

AI!

CS 744: RAY

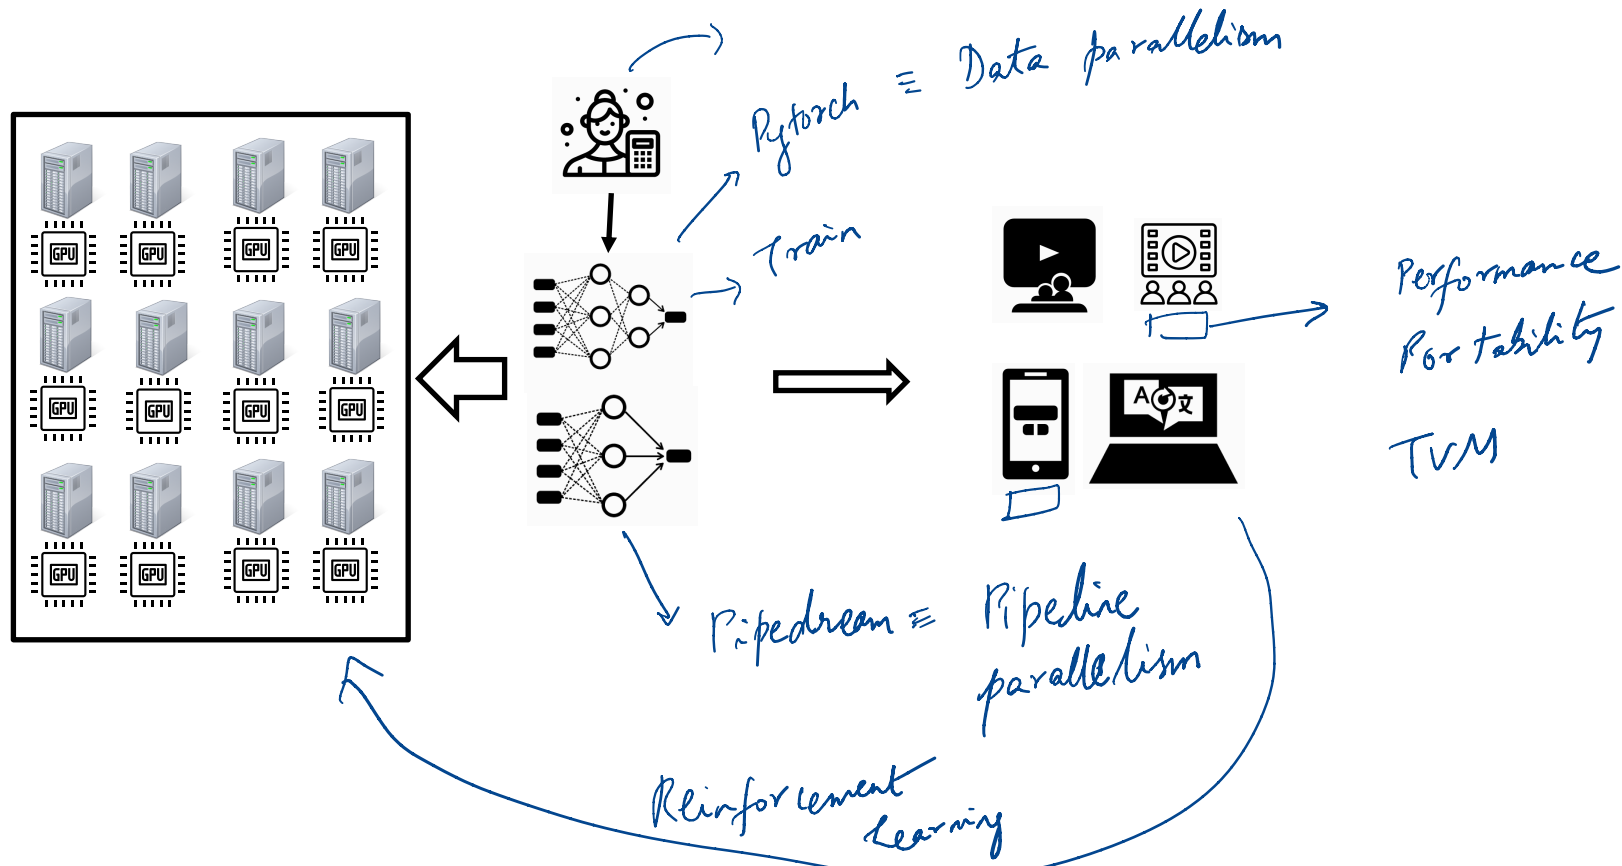
Shivaram Venkataraman

Fall 2020

ADMINISTRIVIA

- Assignment Grades? → *by mid / late next week*
- Project proposal aka Introduction (10/16)
 - Introduction
 - Related Work
 - Timeline (with eval plan)
- Midterm: Oct 22 → *early next week.*

MACHINE LEARNING: STACK



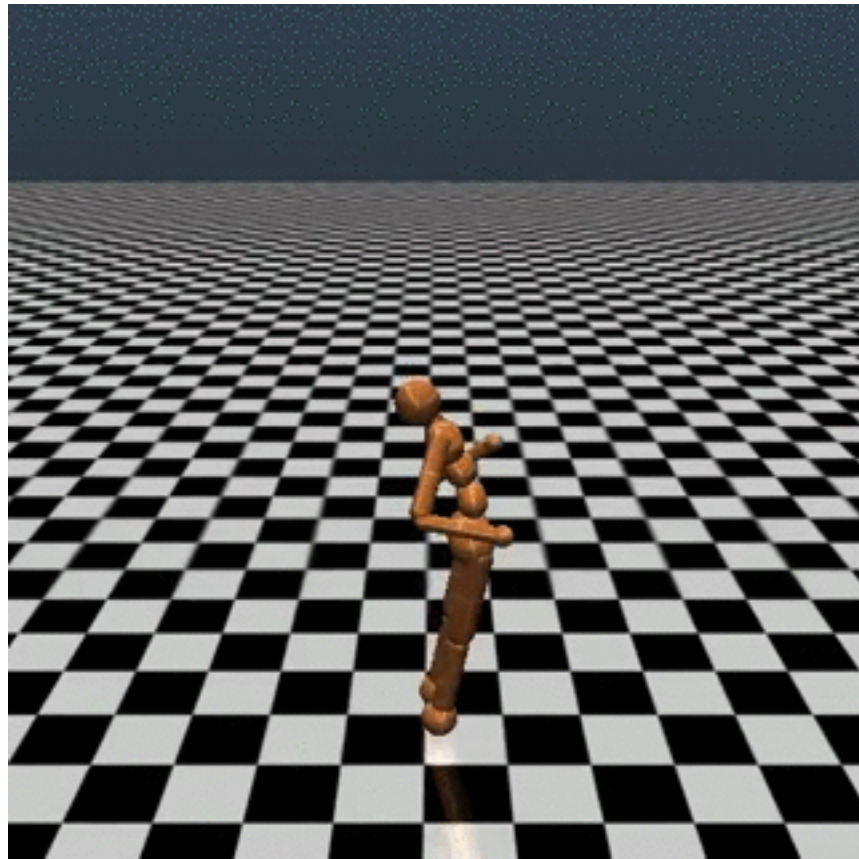
REINFORCEMENT LEARNING

Agent
controlled
by
algorithm



Reward

Not just
supervised learning



RL SETUP

Anything that performs an action

Agent

performs action

road / traffic

Environment

board / game state

Training

Policy improvement
(e.g., SGD)

policy

Serving

Policy evaluation

action (a_i)

state (s_{i+1})
(observation)

reward (r_{i+1})

Simulation

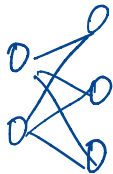


trajectory: $s_0, (s_1, r_1), \dots, (s_n, r_n)$

(history of states, rewards)

tells us the effects of this action

ML model that decides next action



RL REQUIREMENTS

Static execution plan

Simulation

→ fine grained computation → flexibility

↳ each simulation could be ~ ms or hours

Training

→

stateful processing

↳ Model state / simulator state

stateless processing

↳ data pre-processing

Serving

→

very low latency

very high throughput
1M tasks / sec

Dynamic execution

↳ future computation depends on output of past compute

RAY API

stateless and side-effect free?
 same input
 ↓
 same output

Tasks → any function that is run remotely

Erlang ~ 80s/90s

Actors

stateful tasks



Class Actor
<state>
def handle Msg(msg):
 use <state>
 update <state>

futures = f.remote(args)
 ↓
 normal Python function → f(args)
 local variable / future

a handle that you can wait on

actor = Class.remote(args)
futures = actor.method.remote(args)

method name
= actor.method.remote(args 1)
msg

objects = ray.get(futures)
ready = ray.wait(futures, k, timeout)

args will be handled before
args 1

Futures can be arguments to tasks
↳ you can spawn (or wait) for tasks within a task

RAY API

State  run this anywhere

1  Task 

Actors 2 

<internal>

actor = Class.remote(args)

futures = actor.method.remote(args)

for i in 1 to 10:

- ~~ray~~ f.remote(i)

wait()

objects = ray.get(futures)

ready = ray.wait(futures, k, timeout)

trigger to run the computation

Tasks

Nested tasks

futures = f.remote(args)

def f(args):

fo = g.remote(1,2,3)

o = ray.wait(fo)
(args, o)

Lineage!

COMPUTATION MODEL

create_policy.remote()

Dotted lines

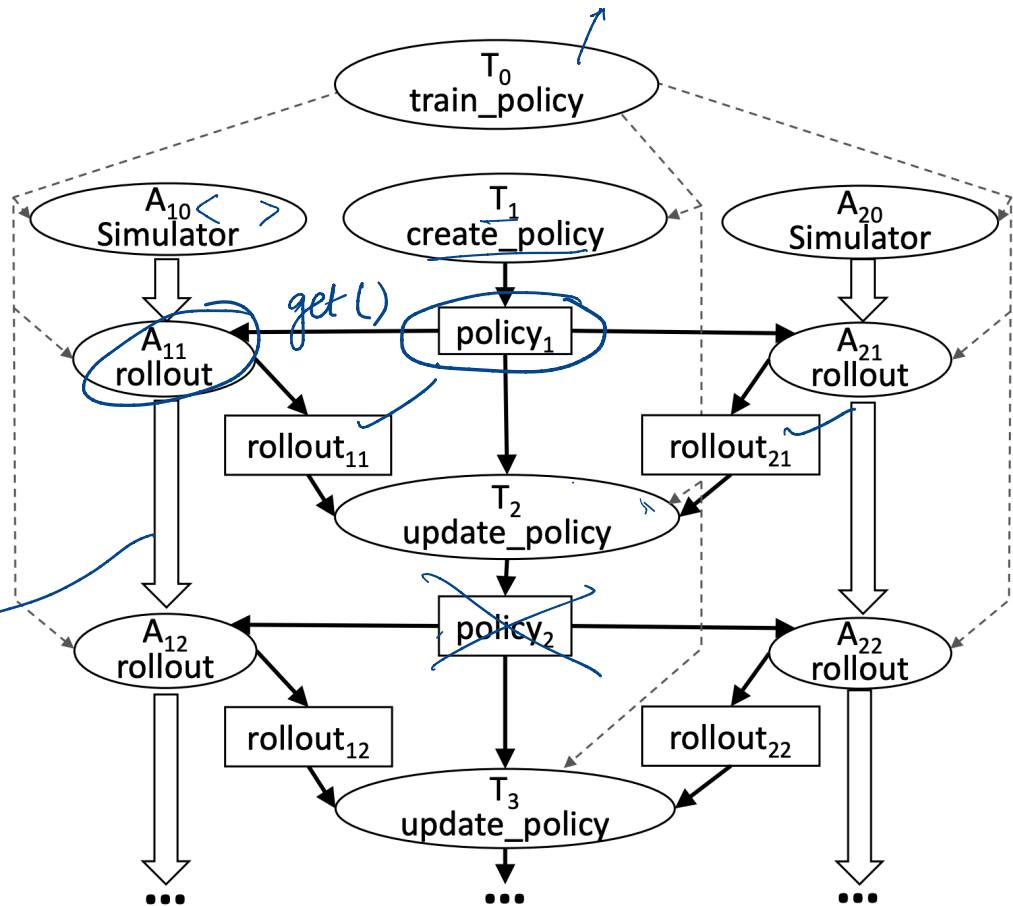
↳ Control edges

spawn a task
spawn an actor

Solid lines

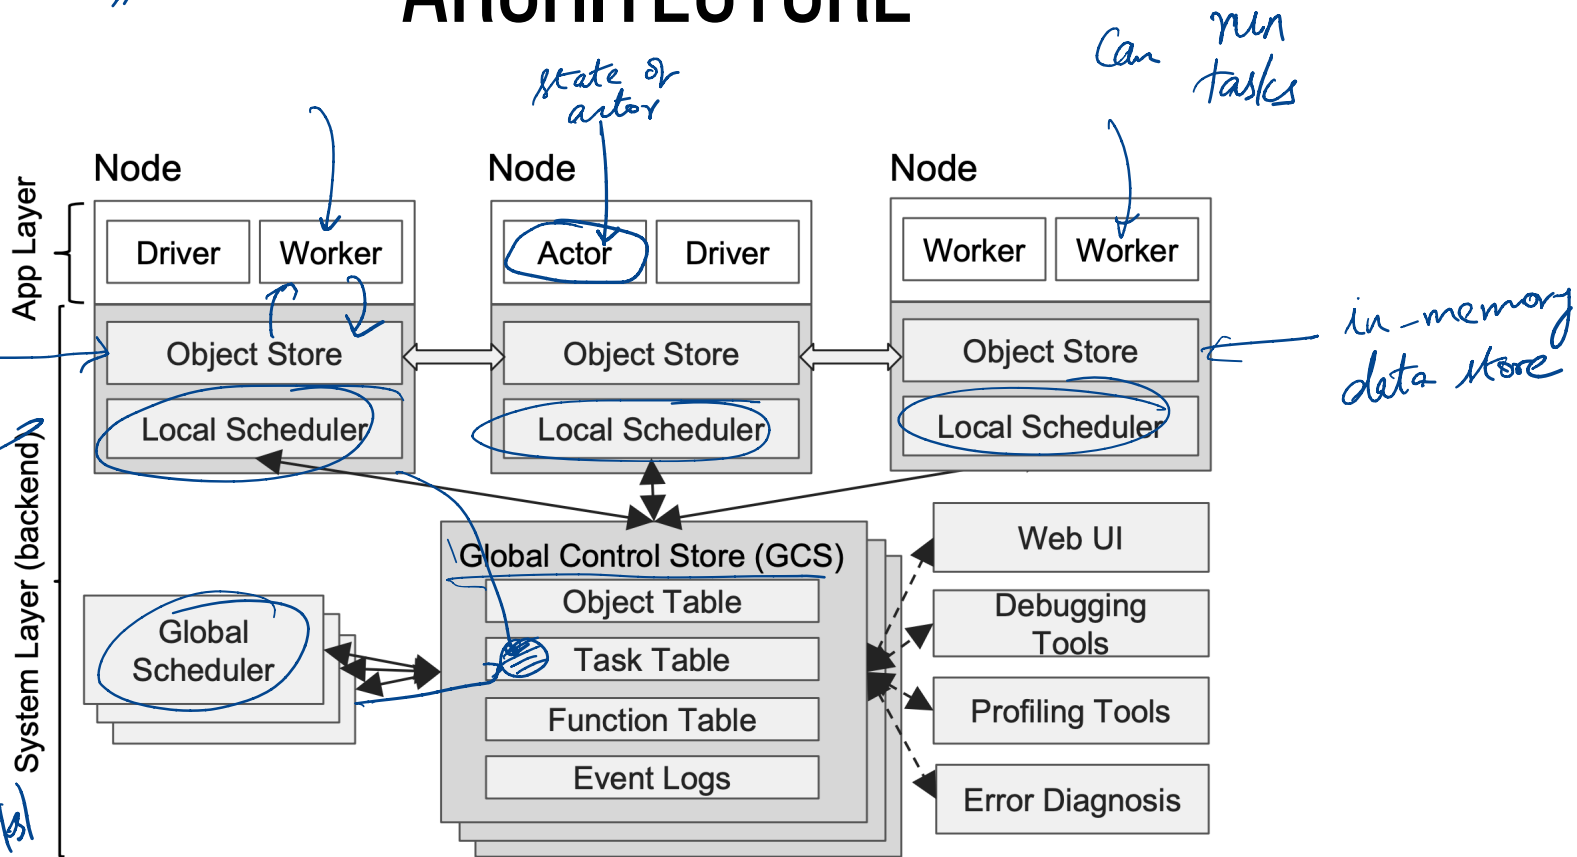
↳ Data edges

⇒ Stateful edges
action on actor
happen sequentially



Deterministic
hash key% hash = machine num

ARCHITECTURE



GLOBAL CONTROL STORE

Sort of a Database

Object table

↳ list of all objects
and their locations } namenode
metadata

Task table

↳ lineage of tasks

Function table

↳ code blocks corresponding
to tasks

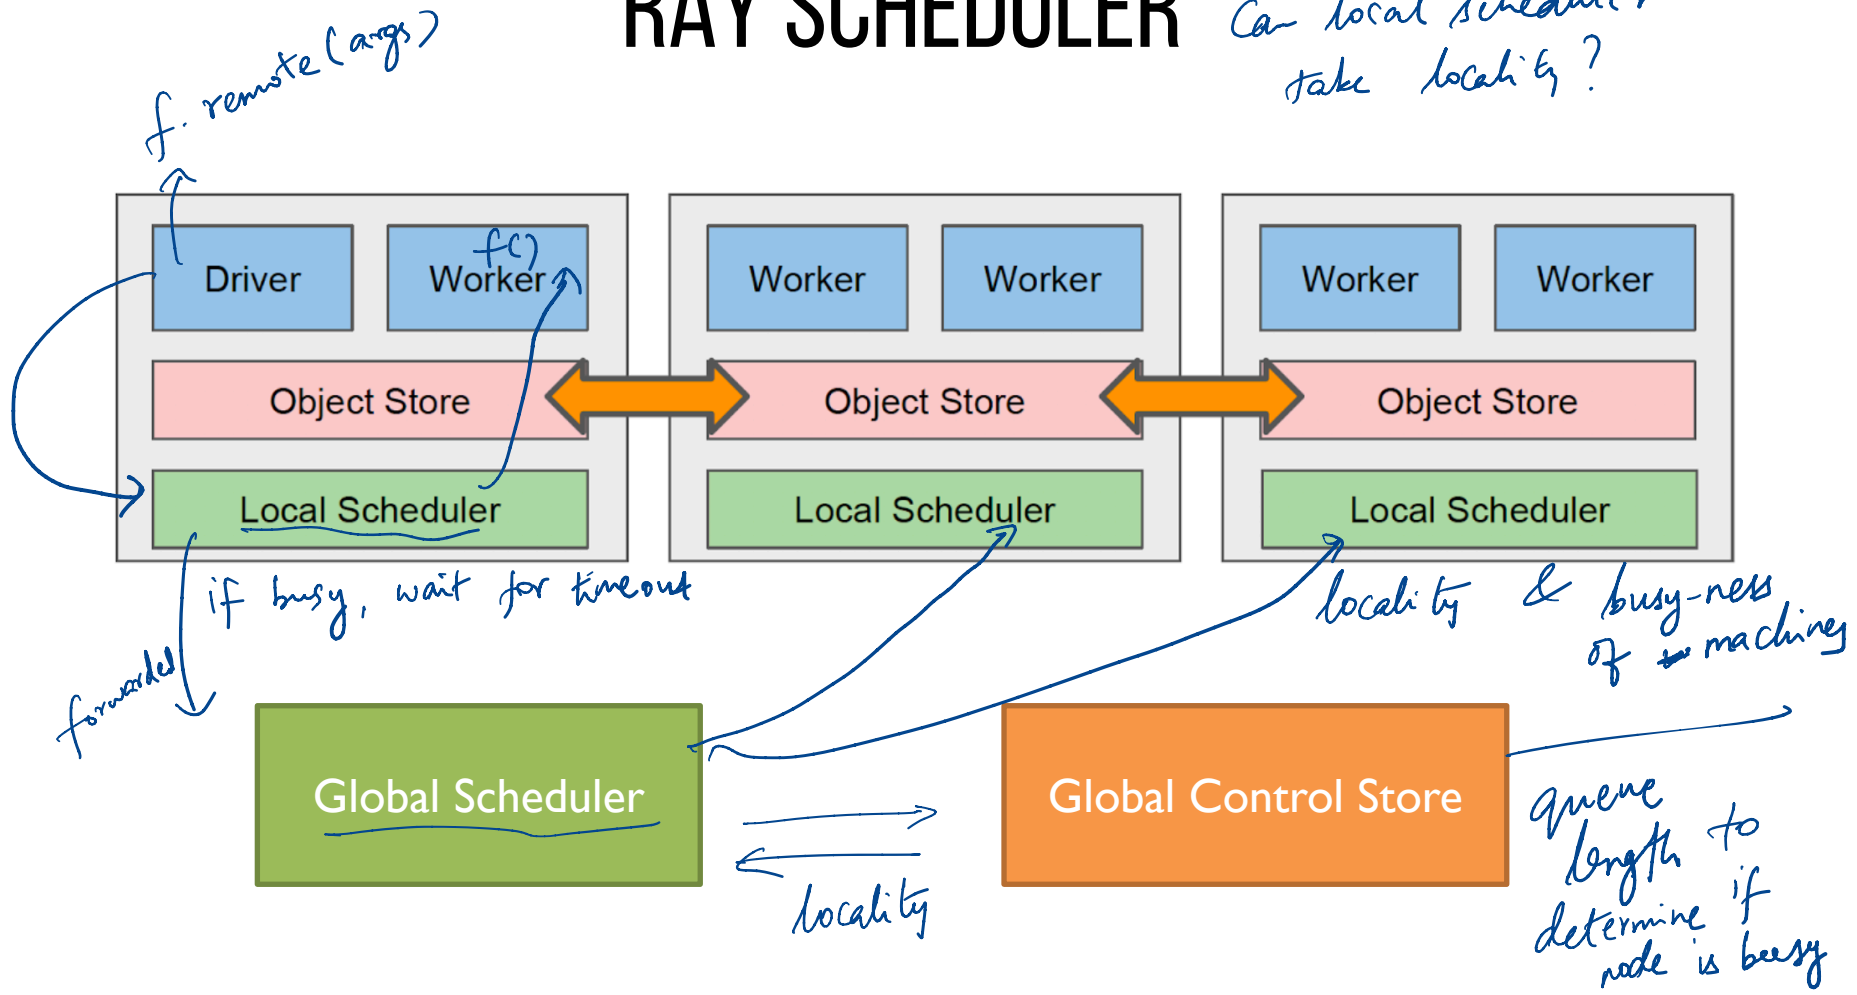
Externalizes
state !

↳ shard
Replicate

scale more
easily, simplify
sched design !
fault tolerance

RAY SCHEDULER

Can local scheduler
take locality?



FAULT TOLERANCE

Tasks → lineage, replay or re-execution of tasks

Actors → periodically checkpoint actors
↔ restore ckpt
replay messages

GCS →  shard | replication chain

Scheduler → Stateless!
Nothing?
Re-spawn or launch a new scheduler

SUMMARY

Ray: Unified system for ML training, serving, simulation

Flexible API with support for

- Stateless tasks

- Stateful Actors

Distributed scheduling, Global control store

DISCUSSION

<https://forms.gle/PN5FSJB6vVkDjoih8>

Consider you are implementing two apps: a deep learning model training and a sorting application. When will use tasks vs actors and why?

Sorting

state
locality

Actors

Does external
sorting \rightarrow state

load-balancing
stateless, Tasks

Deterministic operations

still have
dependencies!

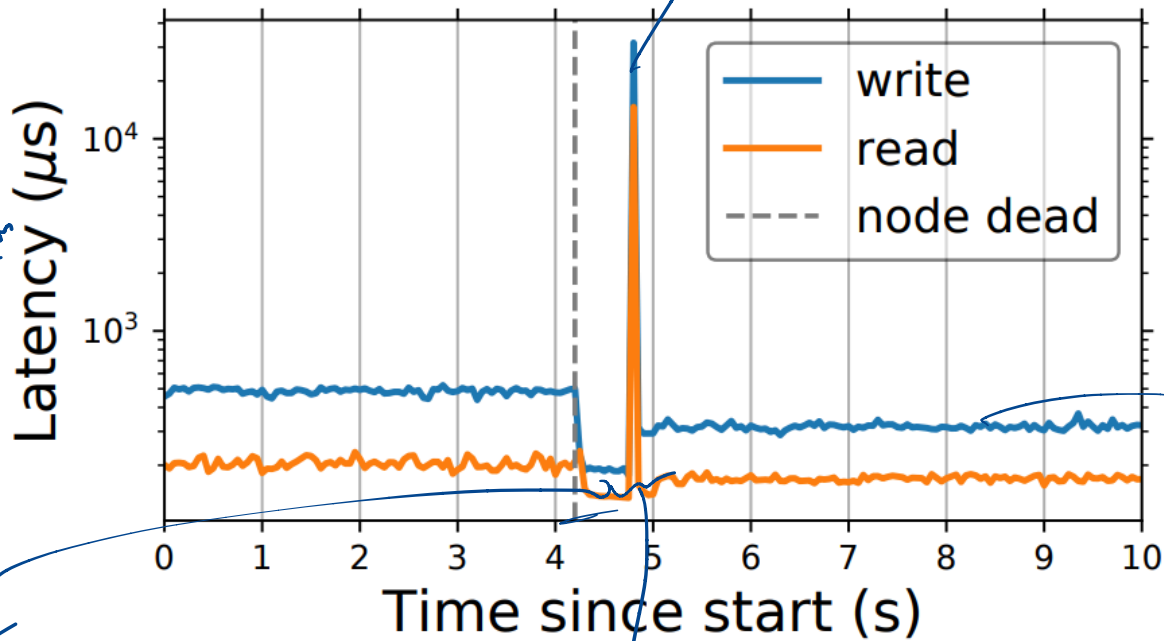
divide into smaller parts?

Model
Training

weights are
state,
Multiple for
data parallel

Can do dependencies
between iterations?

fine-grained
recovery



one replica
no overhead
from replication

during failure
goes down here?

warm-up
this new replica

latency spike lasts for 30ms

new node
has better
hardware
better/closer
to replica?
goes down
after recovery?

NEXT STEPS

Next class: Clipper
Last lecture on ML!

