Lyra: Elastic Scheduling for Deep Learning Clusters

Jiamin Li^{*}, Hong Xu⁺, Yibo Zhu[‡], Zherui Liu[‡], Chuanxiong Guo[‡], Cong Wang^{*}, *City University of Hong Kong, +The Chinese University of Hong Kong, +ByteDance Inc.

Introduction [Separate management of training and inference cluster.]

Inference requires less computation and GPU memory than training, therefore using weaker GPUs like Nvidia T4, with a fraction of the resources of the training GPUs, such as Nvidia V100 and A100.



1157 22

3204



[Elastic scaling of distributed training jobs.]

Jobs can take a **variable** number of workers according to resource availability. One can even adjust the number of workers on-the-fly when the job is running.

[Limited elasticity] Some model families enjoy a linear scaling efficiency within a range.

 \sim 5% of all jobs (account for 36% of training) cluster resources)

Which on-loan servers should be reclaimed to minimise

For a 2-server reclaim request...

Naively selecting servers leads to unnecessary preemption.

and servers have inter-dependency

[Knapsack with dependent item values]

A new value definition: sum of job's server fraction

J	ob	Co-hosted by # of servers	Per Server Value
	а	2	0.5
	b	1	1
	С	2	0.5
	d	2	0.5

Sal	Initial Allocation		JCT		Average
501.	А	В	Α	В	JCT
1	6	2	50	53.33	51.67
2	2	6	63.33	20	41.67
3	4	4	60	30	45
	Sol. 1 2 3	Initial All Sol. A 1 6 2 2 3 4	Sol.Initial AllocationAB16222234	Initial Allocation JC Sol. A B A 1 6 2 50 2 2 6 63.33 3 4 4 60	Initial Allocation JCI A B A B 1 6 2 50 53.33 2 2 6 63.33 20 3 4 4 60 30

			Sal	Initial Al	itial Allocation		JCT		
Job	W ^{min}	W ^{max}	Min. running time	501	A	В	А	В	JCT
A	2	3	100	1	3	5	100	24	62
В	2	6	20	2	2	6	106.67	20	63.33

Elastic job = Base (first-class citizen) + Flexible Demand [Two-phase resource allocation]

- a. Prioritise Base Demand SJF to minimise queuing time,
- b. Allocate the remaining resources to fulfill the Flexible Demand to minimise running time.

JCT (s)			GPU	Preemption	
Mean	Median	95%ile	Training	Overall ¹	Ratio ²
16610	791	82933	0.72	0.52	0
11236	568	56477	0.86	0.65	12.24%
10434	525	56553	0.86	0.68	7.35%
8891	422	41146	0.93	0.72	5.72%



字节跳动 ByteDance

University of Hona Kona



香港中文大學 Chinese University of Hong Kong





Elastic Scaling Shortest-job-first is not always optimal.

Group	ltem	Weight	Value	
А	1	2	50	
	1	1	20	
D	2	2	30	
D	3	3	36	
	4	4	40	

Multiple-choice Knapsack problem