CS 744: DRF

Shivaram Venkataraman

Fall 2020

ADMINISTRIVIA

- Assignment 2 out!
- Course Project
  - Form groups?
  - Project list by Monday (9/28)
  - Submit project bids by Thursday (10/1)
  - Assigned project by Friday (10/2)

SETTING: FAIR SHARING

Equal Share

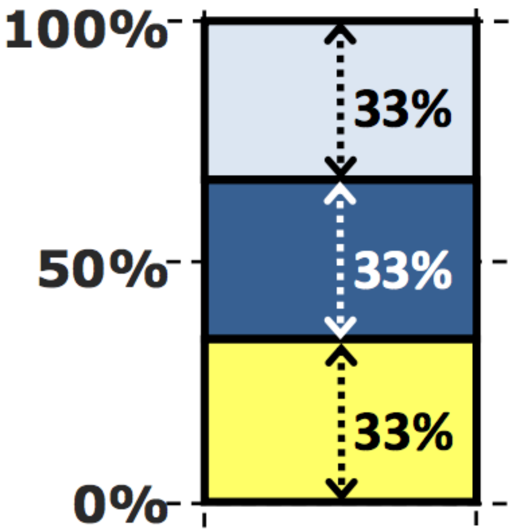

Max-Min Share

Maximize the allocation for most poorly treated users

Maximize the minimum

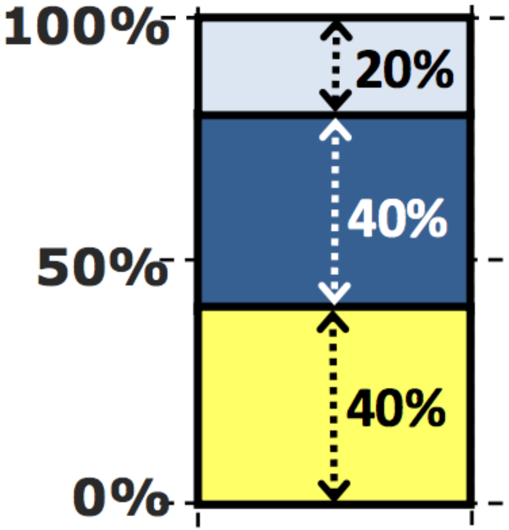

MOTIVATION: MULTI RESOURCES

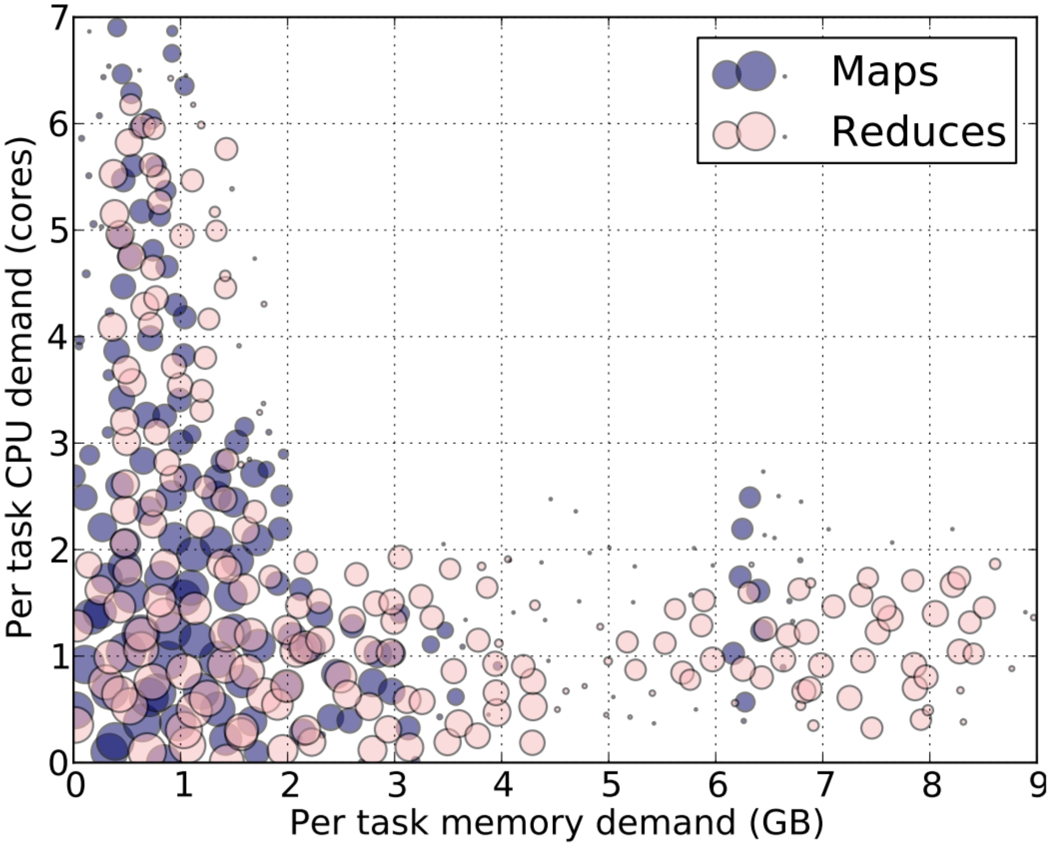

DRF: MODEL

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>

PROPERTIES

Sharing Incentive

Strategy Proof

Pareto Efficiency

Envy free

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: 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> 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

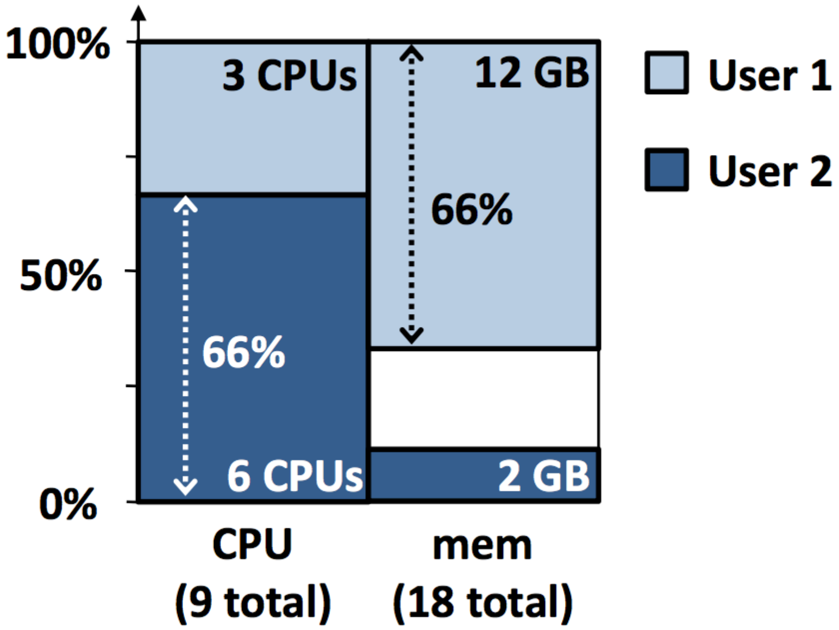

DRF ALGORITHM

Whenever there are available resources:

Schedule a task to the user with smallest dominant share

DRF ALGORITHM

Algorithm 1 DRF pseudo-code

$R = \langle r_1, \cdots, r_m \rangle$ ▷ total resource capacities
$C = \langle c_1, \cdots, c_m \rangle$ ▷ consumed resources, initially 0
$s_i \ (i = 1..n)$ ▷ user i's dominant shares, initially 0
$U_i = \langle u_{i,1}, \cdots, u_{i,m} \rangle \ (i = 1..n)$ ▷ resources given to user i, initially 0

pick user i with lowest dominant share s_i
$D_i \leftarrow$ demand of user i's next task
if $C + D_i \leq R$ **then**
    $C = C + D_i$ ▷ update consumed vector
    $U_i = U_i + D_i$ ▷ update i's allocation vector
    $s_i = \max_{j=1}^{m}\{u_{i,j}/r_j\}$
else
    return ▷ the cluster is full
end if

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:
U2:

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)

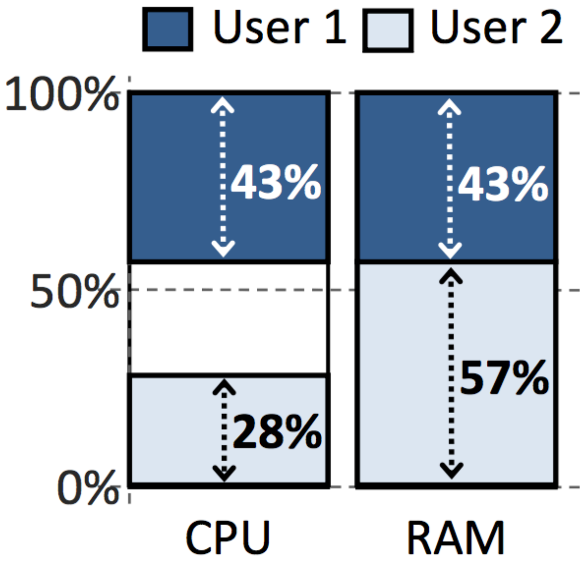

COMPARISON: CEEI

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>

$$\max\,(x \cdot y)$$

subject to

$$\begin{aligned} x + 3y &\leq 9 \\ 4x + y &\leq 18 \end{aligned}$$

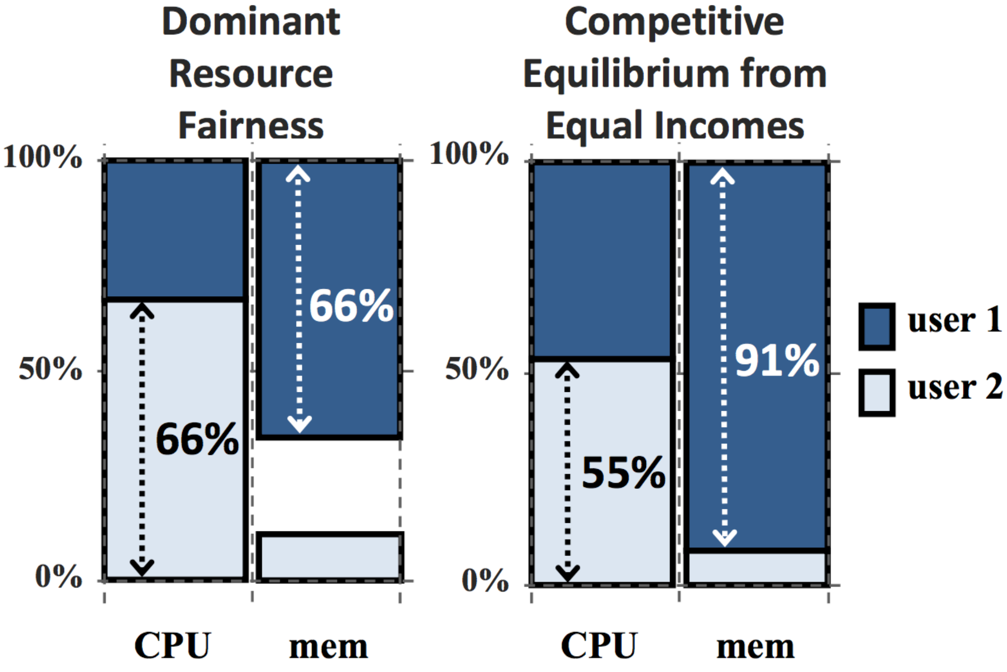

CEEI: STRATEGY PROOFNESS

Total: <9 CPU, 18 GB>

User2 Before:
CEEI: 55% CPU, 9% mem

Total: <9 CPU, 18 GB>

User1: <1 CPU, 4 GB>
User2: <3 CPU, 2 GB>

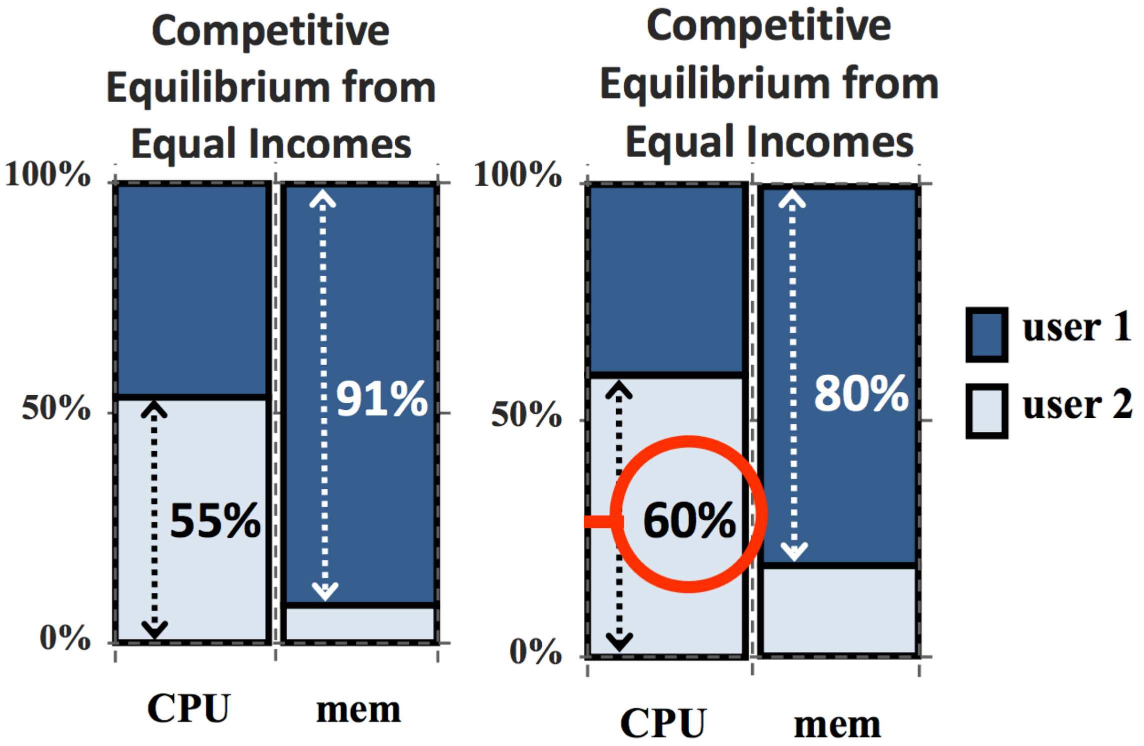

COMPARISON

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.

SUMMARY

DRF: Dominant Resource Fairness

Allocation policy for scheduling

Provides multi-resource fairness

Ensures sharing incentive, strategy proofness

DISCUSSION

https://forms.gle/i7m7xXxKhtfvL9UD9

Consider a system with 100 units of CPU, 50 units of memory and 200 units of disk. Consider three users with the following requirements

Alice (4 CPU, 1 memory, 1 disk)

Bob (1 CPU, 4 memory and 4 disk)

Carol (1 CPU, 2 memory and 16 disk)

List the dominant resource as defined in DRF for Alice, Bob and Carol

What would be the final task allocation in the given cluster for Alice, Bob and Carol ?

What could be one workload / cluster scenario where DRF implemented on Mesos will NOT be optimal?

NEXT STEPS

Next Week: Machine Learning

Assignment 2 out!