{G_c})
6 | | | $\mathcal{F}_x = \mathcal{F}_x - \{\mathcal{G}_p, \mathcal{G}_c\}$; $\mathcal{F}_x = \mathcal{F}_x \cup \mathcal{G}_a$
7 | | | $IR_{F_x} = IR_{F_x} - \{IR_{G_p}, IR_{G_c}\}$; $IR_{F_x} = IR_{F_x} \cup IR_{G_a}$
  | | | **if** sizes$(e.tensors) \leq$ sizes$(local\_mem)$ **then**
8 | | | | allocate$(e.tensors, local\_mem)$
  | | | **else**
9 | | | | allocate$(e.tensors, global\_mem)$
Output: IR_{F_x}

Algo.3 describes the fusion approach of Layer III (§4.3). It extracts all sets of parallelizable candidates (line 2) by inspecting each pattern of $\mathcal{P}$ in order (line 1) and tries to compose the independent ops within each s (lines 3-11). Note that we add a virtual common tail/head for at most 7 independent branches for fast compilation purpose. In addition, we prefer to add a common tail when a head and a tail are both legal.

Algorithm 3 Parallelism Stitching Algorithm

Input: $IR_P = \{IR_{F_x}\}$
1   $Patterns$ ={multi-head, multi-tail, independent}; $Sets$=$\varnothing$
  foreach $p \in Patterns$ **do**
2 | $Sets$ = extract_sets_of_parallelizable_candidates(p)
  foreach $s \in Sets$ **do**
3 | s = mark_each_op_unvisited(s)
  | **while** num_of_branches_with_unvisited$(s, 1) \geq 2$ **do**
4 | | $group$ = extract_one_op_from_each_branches(s)
5 | | $group$ = sort_by_cost_of_ops$(group)$
6 | | $m = 1$, $k = n$ = num_of_op$(group)$
  | | **while** $m \leq n \wedge k \geq 2$ **do**
  | | | **if** is_positive(cost model (3), m, k) **then**
7 | | | | IR_P =parallel_stitch_ops$(\{IR_{F_x}\}, m, k)$
8 | | | | $m = m + k$; $k = n - m + 1$
  | | | **else**
9 | | | | $k = k - 1$
  | | | **if** $m \neq n \wedge k = 1$ **then**
10 | | | | $m = m + k$; $k = n - m + 1$
11 | | mark_each_op_visited$(group)$
Output: IR_P

The algorithm first marks each *op* within *s* as unvisited (line 3) and iterates over the **while** loop between lines 4 and 11 when there still exist at least two branches within *s*, each of which has at least one unvisited *op*, determined using the *num_of_branches_with_unvisited* function. *group* represents the collection of *op*s extracted from different branches of *s* (line 4); it is then sorted in a descending order by the cost of its *op*s (line 5). The inner **while** loop is used to perform parallelism stitching by compose *k* independent *op*s starting from *m* (line 7), which is triggered when positive performance gain can be estimated by Cost Model (3). *m* and *k* are initialized at line 6 and updated when the **if** conditional is satisfied (line 8) or violated (line 9), guaranteeing the greedy attempts of the algorithm. The updates of *m* and *k* at line 10 are used to ensure that the remaining unvisited *op*s can be evaluated when the leading ones are overweight. Each *op* of *group* is marked as visited (line 11) and the algorithm outputs an updated IR_P.

B ARTIFACT

This section offers the description to those who are interested in reproducing our results. A publicly accessible DOI of this section is https://doi.org/10.6084/m9.figshare.19383890.v1.

B.1 Artifact check-list (meta-information)

- **Compilation: CUDA toolkits.**

- **Data set: Available online.**

- **Hardware: V100 GPUs and Huawei Ascend 910 chips.**

- **Output: Execution times.**

- **How much time is needed to prepare workflow (approximately)?: 30 minutes to two hours.**

- **How much time is needed to complete experiments (approximately)?: about one hour.**

- **Publicly available?: Yes.**

- **Code licenses (if publicly available)?: Apache 2.0.**

- **Data licenses (if publicly available)?: Apache 2.0.**

- **Archived (provide DOI)?: Available online.**

B.2 Description

B.2.1 How delivered

The data sets and workloads used in the paper can be accessed from the repository of MindSpore. A patch is also offered online to ease the experiment workflow, the path to which will be described later in §B.3.

B.2.2 Hardware dependencies

Our work depends on two kinds of hardware environments. For the GPU environments, please refer to Table 9; and for the Ascend environment, please follow the installation guide at https://

www.mindspore.cn/install/en (choose the Ascend 910 section) to configure the Ascend 910 environment.

Table 9: GPU configurations.

Hardware	Requirements
GPU	eight Tesla V100-SXM2-16GB
Network Interface Controller	Mellanox Technologies MT28800 Family[ConnextX-5] 100Gb/sec[4X EDR]
Processor	dual-socket Intel(R) Xeon(R) Platinum 8160 CPU @ 2.10GHz
Memory	256GB
Operating System	Ubuntu 16.04.4 LTS (GNU/Linux 4.4.0-116-generic x86_64)

B.2.3 Software dependencies

We assume the readers have installed git, or please first install git if it has not yet been installed. To execute the script, please install Python version 3. All network models used in the experiment were written using the MindSpore framework (version 1.3.0) developed by Huawei. One can install the correct version of MindSpore from https://www.mindspore.cn/install/en. The CUDA toolkit version 11.1 is preferred, or one can install a higher version. We recommend the pip or docker installation method of the CUDA toolkit. The CUDA profiler is needed to validate the results of sub-graph case study.

B.2.4 Data sets

The data sets considered during the training of each model in the experiments can also be downloaded from the open access website. In particular, one can download the data sets for individual models as follows.

- For BERT, please download its data set by following the online document https://gitee.com/mindspore/mindspore/blob/r1.3/model_zoo/official/nlp/bert/README.md

- For Transformer, the publicly accessible website is https://gitee.com/mindspore/mindspore/blob/r1.3/model_zoo/official/nlp/transformer/README.md

- For Wide&Deep, one can download the data set from https://gitee.com/mindspore/mindspore/blob/r1.3/model_zoo/official/recommend/wide_and_deep/README.md

- For Yolo-v3, the readme file can be accessed from https://gitee.com/mindspore/mindspore/blob/r1.3/model_zoo/official/cv/yolov3_darknet53/README.md

- And for DeepFM, one can refer to https://gitee.com/mindspore/mindspore/blob/r1.3/model_zoo/official/recommend/deepfm/README.md

If the data of Wide&Deep and DeepFM is not accessible, one can also fetch them from https://gitee.com/link?target=http%3A%2F%2Fgo.criteo.net%2Fcriteo-research-kaggle-display-advertising-challenge-dataset.tar.gz

B.3 Installation

Each model has been well configured in the repository of Mind-Spore. A model, however, has to be switched between different configurations as described in the paper. To ease the switching between different configurations, we offer a patch which can be downloaded and applied as follows.

```
$ git clone -b r1.3 https://gitee.com/mindspore/
  mindspore.git
$ cd ./mindspore
$ git clone https://gitee.com/yaozhujia/apollo-ae-patch.git
$ git apply apollo-ae-patch/apollo_bert.diff
$ git apply apollo-ae-patch/apollo_transformer.diff
$ git apply apollo-ae-patch/apollo_widedeep.diff
$ git apply apollo-ae-patch/apollo_deepfm.diff
$ git apply apollo-ae-patch/apollo_yolov3.diff
```

In case the gitee repository is not convenient for some users, please use the pip installation method described at https://gitee.com/mindspore/docs/blob/r1.3/install/mindspore_gpu_install_pip_en.md. Note that the implementations of our work is embedded into the repository of MindSpore. Please pay attention to the notifications of MindSpore's official website https://www.mindspore.cn/install/en provided some links do not work.

B.4 Evaluation and expected results

The evaluation results will be printed on screen during the experiment workflow. In most cases, the execution time will be reported. Please note that, we use "/path/to/dataset" as an abstraction of the real path of the datasets. Please replace it depending on your environment. For the first commands of each case like "cd top/path", "top/path" refers to the top directory of MindSpore code. That is to say, one is advised to go back to the top directory of MindSpore after each evaluation.

B.5 Methodology

B.5.1 Sub-graph case study

Each sub-graph is a small-scale benchmark when compared to the models used in the experiments. In addition, multiple DL operators are executed in a sequential streaming manner by the MindSpore framework. CUDA nvprof is thus used to collect the execution time of a sub-graph.

The command to reproduce the result of sub-graph case study is

```
$ nvprof python3 <test_case.py> -t <test_type> -d <device_id>
```

where test_type can be instantiated using 0, 1 or 2, which represents the MindSpore, partial and full versions reported in Table 3. device_id is the ID of the used device, which can be specified as 0 by default.

For instance, one can execute the example in Fig.6(c) using

```
$ nvprof --print-gpu-summary python3 6c.py -t 2 -d 0
```

which will output the profiling of the result when APOLLO is enabled. The overall execution time, which we reported in Table 3, is the sum of the last three rows, *i.e.*, 5.3440 + 4.2240 + 1.7600 = 11.328 μs.

Or one can type the following instructions in the terminal to reproduce the results.

```
$ cd 6.1
# 5a, MindSpore
$ nvprof --print-gpu-summary python3 5a.py -t 0 -d 0
# 5a, partial
$ nvprof --print-gpu-summary python3 5a.py -t 1 -d 0
```

```
# 5a, full
$ nvprof --print-gpu-summary python3 5a.py -t 2 -d 0
# 5b, MindSpore
$ nvprof --print-gpu-summary python3 5b.py -t 0 -d 0
# 5b, partial
$ nvprof --print-gpu-summary python3 5b.py -t 1 -d 0
# 5b, full
$ nvprof --print-gpu-summary python3 5b.py -t 2 -d 0
# 5c, MindSpore
$ nvprof --print-gpu-summary python3 5c.py -t 0 -d 0
# 5c, partial
$ nvprof --print-gpu-summary python3 5c.py -t 1 -d 0
# 5c, full
$ nvprof --print-gpu-summary python3 5c.py -t 2 -d 0
# 6a, MindSpore
$ nvprof --print-gpu-summary python3 6a.py -t 0 -d 0
# 6a, partial
$ nvprof --print-gpu-summary python3 6a.py -t 1 -d 0
# 6a, full
$ nvprof --print-gpu-summary python3 6a.py -t 2 -d 0
# 6b, MindSpore
$ nvprof --print-gpu-summary python3 6b.py -t 0 -d 0
# 6b, partial
$ nvprof --print-gpu-summary python3 6b.py -t 1 -d 0
# 6b, full
$ nvprof --print-gpu-summary python3 6b.py -t 2 -d 0
# 6c, MindSpore
$ nvprof --print-gpu-summary python3 6c.py -t 0 -d 0
# 6c, partial
$ nvprof --print-gpu-summary python3 6c.py -t 1 -d 0
# 6c, full
$ nvprof --print-gpu-summary python3 6c.py -t 2 -d 0
```

B.5.2 Results on a single GPU

To reproduce the results of MindSpore and APOLLO, one can use the following commands.

- BERT-base.
  ```
  $ cd mindspore/model_zoo/official/nlp/bert
  # BT-base, batchsize = 32, MS only
  $ bash scripts/run_standalone_pretrain_for_gpu.sh 0 1
    /path/to/cn-wiki-128 base 32
  # BT-base, batchsize = 32, MS with apollo
  $ bash scripts/run_standalone_pretrain_for_gpu.sh 0 1
    /path/to/cn-wiki-128 base 32 1
  # BT-base, batchsize = 64, MS only
  $ bash scripts/run_standalone_pretrain_for_gpu.sh 0 1
    /path/to/cn-wiki-128 base 64
  # BT-base, batchsize = 64, MS with apollo
  $ bash scripts/run_standalone_pretrain_for_gpu.sh 0 1
    /path/to/cn-wiki-128 base 64 1
  ```

  Specifically, "/path/to/cn-wiki-128" should be replaced with the path location where you store the data set for BERT-base. The execution time of each training epoch is cached in a log file during the execution of a model written using the MindSpore framework. The throughputs reported in the paper are calculated according to these execution times.

- Transformer.
  ```
  $ cd mindspore/model_zoo/official/nlp/transformer
  # TR, batchsize=8, MS only
  $ bash scripts/run_standalone_train.sh GPU 0 2 8
    /path/to/ende-l128-mindrecord 8
  # TR, batchsize=8, MS with apollo
  $ bash scripts/run_standalone_train.sh GPU 0 2 8
    /path/to/ende-l128-mindrecord 8 1
  # TR, batchsize=16, MS only
  $ bash scripts/run_standalone_train.sh GPU 0 2 8
    /path/to/ende-l128-mindrecord 16
  # TR, batchsize=16, MS with apollo
  $ bash scripts/run_standalone_train.sh GPU 0 2 8
    /path/to/ende-l128-mindrecord 16 1
  ```

- Wide&Deep.
  ```
  $ cd mindspore/model_zoo/official/recommend/wide_and_deep
  # WD, batchsize=16000, MS only
  $bash script/run_standalone_train_for_gpu.sh 2
    /path/to/mindrecord 16000
  # WD, batchsize=16000, MS with apollo
  $bash script/run_standalone_train_for_gpu.sh 2
    /path/to/mindrecord 16000 1
  # WD, batchsize=32000, MS only
  ```

```
$bash script/run_standalone_train_for_gpu.sh 2
  /path/to/mindrecord 32000
# WD, batchsize=32000, MS with apollo
$bash script/run_standalone_train_for_gpu.sh 2
  /path/to/mindrecord 32000 1
```

- Yolo-v3.
  ```
  $ cd mindspore/model_zoo/official/cv/yolov3_darknet53/scripts
  # YO, batchsize=4, MS only
  $ bash run_standalone_train_gpu.sh /path/to/coco2014
    /path/to/backbone_darknet53.ckpt 4
  # YO, batchsize=4, MS with apollo
  $ bash run_standalone_train_gpu.sh /path/to/coco2014
    /path/to/backbone_darknet53.ckpt 4 1
  # YO, batchsize=8, MS only
  $ bash run_standalone_train_gpu.sh /path/to/coco2014
    /path/to/backbone_darknet53.ckpt 8
  # YO, batchsize=8, MS with apollo
  $ bash run_standalone_train_gpu.sh /path/to/coco2014
    /path/to/backbone_darknet53.ckpt 8 1
  ```

- DeepFM.
  ```
  $ cd mindspore/model_zoo/official/recommend/deepfm
  # FM, batchsize=8192, MS only
  $ bash scripts/run_standalone_train.sh 0 GPU
    /path/to/dataset 8192
  # FM, batchsize=8192, MS with apollo
  $ bash scripts/run_standalone_train.sh 0 GPU
    /path/to/dataset 8192 1
  # FM, batchsize=16384, MS only
  $ bash scripts/run_standalone_train.sh 0 GPU
    /path/to/dataset 16384
  # FM, batchsize=16384, MS with apollo
  $ bash scripts/run_standalone_train.sh 0 GPU
    /path/to/dataset 16384 1
  ```

To reproduce the results of TensorFlow and XLA, one can follow the instructions below.

- For BERT-base, please access it from https://github.com/NVIDIA/DeepLearningExamples/tree/master/TensorFlow/LanguageModeling/BERT and execute the model using the configurations and datasets specified on the official website.

- For Transformer, one can access the model from https://github.com/NVIDIA/DeepLearningExamples/tree/master/TensorFlow/LanguageModeling/Transformer-XL and execute the model using the configurations and datasets specified on the official website.

- For Wide&Deep, please visit it at https://github.com/NVIDIA/DeepLearningExamples/tree/master/TensorFlow/Recommendation/WideAndDeep and execute the model using the configurations and datasets specified on the official website.

- For Yolo-v3, one configure it using the configurations and datasets specified at https://github.com/wizyoung/YOLOv3_TensorFlowandexecutethemodel.

- For DeepFM, one can obtain it from https://github.com/ChenglongChen/tensorflow-DeepFM.

B.5.3 Results on multiple GPUs

To reproduce the results of MindSpore and APOLLO, one can use the following commands.

- BERT (both base and large versions).
  ```
  $ cd mindspore/model_zoo/official/nlp/bert
  # BT-base(32), device_num = 8, MS only
  $ bash scripts/run_distributed_pretrain_for_gpu.sh 8 1
    /path/to/cn-wiki-128 base 32
  # BT-base(32), device_num = 8, MS with apollo
  ```

```
$ bash scripts/run_distributed_pretrain_for_gpu.sh 8 1
  /path/to/cn-wiki-128 base 32 1
# BT-base(64), device_num = 8, MS only
$ bash scripts/run_distributed_pretrain_for_gpu.sh 8 1
  /path/to/cn-wiki-128 base 64
# BT-base(64), device_num = 8, MS with apollo
$ bash scripts/run_distributed_pretrain_for_gpu.sh 8 1
  /path/to/cn-wiki-128 base 64 1
# BT-large(4), device_num = 4, MS only
$ bash scripts/run_distributed_pretrain_for_gpu.sh 4 1
  /path/to/cn-wiki-128 large 4
# BT-large(4), device_num = 4, MS with apollo
$ bash scripts/run_distributed_pretrain_for_gpu.sh 4 1
  /path/to/cn-wiki-128 large 4 1
```

- Wide&Deep.
  ```
  $ cd mindspore/model_zoo/official/recommend/wide_and_deep
  # WD(16000), device_num=8, MS only
  $ bash script/run_multigpu_train.sh 8 2
    /path/to/mindrecord 16000
  # WD(16000), device_num=8, MS with apollo
  $ bash script/run_multigpu_train.sh 8 2
    /path/to/mindrecord 16000 1
  ```

- DeepFM.
  ```
  $ cd mindspore/model_zoo/official/recommend/deepfm
  # FM(16384), device_num=4, MS only
  $ bash scripts/run_distribute_train_gpu.sh 4
    /path/to/dataset 16384
  # FM(16384), device_num=4, MS wih apollo
  $ bash scripts/run_distribute_train_gpu.sh 4
    /path/to/dataset 16384 1
  ```

To reproduce the results of TensorFlow and XLA, one can follow the instructions below.

- For BERT-base and BERT-large, one can access the model from https://github.com/NVIDIA/DeepLearningExamples/tree/master/TensorFlow/LanguageModeling/BERT and execute the model using the configurations and datasets specified on the official website.

- For Wide&Deep, please visit the open access address https://github.com/NVIDIA/DeepLearningExamples/tree/master/TensorFlow/Recommendation/WideAndDeep.

- For DeepFM, the model can be retrieved from https://github.com/ChenglongChen/tensorflow-DeepFM.

B.5.4 Results on Ascend 910 chips

- BERT-large.

  To reproduce the results of MindSpore and APOLLO, one can use the following commands.
  ```
  $ cd mindspore/model_zoo/official/nlp/bert
  # BERT(24), batchsize = 24, MS only
  $ bash scripts/run_standalone_pretrain_ascend.sh 0 1
    /path/to/cn-wiki-128 large 24
  # BERT(24), batchsize = 24, MS with apollo
  $ bash scripts/run_standalone_pretrain_ascend.sh 0 1
    /path/to/cn-wiki-128 large 24 1
  ```

  To reproduce the results of TensorFlow and XLA, please visit https://github.com/NVIDIA/DeepLearningExamples/tree/master/TensorFlow/LanguageModeling/BERT and execute the model using the configurations and datasets specified on the official website.

- Pangu-α.

  As the architecture of the Pangu-alpha model has not been released yet, we feel very sorry that the architectures cannot be provided.

B.5.5 Compilation overhead

To reproduce the compilation overhead, one has to compile each model written in MindSpore 1.3.0, with necessary compilation options required to be turned on. To achieve this, one needs to select the installation mode after downloading the MindSpore framework, and compilation options "-p on" should be turned on, which enables the pipelined profiling functionality of the framework. Specifically, one can type

```
$ bash build.sh -e gpu -p on
```

on the terminal. The cached information, including the compilation overhead, will be printed onto stdout. The "TotalTime" indicates the compilation overhead of a model in seconds.

B.5.6 Results of inference workloads

To reproduce the results of MindSpore and APOLLO, one can follow the commands below.

- Wide&Deep.
```
$ cd mindspore/model_zoo/official/recommend/wide_and_deep
$ export WD_EVAL_ONLY=1
# WD, batchsize=16000, MS only
$ bash script/run_standalone_train_for_gpu.sh 2
  /path/to/mindrecord 16000
# WD, batchsize=16000, MS with apollo
$ bash script/run_standalone_train_for_gpu.sh 2
  /path/to/mindrecord 16000 1
$ unset WD_EVAL_ONLY
```

- Yolo-v3.
```
$ cd mindspore/model_zoo/official/cv/yolov3_darknet53/scripts
# Yolo-v3, batchsize=32, MS only
$ bash run_eval_gpu.sh
  </path/to/dataset> <checkpoint_path> 32
# Yolo-v3, batchsize=32, MS with apollo
$ bash run_eval_gpu.sh
  </path/to/dataset> <checkpoint_path> 32 1
```

- EPPMVSNet.

  EPPMVSNet is a small-scale workload developed within Huawei, used for the 3D reconstruction of real objects. This workload can be cloned using
```
$ git clone -b r1.4 ttps://gitee.com/mindspore/models.git
```

  In particular, this workload is developed under the version 1.4 of MindSpore's Model Zoo. To reproduce the results of MindSpore and APOLLO, one can follow the commands below.
```
$ cd model/research/cv/eppmvsnet
# EPPMVSNet, batchsize=1, MS with apollo
$ bash eval.sh /path/to/dataset 0
# disable apollo
$ sed 's/enable_graph_kernel=True/enable_graph_kernel=False/g'
  ./validate.py > ./validate_new.py
$ mv validate_new.py validate.py
# EPPMVSNet, batchsize=1, MS only
$ bash eval.sh /path/to/dataset 0
```

  The evaluation results will be stored in "./results/blended-mvs/val/metrics.txt". One can find the results in the log file.

B.5.7 Ablation study of BERT on a single GPU

To reproduce the results the ablation study of BERT on a single GPU, one can follow the commands below.
```
$ cd mindspore/model_zoo/official/nlp/bert
# BT-base, batchsize = 32, MS only
$ bash scripts/run_standalone_pretrain_for_gpu.sh 0 1
  /path/to/cn-wiki-128 base 32
# BT-base, batchsize = 32, MS with apollo L1
$ bash scripts/run_standalone_pretrain_for_gpu.sh 0 1
  /path/to/cn-wiki-128 base 32 L1
# BT-base, batchsize = 32, MS with apollo L2(Memory Stitch)
$ bash scripts/run_standalone_pretrain_for_gpu.sh 0 1
  /path/to/cn-wiki-128 base 32 L2
# BT-base, batchsize = 32, MS with apollo All
# (Memory Stitch + Parallel Stitch)
$ bash scripts/run_standalone_pretrain_for_gpu.sh 0 1
  /path/to/cn-wiki-128 base 32 ALL
```