

CS 744: RAY

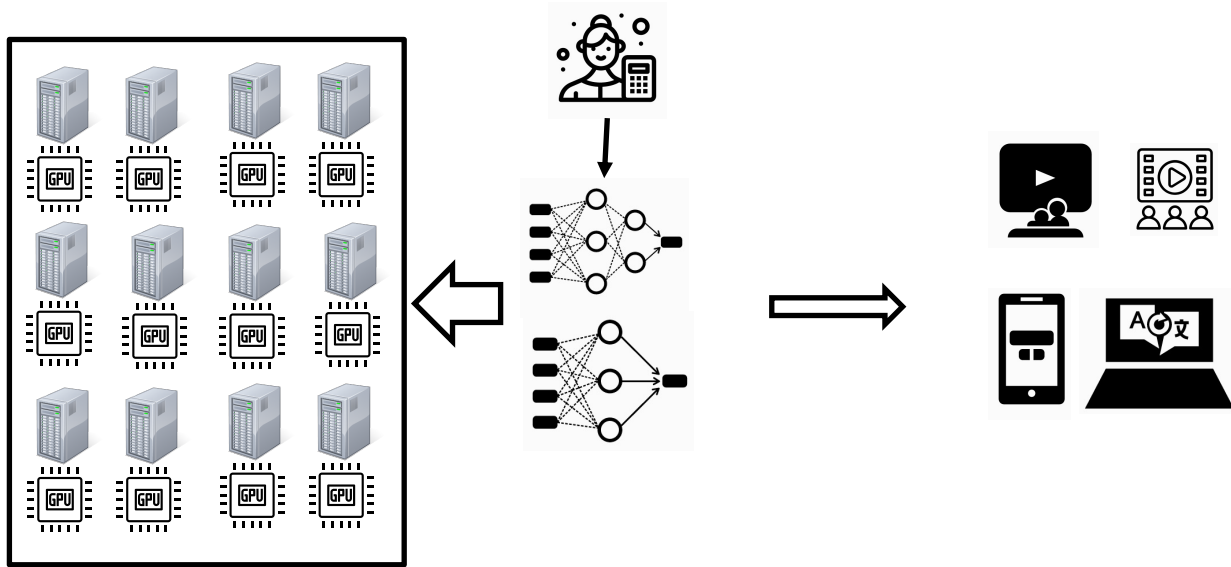
Shivaram Venkataraman

Fall 2021


ADMINISTRIVIA

- Assignment Two: Due tonight!
- Project proposal aka Introduction (10/25)
 - Introduction
 - Related Work
 - Timeline (with eval plan)

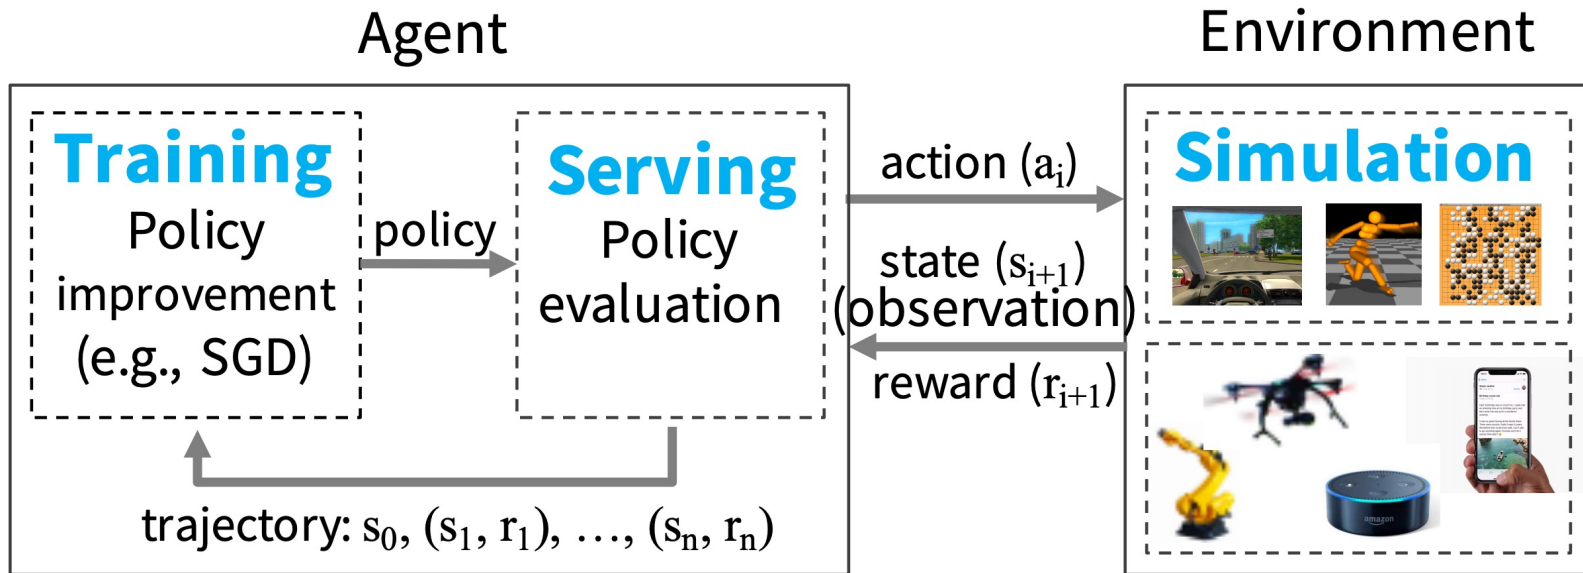
MACHINE LEARNING: STACK



REINFORCEMENT LEARNING: APPLICATIONS



RL SETUP



RL REQUIREMENTS

Simulation

Training

Serving

RAY API

Tasks

```
futures = f.remote(args)
```

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

Actors

```
Class.remote  
actor.method.remote
```

RAY API

Tasks


```
futures =  
    f.remote(args)
```

Actors

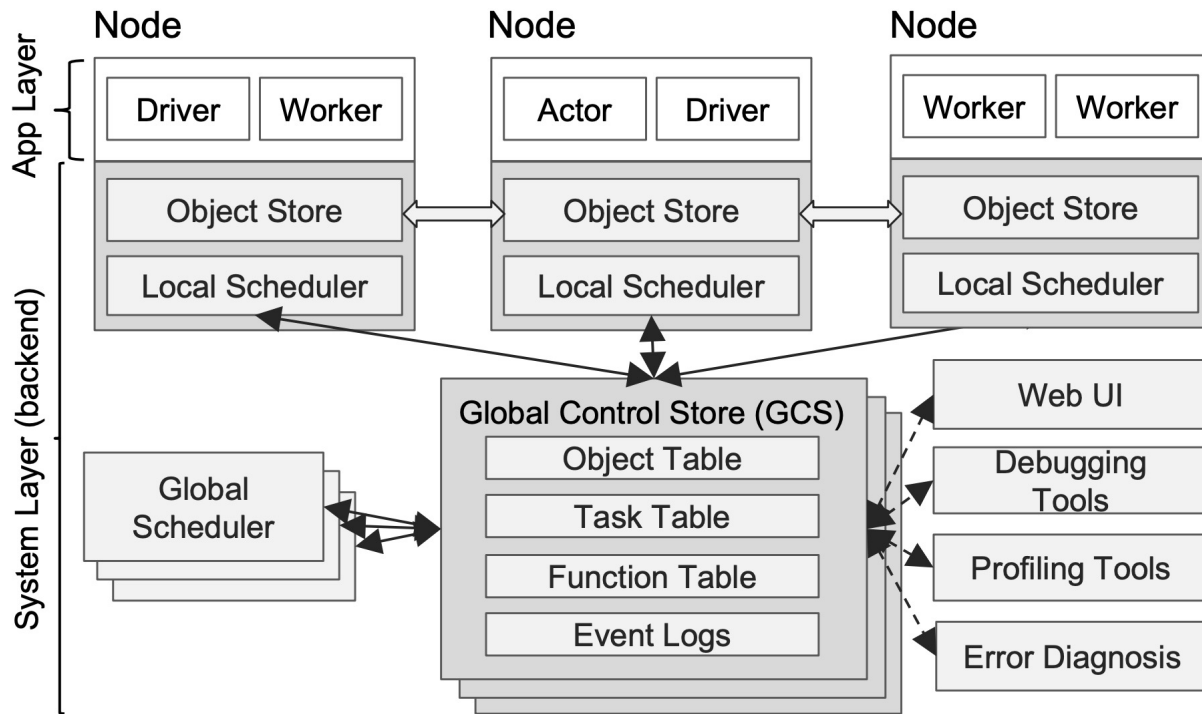
```
actor = Class.remote(args)  
futures = actor.method.remote(args)
```

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


COMPUTATION MODEL



ARCHITECTURE



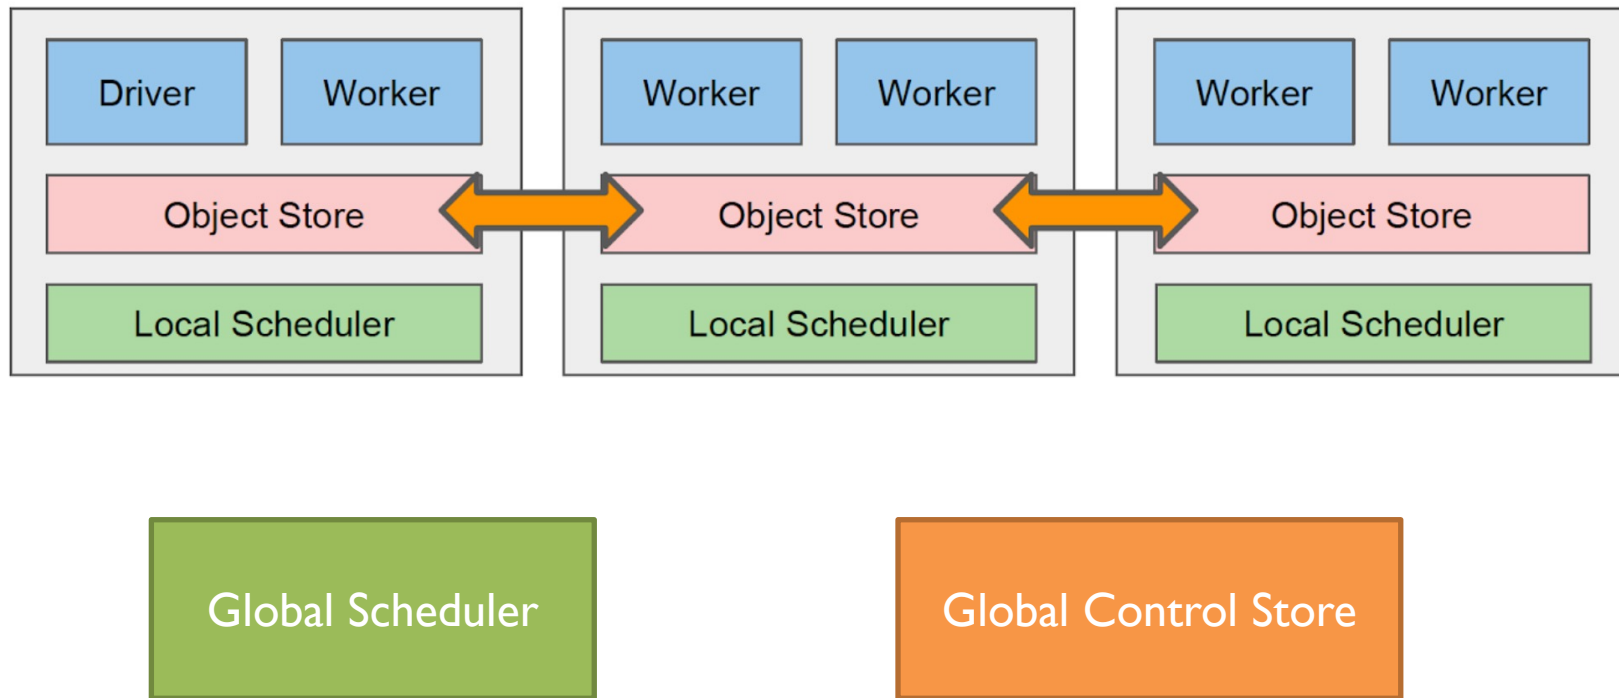
GLOBAL CONTROL STORE

Object table

Task table

Function table

RAY SCHEDULER



FAULT TOLERANCE

Tasks

Actors

GCS

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/MnsCJA87CVhMmShs8>

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



(a) Task reconstruction



(b) Actor reconstruction

NEXT STEPS

Next class: Pollux