CS 744: BIG DATA SYSTEMS

Shivaram Venkataraman

Fall 2018

ADMINISTRIVIA

- Assignment 2
- Project meetings

SETTING: FAIR SHARING

Equal Share

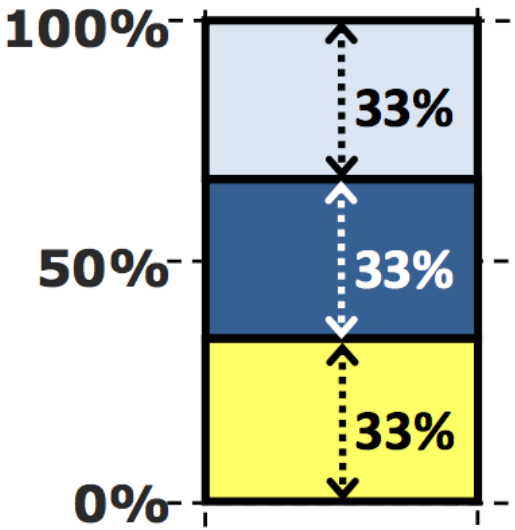

Max-Min Share

Maximize the allocation for most poorly treated users

Maximize the minimum

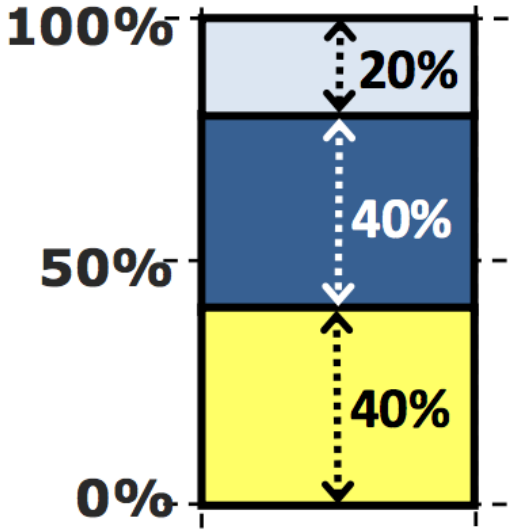

MOTIVATION: MULTI RESOURCES

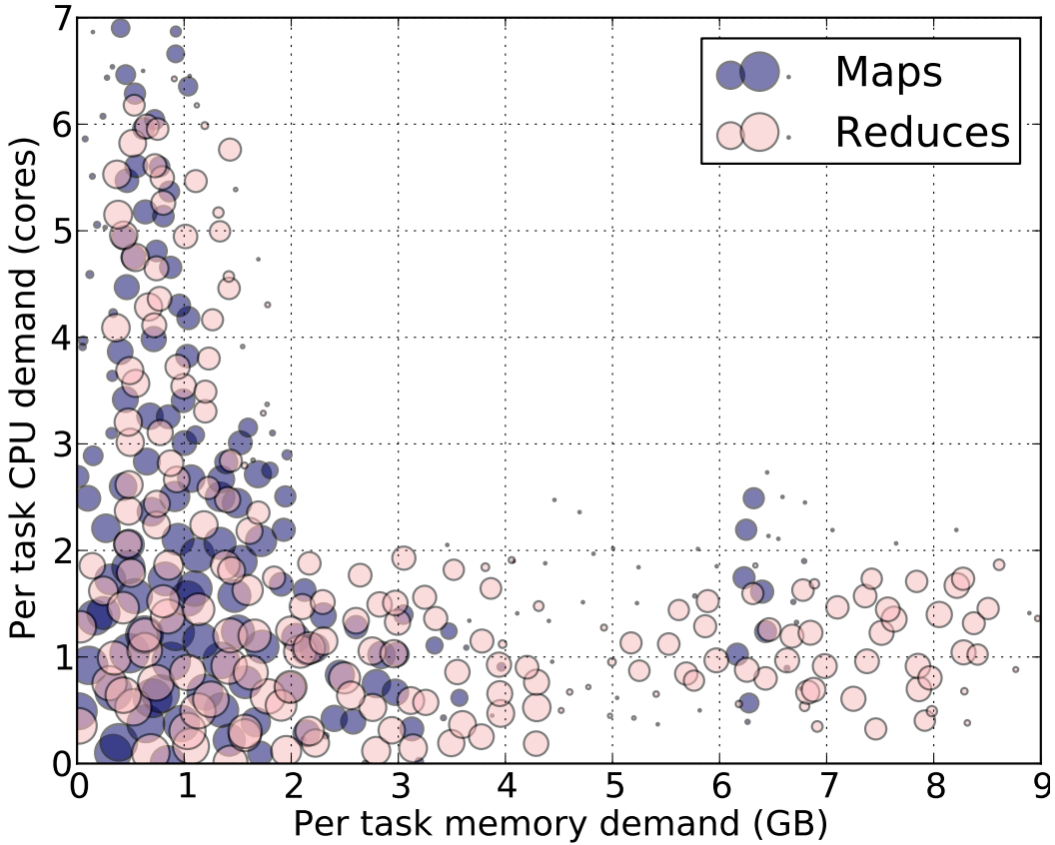

PROPERTIES

Sharing Incentive

User is no worse off than a cluster with 1/n resources

Strategy Proof

User should not benefit by lying about demands

Pareto Efficiency

Not possible to increase one user without decreasing another

Envy free

User should not desire the allocation of another user

DRF: MODEL

Users have a demand vector

<2, 3, 1> means user's task needs 2 R1, 3 R2, 1 R3

Resources given in multiples of demand vector

i.e., users might get <4,6,2>

DRF: APPROACH

Dominant Resource

Resource user has the biggest share of

Total: <10 CPU, 4 GB>

User 1: <1 CPU, 1 GB>

Dominant resource is memory

Dominant Share

Fraction of the dominant resource user is allocated

E.g., for User 1 this is 25% or 1/4

DRF: APPROACH

Equalize the dominant share of users

Total: <9 CPU, 18 GB>

User1: <1 CPU, 4 GB>
dom res: mem
User2: <3 CPU, 1 GB>
dom res: CPU

User	Allocation	Dominant Share
User1	<0 CPU, 0 GB>	0
User2	<0 CPU, 0 GB>	0

DRF: APPROACH

Total: <9 CPU, 18 GB>

User1: <1 CPU, 4 GB>
dom res: mem
User2: <3 CPU, 1 GB>
dom res: CPU

User	Allocation	Dominant Share
User1	<1 CPU, 4 GB>	2/9
User2	<0 CPU, 0 GB>	0

DRF: APPROACH

Total: <9 CPU, 18 GB>

User1: <1 CPU, 4 GB>
dom res: mem
User2: <3 CPU, 1 GB>
dom res: CPU

User	Allocation	Dominant Share
User1	<1 CPU, 4 GB>	2/9
User2	<3 CPU, 1 GB>	1/3

DRF: APPROACH

Total: <9 CPU, 18 GB>

User1: <1 CPU, 4 GB>
dom res: mem
User2: <3 CPU, 1 GB>
dom res: CPU

User	Allocation	Dominant Share
User1	<2 CPU, 8 GB>	8/18 = 4/9
User2	<3 CPU, 1 GB>	1/3

DRF: APPROACH

Total: <9 CPU, 18 GB>

User1: <1 CPU, 4 GB>
dom res: mem
User2: <3 CPU, 1 GB>
dom res: CPU

User	Allocation	Dominant Share
User1	<2 CPU, 8 GB>	4/9
User2	<6 CPU, 2 GB>	6/9 = 2/3

DRF: APPROACH

Total: <9 CPU, 18 GB>

User1: <1 CPU, 4 GB>
dom res: mem
User2: <3 CPU, 1 GB>
dom res: CPU

User	Allocation	Dominant Share
User1	<3 CPU, 12 GB>	12/18 = 2/3
User2	<6 CPU, 2 GB>	2/3

DRF: APPROACH

Total: <9 CPU, 18 GB>

User1: <1 CPU, 4 GB> per task
<3 CPU, 12 GB> for 3 tasks
dom res: mem
dom share: 12/18 = 2/3

User2: <3 CPU, 1 GB>
<6 GPU, 2 GB> for 2 tasks
dom res: CPU
dom share: 6/9 = 2/3

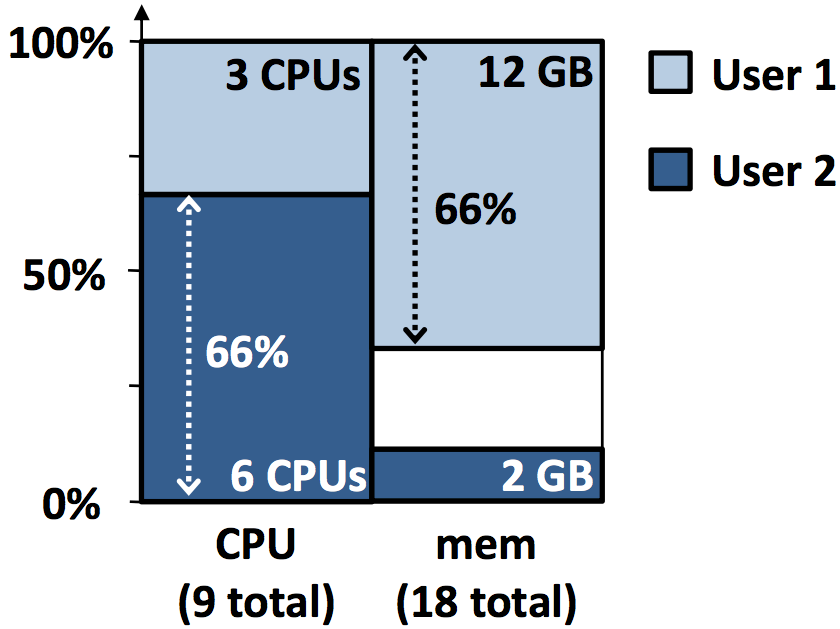

DRF SCHEDULING ALGORITHM

Whenever there are available resources:

Schedule a task to the user with smallest dominant share

COMPARISON: ASSET FAIRNESS

Asset Fairness: Equalize each user's sum of resource shares

Violates Sharing Incentive

Consider total of 70 CPUs, 70 GB RAM
U1 needs <2 CPU, 2 GB RAM> per task
U2 needs <1 CPU, 2 GB RAM> per task

Asset Fair Allocation:
U1: 15 tasks: 30 CPU, 30 GB (Sum = 60)
U2: 20 tasks: 20 CPU, 40 GB (Sum = 60)

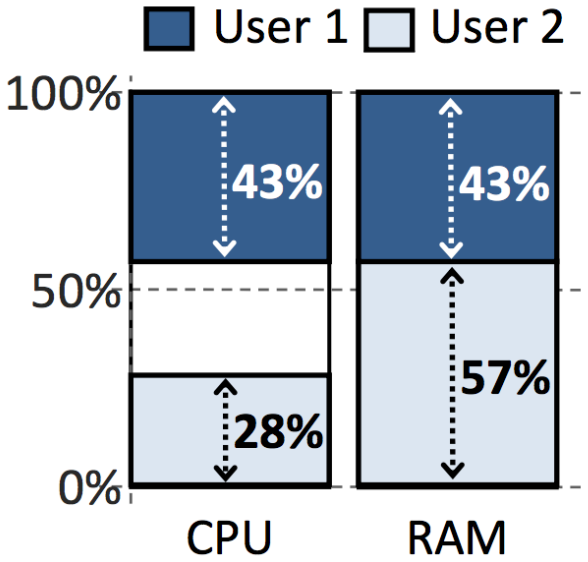

COMPARISON: CEEI

CEEI: Competitive Equilibrium from Equal Incomes

- Each user receives initially 1/n of every resource,
- Subsequently, each user can trade resources with other users in a perfectly competitive market
- Computed by maximizing product of utilities across users

COMPARISON: CEEI

Total: <9 CPU, 18 GB>

User1: <1 CPU, 4 GB>
User2: <3 CPU, 1 GB>

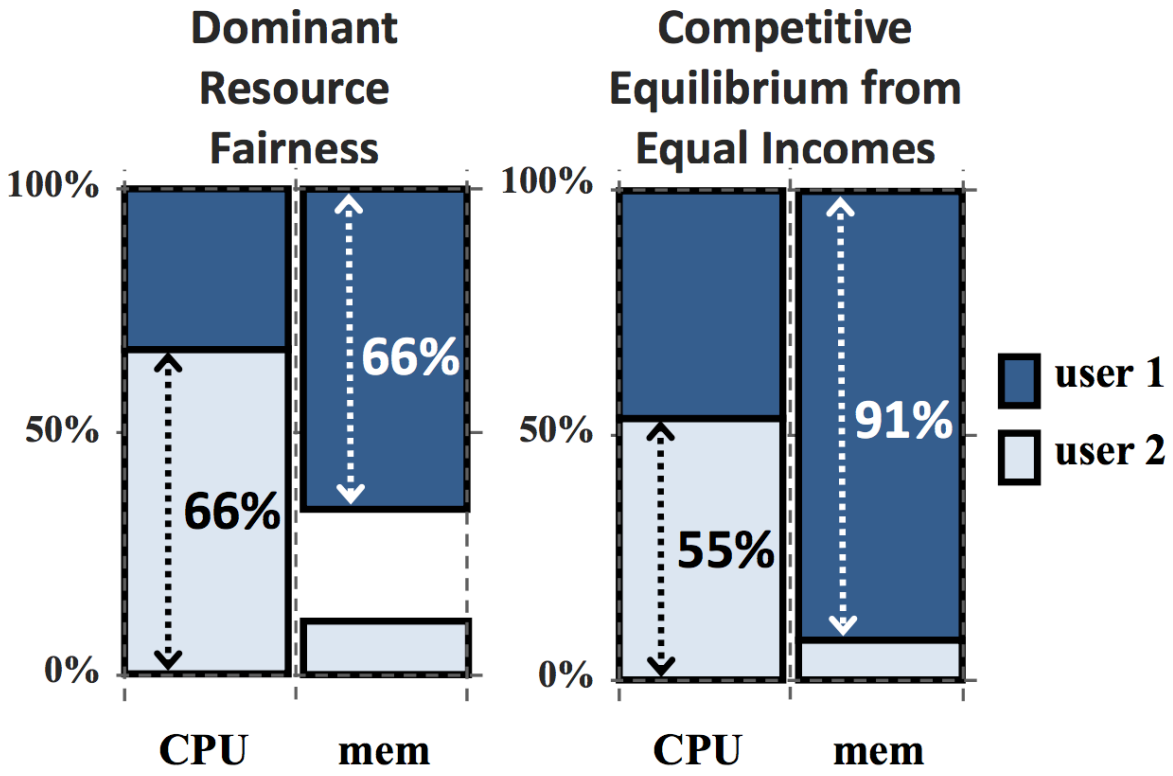

COMPARISON: CEEI

Total: <9 CPU, 18 GB>

User2 Before:

DRF: 66% CPU, 16% mem

CEEI: 55% CPU, 9% mem

Total: <9 CPU, 18 GB>

User1: <1 CPU, 4 GB>

User2: <3 CPU, 2 GB>

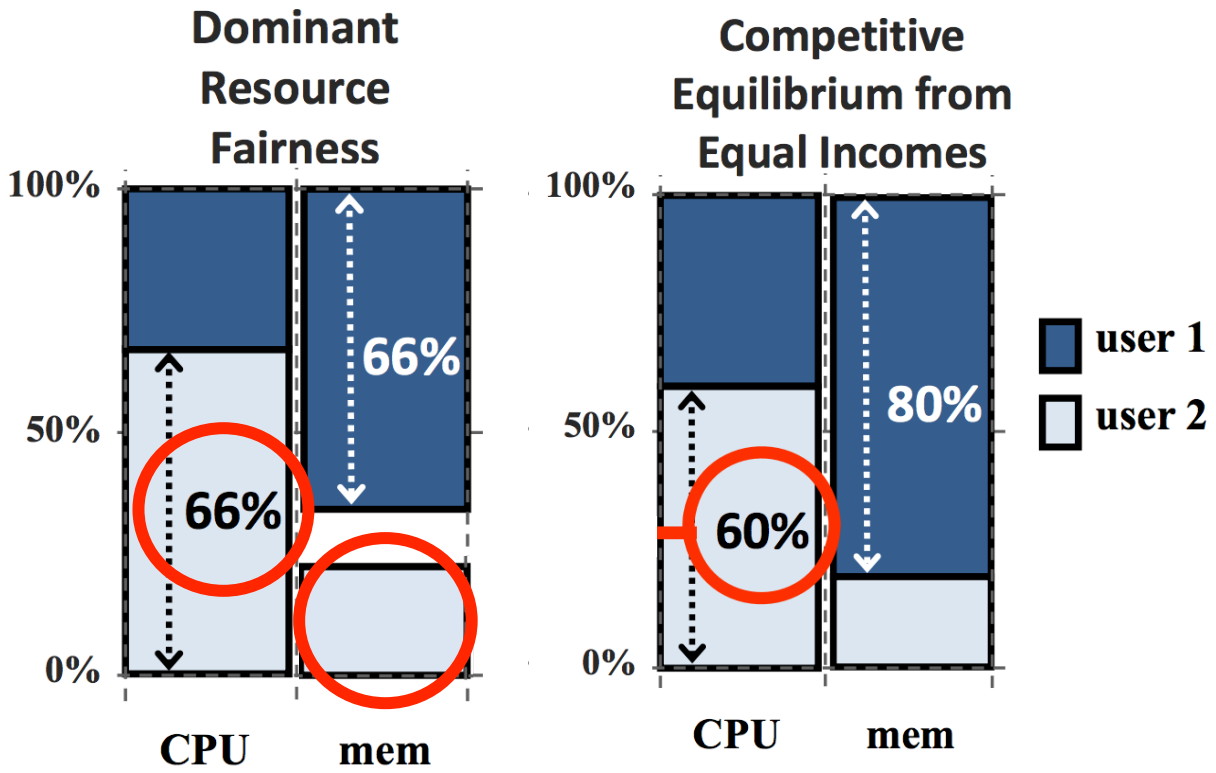

COMPARISON

Property	Allocation Policy		
	Asset	CEEI	DRF
Sharing Incentive		✓	✓
Strategy-proofness	✓		✓
Envy-freeness	✓	✓	✓
Pareto efficiency	✓	✓	✓
Single Resource Fairness	✓	✓	✓
Bottleneck Fairness		✓	✓
Population Monotonicity	✓		✓
Resource Monotonicity			

Table 2: Properties of Asset Fairness, CEEI and DRF.

QUESTIONS / DISCUSSION ?