

# Multi-Resource Allocation: Fairness-Efficiency Tradeoffs in a Unifying Framework

Carlee Joe-Wong, Soumya Sen, Tian Lan, and Mung Chiang

**Abstract**—Quantifying the notion of fairness is under-explored when there are multiple types of resources and users request different ratios of the different resources. A typical example is datacenters processing jobs with heterogeneous resource requirements on CPU, memory, network, bandwidth, etc. This paper develops a unifying framework addressing the fairness-efficiency tradeoff in light of multiple types of resources. We develop two families of fairness functions that provide different tradeoffs, characterize the effect of user requests’ heterogeneity, and prove conditions under which these fairness measures satisfy the Pareto efficiency, sharing incentive, and envy-free properties. Intuitions behind the analysis are explained in two visualizations of multi-resource allocation. We also investigate people’s fairness perceptions through an online survey of allocation preferences and provide a brief overview of related work on fairness.

## I. INTRODUCTION

### A. Motivation

Comparing fairness of different allocations of a *single* type of resource has been extensively studied. Fairness can be quantified with a variety of metrics, such as Jain’s index [1]. Alternatively, different notions of fairness, including proportional and max-min fairness, can be achieved through maximization of  $\alpha$ -fair or isoelastic utility functions [2]. These approaches, as well as others from economics and sociology, have recently been unified as the unique family of functions satisfying four axioms for fairness metrics, as summarized in [3], [4]. The tradeoff between fairness and efficiency has also been studied in [5]–[7].

When it comes to allocating *multiple* types of resources, however, there has been much less systematic study, the recent paper [8] being a notable exception. Indeed, it is unclear what it means to say that a multi-resource allocation is “fair.” Each user in a network requires a certain *combination* of different resource types to process one job, and this combination may differ from user to user. For example, datacenters allocate different resources (memory, CPUs, storage, bandwidth, etc.) to competing users with different requirements. One user might




Fig. 1. An example of multi-resource requirements in datacenters.

have computational jobs requiring more CPU cycles than memory, while another might have the opposite requirements.

The need for multi-resource fairness functions can be illustrated with a very simple example, as shown in Fig. 1. In this example, two users require CPUs and memory in order to perform some jobs. User 1 requires 2 GB of memory and 3 CPUs per job, while user 2 needs 2 GB of memory and 1 CPU per job. There is a total of 6 GB of memory and 4 CPUs.

Many allocations might be considered “fair” in this example: should users be allocated resources in proportion to their resource requirements? Or should they be allocated resources so as to process equal numbers of jobs? The fairness measure proposed recently in [8], called **Dominant Resource Fairness (DRF)**, allocates resources according to max-min fairness on dominant resource shares. In this example, DRF would allocate 0.76 jobs to user 1 and 1.71 jobs to user 2, for a total of 2.47 jobs processed. But this allocation brings about a significant loss in system efficiency; e.g., a more unequal allocation of 0.17 jobs to user 1 and 2.83 jobs to user 2 yields a total of 3 jobs. An in-between allocation can be realized if another well-known fairness metric,  $\alpha$ -fairness, is adapted for multiple resources following our methods in Section III-B. For  $\alpha = 0.5$ , user 1 has 0.57 jobs and user 2 has 2.29 jobs, for a total of 2.86 jobs. Each of these allocations represents one point of the fairness-efficiency tradeoff. This paper develops a unifying framework for studying this tradeoff in light of multiple types of resources and heterogeneity in users’ resource requirements.

Multi-resource allocation problems arise in increasingly many applications. Datacenters that sell *bundles* of CPUs, memory, storage, and network bandwidth are just one example. In fact, even the classical problem of bandwidth allocation in a congested network can be viewed as a special case of multi-resource allocation. Given a network and its topology, we can view each link as a separate resource with a distinct capacity. Each user is represented by a network flow, which uses a pre-defined subset of links. In this special case, resource requests on all the links must be the same for each user.

In general, multi-resource allocation *cannot* be trivially turned into single-resource allocation by assuming different resources are interchangeable. For example, if a cloud client needs 2 units of CPU and 5 units of networking bandwidth to

C. Joe-Wong is with the Program in Applied and Computational Mathematics at Princeton University (email:cjoe@princeton.edu). S. Sen and M. Chiang are with the Department of Electrical Engineering, Princeton University (emails: {soumyas, chiangm}@princeton.edu). T. Lan is with the Department of Electrical and Computer Engineering at George Washington University (email: tlan@gwu.edu). An earlier version of part of this work appeared at the 2012 IEEE International Conference on Computer Communications (Infocom).

This work was partly supported by NSF grant CNS-0905086, DARPA grant FA8750-11-C-0254, and AFOSR MURI grant FA9550-09-1-0643. C. Joe-Wong acknowledges support from the NDSEG fellowship.

Manuscript received Month Day 2012.

finish 1 unit of job, adding more does *not* reduce the need for 5 units of bandwidth.

### B. Unique Challenges of Multi-Resource Fairness

The following new challenges on fairness arise due to the presence of multiple types of resources:

- In a single-resource scenario, users' resource requirements can be represented with a scalar. With multiple resources, users have vectors of resource requirements, which may all look different and must be scalarized before fairness can be evaluated. We present two ways to visualize user heterogeneity in Section III-A and two methods for this scalarization in Section III-B, yielding parametrized families of multi-resource fairness measures that satisfy the axioms of [4].
- In a single-resource scenario, the most efficient allocation will clearly use the entire resource. In a multi-resource scenario, however, users' heterogeneous resource requirements may not allow each resource to be completely used. Even how to measure efficiency is unclear: should we use the total number of jobs allocated?<sup>1</sup> Or the amount of leftover resource capacity? Section V numerically examines both of these efficiency metrics, while Props. 1 and 2 and their corollaries examine the impact of user heterogeneity on the number of jobs processed.
- The extension of max-min fairness to multiple resources is shown in [8] to satisfy such properties as Pareto-efficiency for certain parameter values. We characterize the parameterizations under which our multi-resource fairness functions satisfy Pareto-efficiency, sharing incentive, and envy-freeness (Props. 3-5 and their corollaries).
- The existence of a fairness-efficiency tradeoff depends on both the scalarization of users' resource requirements and the subsequent evaluation of fairness. We show that a greater emphasis on equity or fairness need not always decrease efficiency (Prop. 6) and give analytical conditions on when the fairness-efficiency tradeoff exists (Props. 7 and 8 and their corollaries).
- When a fairness-efficiency tradeoff exists, the "best" operating point along this tradeoff depends on the operator's exogenously determined preferences. We characterize this *psychological* component to fairness by conducting a human subject experiment in which participants are asked to rank possible allocation choices given in an online survey. Our results indicate that people tend to cluster into two different groups—one preferring efficiency over fairness and one fairness over efficiency.

After further discussion of related work in Section II, Section III develops our two new families of fairness functions, which we call **Fairness on Dominant Shares (FDS)** and **Generalized Fairness on Jobs (GFJ)**. FDS includes the max-min fairness measure DRF proposed in [8] as a special case. We investigate key properties of these functions in Section IV and characterize conditions under which they are satisfied by FDS and GFJ. Section V then applies our fairness

functions to numerical examples of datacenters. We examine the relationship between the fairness-efficiency tradeoff and FDS and GFJ parameterizations. In Section VI, we experiment with characterizing the parameter values consistent with real people's fairness judgements, analyzing results from an online survey of 143 participants who were asked to rank different possible resource allocations for an example datacenter. All proofs can be found in the technical report [3].

## II. RELATED WORK

Much of the existing theory on the fairness of resource allocations is devoted to allocations of a single resource [4], [9]–[11] (e.g. allocating available link bandwidth to network flows [12]–[15]). The recent work [4] develops the following family of fairness functions for a single resource, unifying previously developed fairness measures. It was proven that this family, parametrized by two numbers, is the *only* family of functions satisfying four simple axioms of fairness metrics:

$$f_{\beta,\lambda}(\mathbf{x}) = \text{sgn}(1-\beta) \left( \sum_{i=1}^n \left( \frac{x_i}{\sum_{j=1}^n x_j} \right)^{1-\beta} \right)^{\frac{1}{\beta}} \left( \sum_{i=1}^n x_i \right)^{\lambda}, \quad (1)$$

where  $\beta \in \mathbb{R}$  and  $\lambda \in \mathbb{R}$  are parameters. The parameter  $\beta$  gives the "type" of fairness measured by (1), and the parameter  $\lambda$  gives the emphasis on efficiency. A larger  $|\lambda|$  indicates greater emphasis on efficiency over fairness. If we take  $\lambda = \frac{1-\beta}{\beta}$  and  $\beta > 0$ , we recover  $\alpha$ -fairness for  $\alpha = \beta$ . In particular, taking the limit as  $\beta \rightarrow 1$  yields proportional fairness.

Even multi-resource allocation problems, such as scheduling jobs in a datacenter, are often simply treated as a single resource problem (e.g. the Hadoop and Dryad schedulers [16]). A recent paper [8] generalizes the max-min fairness measure to multiple resource settings. Our work develops a unified analytical framework for fairness of multi-resource allocations. In particular, in contrast to [8], we incorporate the tradeoff between fairness and efficiency in multi-resource settings.

In the technical report [3], we provide a more comprehensive survey of other work on fairness. In addition to further discussion on fairness in engineering frameworks, we summarize theories of fairness from computer science, economics, political philosophy, and sociology.

## III. FAIRNESS-EFFICIENCY OF MULTI-RESOURCE ALLOCATIONS

We first present "dual" visualizations of heterogeneity among users' requirements for multiple resources in Section III-A. Section III-B then develops two new families of fairness functions, which scalarize these heterogeneous resource requirement vectors and use them to evaluate the fairness of multi-resource allocations. These two families are Fairness on Dominant Shares (FDS) and Generalized Fairness on Jobs (GFJ). FDS measures the fairness of users' resource allocations by accounting for both the number of jobs allocated to each user (a function of the resources available) and the heterogeneity in different resource requirements across users. GFJ, on the other hand, assumes that users' utility depends solely on the number of jobs they are allocated, irrespective of their differing resource needs.

<sup>1</sup>The phrases "jobs allocated" and "jobs processed" are used interchangeably throughout the paper.

### A. Visualizing User Heterogeneity

A major challenge of multi-resource fairness is incorporating the heterogeneity of different users' requirements for different resources into the assessment of its fairness. Visualizing this heterogeneity can yield useful insights. Moreover, Section V examines in detail how heterogeneity affects the optimal allocation and achieved efficiency.

Figure 2 provides two ways to visualize user heterogeneity. Each user  $j$  requires  $R_{ij}$  of resource type  $i$  for each job.

The first (top) visualization has as many dimensions as there are different types of resources. The axes correspond to the resources (two types of resources here for visual simplicity), with the box representing the resource constraints. The slope  $\sigma_i$  of the line corresponding to each user  $i$  is the ratio of that user's requirements for the two resources. The heterogeneity of users' resource requirements can be captured with the variance of the  $\{\sigma_i\}$ :<sup>2</sup> homogeneity occurs at 0 variance (all users have the same resource requirements) and the dashed line becomes straight. Heterogeneity increases with the variance of  $\sigma$ .

The second (bottom) visualization has as many dimensions as there are different users. The axes correspond to the jobs allocated to each user (two users here for simplicity of drawing), with feasible allocations shown as shaded regions bounded by linear resource constraints. The slopes  $\tau_i$  of constraint line  $i$  reflect the ratio of user 1's and user 2's requirements for resource  $i$ . Again, the heterogeneity of users' resource requirements can be captured in the variance of the  $\tau_i$ . Homogeneity occurs when the variance is 0; in that case the resource constraints have the same slope and reduce to one constraint. Heterogeneity increases with the variance of  $\tau$ .

### B. Defining Multi-Resource Fairness

1) *Fairness on Dominant Shares (FDS)*: As defined in [8], a user's **dominant share** is the maximum share of any resource allocated to that user.

Let  $x_j$  denote the number of jobs allocated to each user  $j$  and  $C_i$  the capacity of each resource  $i$ . Then we have the resource constraints  $\sum_{j=1}^n R_{ij}x_j \leq C_i$  for all resources  $i$ , where  $R_{ij}$  is the amount of resource  $i$  which user  $j$  requires for one job, and there are  $n$  users. For ease of notation, we define  $\gamma_{ij} = R_{ij}/C_i$  as the share of resource  $i$  required by user  $j$  to process one job. We let

$$\mu_j = \max_i \left\{ \frac{R_{ij}}{C_i} \right\} \quad (2)$$

denote the maximum share of a resource required by user  $j$  to process one job; then  $\mu_j x_j$  is user  $j$ 's dominant share.

We introduce the fairness measures  $f_{\beta,\lambda}^{\text{FDS}}$ :

$$\text{sgn}(1 - \beta) \left( \sum_{j=1}^n \left( \frac{\mu_j x_j}{\sum_{k=1}^n \mu_k x_k} \right)^{1-\beta} \right)^{\frac{1}{\beta}} \left( \sum_{j=1}^n \mu_j x_j \right)^{\lambda} \quad (3)$$

These fairness measures extend those developed in [4] for a single resource; details on their derivation are given in that work. Here  $\beta \neq 1$  and  $\lambda$  are pre-specified parameters.

<sup>2</sup>We assume that the  $\sigma_i$  are realizations of a random variable  $\sigma$ .

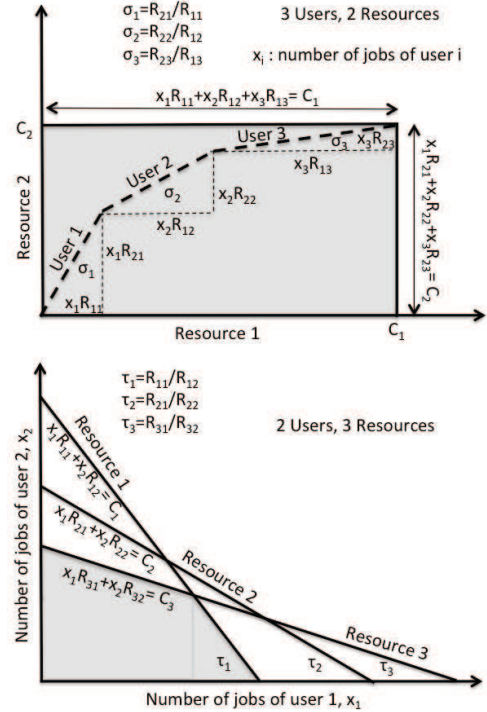


Fig. 2. Two visualizations of user heterogeneity. The lines in the top graph show the ratio of users' requirements for two different resources, while the lines in the bottom graph show the feasible allocation region. The slopes of those lines reflect the ratio of two users' requirements for each resource.

Note that  $\beta = 1$  is a trivial case, since (3) then reduces to  $n \left( \sum_{j=1}^n \mu_j x_j \right)^{\lambda}$ , so that each allocation gives equal fairness. We make a standard assumption that all resources and all jobs are infinitely divisible, which is typical of many multi-resource settings [17], [18]. An illustrative example of FDS is given in Section III-B3.

The fairness function (3) may be divided into two components, one representing fairness and one efficiency. The sum of the dominant shares raised to the power  $\lambda$  represents efficiency; thus,  $\lambda$  parametrizes efficiency's relative importance.

The remainder of (3) is parametrized by  $\beta$  and represents the fairness of the allocation. It is easily seen that for any value of  $\beta \neq 1$ , this component of (3) is maximized at an equal allocation. However, different values of  $\beta$  will yield different orderings of unequal allocations. One allocation may be more fair than another when  $\beta = \beta_1$  is used to parametrize fairness, but the second allocation may be more fair than the first when  $\beta = \beta_2 \neq \beta_1$  is used.

Though different values of  $\beta$  give different types of fairness, we can generally say that "larger  $\beta$  is more fair." As  $\beta \rightarrow \infty$ , we obtain max-min fairness on the ratio of each user's dominant share to the sum of all the dominant shares.

As  $\beta \rightarrow \infty$  and  $\lambda = \frac{1-\beta}{\beta}$ , the fairness function  $f_{\beta,\lambda}$  approaches max-min fairness on the dominant shares. Dominant resource fairness (DRF), proposed in [8], is thus a special case of FDS. Again letting  $\mu_j x_j$  denote the dominant share of user  $j$ , DRF can be expressed as

$$\min \{ \mu_1 x_1, \mu_2 x_2, \dots, \mu_n x_n \}. \quad (4)$$

Maximizing this equation subject to the constraints  $\sum_{j=1}^n R_{ij}x_j \leq C_i, \forall i$ , yields the DRF-optimal allocation.

FDS is therefore a generalization of DRF, in which choosing the parameters  $\beta$  and  $\lambda$  allows one to achieve different tradeoffs between fairness and efficiency.

FDS also includes the well-known  $\alpha$ -fairness family of functions as a special case. This fact easily follows from the relationship of the single-resource functions in [4] to  $\alpha$ -fairness, which is generally used to measure fairness in bandwidth allocation (see references in Section II). Taking  $\alpha = \beta \geq 0$  and  $\lambda = \frac{1-\beta}{\beta}$ , the FDS function (3) becomes

$$\text{sgn}(1 - \beta) \left( \sum_{i=1}^n (\mu_i x_i)^{1-\beta} \right)^{\frac{1}{\beta}}; \quad (5)$$

optimizing this function is equivalent to optimizing the  $\alpha$ -fairness function on dominant shares

$$\sum_{j=1}^n \frac{(\mu_j x_j)^{1-\alpha}}{1-\alpha}. \quad (6)$$

2) *Generalized Fairness on Jobs (GFJ)*: Since some users require more resources per job than others, it might be more fair for those who require more resources to be allocated fewer jobs, thus increasing efficiency across all users. FDS captures this perspective. However, an individual user often cares only about the number of jobs processed (without accounting for heterogeneous resource requirements), and hence each user's notion of fairness may be based only on the number of jobs she is allocated. This motivates us to introduce another fairness measure called Generalized Fairness on Jobs (GFJ), which uses the number of jobs allocated (instead of dominant shares) in the fairness function.

GFJ can be further motivated with bandwidth allocation examples. The utility function used in these scenarios is generally  $\alpha$ -fairness applied to the bandwidth allocated to each flow. These functions are therefore a special case of GFJ, a family of functions given by

$$f_{\beta,\lambda}^{\text{GFJ}} = \text{sgn}(1 - \beta) \left( \sum_{j=1}^n \left( \frac{x_j}{\sum_{k=1}^n x_k} \right)^{1-\beta} \right)^{\frac{1}{\beta}} \left( \sum_{k=1}^n x_k \right)^{\lambda}. \quad (7)$$

Here  $\beta$  and  $\lambda$  are two parameters (just as in FDS) and  $x_j$  is the number of jobs processed for user  $j$ . As for FDS, we have the resource constraints  $\sum_{j=1}^n R_{ij} x_j \leq C_i$  for each resource  $i$ . An illustrative example is given in the next section.

For  $\beta > 0$  and  $\lambda = \frac{1-\beta}{\beta}$ , GFJ reduces to  $\alpha$ -fairness on the number of jobs allocated to each user.

3) *Differences between FDS and GFJ*: We can summarize FDS' and GFJ's approaches as follows:

- FDS measures fairness in terms of the relative size of the dominant shares, explicitly accounting for heterogeneous resource requirements in both the objective function and the constraints. As a limiting case of FDS, DRF also follows this approach.
- On the other hand, GFJ measures fairness only in terms of the number of jobs allocated to each user; the heterogeneity in resource requirements only appears in the resource constraints. Users requiring more resources are thus treated equally, a result observed in Section V.




Fig. 3. Overall schematic of our multi-resource fairness approach.

When  $\mu_j = \mu$  for all  $j$ , FDS and GFJ are equivalent.

Revisiting the example in the Introduction, we have the resource constraints  $2x_1 + 2x_2 \leq 6$  and  $3x_1 + x_2 \leq 4$ . Thus, the dominant share of user 1 is  $\frac{3}{4}x_1$ , since user 1 requires  $\frac{3}{4}$  of the available CPUs and  $\frac{1}{3}$  of the available memory for each job. Similarly, the dominant share of user 2 is  $\frac{1}{3}x_2$ , since user 2 requires  $\frac{1}{3}$  of the available memory and  $\frac{1}{4}$  of the available CPUs for each job. FDS and GFJ can then be expressed as

$$\begin{aligned} \max_{x_1, x_2} f(x_1, x_2) \\ \text{s.t. } 2x_1 + 2x_2 \leq 6, 3x_1 + x_2 \leq 4, \end{aligned} \quad (8)$$

where the fairness function is

$$f = \text{sgn}(1 - \beta) \left( \frac{\left(\frac{3x_1}{4}\right)^{1-\beta} + \left(\frac{x_2}{3}\right)^{1-\beta}}{\left(\frac{3x_1}{4} + \frac{x_2}{3}\right)^{1-\beta}} \right)^{\frac{1}{\beta}} \left( \frac{3x_1}{4} + \frac{x_2}{3} \right)^{\lambda}$$

for FDS and

$$f = \text{sgn}(1 - \beta) \left( \frac{x_1^{1-\beta} + x_2^{1-\beta}}{(x_1 + x_2)^{1-\beta}} \right)^{\frac{1}{\beta}} (x_1 + x_2)^{\lambda}$$

for GFJ.

Figure 3 illustrates the approaches to multi-resource fairness. We transpose the matrix  $\mathbf{R}$  to capture users' resource requirements; each row represents one user's requirements. One simplistic approach would assume perfectly substitutable resources; in that case, this matrix immediately collapses into a vector of users' single resource requirements. However, this substitutability often does not hold. For example, CPUs and memory are not directly substitutable.

FDS and GFJ represent alternative approaches to the scalarization of each row in Fig. 3's matrix. FDS and its limiting case DRF choose a dominant entry from the row vector of users' requirements. GFJ, on the other hand, scalarizes each row by the number of jobs processed with a bundle of different resources. These row-by-row scalarizations then yield another vector of users' scalars; evaluating fairness with  $f_{\beta,\lambda}^{\text{FDS}}$  or  $f_{\beta,\lambda}^{\text{GFJ}}$  further reduces this vector to a final scalar quantifying fairness.

#### IV. PROPERTIES OF FDS AND GFJ

In this section, we prove key properties of the FDS and GFJ functions introduced above. Section IV-A characterizes the optimal fairness values in certain special cases, while Section IV-B examines the conditions of  $\beta$  and  $\lambda$  under which FDS and GFJ satisfy important properties relevant to fairness quantification and fairness-efficiency tradeoffs:

- What happens to the optimal allocations when users have the same resource requirements?

- What fairness properties do FDS and GFJ satisfy? For instance, are their optimal allocations Pareto-efficient? Sharing incentive compatible? Envy-free?
- Does there always exist a fairness-efficiency tradeoff?

Finally, Section IV-C examines the conditions under which a fairness-efficiency tradeoff exists.

We consider  $n$  users and  $m$  different resources. Users have the same resource requirements when they are homogeneous, i.e., their heterogeneity is zero. In the special cases  $n = 2$  or  $m = 2$ , user heterogeneity may be easily visualized as in Fig. 2 in Section III-A. We use the term *user-resource system* to refer to a given set of resources and users with associated resource requirements and capacities.

#### A. Values of FDS and GFJ

Heterogeneity is measured by the variance in the slopes  $\sigma_i$  or  $\tau_i$  of Fig. 2. When all users have the same ratios of multi-resource requirements (i.e., the variance of the  $\{\sigma_i\}$  and  $\{\tau_i\}$  is zero), the problem reduces to that of a single resource:

*Proposition 1 (Reduction to Single-Resource Case):*

Suppose that the resource constraints may be written as

$$\eta_i (\mu_1 x_1 + \mu_2 x_2 + \dots + \mu_n x_n) \leq 1, \quad (9)$$

$i = 1, 2, \dots, m$ . Let  $\eta_{\max} = \max_i \eta_i$ . Then the problem reduces to single-resource fairness on resource 1. Moreover, FDS and DRF both yield the allocation  $x_j = \frac{1}{\eta_{\max} \mu_j n}$ . GFJ

yields the allocation  $x_j = \frac{\mu_j^{-\frac{1}{\beta}}}{\eta_{\max} \sum_{i=1}^n \mu_i^{\frac{\beta-1}{\beta}}}$ .

*Definition 1 (Efficiency):* Let  $X = x_1 + x_2 + \dots + x_n$  denote the allocation efficiency.

In this special case, we also have the following corollary:

*Corollary 1:* For allocations that maximize DRF and FDS,

$$\frac{\partial X}{\partial \mu_j} = \left( \frac{-1}{n \eta_{\max}} \right) \left( \frac{1}{\mu_j^2} \right)$$

and the efficiency of these allocations increases the fastest if  $\min_j \mu_j$  is decreased. For allocations that maximize GFJ,

$$\frac{\partial X}{\partial \mu_j} = \frac{-\mu_j^{-\frac{1+\beta}{\beta}}}{\eta_{\max} \beta \sum_{i=1}^n \mu_i^{\frac{\beta-1}{\beta}}} + \frac{(1-\beta) \mu_j^{-\frac{1}{\beta}} \sum_{i=1}^n \mu_i^{-\frac{1}{\beta}}}{\eta_{\max} \beta \left( \sum_{i=1}^n \mu_i^{\frac{\beta-1}{\beta}} \right)^2}.$$

In other words, the system's efficiency will increase if the user with the lowest  $\mu_j$  gives up some resources.

We now consider heterogeneous users, and assume that their resource requirements  $R_{ij}$  are uniformly distributed in  $[0, \nu C_i]$ ,  $\nu$  a given positive constant. Then, as the number of users  $n$  goes to infinity, the optimal FDS and GFJ values converge as follows:

*Proposition 2 (Optimal FDS and GFJ Values):* The optimal FDS value converges in probability as

$$\lim_{n \rightarrow \infty} (\max f_{\infty, -1}^{\text{FDS}})^{-1} \cdot \frac{2m}{n(m+1)} = 1. \quad (10)$$

Thus, users' asymptotic dominant share is  $\frac{1}{n} \cdot \frac{2m}{m+1}$ . In contrast, the optimal GFJ value converges in probability as

$$\lim_{n \rightarrow \infty} (\max f_{\infty, -1}^{\text{GFJ}})^{-1} \cdot \frac{2}{\nu (\sqrt{mn/3} + n)} = 1. \quad (11)$$

Users are asymptotically allocated resources for  $\frac{2}{\nu n}$  jobs. We note that  $\nu$  appears in (11) but not (10), since the dominant shares, not the number of jobs, appear in the FDS objective function. Scaling the resource requirements  $R_{ij}$  by  $\nu$  is equivalent to scaling the optimal allocations  $x_j$  by  $\nu^{-1}$ ; these cancel in calculating the dominant shares  $\mu_j x_j$ .

We thus see that in the limit of a large number of heterogeneous users, with  $\beta = \infty$  and  $\lambda = -1$ , the optimal FDS value increases while the optimal GFJ value decreases as more resources are added to the system. This proposition highlights the fundamental difference between FDS and GFJ: in the limit, they yield very different allocations.

#### B. Three Key Properties of Fairness

We next turn our attention to fairness and its relationship with efficiency, using three widely-used properties of fairness functions (see e.g., [8] and the many references therein):

*Definition 2:* A function  $f$  is **Pareto-efficient** if, whenever  $\mathbf{x}$  Pareto-dominates  $\mathbf{y}$  (i.e.,  $x_i \geq y_i$  for each index  $i$  and  $x_j > y_j$  for some  $j$ ),  $f(\mathbf{x}) > f(\mathbf{y})$ .

*Definition 3:* **Sharing incentive** is the property that no user's dominant share is less than  $\frac{1}{n}$ ; each user has an incentive not to simply split the resources equally.

*Definition 4:* **Envy-freeness** holds if and only if no user envies another user's allocation. Mathematically, let  $r_{ij}$  denote the amount of resource  $i$  allocated to user  $j$ . User  $j$  can then process  $\max_i r_{ij}/R_{ij}$ . Envy-freeness is defined as the property that  $\max_i r_{ij}/R_{ij} > \max_i r_{ik}/R_{ij}$  for any  $j \neq k$ . In words, no other user's allocation would enable a user to process more jobs than her allocation would.

We investigate if and when these properties are satisfied by FDS and GFJ. Our results show that the answer depends on several factors, e.g. the values of the parameters  $\beta$  and  $\lambda$ . Table I summarizes our findings.

We first consider Pareto-efficiency. Evidently, this property holds for large  $\lambda$ . Based on [4], we can in fact specify a threshold for  $\lambda$  above which Pareto-efficiency holds:

*Proposition 3 (Pareto-efficiency of FDS and GFJ):* The fairness functions (3) and (7) are Pareto-efficient when  $\beta > 0$  if and only if  $|\lambda| \geq \left| \frac{1-\beta}{\beta} \right|$ .

The absolute value signs are necessary, as for  $\beta > 1$ , (3) and (7) are negative. For this range of  $\beta$ , a more negative  $\lambda$  therefore emphasizes efficiency. As Pareto-efficiency is a highly desirable property for fairness functions (both single and multi-resource), the following analysis considers only values of  $\lambda$  satisfying  $|\lambda| \geq \left| \frac{1-\beta}{\beta} \right|$ .

*Proposition 4 (Sharing Incentive of FDS):* Suppose  $\beta > 0$ . Then we can prove the following:



- (a) Sharing incentive is satisfied by the FDS-optimal allocation when  $\lambda = \frac{1-\beta}{\beta}$  and  $\beta > 1$ .
- (b) For  $0 < \beta < 1$  and  $\lambda = \frac{1-\beta}{\beta}$ , there exists a user-resource system such that the FDS-optimal allocation for this system does not satisfy the sharing incentive property.
- (c) For any  $\beta > 0$ , there exists  $\lambda$  with  $|\lambda|$  sufficiently large so that for some user-resource system, the FDS-optimal allocation need not satisfy the sharing incentive property.
- (d) If  $\lambda = 0$ , then the FDS-optimal allocation always satisfies the sharing incentive property.

We can further bound the allocation efficiency:

**Corollary 2 (Bounds on Allocation Efficiency of FDS):** If  $\beta > 0$  and  $\lambda = \frac{1-\beta}{\beta}$ , the efficiency  $X \geq \frac{1}{\max_j \mu_j}$ .

For  $\lambda = \frac{1-\beta}{\beta}$ , the FDS function becomes equivalent to the isoelastic  $\alpha$ -fair utility in economics;  $\beta$  corresponds to a measure of constant relative risk-aversion for individual users.<sup>3</sup> As  $\beta$  increases, individual risk-averse users find the resource allocation more equitable and become collectively envy-free. The following corollary establishes that this interesting envy-free behavior emerges (for FDS) at a *threshold* of  $\beta > 1$ :

**Corollary 3 (Envy-Freeness of FDS):** For  $\beta > 0$  and  $\lambda = \frac{1-\beta}{\beta}$ , the envy-freeness property holds if  $\beta > 1$ ; if  $\lambda = 0$ , then envy-freeness holds for all user-resource systems and any  $\beta$ . Moreover, there exists a user-resource system whose FDS-optimal allocation does not satisfy envy-freeness under the same conditions (b) and (c) in Prop. 4 for which the sharing incentive property does not always hold.<sup>4</sup>

In contrast to FDS, GFJ need not always satisfy sharing incentive even for  $\beta > 1$ :

**Proposition 5 (Sharing Incentive of GFJ):** Suppose again that  $\beta > 0$ . Then under the conditions enumerated below, there exists a user-resource system whose GFJ-optimal allocation does not satisfy the sharing incentive property:

- (a)  $|\lambda| = |(1-\beta)/\beta|$ ,
- (b)  $|\lambda| > |(1-\beta)/\beta|$  and  $0 < \beta < 1$ ,
- (c)  $|\lambda| < |(1-\beta)/\beta|$  and  $\beta > 1$ ,
- (d)  $|\lambda|$  sufficiently large,
- (e)  $\lambda = 0$ .

Similarly, GFJ-optimal allocations need not be envy-free for any value of  $\beta$ :

**Corollary 4 (Envy-Freeness of GFJ):** Under the conditions specified in Prop. 5, there exists a user-resource system such that envy-freeness does not hold for the GFJ-optimal allocation.

Figure 4 illustrates Props. 4 and 5's results on the sharing incentive property, as well as Corollaries 3 and 4's results on envy-freeness.

<sup>3</sup>**Isoelasticity and relative risk-aversion** in economics are defined as  $\frac{\partial u(x)}{\partial x} \frac{x}{u(x)}$  and  $-\frac{x u''(x)}{u'(x)}$  respectively, where  $u$  is the utility function.

<sup>4</sup>Though it may appear so from this proposition, sharing incentive and envy-freeness are *not* equivalent [8].

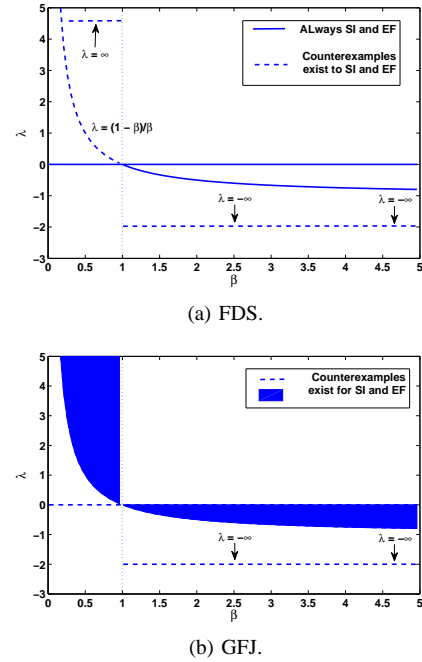


Fig. 4. Conditions under which sharing incentive (SI) and envy-freeness (EF) can be shown either to hold or not to hold (c.f. Props. 4 and 5 and their corollaries 3 and 4).

### C. Fairness-Efficiency Tradeoff

We now consider two ways in which a fairness-efficiency tradeoff does not exist: first, an increased emphasis on fairness need not decrease efficiency. Second, the efficiency-maximizing allocation may also be the “most fair.”

Traditionally, a larger parameter  $\alpha$  in  $\alpha$ -fairness functions is thought to be “more fair” [19], [20]; this statement is made mathematically precise in [4]. In [12], however, it is shown that when a network allocates bandwidth so as to maximize  $\alpha$ -fairness, total throughput in the network will sometimes increase with  $\alpha$ . It may even decrease as capacity increases. These “counter-intuitive” results hold in the general multi-resource problem:

Consider the general family of utility functions  $U(\mathbf{x}, \alpha)$ ; here  $\alpha$  is a parameter indexing the family of utility functions, and the specific functional form of  $U$  is not specified. For instance, we could use the functions in (3), with  $\alpha = \beta$  and  $\lambda = \frac{1-\beta}{\beta}$ , so that the utility function uses “ $\alpha$ -fairness.” We incorporate the resource capacity constraints in the matrix inequality  $\mathbf{R}\mathbf{x} \leq \mathbf{C}$  and assume that  $\mathbf{R}$  is a matrix of full row rank consisting only of those constraints which are tight at the optimal allocation  $\mathbf{x}$  for the given value of  $\alpha$ .

We let  $\mathbf{S}$  be an  $n \times (n - m)$  dimensional matrix whose columns form a basis for the nullspace of  $\mathbf{R}$ , and again let  $X = \sum_{j=1}^n x_j$  denote the total efficiency. The negative of the utility function's Hessian matrix is denoted by  $\mathbf{D}$ , and we define  $\mathbf{b} = \frac{\partial^2 U}{\partial x \partial \alpha}$ ,  $\mathbf{A} = \mathbf{S}^T \mathbf{D} \mathbf{S}$ ,  $\mathbf{v}_j = \mathbf{s}_j^T \mathbf{b}$  and  $\beta_j = -\mathbf{1}^T \mathbf{s}_j$ , where the  $\mathbf{s}_j$  are the columns of the matrix  $\mathbf{S}$ . Let  $\bar{\mathbf{A}}_i$  denote the matrix  $\mathbf{A}$  with the  $i$ th row replaced by  $\beta = [\beta_1 \ \beta_2 \ \cdots \ \beta_n]$ . We use  $\delta$  to denote a direction of perturbation of the capacity vector  $\mathbf{C}$  and  $\mathbf{D}X(\delta)$  to denote the derivative of  $X$  in the

Fairness	Sharing Incentive		Envy-Freeness	
FDS	$\lambda = \frac{1-\beta}{\beta}, 0 < \beta < 1$	$\lambda = \infty, \text{ any } \beta$	$\lambda = \frac{1-\beta}{\beta}, 0 < \beta < 1$	$\lambda = \infty, \text{ any } \beta$
GFJ	$\lambda = \frac{1-\beta}{\beta}, \beta > 0$	$\lambda = \infty \text{ or } 0, \text{ any } \beta$	$\lambda = \frac{1-\beta}{\beta}, \beta > 0$	$\lambda = \infty \text{ or } 0, \text{ any } \beta$
	$ \lambda  < \frac{ 1-\beta }{\beta}, \beta > 1$	$ \lambda  > \frac{ 1-\beta }{\beta}, 0 < \beta < 1$	$ \lambda  < \frac{ 1-\beta }{\beta}, \beta > 1$	$ \lambda  > \frac{ 1-\beta }{\beta}, 0 < \beta < 1$

TABLE I  
CONDITIONS UNDER WHICH PROPERTIES DO NOT HOLD FOR ALL USER-RESOURCE SYSTEMS.

direction of  $\delta$ . From [12], we have

$$\frac{\partial X}{\partial \alpha} = \mathbf{1}^T \mathbf{S} \mathbf{A}^{-1} \mathbf{S}^T \mathbf{b} \quad (12)$$

$$\mathbf{D}X(\delta) = \mathbf{1}^T \frac{\partial x}{\partial \mathbf{C}} \delta = \mathbf{1}^T \mathbf{D}^{-1} \mathbf{R}^T (\mathbf{R} \mathbf{D}^{-1} \mathbf{R}^T)^{-1} \delta. \quad (13)$$

We can further prove the following proposition:

**Proposition 6 (Efficiency Non-Monotonicity):** Efficiency increases with  $\alpha$  if and only if

$$\sum_{i=1}^{N-L} v_i \det \bar{\mathbf{A}}_i \geq 0. \quad (14)$$

Moreover, efficiency may decrease with an increase in the capacity vector  $\mathbf{C}$ . If capacity increases proportionally, i.e.,  $\delta = \epsilon \mathbf{C}$  for some small  $\epsilon$ , then  $\mathbf{D}X(\delta) \geq 0$ .

As a special case, when only one capacity constraint is tight (e.g., one resource), efficiency always increases with capacity. The technical report [3] contains a numerical example in which efficiency increases with  $\beta$ .

We next examine the conditions under which an equal allocation (equal dominant shares for FDS or an equal number of jobs for GFJ) maximizes efficiency. In these situations, there is no fairness-efficiency tradeoff; the most fair allocation maximizes the total number of jobs processed. As this property is an ideal case, it will likely be satisfied only under rather stringent conditions. Indeed, our results show that this ideal case occurs only when the resource constraints “line up” exactly.

We again express the resource constraints in matrix form as  $\mathbf{R}\mathbf{x} \leq \mathbf{C}$ , and simplify them to  $\gamma\mathbf{x} \leq \mathbf{1}_m$ , where  $\mathbf{1}_m$  is a vector of  $m$  1's and  $\gamma_{ij} = \frac{R_{ij}}{C_i}$ .

**Proposition 7 (Maximizing Fairness and Efficiency (I)):** Suppose that  $m = n$  constraints are tight at the maximum-efficiency allocation. Then this allocation equalizes the dominant shares (FDS has no fairness-efficiency tradeoff) if and only if

$$\sum_{j=1}^n \frac{\gamma_{ij}}{\mu_j} = \rho \quad (15)$$

for some constant  $\rho$  and all resources  $i$ . The number of jobs per user is equalized (GFJ has no fairness-efficiency tradeoff) if

$$\sum_{j=1}^n \gamma_{ij} = r \quad (16)$$

for some constant  $r$  and all resources  $i$ .




Fig. 5. Illustration of Prop. 7.




Fig. 6. Illustration of Prop. 8 in two dimensions. In the top graph, exactly one resource constraint is tight at the unique efficiency-maximizing allocation, and  $x_2 = 0$ . In the bottom graph, exactly one resource constraint is tight at any of the multiple efficiency-maximizing allocations.

Looking back at Fig. 2, we see that the number of jobs per user is equal at the efficiency-maximizing allocation if  $\sigma_1 = \dots = \sigma_n$  for  $n$  users and two resources. For two users and  $m$  resources, the number of jobs per user is equal at the efficiency-maximizing allocation if  $\sum_{j=1}^m \tau_j R_{j2} = \sum_{j=1}^m R_{j2}$ .

Our conclusions are more subtle when  $m < n$  constraints are tight at an efficiency-maximizing allocation:

**Proposition 8 (Maximizing Fairness and Efficiency (II)):** Suppose that  $m < n$  constraints are tight at an efficiency-maximizing allocation  $\mathbf{x}^*$ . If this allocation is the unique allocation maximizing efficiency, then at least one of the  $x_j^* = 0$  and one user is allocated no jobs. If other allocations also maximize efficiency, an allocation equalizing either the dominant shares or number of jobs processed maximizes efficiency if and only if at the equal allocation, the constraint set intersects the hyperplane  $\sum_{j=1}^n x_j = \sum_{j=1}^n x_j^*$  on a set of dimension at least 1.

Figure 6 shows the two-dimensional illustration of this theorem's statements. The top graph shows a unique efficiency-maximizing allocation when exactly one resource constraint is tight, and the bottom graph shows a set of multiple efficiency-maximizing allocations.

We can use this proposition to derive a sufficient condition for the efficiency-maximizing allocation to equalize the dominant shares or number of jobs for each user:

*Corollary 5:* Suppose  $m < n$  resource constraints hold at the efficiency-maximizing allocation. Then if  $R_{ij} > R_{ik}$  for some users  $j$  and  $k$  and all resources  $i$ ,  $x_j = 0$  (user  $j$  is allocated no jobs) at any efficiency-maximizing allocation.

If  $m = 1$  (the single-resource case), this result implies the following:

*Corollary 6:* The maximum efficiency allocation equalizes the dominant shares (FDS) or jobs per user (GFJ) if and only if  $\mu_j = \mu \forall$  users  $j$ . In other words, each user needs the same amount of the single resource to process one job.

## V. APPLICATIONS AND ILLUSTRATIONS

We consider an illustrative example of a datacenter with CPU and RAM constraints. There are two users, each of whom requires a fixed amount of each resource to accomplish a job. Jobs are assumed to be infinitely divisible [17], [18]. In order to benchmark performance, we use the same parameters as [8]: user 1 requires 1 CPU and 4 GB of RAM for each job, and user 2 requires 3 CPUs and 1 GB of RAM for each job. There are 9 CPUs and 18 GB of RAM at first. We then vary these constraint values to observe their impact on fairness.

Suppose that the fairness function is given by  $f$  (e.g. FDS (3), DRF (4), GFJ (7)). Then the allocation problem is

$$\max_{x,y} f(x,y) \quad (17)$$

$$\text{s.t. } x + 3y \leq 9, 4x + y \leq 18 \quad (18)$$

where  $x$  and  $y$  are the number of jobs allocated to users 1 and 2 respectively.

We use DRF as the benchmark fairness to compare the performance of our FDS and GFJ functions. We define **percent fairness** as the percentage difference between the optimal DRF fairness value (i.e., the minimum dominant share) and the DRF fairness value of the allocation obtained from FDS or GFJ. The **percent efficiency** is defined as the percentage difference between the total number of jobs processed in the given allocation and the maximum number of jobs that can be processed, given the same capacity constraints. We also introduce another efficiency measure, the **leftover capacity** (i.e., the amount of unused resources).

We investigate the outcomes of the proposed fairness measures along two dimensions:

- Comparing the achieved efficiency when user heterogeneity and resource capacity are varied.
- Examining the range of attainable fairness-efficiency tradeoffs for different values of the parameters  $\beta$  and  $\lambda$ .

### A. Efficiency

We first use our two efficiency measures—leftover capacity and percent efficiency—to investigate user heterogeneity’s effect on achieved efficiency. Heterogeneity is measured by the variance in the slopes  $\tau_i$  and in the slopes  $\sigma_i$  of users’ resource requirements, as introduced in Fig. 2 in Section III-A. If two users have identical resource requests, they become




Fig. 7. Too much or too little variance in  $\tau$  leads to inefficiency from leftover capacity: Leftover capacity versus variance in user heterogeneity in a datacenter example. Variances below 0.5 have only leftover CPUs; variances above 0.5 have only leftover RAM.

homogeneous, and both variances are 0. At the other extreme, the users do not share any resource requirements; they become decoupled, with infinite variances.

We calculate the optimal FDS, GFJ and DRF allocations for  $\beta = 2$ ,  $\lambda = -0.5$ . First, Fig. 7 examines the leftover capacity as a function of the variance in  $\tau$ . The heterogeneity was varied by changing the RAM requirement of user 2 from 1 GB to 13 GB. Thus, the RAM constraint line in Fig. 2’s representation tilts from very steep to very flat. This tilting geometrically explains the overall “V” trend in Figs. 7 and 8. When the RAM requirement is below 3 GB (a steep constraint line), the variance of  $\tau$  is over 0.5 and the variance of  $\sigma$  is over 4.5: only RAM is leftover. When the RAM requirement is above 3 GB (a flatter line), the variance of  $\tau$  is less than 0.5 and the variance of  $\sigma$  less than 4.5: only CPUs are leftover. The change in the leftover resource is due to the changing shape of the feasible region.

In this example, we see that for low heterogeneity in users’ resource requirements, FDS, GFJ, and DRF have similar efficiency values. In fact, Prop. 1 states that at zero heterogeneity, DRF and FDS are optimized at the same allocation, predicting part of the observed behavior. As the heterogeneity increases, DRF has a lot of leftover capacity compared to GFJ and FDS, especially for a variance larger than 1 in Fig. 7 and larger than 5 in Fig. 8. DRF trades off efficiency significantly to preserve users’ minimum dominant share with increasingly heterogeneous resource requirements. Even GFJ performs worse than FDS, which yields the lowest leftover capacity. As FDS includes resource requirements in its fairness function, we intuitively expect such a result.

We next examine the percent efficiency in jobs processed as a function of the variances in  $\tau$  and  $\sigma$  in Figs. 9 and 10. As in the previous figures, for low heterogeneity across users’ resource requirements, FDS, GFJ, and DRF perform at similar efficiency levels. All three achieve full efficiency for a  $\tau$  variance near 0.5 and  $\sigma$  variance near 4.5. Again, the efficiency attained is also much higher (about 15%) for FDS and GFJ than for DRF as the variance increases.

In summary, enforcing DRF can significantly reduce efficiency as measured by either leftover capacity or percent efficiency. This is also the case when the number of users grows; Fig. 11 shows the leftover capacity versus the number






Fig. 8. Too much or too little variance in  $\sigma$  leads to inefficiency from leftover capacity: Leftover capacity versus variance in user heterogeneity in a datacenter example. Variances below 4.5 have only leftover CPUs; variances above 4.5 have only leftover RAM.




Fig. 9. Greater variance in  $\tau$  leads to DRF inefficiency in the number of jobs processed: Percentage efficiency versus variance in user heterogeneity in a datacenter example.

of users in the system. Only RAM capacity was leftover; in all scenarios, all of the CPUs were used. For a large number of users, we see that FDS and GFJ both use more capacity than DRF. Users' CPU requirements were fixed at 2 CPUs; their RAM requirements were drawn from a uniform distribution. Other randomly chosen RAM requirements yield similar plots.

Finally, we examine the impact of changing RAM capacity on the attainable efficiency levels. Figure 12 shows how varying this capacity affects the efficiency attained at the optimal allocation. We see that when the dominant shares for both users are equal, at 12 GB of RAM capacity, GFJ and FDS have the same range of achievable efficiency. Moreover,




Fig. 10. Greater variance in  $\sigma$  leads to DRF inefficiency in the number of jobs processed: Percentage efficiency versus variance in user heterogeneity in a datacenter example.




Fig. 11. Even with a large number of users, DRF uses less available capacity than FDS and GFJ: Leftover capacity versus the number of users in a datacenter example.

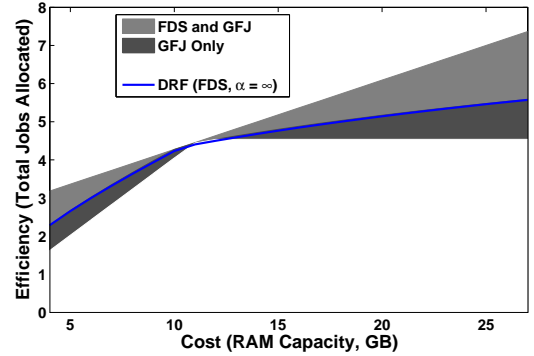


Fig. 12. Capacity expansion can increase the range of operating efficiencies for FDS and GFJ over DRF: Attainable efficiency for varying capacity constraint, given different implicit realizations of  $\beta \in (-5, 5)$  and  $\lambda \in (0.01, 1.91)$  for  $\beta < 0$ ,  $\lambda \in \left(0.005 \left(\frac{1}{\beta} - 2\right), 0.955 \left(\frac{1}{\beta} - 2\right)\right)$  for  $\beta > 0$  values. The region labels refer to the fairness functions that attain those efficiencies.

$\beta$  and  $\lambda$  can be chosen to achieve higher efficiency in FDS and GFJ. The DRF function serves as a “lower bound” to the efficiency values attainable with the FDS functions.

The impact of capacity expansion also highlights an interesting dimension of the *economy of scale* in large networks. The standard view is that a large scale helps smooth out temporal fluctuations of demands through statistical multiplexing, e.g., at any aggregation point in a broadband access network. In addition to temporal “heterogeneity” (bursting at different times), network users may have *resource type heterogeneity*: some applications need more CPU processing while others need more storage or bandwidth. Can this heterogeneity be exploited to utilize different types of resources more efficiently? The answer depends on how these different resources are allocated among the users. If DRF is used, for example, efficiency can be quite low. However, by using the appropriate FDS parametrization, resource request heterogeneity can indeed be leveraged along with increases in resource capacity and turned into another type of economy of scale.

### B. Fairness-Efficiency Tradeoffs

The previous section established that when users are very heterogeneous, FDS and GFJ outperform DRF, achieving a much greater efficiency. However, we expect that this larger efficiency comes at a cost of decreased fairness. This section




Fig. 13. Larger  $\beta$  values lead to more equitable allocations: Optimal allocations for various fairness measures in a datacenter example, using  $\alpha = \beta$  fairness for FDS and GFJ.




Fig. 14. The fairness-efficiency tradeoff can be tuned by changing  $\beta$ : Percentage of fairness and efficiency achieved for various fairness measures in a datacenter example, using  $\alpha = \beta$  fairness for FDS and GFJ. Notice that an increased emphasis on fairness (i.e. larger  $\beta$ ) need not always decrease the efficiency of the allocation, as seen for  $\beta < 2.6$  for GFJ measure.

examines the general behavior of fairness when a larger efficiency is achieved. Here we measure fairness as percent fairness with the DRF metric and efficiency as percent efficiency on the number of jobs processed.

Figure 13 shows the optimal allocations of jobs for different values of  $\beta$ ,  $\lambda = \frac{1-\beta}{\beta}$ . Both FDS and GFJ become  $\alpha$ -fair on the dominant shares of and jobs allocated to each user, respectively, for  $\alpha = \beta$ . As  $\beta$  increases,  $\lambda$  decreases, so that fairness is emphasized more than efficiency and FDS asymptotes to DRF. For small  $\beta$  (i.e., more relative emphasis on efficiency than fairness), the optimal FDS allocation maximizes efficiency. In the case of GFJ, which emphasizes the fairness on jobs allocated, larger  $\beta$  values produce a more fair allocation of jobs across users than FDS, as expected. Consequently, the total number of jobs processed (i.e., efficiency) is lower for GFJ than for FDS.

Figure 14 gives a representative plot of how this tradeoff varies with  $\beta$  and  $\lambda = \frac{1-\beta}{\beta}$ . As  $\beta$  grows larger, the percent efficiency from the FDS measure drops, approaching DRF in the limit  $\beta \rightarrow \infty$ . The GFJ fairness increases until  $\beta = 2.6$ , at which point the GFJ-optimal allocation is also DRF-optimal. (We see in Fig. 13 that the GFJ allocation “crosses” the DRF allocation line at this value of  $\beta$ ). For larger values of  $\beta$ , GFJ quickly converges to an allocation with a more equal number of jobs per user; thus, its efficiency decreases. But efficiency in FDS decreases more slowly since FDS attempts to make the dominant shares, not the number of jobs, more equitable.




Fig. 15. Capacity expansion allows different FDS fairness-efficiency tradeoff contours: Attainable efficiency vs. fairness tradeoffs from different implicit realizations of  $\beta \in (-5, 5)$  and  $\lambda \in (0.01, 1.91)$  for  $\beta < 0$ ,  $\lambda \in \left(0.005 \left(\frac{1}{\beta} - 2\right), 0.955 \left(\frac{1}{\beta} - 2\right)\right)$  for  $\beta > 0$  values. DRF is used as the fairness benchmark and metric.

Finally, we show the interaction between capacity constraints and the range of fairness-efficiency tradeoffs achieved. The shaded region in Fig. 15 shows the attained tradeoffs for a large range of  $\beta$  and  $\lambda$  values; each point corresponds to some  $\beta$  and  $\lambda$  values in the FDS function that achieve the shown operating tradeoff. This achieved tradeoff depends on the available capacity, with contour lines for various RAM capacities shown in the figure. As RAM capacity increases from 4 GB to  $6\sqrt{3}$  GB, the tradeoff stops: one can increase both fairness and efficiency. At a RAM capacity of  $6\sqrt{3}$  GB, the conditions of Prop. 7 are satisfied, and efficiency is maximized when the dominant shares are equal. When the RAM capacity goes above  $6\sqrt{3}$  GB up to 25 GB, user 1’s dominant share  $\frac{4x_1}{\text{RAM capacity}}$  decreases. Thus, an increase in fairness requires an increase in  $x_1$  and user 1’s CPU allocation. User 2 is then allocated fewer jobs, decreasing efficiency. In this figure, one can achieve 100% efficiency and fairness when RAM capacity is  $6\sqrt{3}$  GB, but such an ideal operating point does not always exist.

Figure 16 shows the analogue of Fig. 15 for GFJ functions. In this case, the range of attainable efficiency at the maximum allocation decreases as the fairness value increases. Thus, one can increase both fairness and efficiency as RAM capacity goes from 4 GB to 25 GB. Moreover, the contour lines “bend back” on themselves, indicating that for different  $\beta$  and  $\lambda$  parameters, the same fairness value can result in many efficiency values at the optimal allocation. When RAM capacity equals 11.25 GB, the conditions of Prop. 7 are satisfied and there is no tradeoff between fairness and efficiency.

## VI. SURVEY ON FAIRNESS PARAMETERS

In this section, we provide results from a simple survey to complement the proposed theoretical framework with a demonstration of how the typical values of fairness function parameters can be estimated from large scale consumer surveys. We note that our survey methodology and results should be considered as a demonstration of one out of many feasible approaches rather than a prescription of what exact parameter values to choose in a given real world scenario. In particular, this survey provides a systematic way of inferring an initial estimate for  $(\beta, \lambda)$  values, visualizes participant clusters in




Fig. 16. Capacity expansion allows different GFJ fairness-efficiency tradeoff contours: Attainable efficiency vs. fairness tradeoffs from different implicit realizations of  $\beta \in (-5, 5)$  and  $\lambda \in (0.01, 1.91)$  for  $\beta < 0$ ,  $\lambda \in (0.005(\frac{1}{\beta} - 2), 0.955(\frac{1}{\beta} - 2))$  for  $\beta > 0$  values. DRF is used as the fairness benchmark and metric.

the fairness-efficiency space, and connects the FDS and GFJ functions with participants' responses.

#### A. Survey Methodology

We conducted an online survey in January-February 2012, which received 143 responses, mostly from the U.S. Out of these responses, 110 were complete and were used in the subsequent analysis. The participants were given six questions, each with a simplified 'toy' scenario of resource allocation in a datacenter, where jobs from two different clients had heterogeneous resource requirements over multiple resources (CPU and storage). Our online survey participants were faculty, students, and staff primarily from the EE and CS departments of Princeton and George Washington University. They all were familiar with everyday computer use, and hence intuitively understood the two resources considered (processing power and storage capacity). The survey questionnaire further explained the context to ensure participants' understanding.

We limited our question scenarios to only two types of resources in order to ease participants' understanding of the questions, although more sophisticated methods using conjoint analysis can be used on data with more resources [21]. In the last question, we increase the number of resources to three: clients' jobs required CPU, storage, and bandwidth. Each of the six questions offered five different options of distributing resources among the two clients, with each option resulting in a particular outcome. For each question, the survey participants were asked to rank the five allocation options in decreasing order of preference, as shown in Fig. 17.

In four of the questions, the five options that the survey participants were asked to rank were reported in terms of the number of jobs completed for each datacenter client under that option's resource allocation. In the other questions, the options were reported in terms of the leftover (unused or wasted) capacity resulting from that resource allocation option. The questions had either the same set of allocation choices or a scalar multiple, thus permitting a sanity check on whether participants made consistent choices when the outcomes were

Now consider a similar scenario as before, that is:

- Both clients want to complete as many jobs as possible (including fractional completion of jobs)
- Client A needs 1 CPU and 4 TB per job, while Client B needs 3 CPU and 1 TB per job.
- Both clients would like to be treated more or less fairly in the amount of resources allocated to them
- Both clients pay you the same amount of \$ per jobs completed. Thus, your revenue depends only on the "Total number of jobs completed"

But now you have a total of 108 CPUs and 180 TB (Terabytes) of storage available to allocate.

- The following table shows five "Allocation" options at your disposal to distribute your available resources among these two clients, and the corresponding in number of jobs completed for each client and a sum total number of jobs completed from that allocation.

Allocation Options	Allocated to Client A			Allocated to Client B			Total no. of Jobs Completed
	CPU	TB	No. of Jobs Completed for Client A	CPU	TB	No. of Jobs Completed for Client B	
Allocation 1	24	96	24	84	28	28	52
Allocation 2	12	48	12	96	32	32	44
Allocation 3	36	144	36	72	24	24	60
Allocation 4	45	180	45	0	0	0	45
Allocation 5	27	108	27	81	27	27	54

Fig. 17. Question 2 of our online fairness survey. Client A required 1 CPU and 4 TB per job, while client B required 3 CPU and 1 TB per job. The datacenter had a total of 108 CPUs and 180 TB to allocation.

reported in different metrics ('total number of jobs completed' and 'leftover resources') or were scaled by a constant factor. To avoid influencing the participant's decisions, we did not explicitly inform them of the survey's purpose, i.e., evaluating their fairness-efficiency tradeoff.

The full survey is available in Appendix C of [3]. The results obtained from analyzing the survey responses are reported in the next subsection.

#### B. Results

Our analysis of the survey results focuses on three goals:

- Evaluate consistency of the results across users with the fairness axioms in [3], [4].
- Cluster participants based on the fairness and efficiency values inferred from their preferences in their rankings of resource allocations.
- Determine the different  $\beta$  and  $\lambda$  heat maps of compatible parameter values for participants in each cluster.

We address these sequentially below.

1) *Axiom Validation*: We first use the survey results to examine our construction of the fairness functions, evaluating the consistency of the results with three of the four axioms from which these functions are constructed (see Appendix A of [3] for a full list of the axioms). To keep the survey simple, we were unable to evaluate the Axiom of Continuity, which, however, is quite intuitive.

Figure 18 shows the number of participants ranking each allocation first, second third, etc. in each question of the survey. We see that a clear consensus emerges across the participant pool: for instance, for question 2 most people rank the allocations from best to worst as 3, 5, 1, 2, 4. It is interesting to note that the fourth allocation, under which client B had no jobs done, has the lowest rank. In fact, allocation 2, which is less efficient than allocation 4, was more preferred. This result is thus consistent with the Axiom of Starvation: participants generally dislike starvation allocations, even if they are more efficient.




Fig. 18. Allocation rankings for each survey question. The responses to each question are shown in a row at each allocation, e.g. the first six dots correspond to rankings of allocation 1 in questions 1-6, the second six dots correspond to rankings of allocation 2 in questions 1-6, etc. The size of the bubble is proportional to the number of people choosing a particular allocation for a particular allocation.

We implicitly evaluate the remaining two axioms (those of Saturation and Partition) by examining the consistency of participants' responses when the allocations are scaled up or down. Our fairness functions predict that a person's rankings of different allocations should not change with this scaling. Figure 18 shows that for each question, a clear consensus ranking emerges; moreover, this ordering of allocations is consistent across all questions<sup>5</sup>. This observation is especially significant since questions 3 and 5 report the leftover capacity as a metric instead of the total number of jobs processed—thus, even when the efficiency metric changes, participants' answers are consistent across the different survey questions.

2) *Participant Clustering*: We now evaluate the consistency of different people's responses by calculating the *average preferred fairness and efficiency values* for each person and each question. These are calculated by taking a weighted average of the efficiency and fairness values for each allocation; the weights are determined by the participant's ranking of that allocation. The fairness metric is defined to be the negative of the difference between the numbers of jobs processed for clients *A* and *B*, while the efficiency metric is taken from the survey as the total number of jobs processed (or the leftover capacity). The leftover capacity is measured by the negative of the percentage of leftover capacity for each resource, to facilitate comparison of leftover CPUs with leftover GB. We use negatives for the fairness value and leftover capacity metric so that an increase in the fairness or leftover capacity value indicates a more fair or more efficient allocation.

We see from Fig. 19 that for all questions, participants tend to fall into two distinct groups, one of which puts more emphasis on efficiency, and one which puts more emphasis on fairness. The two groups have approximately equal numbers of participants (e.g., 52 in each for question 1). Moreover, these groups are consistent across questions. While the numerical fairness and efficiency values vary depending on the allocation scalarization and efficiency metric used in a question, we see that both clusters lie in approximately the same position in the graph for each question.

<sup>5</sup>This ranking consensus is simply in terms of majority agreement on the rank of the allocations, but does not mean that the individual participants'  $(\beta, \lambda)$  values agree.

3) *Parameter Choices*: We next determine  $\beta$  and  $\lambda$  values compatible with the answers in Fig. 19's clusters. The results for participants in both clusters were the same for all questions; thus, we only show the  $\beta$  and  $\lambda$  values for question 2.

We use exhaustive search for discretized  $\beta$  and  $\lambda$  values to determine whether a given person's allocation ranking is compatible with that obtained using the  $(\beta, \lambda)$  fairness function. Figure 20 shows the heat map of compatible  $\beta$  and  $\lambda$  values for a person in each of the two dominant clusters; the intensity of the color corresponds to the number of times an answer is compatible with the given  $(\beta, \lambda)$  value. A darker color indicates a larger number of compatible answers across users. In this figure, we assume that participants use a GFJ fairness function. Though no single  $(\beta, \lambda)$  value is compatible with all participants (the single black squares represent a maximum number of compatible answers), a majority of responses were compatible with some  $(\beta, \lambda)$  value: 50% of cluster 1 and 60% of cluster 2 participants agreed on at least one  $(\beta, \lambda)$  pair.<sup>6</sup>

As expected, the compatible  $\lambda$  values for cluster 1 (Fig. 20a) are higher in absolute value than those in cluster 2 (Fig. 20b), as is consistent with cluster 2 participants' preferring fairness over efficiency (Fig. 19). The reference lines in the figure show the Pareto-efficient frontier. For  $\beta > 1$ , most of the compatible  $(\beta, \lambda)$  values are below the Pareto-efficient frontier, i.e., not Pareto-efficient. This does not happen for cluster 1 participants, as might be expected since they emphasize efficiency. However, as  $\beta$  increases, more Pareto-efficient  $(\beta, \lambda)$  values are compatible with at least some answers.

Figure 21 shows the  $(\beta, \lambda)$  heat graphs for both participant clusters when FDS-fairness is used. Only the heat graphs for question 2 are shown; the other questions give similar results. We see that all of the  $(\beta, \lambda)$  values tested in Fig. 21a are compatible with the cluster 1 responses (50% of responses agree on these values). We may partially explain these results by the fact that cluster 1 participants all favor allocation 3 over allocation 5: calculating the dominant shares of each client, we see that allocation 5 actually gives clients *less equitable* dominant shares, and that the sum of dominant shares for allocation 3 is also larger than that for allocation 5. Thus, no matter which  $\beta$  and  $\lambda$  are considered, allocation 3 will be ranked above allocation 5. All  $(\beta, \lambda)$  pairs are therefore consistent with this ranking. Most participants rank the other allocations in a manner consistent with ranking 3 above 5; those participants whose additional rankings are inconsistent do not show any compatible  $(\beta, \lambda)$  values.

In contrast to cluster 1, *all* of the  $(\beta, \lambda)$  values tested are inconsistent with cluster 2's allocation preferences. We can account for this result by noting that all cluster 2 participants prefer allocation 5 (processing an equal number of jobs for each client) over allocation 3. However, allocation 3 is both more more efficient (under FDS) than allocation 5, and hence is inconsistent with cluster 2's answers if they used FDS.

<sup>6</sup>This result may be due to our discretization; for instance, using a  $\lambda$  closer to zero may improve the compatibility with cluster 2 participants, who emphasize fairness over efficiency. Using a larger  $\lambda$  may improve compatibility with cluster 1 participants. It is also possible that a minority of participants provided inconsistent responses compatible with no  $(\beta, \lambda)$  values, e.g. preferring efficiency to fairness in ranking two allocations, and fairness over efficiency in another two allocations.






Fig. 19. Average preferred fairness and efficiency values for each survey question. The size of the circle is proportional to the number of people with those particular fairness-efficiency values. Participants tend to fall into two distinct groups, with some other extraneous points.




Fig. 20. Heat map of compatible  $(\beta, \lambda)$  values for clusters 1 and 2 participants in Fig. 19, GFJ fairness. The reference line is the Pareto-efficient boundary  $|\lambda| = |(1 - \beta)/\beta|$ , and the black dot at  $(\beta, \lambda) = (2, 2)$  represents a maximum number of compatible answers. Only question 2 results are shown; those for all other questions are similar.




Fig. 21. Heat map of compatible  $(\beta, \lambda)$  values for clusters 1 and 2 participants in Fig. 19, FDS fairness. The reference line is the Pareto-efficient boundary  $|\lambda| = |(1 - \beta)/\beta|$ , and the black dot at  $(\beta, \lambda) = (2, 2)$  represents a maximum number of compatible answers. Only question 2 results are shown; results for the other questions are similar.

We thus conjecture that inconsistency arises because GFJ is a “more natural” fairness function: for certain  $\beta$  and  $\lambda$  values, most of cluster 1 and cluster 2 participants exhibit preferences consistent with GFJ fairness. While it is intuitive that most people find it natural to understand fairness in terms of jobs completed rather than dominant share, this is an interesting direction to explore through repeated and controlled behavioral experiments.

The fact that participants generally seem to follow GFJ rather than FDS fairness has interesting implications, as Props. 4 and 5 show that sharing incentive and envy-freeness are more

likely to hold when FDS is used instead of GFJ. Participants thus pay more attention to the number of jobs allocated to each client, rather than each client’s share of the resources allocated; more generally, we can say that in making allocation decisions, many participants did not fully internalize the heterogeneity in the clients’ different resource requirements. Intuitively, this might be expected, since the number of jobs allocated is a more “natural” measure of fairness than the proportion of different resources allocated. However, this observation, if validated in a larger survey, can provide useful guidelines for datacenter operators in that they need to educate their clients about the externality imposed on others by each client’s unique heterogeneous resource requirements.

## VII. FUTURE WORK

Initial exploration suggests that both FDS and GFJ can be unified into a single framework. The idea is to use a  $p$ -norm function  $g(\gamma_{1,j}, \dots, \gamma_{n,j}) = (\sum_i \gamma_{i,j}^p)^{\frac{1}{p}}$  to scalarize the resource requirement vector of user  $j$ , and then evaluate the resulting fairness by  $f_{\beta,\lambda}$ . This method leads to a new family of fairness measures, parameterized  $p$ ,  $\beta$ , and  $\lambda$ , i.e.,

$$f_{p,\beta,\lambda} = \text{sgn}(1 - \beta) \left( \sum_{j=1}^m \left( \sum_{k=1}^n R_{kj}^p \right)^{\frac{1-\beta}{p}} x_j^{1-\beta} \right)^{\frac{1}{\beta}} \times \left( \sum_{j=1}^m \left( \sum_{k=1}^n R_{kj}^p \right)^{\frac{1}{p}} x_j \right)^{\lambda+1-\frac{1}{\beta}}. \quad (19)$$

Fairness  $f_{p,\beta,\lambda}$  includes many fairness measures as special cases. For instance,  $f_{0,\beta,\lambda} = f_{\beta,\lambda}^{GFJ}$  and  $f_{\infty,\beta,\lambda} = f_{\beta,\lambda}^{FDS}$ , while  $f_{1,\beta,\lambda}$  gives the total resource usage in the system.

This function again satisfies the four axioms of [4], as do FDS and GFJ. Moreover, Pareto-efficiency is satisfied for  $|\lambda| \geq \left| \frac{1-\beta}{\beta} \right|$ ,  $\beta > 0$ . We expect that, in analogy with Props. 4 and 5 and their corollaries, threshold values of  $p$  and  $\beta$  can be



found, above which sharing incentive and envy-freeness are satisfied if  $\beta > 0$  and  $\lambda = \frac{1-\beta}{\beta}$ .

In addition to the functional unification proposed in (19), a number of extensions to the current framework are possible. First, we have assumed that both resources and jobs are infinitely divisible. However, in practice a job may require a minimum, indivisible bundle of resources, e.g., 2 GB of memory and 1 CPUs, to run one instance of the job, whereas allocating 1 GB of memory and 1/2 CPUs offers no more benefit than allocating nothing at all. Second, our fairness measures are assumed to be irrelevant to the feasible region of resources. Adding a feasible region and indivisible resources would lead to a fairness version of the knapsack problem, which has no known solution. Some approaches to the knapsack problem are summarized in Appendix D of [3].

Another interesting direction to explore is to extend our multi-resource fairness theory to account for job deadlines, scheduling, and user utility from allocated resources. Finally, our fairness analysis is based on a model of static jobs whose resource demands follow a constant pattern. Many applications not only have time elasticity of demand, but also allow jobs to dynamically change the composition of a bundle of different types of resources. These are all challenging problems that can be explored as future work.

### VIII. CONCLUDING REMARKS

In this paper, we introduce FDS and GFJ, two families of fairness functions for multi-resource allocations. FDS also includes as a special case the recently-proposed generalization of the max-min fairness measure for multiple resources. Different parameterizations of these functions generate a range of fairness-efficiency tradeoffs, thus allowing for different degrees of emphasis on fairness and efficiency that suit different network operation needs.

We consider three key properties of fairness functions: Pareto-efficiency, sharing incentive, and envy-freeness. FDS and GFJ are both Pareto-efficient if  $|\lambda| \geq \frac{1-\beta}{\beta}$ ,  $\beta > 0$ . FDS satisfies the sharing incentive property and is envy-free for  $\beta > 1$  and  $\lambda = \frac{1-\beta}{\beta}$ ; if  $0 < \beta < 1$  and  $\lambda = \frac{1-\beta}{\beta}$ , then sharing incentive and envy-freeness are only sometimes satisfied. GFJ may or may not be sharing-incentive compatible or envy-free for any  $\beta > 0$ ,  $\lambda = \frac{1-\beta}{\beta}$ .

We also explore the estimation of the  $\beta$  and  $\lambda$  values which correspond to people's preferences. Preliminary results along these lines are given in Section VI, though one can easily imagine extensions of both the results analysis and the questions asked to participants. Given the limited set of allocations ranked by the participants, reverse-engineering *unique*  $(\beta, \lambda)$  values compatible with each response was not feasible, but it would be interesting to determine if such unique parameters exist given the rankings of more allocations. Moreover, our current sample size consists primarily of students and others in the academic community who are familiar with computers; with a more diverse demographic of participants, we could examine the impact of various demographic factors on participants' responses. In particular, we could investigate whether participants naturally group themselves into more

than two clusters, and whether these have any demographic correlations.

### ACKNOWLEDGMENT

The authors wish to thank Augustin Chaintreau and Chee-Wei Tan for their comments and assistance with recruiting survey participants.

### REFERENCES

- [1] R. Jain, D. M. Chiu, and W. R. Hawe, *A quantitative measure of fairness and discrimination for resource allocation in shared computer system*. Eastern Research Laboratory, Digital Equipment Corp., 1984.
- [2] K. J. Arrow, "The theory of risk aversion," *Essays in the theory of risk-bearing*, pp. 90–120, 1971.
- [3] C. Joe-Wong, S. Sen, T. Lan, and M. Chiang, "Multi-resource allocation: Fairness-efficiency tradeoffs in a unifying framework," Princeton University, [http://www.princeton.edu/~cjoe/Fairness\\_ToN.pdf](http://www.princeton.edu/~cjoe/Fairness_ToN.pdf), Tech. Rep., 2012.
- [4] T. Lan, D. Kao, M. Chiang, and A. Sabharwal, "An axiomatic theory of fairness in network resource allocation," in *Proceedings of IEEE INFOCOM*. IEEE, 2010, pp. 1–9.
- [5] H. Varian, "Equity, envy, and efficiency," *Journal of Economic Theory*, vol. 9, no. 1, pp. 63–91, 1974.
- [6] A. Odlyzko, "Network neutrality, search neutrality, and the never-ending conflict between efficiency and fairness in markets," in *Review of Network Economics*, vol. 8, no. 1, March 2009, pp. 40–60.
- [7] M. Zukerman, L. Tan, H. Wang, and I. Ouyeyi, "Efficiency-fairness tradeoff in telecommunication networks," in *IEEE Communications Letters*. IEEE, 2005, pp. 643–645.
- [8] A. Ghodsi, M. Zaharia, B. Hindman, A. Konwinski, S. Shenker, and I. Stoica, "Dominant resource fairness: Fair allocation of multiple resource types," in *Proceedings of the 8th USENIX conference on networked systems design and implementation*. USENIX Association, 2011, pp. 24–37.
- [9] S. K. Baruah, N. K. Cohen, C. G. Plaxton, and D. A. Varvel, "Proportionate progress: A notion of fairness in resource allocation," *Algorithmica*, vol. 15, no. 6, pp. 600–625, 1996.
- [10] J. F. Nash, "The bargaining problem," *Econometrica*, vol. 18, no. 2, pp. 155–162, 1950. [Online]. Available: <http://www.jstor.org/stable/1907266>
- [11] R. Mazumdar, L. G. Mason, and C. Douligeris, "Fairness in network optimal flow control: Optimality of product forms," *IEEE Transactions on Communications*, vol. 39, no. 5, pp. 775–782, 1991.
- [12] A. Tang, J. Wang, and S. H. Low, "Counter-intuitive throughput behaviors in networks under end-to-end control," *IEEE/ACM Transactions on Networking*, vol. 14, no. 2, pp. 355–368, 2006.
- [13] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan, "Rate control for communication networks: Shadow prices, proportional fairness and stability," *The Journal of the Operational Research Society*, vol. 49, no. 3, pp. 237–252, 1998.
- [14] R. Srinivasan and A. K. Somani, "On achieving fairness and efficiency in high-speed shared medium access," *IEEE/ACM Transactions on Networking*, vol. 11, no. 1, pp. 111–124, 2003.
- [15] M. Dianati, X. Shen, and S. Naik, "A new fairness index for radio resource allocation in wireless networks," in *Proceedings of the 2005 IEEE Wireless Communications and Networking Conference*, vol. 2. IEEE, 2005, pp. 712–717.
- [16] M. Zaharia, D. Borthakur, J. Sen Sarma, K. Elmeleegy, S. Shenker, and I. Stoica, "Delay scheduling: A simple technique for achieving locality and fairness in cluster scheduling," in *Proceedings of the 5th European Conference on Computer Systems*. ACM, 2010, pp. 265–278.
- [17] Y. Yang, "Rumr: Robust scheduling for divisible workloads," in *Proceedings of the 12th IEEE Symposium on High Performance and Distributed Computing*. IEEE, 2003.
- [18] M. Drozdowski. (2011, Jul.) Introduction to divisible tasks. [Online]. Available: [http://www.cs.put.poznan.pl/mdrozdowski/divisible/divisible\\_intro/divisible\\_intro.html](http://www.cs.put.poznan.pl/mdrozdowski/divisible/divisible_intro/divisible_intro.html)
- [19] T. Bonald and L. Massoulié, "Impact of fairness on Internet performance," in *ACM SIGMETRICS Performance Evaluation Review*, vol. 29, no. 1. ACM, 2001, pp. 82–91.
- [20] L. Massoulié and J. Roberts, "Bandwidth sharing: Objectives and algorithms," *IEEE/ACM Transactions on Networking*, vol. 10, no. 3, pp. 320–328, 2002.
- [21] P. Green and V. Srinivasan, "Conjoint analysis in consumer research: Issues and outlook," *Journal of Consumer Research*, pp. 103–123, 1978.