

CS 744: MESOS

Shivaram Venkataraman

Fall 2020

ADMINISTRIVIA

- Assignment 1: How did it go?
- Assignment 2 out tonight
- Project details
 - Create project groups
 - Bid for projects/Propose your own
 - Work on Introduction

COURSE FORMAT

Paper reviews

“Compare, contrast and evaluate research papers”

Discussion

Applications

Machine Learning

SQL

Streaming

Graph

Computational Engines

Scalable Storage Systems

Resource Management



Datacenter Architecture






MapReduce


GFS

Spark

BACKGROUND: OS SCHEDULING



CLUSTER SCHEDULING



TARGET ENVIRONMENT

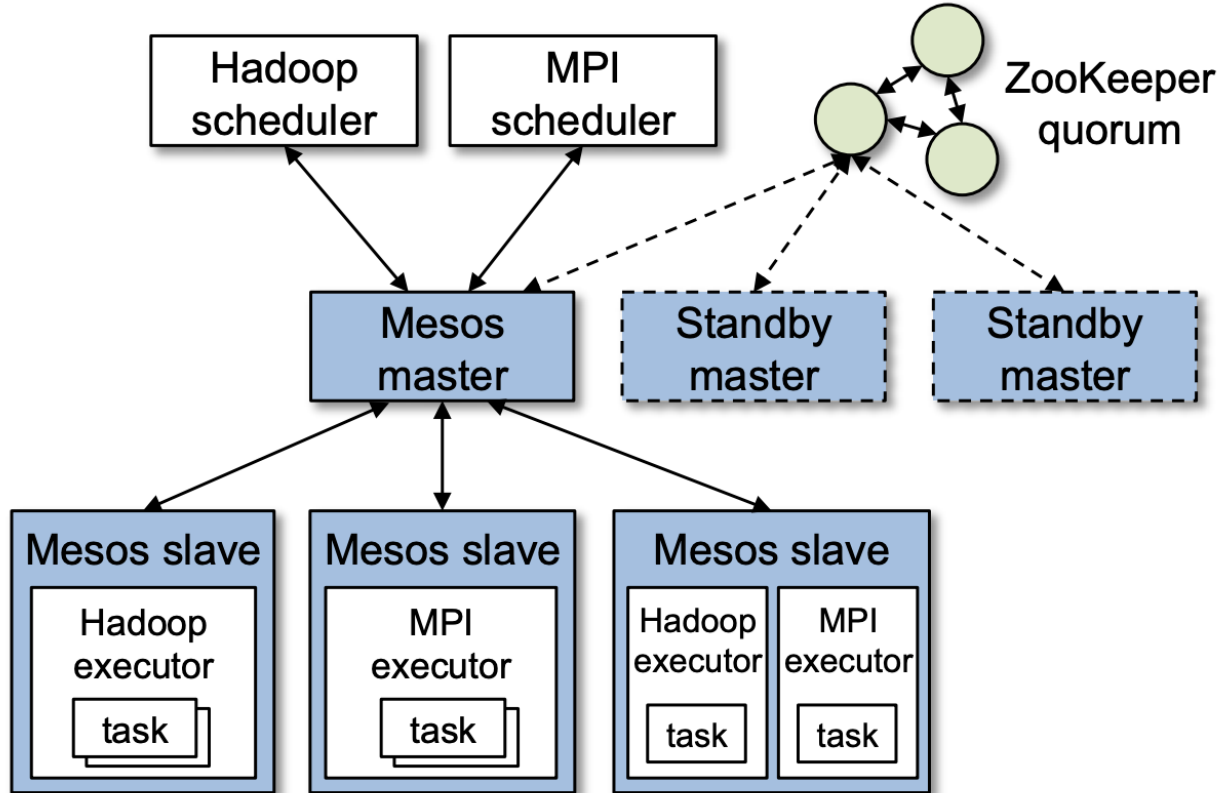
Multiple MapReduce versions

Mix of frameworks: MPI, Spark, MR

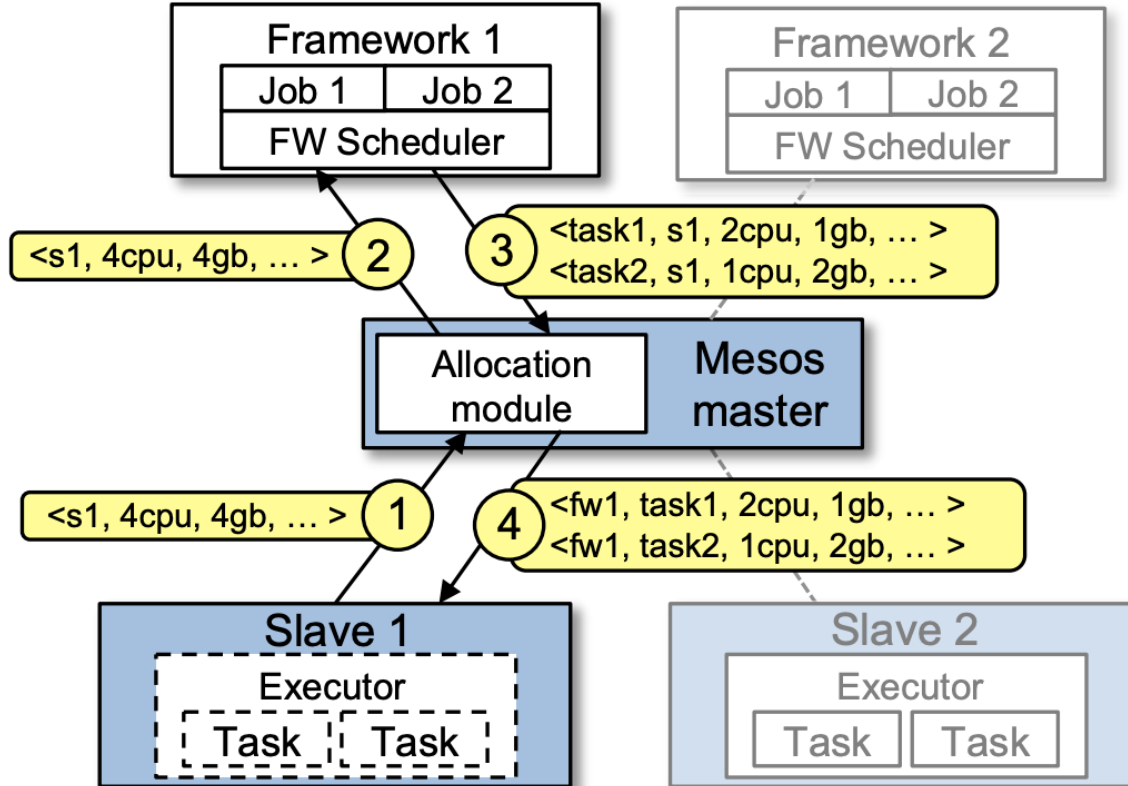
Data sharing across frameworks

Avoid per-framework clusters

DESIGN



RESOURCE OFFERS



CONSTRAINTS

Examples of constraints

Constraints in Mesos:

DESIGN DETAILS

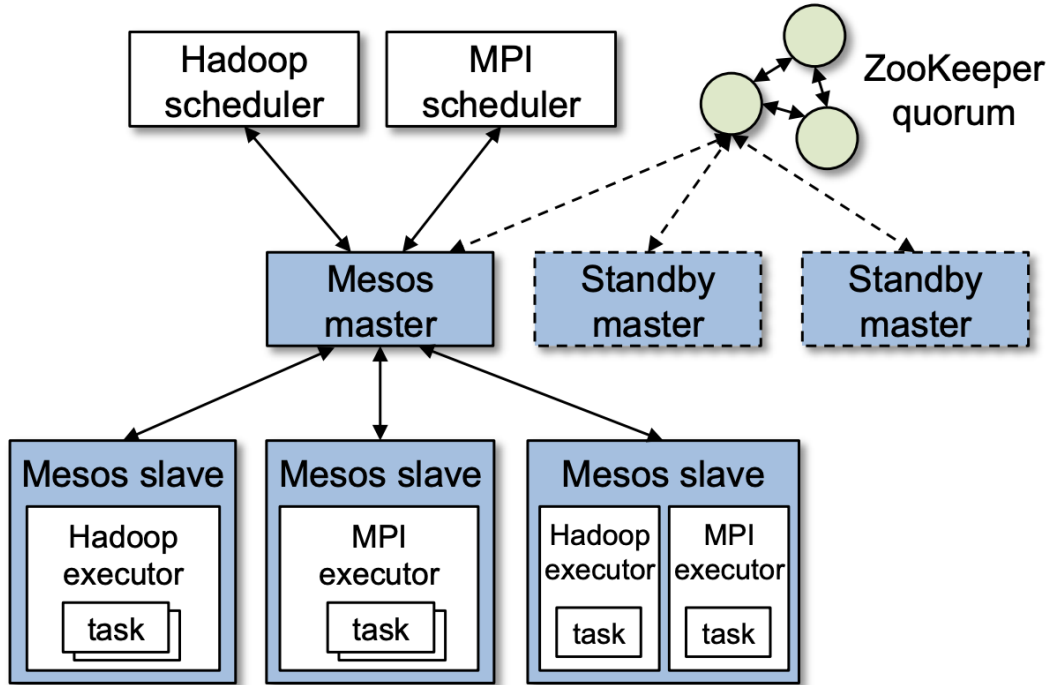
Allocation:

Guaranteed allocation, revocation

Isolation

Containers (Docker)

FAULT TOLERANCE



PLACEMENT PREFERENCES

What is the problem?

How do we do allocations?

CENTRALIZED VS DECENTRALIZED

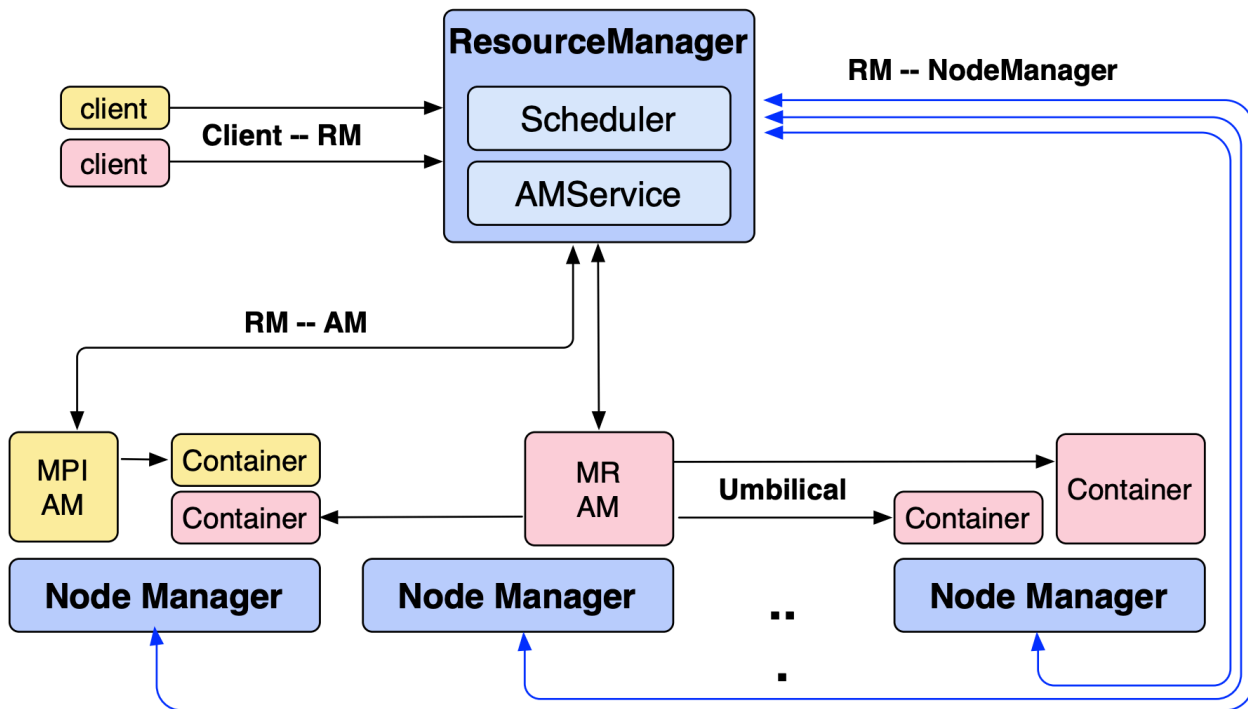
CENTRALIZED VS DECENTRALIZED

Framework complexity

Fragmentation, Starvation

Inter-dependent framework

COMPARISON: YARN



Per-job scheduler

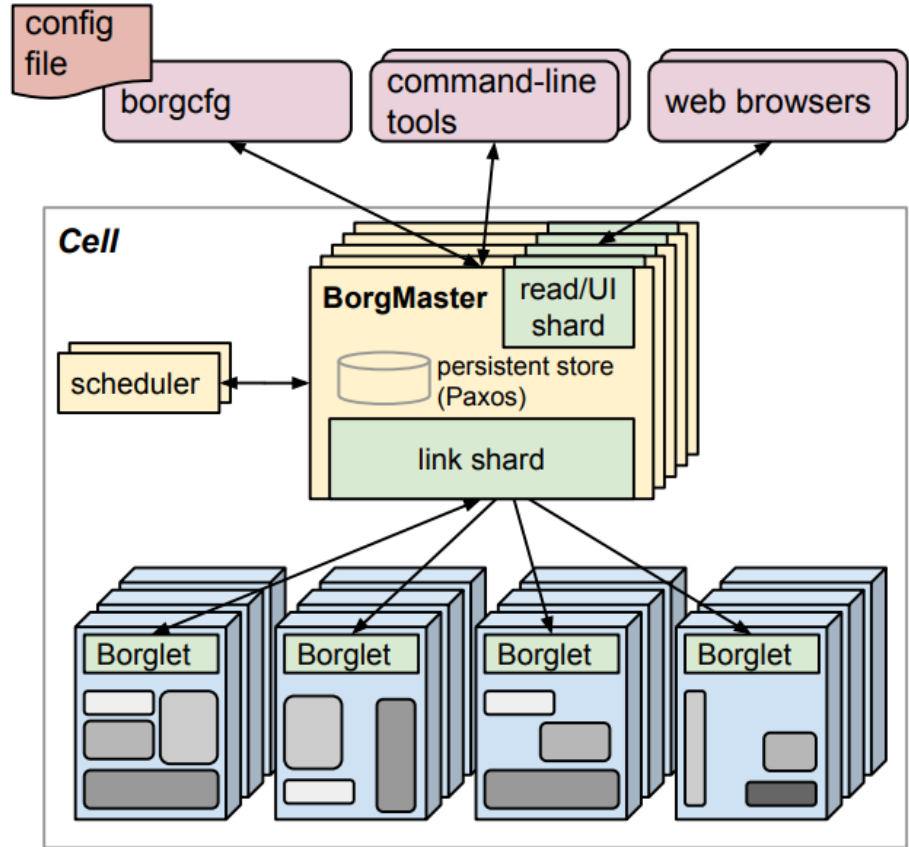
AM asks for resource
RM replies

COMPARISON: BORG

Single centralized scheduler

Requests mem, cpu in cfg
Priority per user / service

Support for quotas / reservations



SUMMARY


- Mesos: Scheduler to share cluster between Spark, MR, etc.
- Two-level scheduling with app-specific schedulers
- Provides scalable, decentralized scheduling
- Pluggable Policy ? Next class!

DISCUSSION


<https://forms.gle/urHSeukfyipCKjue6>

What are some problems that could come up if we scale from 10 frameworks to 1000 frameworks in Mesos?


(a) Facebook Hadoop Mix




(b) Large Hadoop Mix



(c) Spark



(d) Torque / MPI



List any one difference between an OS scheduler and Mesos

NEXT STEPS

Next class: Scheduling Policy

Further reading

- <https://www.umbrant.com/2015/05/27/mesos-omega-borg-a-survey/>
- <https://queue.acm.org/detail.cfm?id=3173558>