

REGULAR ARTICLE

Multi-resource Allocation in Cloud Data Centers: A Trade-off on Fairness and Efficiency

Suhan Jiang | Jie Wu*

Department of Computer and Information Sciences, Temple University, Pennsylvania, USA

Correspondence

*Jie Wu, Email: jiewu@temple.edu

Present Address

1925 N. 12th Street, Philadelphia, PA, 19122

Summary

Fair allocation has been studied intensively in both economics and computer science. Many existing mechanisms that consider fairness of resource allocation focus on a single resource. With the advance of cloud computing that centralizes multiple types of resources under one shared platform, multi-resource allocation has come into the spotlight. In fact, fair/efficient multi-resource allocation has become a fundamental problem in any shared computer system. The widely-used solution is to partition resources into bundles that contain fixed amounts of different resources, so that multiple resources are abstracted as a single resource. However, this abstraction cannot satisfy different demands from heterogeneous users, especially on ensuring fairness among users competing for resources with different capacity limits. A promising approach to this problem is dominant resource fairness (DRF), which tries to equalize each user's dominant share (share of a user's most highly demanded resource, *i.e.*, the largest fraction of any resource that the user has required for a task), but this method may still suffer from significant loss of efficiency (*i.e.*, some resources are underused). This paper develops a new allocation mechanism based on DRF aiming to balance fairness and efficiency. We consider fairness not only in terms of a user's dominant resource, but also in another resource dimension which is secondarily desired by this user. We call this allocation mechanism 2-dominant resource fairness (2-DF). Then, we design a non-trivial on-line algorithm to find a 2-DF allocation and extend this concept to k -dominant resource fairness (k -DF).

KEYWORDS:

Allocation, efficiency, fairness, multi-resource

1 | INTRODUCTION

Resource allocation is one of the central topics in the field of computer science. There are many policies that govern resource allocation to achieve fairness among users. Max-min fairness, one of the most popular allocation policies, tries to maximize the allocation for the most poorly treated users, *i.e.*, maximize the minimum. Weighted max-min fairness adds a new concept called weight based on max-min fairness, and assigns each user with a share of the resources according to a preset weight. However, an obvious limitation in most of the existing works is that they are only devoted to single-resource allocation when quantifying the notion of fairness, *e.g.* by allocating available link bandwidth to network flows.

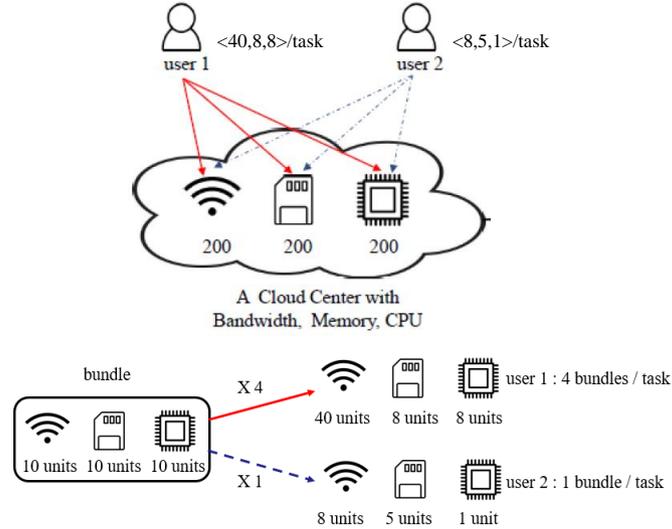


FIGURE 1 A multi-resource allocation setting using resource abstraction.

Fair (or efficient) multi-resource allocation is a fundamental problem in any shared computer system. A typical example is data centers that process numerous jobs with heterogeneous resource requirements on bandwidth, memory, and CPU etc. Both Hadoop and Dryad employed a simple solution based on resource abstraction. As is shown in Figure 1, all resources are partitioned into bundles with fixed amounts of different resources, so that multiple resources are abstracted as a single resource. However, this simple resource abstraction ignores the different demands of heterogeneous users, and cannot always match nicely with user demands.

Ghods *et al.*¹ first put forward a compelling approach to this problem, which is known as dominant resource fairness (DRF). In brief, DRF allocates resources according to users' proportional demands, applying max-min fairness to each user's *dominant share*. Dominant share is the maximum share that a user has been allocated of any resource. Such a resource is then called a *dominant resource*. Although DRF has attracted much attention, this allocation approach has been questioned continuously for the reasons given below: (1) *fairness dispute* - DRF only considers one dimension when allocating all resources. Once the allocation of a user's dominant resource is determined, resources in other dimensions are proportionally assigned according to the user's request. (2) *efficiency loss* - Jin *et al.*² and Bertsimas *et al.*³ respectively showed that proportional fairness, which maximizes the sum of the log of completed tasks of different users, always more efficiently uses resources than DRF does.

1.1 | Motivation

We use the example shown in Figure 1 to illustrate existing problems in the DRF allocation mechanism. Given a system with three resources $\langle \text{Bandwidth}, \text{Memory}, \text{CPU} \rangle$ and two users $\langle \text{user1}, \text{user2} \rangle$, the capacity of each resource is 200 units respectively. User 1 executes each task with the request vector $\langle 40, 8, 8 \rangle$, while user 2 requires $\langle 8, 5, 1 \rangle$ for each task. According to DRF, user 1 gets an allocation of $\langle 100, 20, 20 \rangle$, and user 2 gets $\langle 100, 62.5, 12.5 \rangle$. The resulting allocation is shown in Figure 2.

In the *fairness* domain, it is obvious that the amount of resources allocated to each user is only dependent on his dominant resource request. According to the blue line shown in Figure 4, if user 2's request on memory (which isn't his dominant resource) varies, there is no change on the number of total tasks he can complete. In the *efficiency* domain, if we define efficiency in terms of the aggregate tasks of all users, this example also reflects DRF's deficiency. See the results in Figure 3. In this setting, DRF produces an allocation with 2.5 tasks for user 1 and 12.5 tasks for user 2. This allocation brings about a significant loss in system efficiency. If we assign 2 tasks to user 1 and 15 tasks to user 2, this allocation yields 17 tasks in total, which is more efficient than DRF. Consider an extremely unequal allocation, where all resources are allocated to user 2, it will produce 25 tasks completed in total. In addition, we can also measure efficiency based on the amount of leftover resource capacity. In this example, DRF yields 15 tasks, leaving 285 units of resources unused in total. If we assign 22.5 tasks to user 2, it also leads to 285 units of resources left unused. It is obvious, resource utility is higher in the second allocation. In fact, each of the existing mechanisms designed for multi-resource allocations represents one point of the fairness-efficiency tradeoff. Besides, this example motivates us to think

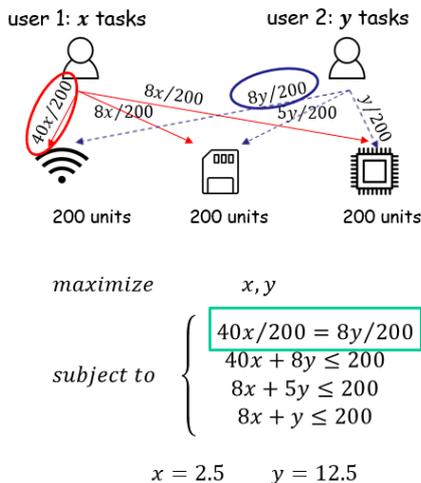


FIGURE 2 An example of dominant resource fairness allocation.

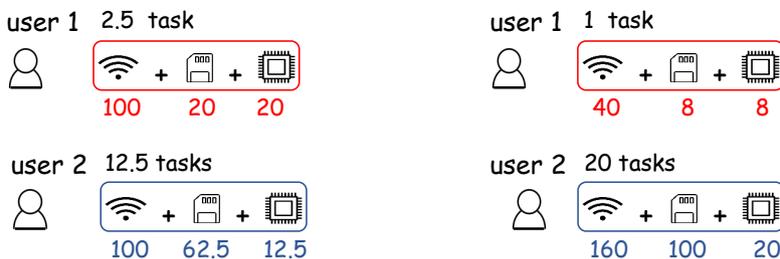


FIGURE 3 Efficiency loss in dominant resource fairness allocation: 15 tasks are completed under DRF (on the left side) while the maximal number pf completed tasks is 21 (on the right side).

about if we could design a new mechanism to reach the following goals: (1) balance the number and the value of resources given to user 1 and user 2, in order to achieve the fairness in the traditional sense; (2) increase total efficiency of resource utility, in order to complete more tasks. While fairness is a basic requirement of different users for resource competition, efficiency is the most desirable property for a system provider, with the efficiency expressed in terms of the aggregate tasks of all users. Hence, a key challenging issue is how to balance these two factors with the desired performance satisfactory. Unfortunately, though generally applicable to multi-resource environments, DRF still performs poorly when efficiency is a concern.

1.2 | Our Result

To some extent, the previous example shows the fundamental tradeoff between fairness and efficiency: fairness and efficiency cannot be achieved simultaneously. In this paper, we seek to answer a fundamental question of resource management: how to allocate multi-type resources among users with heterogeneous demands, in an attempt to balance two opposing factors - efficiency and fairness.

Based on this question, we propose a new allocation mechanism called 2-dominant resource fairness (2-DF). A concept called 2-dominant share (denoted by s) is used. s_i is defined for each user i as the product of first two dominant demand/capacity ratios. Compared with dominant resource share, 2-dominant share is a better reflection of a user’s true demand of all resources. Similarly to DRF, our allocation mechanism tries to equalize each user’s 2-dominant share as much as possible. Further, we extend this 2-dominant resource fairness to k -dominant resource fairness (k -DF). We show our allocation mechanism can improve resource utility in most cases while still keeping some important fairness properties, and we also prove that our allocation mechanism satisfies both strategy-proofness and Pareto-optimality; meanwhile, we could achieve envy-freeness in some scenarios that appear quite frequently.

Our contributions in this paper are summarized as follows:

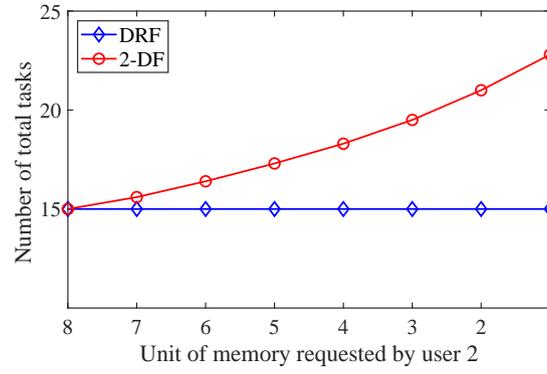


FIGURE 4 User changes his request on a non-dominant resource dimension.

- Motivated by the deficiency of DRF policy, we reconsider the definition of fairness and efficiency, and formulate several basic metrics to measure fairness and efficiency of a specific allocation mechanism in the multi-resource scenario.
- A new allocation mechanism called 2-dominant resource fairness is proposed to fairly and efficiently allocate multiple resources to users with heterogeneous demands.
- An online scheduling algorithm is designed to realize 2-DF allocation mechanism.
- We prove that in 2-DF allocation mechanism, strategy-proofness and Pareto-optimality can be guaranteed, and envy-freeness can be achieved in some scenarios that appear quite frequently.
- We simulate our algorithm in real-world scenarios, and compare the performances of our algorithm with the previous approaches according to our predetermined metrics.

The remainder of the paper is organized as follows. Section 2 briefly reviews the related works, and Section 3 formulates the metrics that are to measure fairness and efficiency in multi-resource allocation settings. Section 4 presents our model and related allocation mechanism. In Section 5, we prove some important fairness properties achieved in our mechanism. We discuss experiment results in Section 6, and conclude our paper in Section 7.

2 | RELATED WORK

Multi-resource allocation problems arise in increasingly many applications. Data centers with multiple resources have often employed a single resource abstraction by partitioning different resources into bundles. However, multi-resource allocation viewed as single-resource allocation inevitably leads to significant inefficiencies because of the heterogeneous user demands. Different approaches have been proposed to deal with multi-resource allocation problems. Many different dimensions have been taken into account, such as desirable allocation characteristics, utility functions used to measure happiness of users, and the step at which a resource allocation approach should be applied⁴.

As discussed earlier, Ghodsi *et al.*¹ proposed the DRF policy, which provides fair allocation of multiple resources in terms of dominant shares. It retains a number of desirable sharing properties and has been widely studied in both theory and practice. In reality, DRF has been modified and adopted by Apache Mesos. As a variation, Mesos doesn't directly allocate resources to tasks. It allocates resources to frameworks based on DRF, and then the frameworks run scheduling algorithms to schedule tasks using the allocated resources. As mentioned in⁵, once Mesos launches tasks, it does not terminate them even if the frameworks' dominant share exceeds the fair share limit of the cluster. This feature can lead other frameworks to starve.

Then, Ghodsi *et al.*⁶ extended DRF to packet networks and proposed DRFQ, the first fair multi-resource queuing algorithm. Gutman and Nisan⁷ considered generalizations of DRF in a more general utility model, such as the Leontief preferences. Joe-Wong *et al.*⁸ paid attention to the tradeoffs between fairness and efficiency, and generalized the DRF policy by designing a unifying multi-resource allocation framework. Parkes *et al.*⁹ extended DRF in several ways, and in particular studied the case of indivisible tasks.

Wang *et al.*¹⁰ and Friedman *et al.*¹¹ extended DRF's all-in-one resource model to distributed systems with heterogeneous machines, while Zeldes and Feitelson¹² proposed an on-line algorithm based on bottlenecks and global priorities. Kash *et al.*¹³

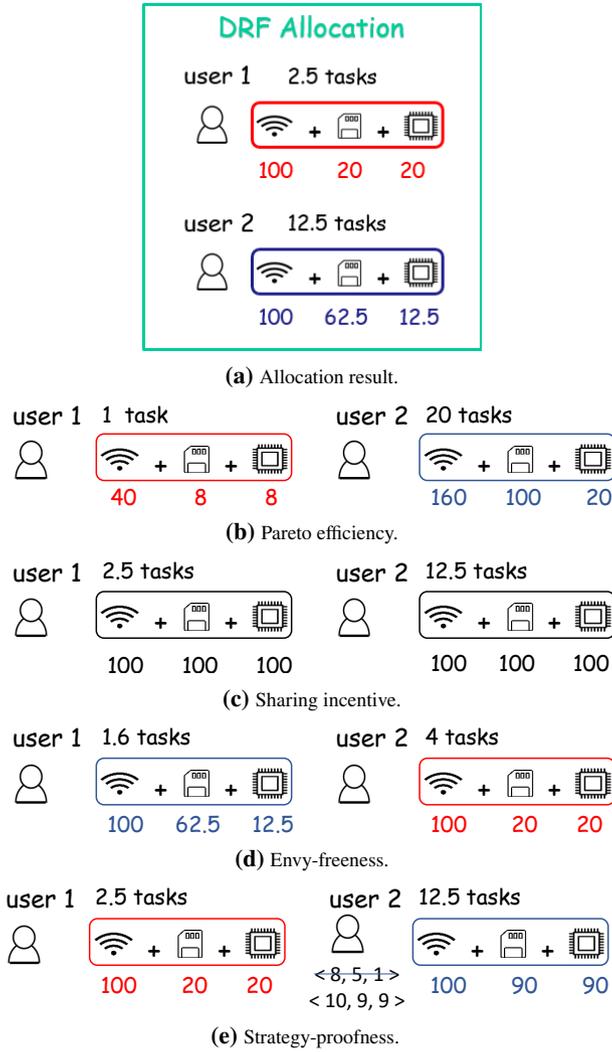


FIGURE 5 Examples of multi-resource fairness properties.

extended DRF to a dynamic setting, where users dynamically arrive over time but never depart. Also, Li *et al.*¹⁴ generalized the dynamic dominant resource fairness mechanism to the bounded case, where each user has a finite number of tasks. In their paper, a linear-time optimal algorithm is presented. Dolev *et al.*¹⁵, on the other hand, suggested a different fairness notion for multi-resource allocation based on fairly dividing a global system bottleneck resource. Zahedi and Lee¹⁶ applied the concept of Competitive Equilibrium from Equal Outcomes (CEEI) in the case of the CobbDouglas utilities to achieve properties similar to DRF.

Recently, a new mechanism called Greediness Metric Fairness has been developed^{17,18,19}, which can be applied to allocate physical resources by periodically adapting the priorities of virtual machines. This mechanism is highly flexible and applicable to scheduling¹, since it has no assumptions on utility functions. Besides, there appears a new trend to take advantage of machine learning to allocate multiple resources to satisfy user requests²⁰.

3 | METRICS ON FAIRNESS AND EFFICIENCY

When assessing the quality of an allocation, we can distinguish two types of indicators of social welfare: fairness and efficiency. While fairness is a basic requirement of different users for resource competition, efficiency is the most desirable property for a system provider. Hence, a key challenging issue is how to balance these two factors with the desired performance and user satisfaction. Before discussing our new allocation mechanism, we start with some widely-accepted definitions of fairness and

efficiency in the multi-resource allocation environments, and formulate some metrics that should be considered to measure a multi-resource allocation mechanism.

In terms of fairness, equal sharing seems to be treated as the conventional idea. However, it cannot always be a good interpretation of fairness even for a single-resource allocation. Given a total of 12-unit bandwidth and 3 network users, *A* needs 1.5 unit bandwidth for web-browsing, *B* needs 4.5 unit bandwidth for watching a film, and *C* needs 6 unit bandwidth to hold a video conference. Egalitarianism will lead to an allocation of 4 units for each user. 4 units for *A* are really a waste while for *B* and *C* are not enough. This equal allocation scheme produces a very low network resource utility. A better allocation can be (1.5, 4.5, 6) for *A*, *B*, *C*, which ensures that all users can be better off or at least no worse off than in the case of equal sharing. A major challenge of multi-resource fairness is incorporating the heterogeneity of different users' requirements for different resources into the assessment of its fairness. As mentioned in⁴, there are two main reasons why more complexity and less agreement are encountered in defining multi-resource fairness: one is that fairness is an intuitive concept, and the other is that the organization of resources also has influence on defining fairness. Instead of being stuck in different definitions of fairness, we list some acknowledged and important properties of a fair multi-resource allocation proposed by Ghodsi *et al.*¹.

Definition 1 (Pareto efficiency). Increasing the allocation of a user will lead to decreasing the allocation of at least another user. This means no user can run more tasks without harming someone else's benefits.

Definition 2 (Sharing incentive). Each user should be at least no worse off by sharing resources compared with exclusively equally partitioning each resource. Given a set of resources and n users, each user should be able to run more tasks if they share resources.

Definition 3 (Envy-freeness). A user should not prefer the allocation of another user. Changing her current allocation with that of anyone else would not improve her total task number.

Definition 4 (Strategy-proofness). Users should not be able to benefit from lying about their resource demands, which means a user cannot run more tasks by lying.

It is obvious that DRF allocation mechanism can satisfy all the properties above. See Figure 5b, if we improve user 2's utility by adding his tasks, then user 1's tasks will definitely decrease. This is so-called Pareto efficiency. Similarly, both user 1 and user 2 will be as good as the condition where all resources are equally allocated to them. This is an example for Sharing incentive. In terms of Envy-freeness, we can find, if user 1 and user 2 exchange their current allocations, both of them will finish fewer tasks, as is shown in Figure 5d. And from Figure 5e, if user 2 cheats on his resource demand, he still only finishes 12.5 tasks as before. Thus, DRF is also subject to Strategy-proofness.

Besides, we want to mention two more important concepts, which are important for an allocation mechanism to achieve so-called fairness in the traditional sense.

- multiple dimensions of resources: when calculating the resource allocation, the mechanism should not consider only one dimension of all resources.
- non-dominant resources: fairness should consider the consumption of non-dominant resources as well. If two different tasks have the same demand of a given dominant resource, then the one having smaller consumption of non-dominant resources should receive some compensation. Further, if a task has a totally lower demand of all given resources compared with other tasks, then it should receive more compensation.

It is also not easy to measure efficiency in the multi-resource allocation setting. In a single-resource scenario, the most efficient allocation will clearly use the entire resources and thus achieve the maximal number of tasks. On this basis, we list two metrics to measure multi-resource allocation efficiency, where an example for the further explanation can be seen in Figure 6.

- the number of total tasks completed (NTT): higher resource efficiency is achieved by more tasks done by all users.
- the amount of unused resources (AUR): after all required resources are allocated, more unused resources allow a datacenter to serve forthcoming users.

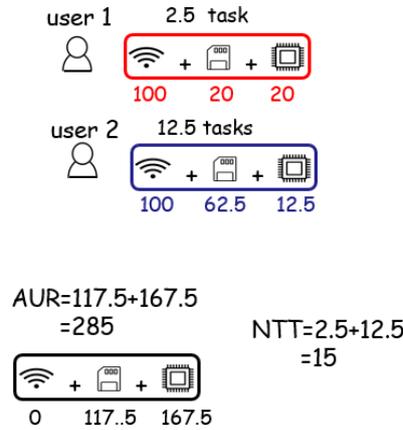


FIGURE 6 Example of efficiency measurements.

4 | 2-DOMINANT RESOURCE FAIRNESS

In this section, we propose a new allocation mechanism called 2-dominant resource fairness (2-DF) for heterogeneous users with different requests for distinct resources.

4.1 | Background and Model

In our model, we follow Ghodsi *et al.*¹ and assume that resources of the same type are assembled in homogeneous pools. Consider r infinitely divisible resources, the capacity of each resource j is C_j . There are n users indexed by i . Each user runs many parallel tasks. Each task is characterized by a request vector $D_i = [a_{i1}, a_{i2}, \dots, a_{ir}]$. This request vector specifies how many units of each type of resource are needed when running such a task. Usually, the tasks from a user are typically the same binary program running on different data blocks of similar sizes, then they require the same amount of resources¹⁰. We therefore assume the same request vector across a user's tasks and that each user i requires an amount of each resource in fixed proportion. User i 's final allocation is defined by a vector $A_i = [\varphi_i a_{i1}, \varphi_i a_{i2}, \dots, \varphi_i a_{ir}]$, where $\varphi_i a_{ij}$ represents the fraction of resource j allocated to user i .

4.2 | Multi-resource Fairshare Function

As discussed in the previous section, dominant share is the core part of DRF. Thus, we consider 'fairshare' which is a generalization of dominant share for each user. When allocating resources, we can easily apply max-min fairness to each user's fairshare. The drawback of dominant share is that it only considers one resource for a user and cannot allow the user to express how important this resource is in comparison with other resources.

In the following part, we propose two major rules that should be considered when defining a fairshare function in any multi-resource allocation mechanism. Then we let these rules guide us to design our allocation mechanism.

- multiple dimensions of resources: when calculating the resource allocation, a fairshare function should not consider only one dimension of all resources. It should reflect the demand of non-dominant resources as well. If two different tasks have the same demand of a given dominant resource, then the one having smaller demand of non-dominant resources should receive some compensation. Further, if a task has a totally lower demand of all given resources compared with other tasks, then it should receive more compensation.
- weights among different resources: for a given resource, a fairshare function should allow users to determine the weight to stress how important this resource is.

4.3 | 2-Dominant Resource Fairness

In this part, we continue the previous model built on our mechanism and introduce a fairshare function called 2-dominant share in our allocation mechanism.

According to user i 's request vector $D_i = [a_{i1}, a_{i2}, \dots, a_{ir}]$, let $d_{i1} = \max \left\{ \frac{a_{ij}}{c_j} \right\}$ where $j \in [1, r]$ be user i 's first dominant request ratio and $d_{i2} = \max \left\{ \frac{a_{ij}}{c_j} \right\} - \{d_{i1}\}$ where $j \in [1, r]$ be i 's second dominant request ratio. The resources corresponding to i 's dominant requests are called her dominant resources. Here, we define that each user i has a 2-dominant share expressed as $\varphi_i \cdot d_{i1} \cdot d_{i2}$.

Definition 5 (2-dominant share). The 2-dominant share of a user i is defined as.

$$s_i = \varphi_i \cdot d_{i1} \cdot d_{i2} \quad (1)$$

In the 2-DF, we are trying to apply max-min fair allocation with respect to the users' 2-dominant share. That is, we always maximize the lowest 2-dominant share first followed by the second lowest, etc. Now, we present an example to illustrate how our 2-DF allocation mechanism works.

An Example. We still use the example from Figure 1 to illustrate how 2-DF allocates resources according to users' different requests. In the above scenario, each task from user 1 consumes $\frac{1}{5}$ of bandwidth, $\frac{1}{25}$ of memory and $\frac{1}{25}$ of CPU, so user 1's first and second dominant requests lie on bandwidth and memory (or CPU), respectively. Similarly, user 2's first and second dominant requests lie on bandwidth and memory, respectively. 2-dominant fairness will equalize users' 2-dominant shares. The allocation can be computed mathematically as follows: Let φ_1 and φ_2 be the number of tasks allocated to user 1 and 2, respectively. Then user 1's allocation vector is $\langle 40\varphi_1, 8\varphi_1, 8\varphi_1 \rangle$, and user 2's allocation vector is $\langle 8\varphi_2, 5\varphi_2, \varphi_2 \rangle$. The total allocated amount of bandwidth is $(40\varphi_a + 8\varphi_b)$, the total allocated amount of memory is $(8\varphi_1 + 5\varphi_2)$, and the total allocated amount of CPU is $(8\varphi_1 + \varphi_2)$. Besides, the 2-dominant share of user 1 and user 2 is $\varphi_1 \cdot \frac{1}{5} \cdot \frac{1}{25} = \frac{1}{125} \cdot \varphi_1$, and $\varphi_2 \cdot \frac{1}{25} \cdot \frac{1}{40} = \frac{1}{1000} \cdot \varphi_2$. The 2-dominant fairness allocation is then given by the solution to the following optimization problem:

$$\text{maximize} \quad \varphi_1, \varphi_2 \quad (2a)$$

$$\text{subject to} \quad \frac{1}{125} \cdot \varphi_1 = \frac{1}{1000} \cdot \varphi_2 \quad (2b)$$

$$40\varphi_1 + 8\varphi_2 \leq 200$$

$$8\varphi_1 + 5\varphi_2 \leq 200$$

$$8\varphi_1 + \varphi_2 \leq 200$$

Solving this problem yields $\varphi_1 = 1.9$, and $\varphi_2 = 15.4$ (See Figure 7). Thus, user 1 gets $\langle 77, 15.4, 15.4 \rangle$, and user 2 gets $\langle 123, 77, 15.4 \rangle$. Recall the result mentioned in Section 1.1, under DRF allocation mechanism, user 1 will receive an allocation of $\langle 100, 20, 20 \rangle$, and user 2 will receive $\langle 100, 62, 5, 12.5 \rangle$. Thus, 2-DF leads to an increase on the number of total tasks from 15 to 17.3. Besides, if user 2's request on memory (which is his second dominant resource) varies, his completes more tasks, as the red increasing line shown in Figure 4.

The intuition behind this allocation mechanism is, if we only consider fairness as equalizing each user's dominant share and satisfying any demand it has of other resource dimensions, it seems unfair for those users whose tasks have a low request on each resource dimension. Besides, it also leads to low efficiency because more tasks come from more resource allocations to users with smaller request vectors. Thus, we take a look at one more resource dimension to know better about a user's real request, and want to give users with lower requests on more resource dimensions more compensation to increase their total allocation shares, thus leading to more tasks done.

4.4 | 2-DF Scheduling Algorithm

We present the 2-DF scheduling pseudocode in Algorithm 1. The algorithm calculates the first dominant request and the second dominant request, thereby obtaining the 2-dominant share, s_i for each user i . The total resources allocated to each user are also be tracked and recorded. At each step, the proposed scheduling algorithm picks a user whose 2-dominant share is the lowest among all users with tasks ready to run, and checks whether that user's task demand can be satisfied. If there are enough resources, then assign the desired resources to the picked user. Otherwise, the user waits until all her desired resources available. According to our assumption, the tasks from a user should have the same request vector, and we use variable D_i to denote the demand vector of user i . Once a launched task finishes, the user releases the resources assigned to her, and our 2-DF mechanism again selects

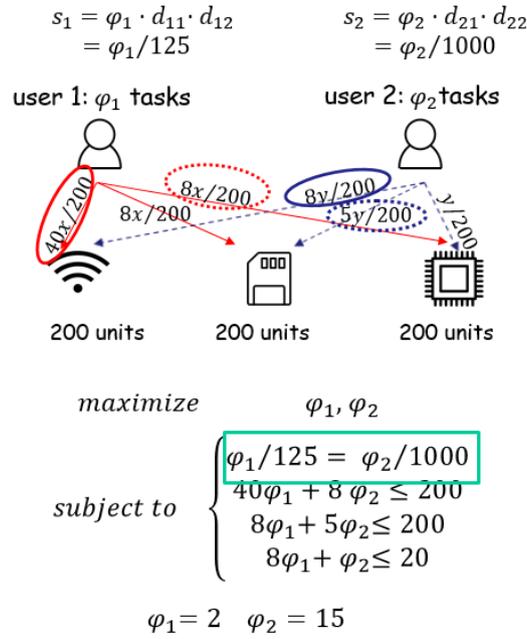


FIGURE 7 Resource allocations under 2-DF.

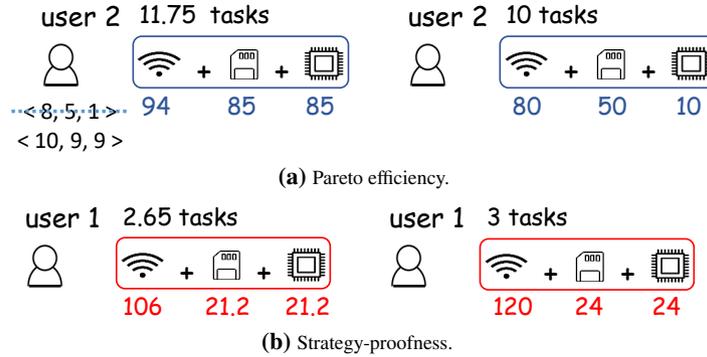


FIGURE 8 Examples of 2-DF fairness properties.

the user with the smallest 2-dominant share and decides if it is possible to run her task. If there are n users simultaneously in the system, the time complexity for each scheduling decision is $O(\log n)$ when we apply the a binary heap data structure to store $2d$ shares for users.

4.5 | Extending from 2-D to k-D with weights

If we evaluate our 2-dominant share using the rules mentioned in Section 4.2, we could see it really takes an additional non-dominant resource into consideration. However, 2-dominant share can be extended to a better fairshare function, which considers k dimensions of resources and allows a user to stress how important each dimension is.

Let $[d_{i1}, d_{i2}, \dots, d_{ik}]$ be the top k largest elements among $\{d_{ij}\}$ where $j \in [1, r]$, and each of the k elements is associated with a weight predefined by user i to express its importance. Then, we define that each user i has a k -dominant weighted share expressed as $\varphi_i \prod_{l=1}^k w_{il} \cdot d_{il}$. Similar to 2-DF, k -DF applies max-min fair allocation with respect to the users' k -dominant weighted share. That is, it always tries to equalize all users' k -dominant weighted shares. Thus, 2-dominant share is a special case of k -dominant weighted share, where $k = 2$ and $w_{i1} = w_{i2} = 1$ for each user i . For simplicity, we assume all weights are defined as 1 in the rest of this paper.

Algorithm 1 2-DF Scheduling Algorithm

```

1:  $R = [C_1, \dots, C_r]$ ; ▷ total resource capacities
2:  $U = [u_1, \dots, u_r]$ ; ▷ total resources used by now, initially 0
3:  $s_i (i = 1 \dots n)$ ; ▷ user  $i$ 's 2-dominant share, initially 0
4:  $d_{i1} = \max \left\{ \frac{a_{ij}}{C_j} \right\} (j = 1 \dots r)$ ; ▷ user  $i$ 's first dominant request
5:  $d_{i2} = \max \left\{ \frac{a_{ij}}{C_j} \right\} - \{d_{i1}\} (j = 1 \dots r)$ ; ▷ user  $i$ 's second dominant request
6:  $A_i = [a_{i1}, \dots, a_{ir}]$ ; ▷ resources allocated to user  $i$ , initially 0
Output 2-DF
7: pick user  $i$  with lowest 2-dominant share  $s_i$ ;
8:  $D_i \leftarrow$  user  $i$ 's demand vector;
9: if  $U + D_i \leq R$  then
10:    $U = U + D_i$ ; ▷ update consumed vector
11:    $A_i = A_i + D_i$ ; ▷ update  $i$ s allocation vector
12:    $s_i + = d_{i1} \cdot d_{i2}$ ;
13: else
14:   return; ▷ the cluster is full
Output Reclaim
15: pick user  $i$  with one task done;
16:  $U = U - D_i$ ; ▷ update consumed vector
17:  $U = U - D_i$ ; ▷ update consumed vector
18:  $A_i = A_i - D_i$ ; ▷ update  $i$ s allocation vector
19:  $s_i - = d_{i1} \cdot d_{i2}$ ;
20:  $s_i + = d_{i1} \cdot d_{i2}$ ;

```

5 | PROPERTIES OF K-DF

Next, we will discuss some of those desirable properties satisfied by k -DF and provide intuitive explanations for our analyses. For simplicity, we normalize the capacities of r infinitely divisible resources to be 1, respectively. We begin with showing that k -DF yields Pareto-efficiency.

Theorem 1. Every k -DF allocation is Pareto-efficient.

Proof. Assume user i can increase her total allocation without decreasing the allocation of anyone else. In fact, every user in a k -DF allocation has at least one saturated resource. If user i is monopolizing her saturated resource, it is impossible to increase i 's allocated fraction on her saturated resource. If the saturated resource is shared by user i and other users, then increasing the allocation of i must lead to decreasing the allocation of at least another user j who shares the same saturated resource, violating the hypothesis. \square

Then, we show that k -DF promises Strategy-proofness.

Theorem 2. The k -dominant fairness is Strategy-proof, *i.e.*, any user cannot increase her allocation of every resource (only increasing fraction on some resources cannot bring about an improved task number) in the k -dominant fairness by boosting some component of her true request vector.

Proof. If user i is monopolizing her saturated resource, it is impossible to increase i 's allocated fraction on her saturated resource; thereby her total allocation cannot be increased no matter what request vector she uses. Then, we discuss the situation where user i 's saturated resource is also shared by other users. Assume user i can increase her total allocation by using a different request vector $D_i' \neq D_i$, which means $\varphi_i' > \varphi_i$. Thus, user i 's k -dominant share $\varphi_i' \cdot \prod_{l=1}^k d_{il}$ is increased. Since we are trying to equalize the k -dominant share of each user, any other user's k -dominant share is also increased, resulting in a larger allocation on every resource. However, in order to increase user i 's total allocation, we must decrease the allocation of at least another user j sharing the same saturated resource, violating the previous analysis. \square

Now, we assume there exists no completely dominant user in the 2-dominant fairness, *i.e.*, for any two users i and j , their first two dominant requests satisfy either $d_{i1} \leq d_{j1}$ and $d_{i2} \geq d_{j2}$, or $d_{i1} \geq d_{j1}$ and $d_{i2} \leq d_{j2}$.

Theorem 3. Under the condition of no completely dominant user in the 2-dominant fairness, every 2-dominant fairness allocation is envy-free.

Proof. Assume by contradiction that there exists envy between user i and user j . Either user i or user j can be an envier. The envied must have a strictly higher fraction of every resource that the envier wants; otherwise, the envier cannot run more tasks under its allocation. Let m and n be i 's two dominant resources and p and q be j 's two dominant resources such that (1) $\varphi_i a_{im} \geq \varphi_i a_{in}$; (2) $\varphi_j a_{jp} \geq \varphi_j a_{jq}$. Now, we prove the theorem based on the following two conditions.

On the first condition, user i is the envier. According to the 2-dominant resource allocation mechanism, we can get the inequality equations 4 below.

$$\text{if } \begin{cases} a_{im} \cdot a_{in} \cdot \varphi_i = a_{jp} \cdot a_{jq} \cdot \varphi_j \\ a_{im} \leq a_{jp}, a_{in} \geq a_{jq} \end{cases} \quad (3)$$

$$\text{then } \begin{cases} \varphi_i a_{im} \leq \varphi_j a_{jp} \\ \varphi_i a_{in} \geq \varphi_j a_{jq} \end{cases} \quad (4)$$

For $\varphi_i a_{in} \geq \varphi_j a_{jq}$, there are three possible conditions:

- if n and q represent the same resource, then $\varphi_i a_{in} \geq \varphi_j a_{jn}$;
- if $a_{jq} \geq a_{jn}$, then $\varphi_i a_{in} \geq \varphi_j a_{jn}$;
- if $a_{jq} \leq a_{jn}$, meaning that resource n is user j 's first dominant resource;

then $\varphi_i a_{im} \geq \varphi_i a_{in} \geq \varphi_j a_{jp} \geq \varphi_j a_{jm}$, namely $\varphi_i a_{im} \geq \varphi_j a_{jm}$. Since $\varphi_i a_{ip} \leq \varphi_i a_{im}$, then $\varphi_i a_{ip} \leq \varphi_j a_{jp}$. Thus, both user i and user j have a higher (at least equal) fraction on one resource than the other does, violating the hypothesis.

On the second condition, user j is the envier. We can get the inequalities 6 below by following the 2-dominant resource allocation mechanism.

$$\text{if } \begin{cases} a_{im} \cdot a_{in} \cdot \varphi_i = a_{jp} \cdot a_{jq} \cdot \varphi_j \\ a_{jp} \leq a_{im}, a_{jq} \geq a_{in} \end{cases} \quad (5)$$

$$\text{then } \begin{cases} \varphi_j a_{jp} \leq \varphi_i a_{im} \\ \varphi_j a_{jq} \geq \varphi_i a_{in} \end{cases} \quad (6)$$

Quite similar to Condition 1, for $\varphi_j a_{jq} \geq \varphi_i a_{in}$, there are three possible conditions:

- if q and n represent the same resource, then $\varphi_j a_{jq} \geq \varphi_i a_{iq}$;
- if $a_{in} \geq a_{iq}$, then $\varphi_j a_{jq} \geq \varphi_i a_{iq}$;
- if $a_{in} \leq a_{iq}$, meaning that resource q is user i 's first dominant resource;

then $\varphi_j a_{jp} \geq \varphi_j a_{jq} \geq \varphi_i a_{im} \geq \varphi_i a_{ip}$, namely $\varphi_j a_{jp} \geq \varphi_i a_{ip}$. Since $\varphi_j a_{jm} \leq \varphi_j a_{jp}$, then $\varphi_j a_{jm} \leq \varphi_i a_{im}$. Quite similar to the first condition, on the second condition where $\varphi_j a_{jq} \geq \varphi_i a_{in}$, we still get the same conclusion, that is, both user i and user j have a higher (at least equal) fraction on one resource than the other does, violating the hypothesis.

Under both conditions, we can achieve the same conclusion that both user i and user j have a higher (at least equal) fraction on one resource than the other does, violating the hypothesis. \square

Next, we will extend this envyfree scenario from 2-d to k -d. Let $[d_{i1}, d_{i12}, \dots, d_{i1k}]$ be the top k largest elements among $\{d_{ij}\}$ where $j \in [1, r]$. We assume there exists no completely dominant user in the k -dominant fairness, *i.e.*, for any two users i and j , their top k dominant requests satisfy:

$$\exists m, \exists n, \begin{cases} \prod_{l \neq m}^k d_{il} \leq \prod_{l \neq m}^k d_{jl} \\ \prod_{l \neq n}^k d_{il} \geq \prod_{l \neq n}^k d_{jl} \end{cases} \quad (7)$$

Theorem 4. Under the condition of having no completely dominant users in the k -dominant fairness, every k -dominant fairness allocation is envy-free.

Capacity	NFC	DRF	2DF
3	1.62	1.34	1.39
5	1.81	1.48	1.56

TABLE 1 Average NTT.

Proof. Assume by contradiction that there exists envy between user i and another user j . Either user i or user j can be an envier. The envied must have a strictly higher fraction of every resource that the envier wants. According to the k -dominant fairness, user i and user j should have equal k -dominant fairness, that is:

$$\varphi_i \cdot \prod_{l=1}^k d_{il} = \varphi_j \cdot \prod_{l=1}^k d_{jl} \quad (8a)$$

$$\text{subject to} \quad \prod_{l \neq m}^k d_{il} \leq \prod_{l \neq m}^k d_{jl} \quad (8b)$$

$$\prod_{l \neq n}^k d_{il} \geq \prod_{l \neq n}^k d_{jl} \quad (8c)$$

Then, equations in 8 can be reduced in the following form:

$$\begin{cases} \varphi_i d_{im} \geq \varphi_j d_{jm} \\ \varphi_i d_{in} \leq \varphi_j d_{jn} \end{cases} \quad (9)$$

Assume d_{im} and d_{in} represent i 's requests on resource u and v , d_{jm} and d_{jn} represent j 's requests on resource x and y . For $\varphi_i d_{im} \geq \varphi_j d_{jm}$, if u and x are the same resource, then $\varphi_i a_{iu} \geq \varphi_j a_{ju}$; if $a_{ju} \leq a_{jx}$, then $\varphi_i a_{iu} \geq \varphi_j a_{jx} \geq \varphi_j a_{ju}$, namely $\varphi_i a_{iu} \geq \varphi_j a_{ju}$. In addition, if $a_{ju} \geq a_{jx}$, there must exist one resource, w , which could satisfy $a_{iw} \geq a_{iu}$ && $a_{ju} \geq a_{jw}$, such that $\varphi_i a_{iw} \geq \varphi_i a_{iu} \geq \varphi_j a_{jx} \geq \varphi_j a_{jw}$, namely $\varphi_i a_{iw} \geq \varphi_j a_{jw}$. Thus, we can always find one resource of which user i has no less fraction than user j does.

For $\varphi_i d_{in} \leq \varphi_j d_{jn}$, we can also achieve a similar conclusion that there always exists one resource of which user j has no less fraction than user i does. Thus, user i has no less fraction than user j does of one resource; meanwhile, user j has no less fraction than user i does of another resource, violating the hypothesis. \square

6 | PERFORMANCE EVALUATION

We consider two multi-resource allocation scenarios in different data centers to evaluate our allocation mechanism. All allocation mechanisms were implemented with MATLAB-R2017b, running on a local machine with an Intel Core 2 Duo E8400 3.0 GHz CPU and 8 GB RAM.

To measure resource efficiency, we use two metrics, which are previously mentioned in Section 3. We formally define the two metrics: **NTT** which represents the number of total tasks completed by all users and **AUR** which represents the amount of unused resource after an allocation. Besides, we also compare our allocation mechanism with some existing allocation works.

6.1 | The first scenario

In the first scenario, there is a data center with three distinct and divisible resources, r_1 , r_2 , and r_3 . There are 3 users, each of whom requires a fixed amount of each resource to accomplish a task. Tasks are assumed to be infinitely divisible. That is, we allow the completed task number to be a decimal. Resource capacity vector is expressed as $\langle C, C, C \rangle$. We conducted two experiments with different values of C where $C \in \{3, 5\}$ to observe how resource capacity would impact fairness and efficiency. In both experiments, each user i 's request vector is $\langle d_{i1}, d_{i2}, d_{i3} \rangle$ where d_{ij} is an integer in the range of 1 to C for any $j = 1, 2, 3$.

6.1.1 | Three comparison allocation mechanisms

We compare the efficiency in terms of **NTT** under three different allocation mechanisms: (1) No Fairness Constraints (NFC) which tries to achieve a maximal number of total tasks without considering fairness, (2) Dominant Resource Fairness (DRF) and (3) 2-Dominant Resource Fairness (2-DF).

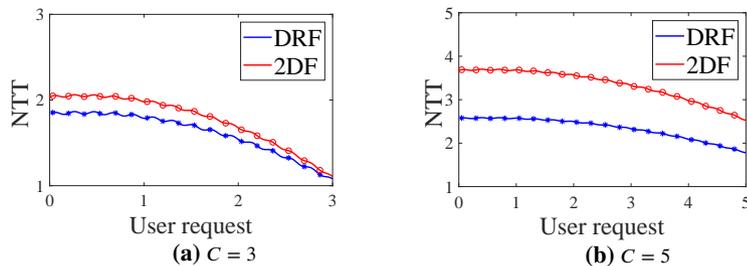


FIGURE 9 Cases where higher NTT is achieved in 2-DF than DRF.

Given the value of C , there are 3^C combinations of i 's request vector. Putting them together, we consider 3^{3C} combinations of all the users' request vectors. For each possible combination, we apply NFC, DRF, and 2-DF, respectively, and then record the corresponding NTTs. After that, we calculate the average NTT of 3^{3C} combinations for each allocation mechanism. The average NTT completed by 3 users with different request vectors under the three allocation mechanisms are shown in Table 1. As can be seen, the number of total tasks obtained by 3 users under NFC is the highest of all allocation mechanisms. This result, to some extent, reflects the fact that fairness is often a conflicting objective against efficiency in the presence of multiple resources. Besides, 2-DF outperforms DRF, and as the resource capacity increases, 2-DF's advantage becomes more evident. Thus, we can conclude that, our allocation mechanism would have good scalability in a data center with large resource capacities.

In the first experiment, among 3^9 combinations of request vectors, 2-DF executes more tasks than DRF does in around 51.7% of total cases. In Figure 9a, we show all unique cases where 2-DF performs better in terms of NTT. In the second experiment, the capacity was changed to 5 units for each resource. Among 5^9 combinations of request vectors, 2-DF executes more tasks than DRF does in around 58.1% of total cases. In Figure 9b, due to the large amount of cases, we only display those unique cases in which NTT obtained by 2-DF is at least 30% more than that obtained by DRF.

6.1.2 | Fairness

In fact, under our 2-DF mechanism, our allocation cannot guarantee sharing incentive (SI) proposed in DRF, which proposes that each user should at least get $\frac{1}{n}$ share on its dominant resource. However, if we compare the efficiency achieved by all 2-DF allocations, which satisfy SI, with that of corresponding DRF allocations, results are still good. In the first experiment, the 2-DF allocations satisfying SI increase by around 0.04 task on average, compared with their DRF counterparts. In the second experiment, the 2-DF allocations satisfying SI increase by around 0.08 task on average, compared with their DRF counterparts.

In Section 3, we mention a scenario where complete envy-freeness can be achieved. We explore the occurrence frequency of this scenario, and the result is not too bad. In the first experiment, the envyfree cases account for 64.0%, and among all cases where 2-DF obtains higher efficiency and envy-free cases occupy around 37.2%. In the second experiment, the envyfree cases account for 58.8%, and among all cases where 2-DF obtains higher efficiency, envyfree cases account for around 38.7%.

6.2 | The second scenario

In the second scenario, we assume a data center with 1000-unit bandwidth, 1000-unit memory and 1000-unit CPU. All resources are shared by two users. Again, each user requires a fixed amount of each resource to accomplish a task, and tasks are assumed to be infinitely divisible.

In this scenario, we consider two users with different request types: heavy and light. A request $D_i = \langle d_{i1}, d_{i2}, d_{i3} \rangle$ is said to be heavy, where $\forall j = [1, 2, 3], d_{ij} \in \{25x_1, 5x_1, x_1\}$. x_1 is a random variable, which is picked randomly while following the normal distribution $\sim N(8, 0)$ in our experiments. Similarly, a request $D_i = \langle d_{i1}, d_{i2}, d_{i3} \rangle$ is said to be light, where $\forall j = [1, 2, 3], d_{ij} \in \{25x_2, 5x_2, x_2\}$ and $x_2 \sim N(1, 0)$.

There are three combinations of user request types, as is shown in Table 2. There are three request levels for both request types: big ($25x$), medium ($5x$), and small (x). Given a specific request type, we design 8 types of request vector with different request levels in different resource dimensions. We list all types here: $\langle 25x, 25x, 25x \rangle$, $\langle 25x, 25x, 5x \rangle$, $\langle 25x, 25x, x \rangle$, $\langle 25x, 5x, 5x \rangle$, $\langle 25x, 5x, x \rangle$, $\langle 5x, 5x, 5x \rangle$, $\langle 5x, 5x, x \rangle$, and $\langle 5x, x, x \rangle$. Given a specific combination of request types, there are $8 \times 8 = 64$ possible pairs of two request vectors. For each pair of request vectors, we want to compare the efficiency of

Combination	User 1	User 2
I	heavy	heavy
II	heavy	light
III	light	light

TABLE 2 Three combinations of user request types

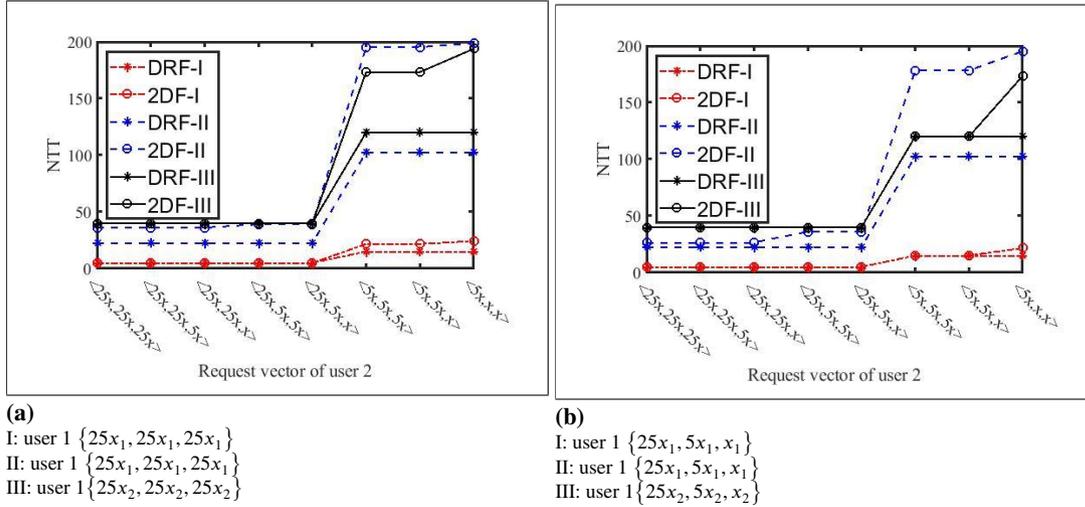


FIGURE 10 Two typical cases.

AUR _{ave}	DRF	2DF
I	418	654
II	418	802
III	418	654

(a)
 I: user 1 $\{25x_1, 25x_1, 25x_1\}$
 II: user 1 $\{25x_1, 25x_1, 25x_1\}$
 III: user 1 $\{25x_2, 25x_2, 25x_2\}$

(b)
 I: user 1 $\{25x_1, 5x_1, x_1\}$
 II: user 1 $\{25x_1, 5x_1, x_1\}$
 III: user 1 $\{25x_2, 5x_2, x_2\}$

TABLE 3 Average AUR.

DRF and 2-DF. To be precise, we conducted 10 experiments over a given pair of request vectors, by randomly choosing values of x_1 and x_2 based on their distributions. Due to the large number of experiments we did, we have confidence in the consistency and reliability of the results. All results present below are the average over their own experiments.

Based on the $3 \times 8 \times 8 \times 10 = 1920$ experiments we did, we conclude two typical cases, which any other case can be induced from. We show these two cases in Figure 10. It is obvious that, when comparing **NTT**, 2-DF outperforms *DRF*, especially when user 1 has big-level requests on each resource dimension and user 2 has small-level requests on each resource dimension, whatever their request types are. The average increase is 45% when comparing **NTT** of 2-DF and DRF among all experiments.

To measure efficiency of DRF and 2-DF in terms of **AUR**, we calculate the average amount of unused resources under DRF and 2-DF, respectively in two cases. The result is shown in Table 3. It is obvious that 2-DF consumes less resources while yielding more tasks, which allows a data center to serve more users in the future.

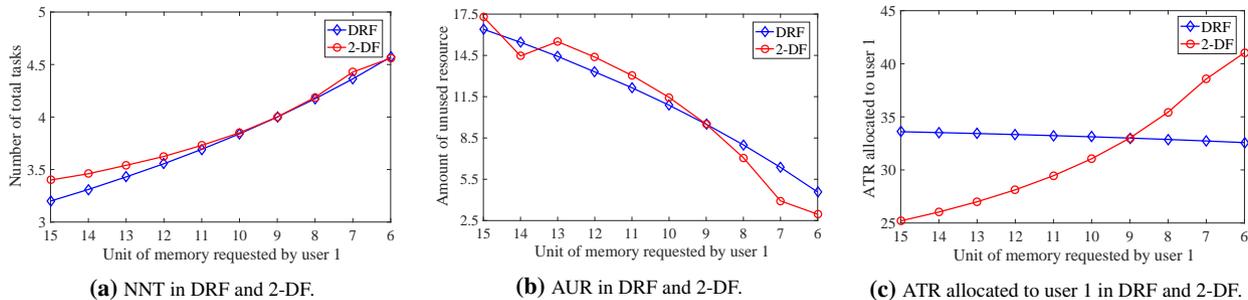


FIGURE 11 Performance changes when user 1's request of memory per task varies.

Machine type	Number of machines	CPUs	Memory	Ratio
0	6732	0.50	0.50	0.5107
1	3863	0.50	0.25	0.2930
2	1001	0.50	0.75	0.0759
3	795	1.00	1.00	0.0603
4	126	0.25	0.25	0.0551
5	52	0.50	0.12	0.0039
6	5	0.50	0.03	0.0004
7	5	0.50	0.97	0.0004
8	3	1.00	0.50	0.0002
9	1	0.50	0.06	0.0001

TABLE 4 Machine configurations summary of Google Trace.

6.2.1 | Sharing incentive

In fact, under our 2-DF mechanism, our allocation cannot guarantee sharing incentive (SI) proposed in DRF, which said that each user should at least get $\frac{1}{n}$ share on its dominant resource. However, if we compare the efficiency achieved by all 2-DF allocations, which satisfy SI, with that of corresponding DRF allocations, results are still good. In the first experiment, the 2-DF allocations satisfying SI increase by around 0.04 task on average, compared with their DRF counterparts. In the second experiment, the 2-DF allocations satisfying SI increase by around 0.08 task on average, compared with their DRF counterparts.

6.2.2 | Complete envy-freeness

In Section 3, we mentioned a scenario where complete envy-freeness can be achieved. We explore the occurrence frequency of this scenario, and the result is not too bad. In the first experiment, the envyfree cases account for 64.0%, and among all cases where 2-DF obtains higher efficiency and envyfree cases occupy around 37.2%. In the second experiment, the envyfree cases account for 58.8%, and among all cases where 2-DF obtains higher efficiency, envyfree cases occupy around 38.7%.

6.2.3 | Traditional fairness

In Section 3, besides discussing the fairness property in the existing works, we also mention two concepts related to the notion of fairness in real life. A fair allocation should take multiple resources into consideration as well, and should give some compensation to users with low requirements in non-dominant resource dimensions. To measure traditional fairness, in the third experiment, we compare the amount of total resources (ATR) allocated to user 1 using DRF and 2-DF, respectively, According to Figure 11, 2-DF allocates obviously more resources user 1, with his decreasing request of non-dominant resource. *i.e.*, memory. However, DRF doesn't show any evidence to compensate to user 1. It is obvious that 2-DF also considers non-dominant resources instead of a single dominant resource when allocating.

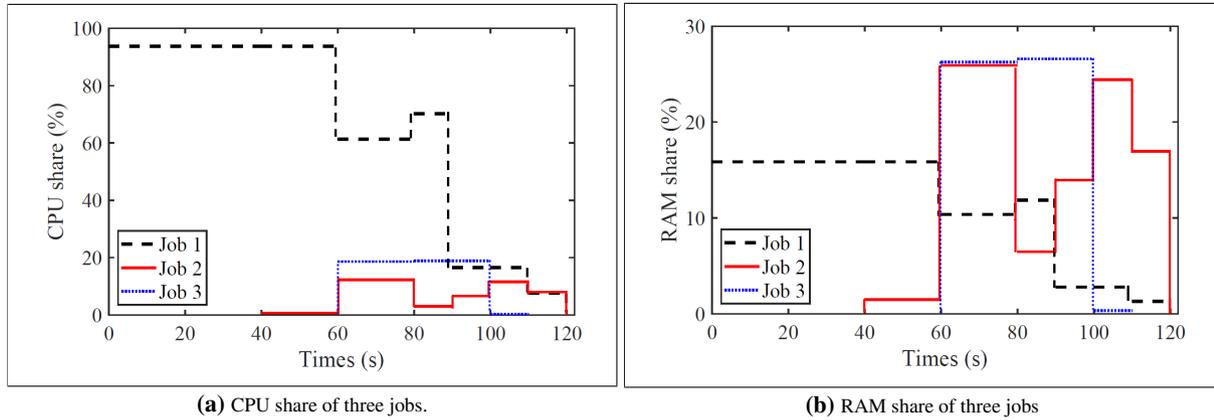


FIGURE 12 Resource shares for three jobs.

6.3 | Trace-Driven online resource allocation

Extensive simulations driven by Google trace²¹ are conducted to further evaluate the performance of the proposed allocation mechanism from a macroscopic perspective. From the trace, we can extract information about machines and user requests. Machines are heterogeneous with different capacities of CPU and memory. However, instead of showing the exact numbers of CPU cores and bytes of memory for each machine, the trace provides machine configurations, where both CPU and memory size measurements are normalized to that of the maximum machines. Table 4 shows the configurations of different types of machines and the corresponding number of each machine type. We also provide the ratio of each machine type among all machines. This information will guide us on how to simulate machines from which the resources are allocated in our experiments. Each user request is a job containing several identical tasks. The request is composed of the task number, the required amount of resources for each task, and task running time, which we will use as the input to the allocation algorithm.

We want to show how the proposed mechanism allocates resources online. We run three jobs on a small cluster of 50 servers. We assume there are 4 types of servers in total, *i.e.*, the top 4 configurations in Table 4 and the maximum server has 16 CPU cores and 128 GB of RAM, *i.e.*, the highest configuration among E2 machines on Google Cloud. The number of each server type is also determined based on the distribution shown in Table 4. This setting leads to a resource pool containing 27 type-1 servers with 8 CPU cores and 64 GB of RAM, 16 type-2 servers with 8 CPU cores and 32 GB of RAM, 4 type-3 servers with 8 CPU cores and 96 GB of RAM, and 3 type-4 servers with 16 CPU cores and 128 GB of RAM. These three jobs launch tasks with different resource demands at different times within a time period of 2 minutes.

Figures 12a and 12b show the CPU and RAM allocations given to each job as a function of time. Job 1 comes at time 0. It contains 450 tasks, and each of its tasks has a request vector of $\langle 2.5 \text{ CPU}, 3 \text{ GB RAM} \rangle$ and an execution time of 30 seconds. Since only Job 1 is active at the beginning, all resources are allocated to Job 1, yielding 159 tasks running in parallel. This allocation lasts 40s, until Job 2 joins with memory-heavy tasks. There are 150 tasks in Job 2, each of which lasts 20 seconds with a request of $\langle 1 \text{ CPU}, 15 \text{ GB RAM} \rangle$. At that time, only type-4 servers can provide enough resources to execute tasks from Job 2. Each type-4 server starts to run a task from Job 2 at the time of 40s.

At 50s, Job 3 becomes active. There are 80 tasks in Job 3. Each task requests $\langle 1 \text{ CPU}, 10 \text{ GB RAM} \rangle$ and its occupation time is 80 seconds. At that time, although memory resources are sufficient, any server can provide at most 0.5 CPU, indicating that no resources can be allocated to Job 3. At 60s, all tasks being executed are finished and the corresponding resources are recycled to be allocated to these three jobs. At that time, 318 tasks from Job 1 and 3 tasks from Job 2 have been completed, respectively.

After applying our proposed allocation mechanism, all type-1 servers are scheduled to execute 56 tasks from Job 1, 28 tasks from Job 2, and 42 tasks from Job 3; all type-2 servers are scheduled to execute 28 tasks from Job 1, 14 tasks from Job 2, and 21 tasks from Job 3; all type-3 servers are scheduled to execute 8 tasks from Job 1, 4 tasks from Job 2, and 6 tasks from Job 3; and all type-4 servers are scheduled to execute 56 tasks from Job 1, 28 tasks from Job 2, and 42 tasks from Job 3. At time 80s, all resources used for Job 2 are released and another 52 tasks from Job 2 are finished. All those free resources are allocated again to execute remaining task. At that time, there remain 28 tasks from Job 1, 95 tasks from Job 2, and 1 tasks from Job 3. Our proposed allocation mechanism assigns the only task from Job 3 to a type-1 server, and all remaining resources are allocated to Job 1 and Job 2, which leads to 15 tasks from Job 1 and 13 tasks from Job 2 scheduled.

At time 90s, the resources allocated to Job 1 since time point 60s are available to further serve Job 1 and Job 2. Since then, another 13 tasks from Job 1 and 15 tasks from Job 2 are executed, respectively. So far, only 34 tasks from Job 2 are left. At time 100s both the resources allocated to Job 1 since time point 60s and the resources allocated to Job 2 since time point 80s are available simultaneously, and can be allocated to Job 2 to execute the remaining 34 tasks.

7 | CONCLUSION

In this paper, we consider the open problem of multi-resource sharing for heterogeneous users. We show that DRF suffers from serious fairness concerns without utility guarantees. We try to mitigate this problem with the new sharing policy known as 2-dominant resource fairness, and we extend this concept to k -dominant resource fairness. In our mechanism, 2-dominant share is defined for each user as the product of first two dominant demand/capacity ratios, and we try to equalize the 2-dominant share of all users as much as possible. 2-dominant resource fairness provides strategy-proofness, in that no user can run more tasks by lying about its demands. Meanwhile, this policy is Pareto-optimal, and envy-freeness can be achieved in certain scenarios. Besides, in the traditional sense of fairness, 2-DF considers more resource dimension when allocating, and results in a more balanced allocation between dominant and non-dominant resources. Compared with dominant resource fairness (DRF), our proposed model achieves better efficiency, in terms of the number of total tasks completed by all users. It is vital for efficiency-needed applications, allowing a small penalty on widely-accepted fairness properties.

8 | ACKNOWLEDGMENT

This research was supported in part by NSF grants CNS 1824440, CNS 1828363, CNS 1757533, CNS 1629746, CNS 1651947, and CNS 1564128.

References

1. Ghodsi A, Zaharia M, Hindman B, Konwinski A, Shenker S, Stoica I. Dominant resource fairness: fair allocation of multiple resource types. *NSDI* 2011.
2. Jin Y, Hayashi M. Efficiency comparison between proportional fairness and dominant resource fairness with two different type resources. *2016 Annual Conference on Information Science and Systems (CISS)* 2016.
3. Bertsimas D, Farias VF, Trichakis N. The price of fairness. *Operations research* 2011.
4. Poullie P, Bocek T, Stiller B. A survey of the state-of-the-art in fair multi-resource allocations for data centers. *IEEE Transactions on Network and Service Management* 2017.
5. Saha P, Beltre A, Govindaraju M. Exploring the fairness and resource distribution in an apache mesos environment. In: *IEEE*. ; 2018: 434–441.
6. Ghodsi A, Sekar V, Zaharia M, Stoica I. Multi-resource fair queueing for packet processing. *Proceedings of the ACM SIGCOMM 2012 conference on Applications, technologies, architectures, and protocols for computer communication* 2012.
7. Gutman A, Nisan N. Fair allocation without trade. *arXiv preprint arXiv:1204.4286* 2012.
8. Joe-Wong C, Sen S, Lan T, Chiang M. Multiresource allocation: Fairness–efficiency tradeoffs in a unifying framework. *IEEE/ACM Transactions on Networking* 2013.
9. Parkes DC, Procaccia AD, Shah N. Beyond dominant resource fairness: Extensions, limitations, and indivisibilities. *ACM Transactions on Economics and Computation (TEAC)* 2015.
10. Wang W, Liang B, Li B. Multi-resource fair allocation in heterogeneous cloud computing systems. *IEEE Transactions on Parallel and Distributed Systems* 2014.

11. Friedman E, Ghodsi A, Psomas CA. Strategyproof allocation of discrete jobs on multiple machines. *Proceedings of the fifteenth ACM conference on Economics and computation* 2014.
12. Zeldes Y, Feitelson DG. On-line fair allocations based on bottlenecks and global priorities. *Proceedings of the 4th ACM/SPEC International Conference on Performance Engineering* 2013.
13. Kash I, Procaccia AD, Shah N. No agent left behind: Dynamic fair division of multiple resources. *Journal of Artificial Intelligence Research* 2014.
14. Li W, Liu X, Zhang X, Zhang X. A further analysis of the dynamic dominant resource fairness mechanism. *International Workshop on Frontiers in Algorithmics* 2017.
15. Dolev D, Feitelson DG, Halpern JY, Kupferman R, Linial N. No justified complaints: On fair sharing of multiple resources. *proceedings of the 3rd Innovations in Theoretical Computer Science Conference* 2012.
16. Zahedi SM, Lee BC. REF: Resource elasticity fairness with sharing incentives for multiprocessors. *ACM SIGPLAN Notices* 2014.
17. Poullie P, Stiller B. Cloud flat rates enabled via fair multi-resource consumption. *IFIP International Conference on Autonomous Infrastructure, Management and Security* 2016.
18. Poullie P, Mannhart S, Stiller B. Virtual machine priority adaption to enforce fairness among cloud users. *2016 12th International Conference on Network and Service Management (CNSM)* 2016.
19. Poullie P, Stiller B. The design and evaluation of a heaviness metric for cloud fairness and correct virtual machine configurations. *International Conference on the Economics of Grids, Clouds, Systems, and Services* 2016.
20. Chen W, Xu Y, Wu X. Deep reinforcement learning for multi-resource multi-machine job scheduling. *arXiv preprint arXiv:1711.07440* 2017.
21. Reiss C, Wilkes J, Hellerstein JL. Google Cluster-Usage Traces. https://github.com/google/cluster-data/blob/master/ClusterData2011_2.md; 2011. Online; accessed 23 May 2020.

