CS 744: POLLUX

Shivaram Venkataraman

Fall 2021

ADMINISTRIVIA

- Course project assignments
  - Emails will go out today (Oct 14th)
  - Introductions due Oct 25th

- Midterm Exam
  - In class on Oct 28th
  - Includes everything from beginning to the end of ML

MACHINE LEARNING: TRAINING

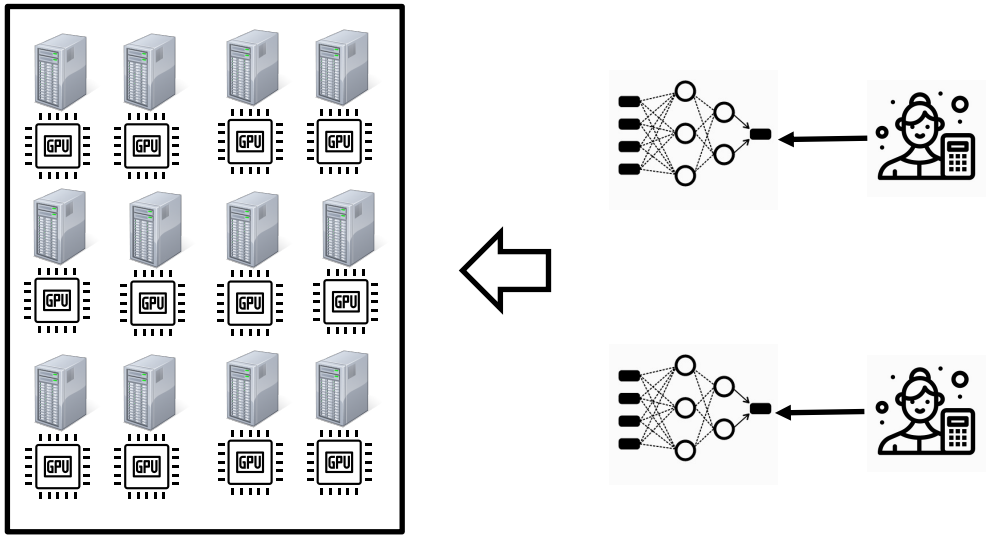

WORKLOAD CHARACTERISTICS

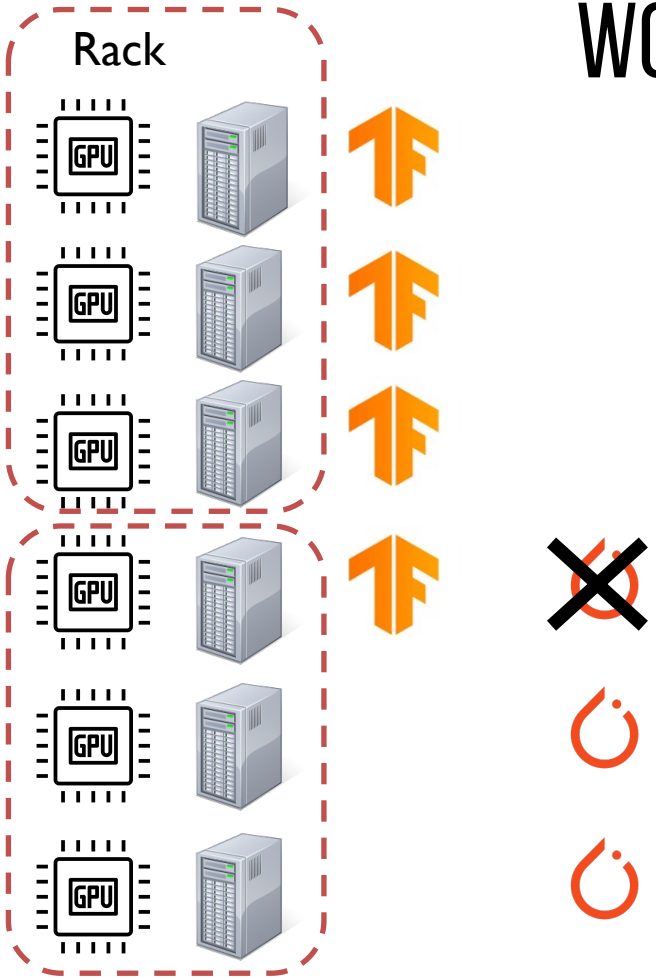

Long running tasks

Gang scheduling

GPU Sharing?

Created by Phonlaphat Thongsriphong
from Noun Project

EXISTING DL SCHEDULER INTERFACE

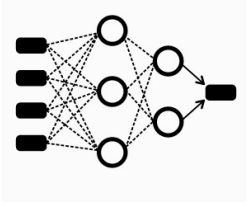

Run job Resnet18
With BatchSize = 64
on Num GPUs = 4

Scheduler

Goals:

- Maximize throughput
- Fairness
- Minimize JCT
- …

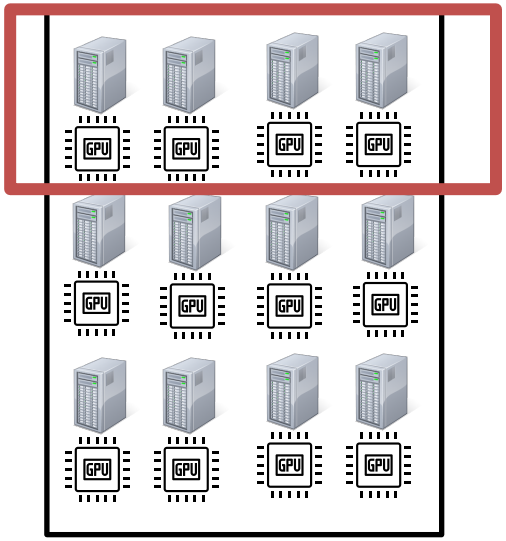

BATCH SIZE VS THROUGHPUT

ResNet18 on CIFAR-10

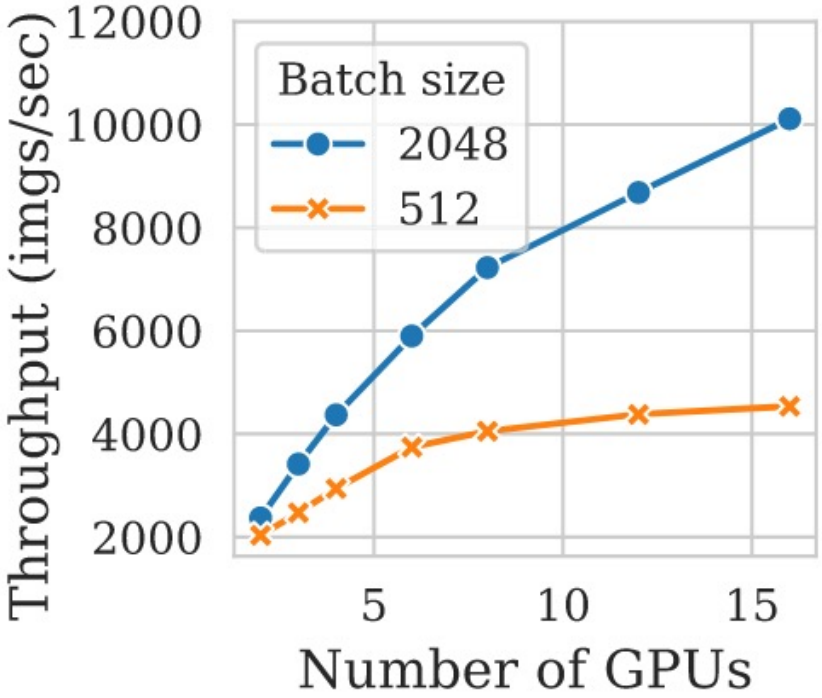

$$\text{Time} = T_{grad} + T_{sync}$$

BATCH SIZE VS STATISTICAL EFFICIENCY

Statistical Efficiency: Amount of progress made *per training example*

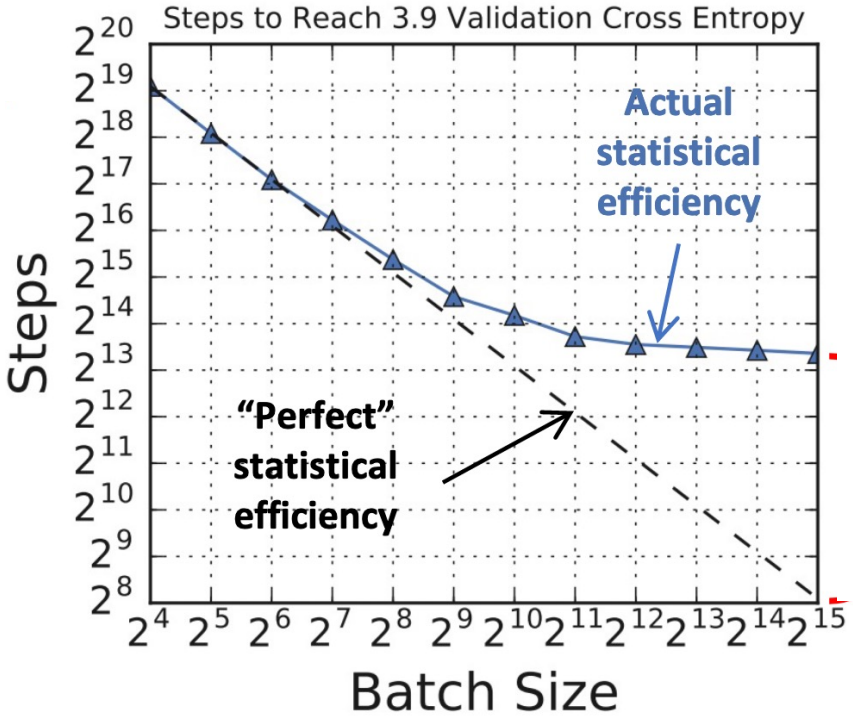

GRADIENT NOISE SCALE (GNS)

Scale batch size *during* training based on gradient noise

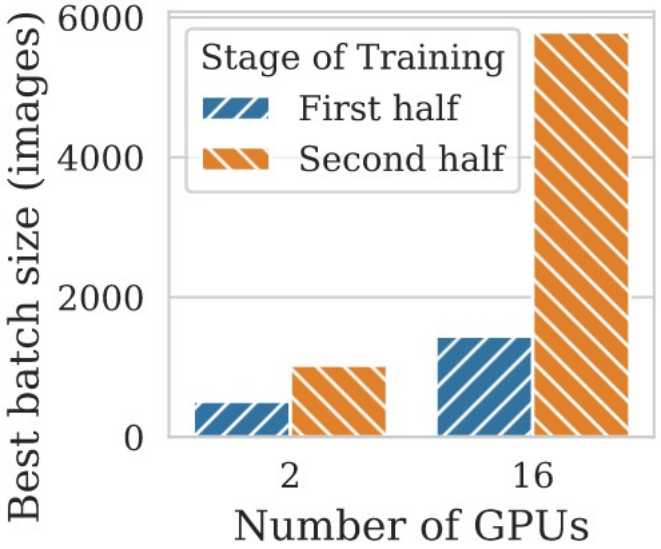

POLLUX: GOAL, APPROACH

Maximize goodput across all jobs in the cluster

$$\texttt{GOODPUT}_t(\star) = \texttt{THROUGHPUT}(\star) \times \texttt{EFFICIENCY}_t(M(\star)),$$

Approach

1. Job submitted with initial batch size, num GPUs
2. Profile jobs during execution to model throughput, efficiency
3. Tune batch size/GPU and num GPUs based on resource availability

MODELING STATISTICAL EFFICIENCY

For a batch size M, initial batch size M_0

$$\texttt{EFFICIENCY}_t(M) = \frac{\varphi_t + M_0}{\varphi_t + M}.$$

φ_t : Computed using gradients noise (GNS) using a **specific batch size**

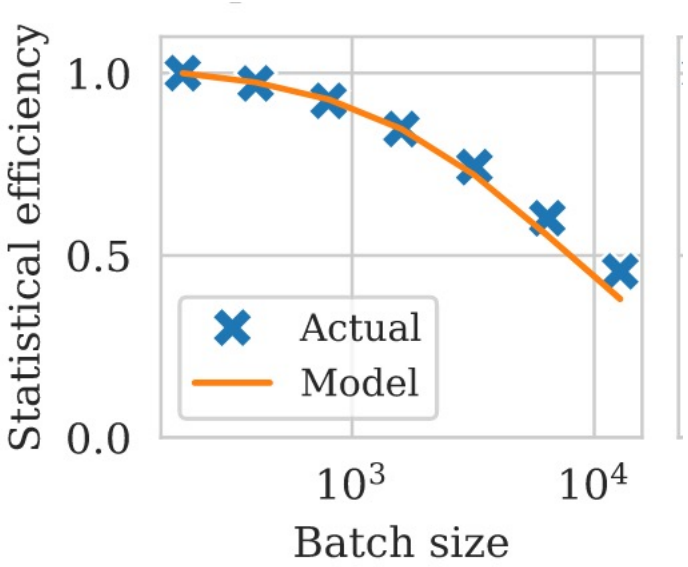

MODELING SYSTEM THROUGHPUT

$$\texttt{THROUGHPUT}(a,m,s) = M(a,m,s)/T_{iter}(a,m,s).$$

$$T_{grad}(m) = \alpha_{grad} + \beta_{grad} \cdot m,$$

$$T_{iter}(a,m,0) = \left(T_{grad}(a,m)^{\gamma} + T_{sync}(a)^{\gamma}\right)^{1/\gamma},$$

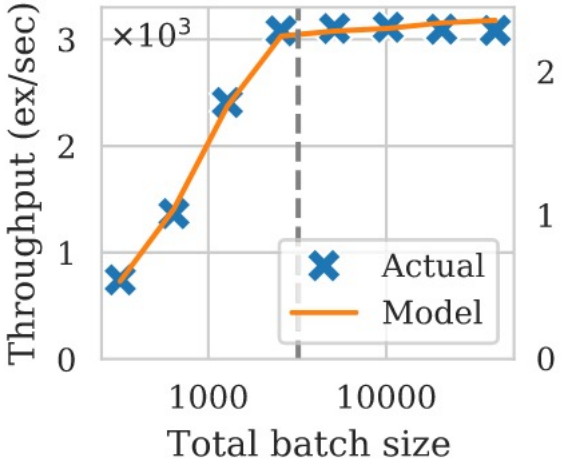

Takeaway : Learn gamma, alpha, beta etc. by running the job
Predict throughputs!

POLLUX: ARCHITECTURE

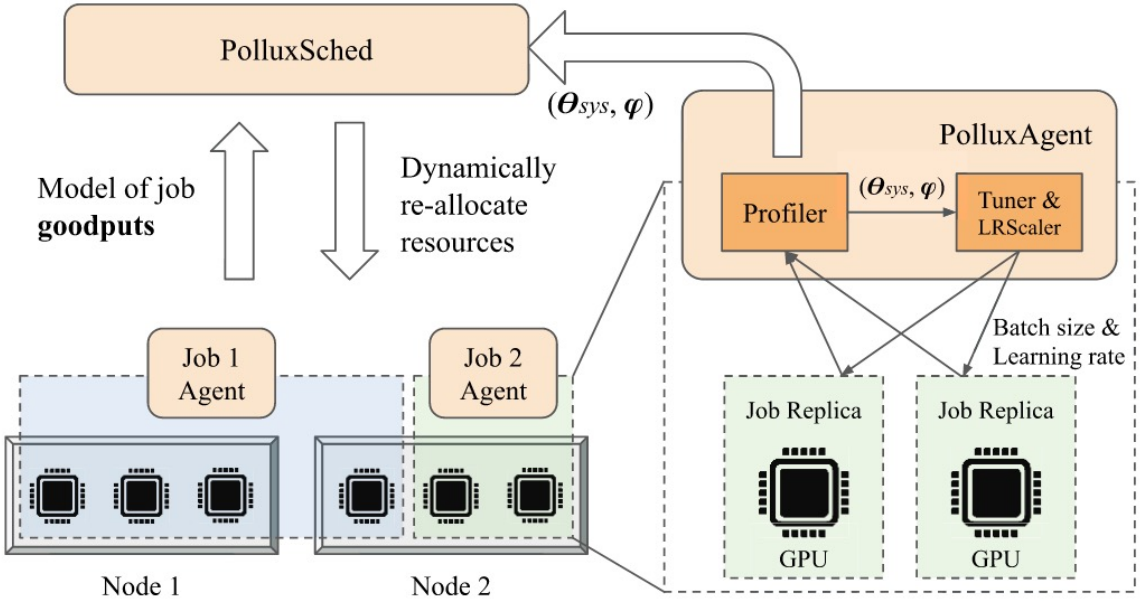

POLLUX: CLUSTER WIDE OPTIMIZATION

Maximize fitness across all jobs

$$\texttt{FITNESS}_p(A) = \left(\frac{1}{J} \sum_{j=1}^{J} \texttt{SPEEDUP}_j(A_j)^p \right)^{1/p}.$$

p → fairness knob

Approach: Use a search algorithm (genetic algorithm) to find allocation

SUMMARY

DL Workloads Scheduling: Batch Size, Num GPUs

Pollux: Optimal batch size, num GPUs varies

Across jobs and during a job's execution

Models for system throughput, statistical efficiency

Cluster-wide optimization

DISCUSSION

https://forms.gle/hQTrk53W3wwkEu9A8

What are some similarities or differences between Mesos and DL schedulers like Pollux?

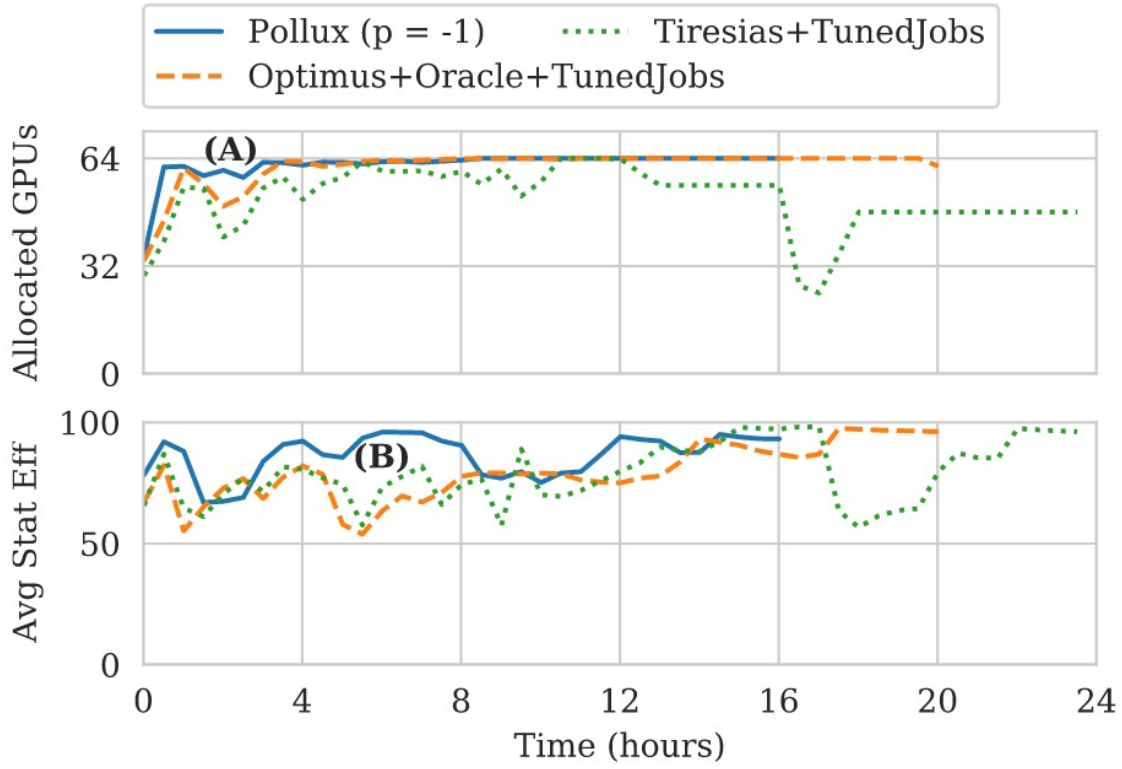
Pollux (p = -1)
Tiresias+TunedJobs
Optimus+Oracle+TunedJobs
(A)
(B)
Allocated GPUs
Avg Stat Eff
64
32
0
100
50
0
0
4
8
12
16
20
24
Time (hours)

NEXT STEPS

Next Class: Nexus

Course Project Introductions!

Midterm after that