

Cuckoo: Deadline-Aware Job Packing on Heterogeneous GPUs for DL Model Training

Yuzheng Zhang¹, Renyu Yang^{1†}, Junhong Liu¹, Weihang Jiang¹, Tianyu Ye¹, Yiqiao Liao², Penghao Zhang², Tiezi Zhang², Kun Shang², Tianyu Wo¹, Chunming Hu¹, Chengru Song², Jin Ouyang²

¹Beihang University

²Kuaishou Inc.

背景：近年来，深度学习（DL）的快速发展推动了智能应用和服务的进步。大规模构建多租户 GPU 集群已经成为训练深度学习作业的事实标准。

论文中的研究现状：Muri 是较早尝试交错利用多维资源的方法之一。它将具有互补资源需求的深度学习作业分组，从而在模型训练的不同阶段充分利用各类空闲资源。然而，Muri 对执行模式进行了过度简化：它将每次训练迭代划分为多个资源互斥的阶段，并假设单个资源在同时被使用时不会产生干扰，并且无法感知作业特定的 QoS 需求。

研究方向和贡献：通过构建一个考虑运行时干扰的估计模型，评估共置作业组的执行效率。基于该估计结果，系统进一步求解一个优化问题，以生成具体的作业打包和执行计划，确保异构 GPU 得到最优利用。




Figure 5: Overview of Cuckoo.

Job Profiler

当用户向 Cuckoo 提交作业时，Job Profiler 会对作业进行若干 step 的试运行。它会记录一次迭代中不同资源的持续时间，例如磁盘 IO、CPU、GPU 和网络通信在不同 GPU 上的执行时长，以及通信与计算之间的重叠时长




Figure 6: Execution timeline. F'_0, B'_0, B'_1 indicate stages with resource interference, where overlapping usage of GPU resources prolongs the actual duration.

Table 3: Job configurations and resource usage time (unit: ms).

	Model	Order	S	F	B	C
Job 0	ResNet	0	10	37	76	98
Job 1	BERT	1	10	72	61	363

Table 4: Resource usage per stage in interleaved execution (unit: ms).

i	f_0	b_0	c_0	f_1	b_1	c_1	R
0	0	0	0	0	0	0	10
1	0	0	0	72	0	0	72
2	74	152	0	0	122	363	363
3	0	0	98	0	0	0	98
4	0	0	0	0	0	0	10
5	0	0	0	72	0	0	72

Job Paker

$$eff_value = \frac{\sum_{j \in jobs} Duration(j)}{InterleavingDuration(jobs)}$$

$$ddl_value = \frac{\min_{j \in jobs} D_j - now}{\max_{j \in jobs} D_j - now}$$

$$total_value = w \cdot eff_value + (1 - w) \cdot ddl_value$$

Algorithm 3 Job Packing Algorithm

```
// Part 1
G ← newGraph()
// J includes all jobs to be packed
for j ∈ J do
    G.addNode(j)
end for
for each pair (u, v) ∈ G.nodes() do
    weight ← ComputeTotalValue(u, v)
    G.addEdge(u, v, weight)
end for
// Part 2
M ← G.ComputeMaximumWeightMatching()
pack_result ← newList()
// Pack each matched job pair into a group
for each pair (u, v) ∈ M do
    p ← PackJob(u, v)
    pack_result.add(p)
end for
// Pack unmatched jobs into individual groups
for unmatched node u ∈ M do
    p ← PackJob(u)
    pack_result.add(p)
end for
return pack_result
```

Resource Scheduling

$$C_1 = (n - 1) \times T + t$$

$$C_2 = \max(0, (n - 1) \times T + t - D)$$




Figure 7: An example of bipartite graph for execution planning.

Table 5: Execution time and deadline for group.

	Time on GPU A	Time on GPU B	Deadline
Group 1	2	3	5
Group 2	3	4	5
Average	2.5	3.5	

Table 6: Weight for each edge.

Group	(GPU, Order)	C_1	C_2	$C_1 + C_2$
Group 1	GPU A, Order 1	2	0	2
	GPU A, Order 2	4.5	0	4.5
	GPU B, Order 1	3	0	3
	GPU B, Order 2	6.5	1.5	8
Group 2	GPU A, Order 1	3	0	3
	GPU A, Order 2	5.5	0.5	6
	GPU B, Order 1	4	0	4
	GPU B, Order 2	7.5	2.5	10

EVALUATION

实验环境:

工作负载:

Table 7: GPU configuration of the experimental cluster.

# of Nodes	GPU Type per Node	# of GPUs per Node
1	NVIDIA A800	2
1	NVIDIA RTX 4090	8
3	NVIDIA A10	2

$$: \frac{deadline - arrival_time}{execution_time} \sim N(8, 2)$$

Table 8: Models and batch sizes used in evaluation.

Model Name	Type	Batch Size(s)
A2C [20]	RL	8, 16, 32
DQN [21]	RL	8, 16, 32
BERT [5]	NLP	2, 4, 8
GPT-2 [24]	NLP	2, 4
VGG19 [27]	CV	8, 16, 32
VGG16 [27]	CV	8, 16, 32
ResNet18 [10]	CV	16, 32, 64
ResNet50 [10]	CV	16, 32, 64





Figure 8: Job submission pattern derived from the Helios trace.


EVALUATION

Table 9: Performance comparison between CUCKOO and Muri.

Method	Norm.JCT	Deadline Satisfaction Rate	Norm.Tail JCT
CUCKOO	1.00	77.16%	1.00
Muri	1.81	32.44%	0.61



(a) Queue length



(b) Number of overtime jobs

Figure 9: Comparisons of queue length and overtime jobs.

EVALUATION




Figure 10: Throughput estimation error comparison.

Table 10: Performance comparison between AM and Muri-P

Method	Norm.JCT	Deadline Satisfaction Rate	Norm.Tail JCT
AM	1.00	77.16%	1.00
Muri-P	1.07	76.58%	1.23

Table 11: Predicted outcome breakdown.

Method	P	TP	FP	N	TN	FN
AM	90.21%	89.06%	1.15%	9.79%	6.82%	2.97%
Muri-P	91.17%	83.69%	7.48%	8.83%	5.10%	3.73%






Figure 11: Precision, recall, and F1-score comparison.

EVALUATION



(a) JCT



(b) Deadline Satisfaction Rate

Figure 12: Impact of varying w on JCT and deadline satisfaction rate.

Table 12: Comparison of different job packing strategies.

Method	Norm.JCT	Deadline Satisfaction Rate	Norm.Tail JCT
DAB	1.00	77.16%	1.00
Blossom	1.32	47.02%	1.11
None	1.79	48.56%	1.77

EVALUATION

Table 13: Comparison of different execution planning strategies.

Method	Norm.JCT	Deadline Satisfaction Rate	Norm.Tail JCT
MCHP	1.00	77.16%	1.00
FIFO	1.30	41.46%	0.79
EDF	1.22	64.88%	1.35

Table 14: Comparison of heterogeneity-aware and homogeneous schedulin

Method	Norm.JCT	Deadline Satisfaction Rate	Norm.Tail JCT
CUCKOO	1.00	77.16%	1.00
CUCKOO-Homo	1.16	71.29%	1.36
Muri-Hetero	1.60	38.20%	1.56
Muri	1.81	32.44%	0.61