



HAL
open science

Cloud Flat Rates Enabled via Fair Multi-resource Consumption

Patrick Poullie, Burkhard Stiller

► **To cite this version:**

Patrick Poullie, Burkhard Stiller. Cloud Flat Rates Enabled via Fair Multi-resource Consumption. 10th IFIP International Conference on Autonomous Infrastructure, Management and Security (AIMS), Jun 2016, Munich, Germany. pp.30-44, 10.1007/978-3-319-39814-3_3. hal-01632743

HAL Id: hal-01632743

<https://inria.hal.science/hal-01632743>

Submitted on 10 Nov 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Cloud Flat Rates

Enabled via Fair Multi-resource Consumption

Patrick Poullie, Burkhard Stiller

Communication Systems Group (CSG),
Department of Informatics (IfI), University of Zürich,
Binzmühlestrasse 14, CH-8050 Zürich, Switzerland
{poullie,stiller}@ifi.uzh.ch

Abstract. Many companies rent Virtual Machines (VM) from cloud providers to meet their computational needs. While this option is also available to end-users, they do not always take advantage of this option. One reason may be that it is common to pay on a per-VM-basis, whereas the telecommunications sector has shown that customers prefer flat rates. A flat rate for cloud services needs to define utilization thresholds, to cap the usage of heavy customers and thereby limit their impact on the flat rate price and the cloud performance. Unfortunately, customers consume multiple heterogenous resources in clouds, *e.g.*, CPU, RAM, disk I/O and space, or network access. This makes the definition of a customer’s fair “cloud share” and according utilization thresholds complex. Backed by a questionnaire among more than 600 individuals, this paper designs the new Greediness Metric (GM) that formalizes an intuitive understanding of multi-resource fairness without access to consumers’ utility functions. This GM enables the introduction of attractive cloud flat rates and fair sharing policies for private/commodity clouds and provides incentive to customers to wisely determine VM configurations.

1 Introduction

Cloud Computing (CC) is a computing paradigm enabled by the growing connectivity provided by modern communication systems combined with virtualization technology [8,27]. CC allows server farms to provide their combined computing power on demand to customers, such as end-users and companies. To process a workload through CC a customer starts Virtual Machines (VM) in the cloud that process the workload. In particular, VMs are defined by Virtual Resources, *e.g.*, virtual CPU and virtual RAM [3,8]. Resources in private/commodity clouds are often managed by quotas, *i.e.*, each user has a quota defining a maximum of VRs that his VMs may have in total.

As opposed to this, in commercial clouds, it is common practice that customers pay on a per-VM-basis [2]. However, the telecommunications sector has shown that customers often prefer flat rates [1,21,22], even if a volume-based tariff would reduce costs [18]. Contrary to a volume-based tariff, a flat rate induces a fixed cost cap. For private users this allows for care-free use and for commercial users an easy budgeting of costs. This suggests high demand for cloud flat

rate schemes. With a cloud flat rate customers pay a monthly fee to get access to a cloud, where they can start VMs. Just as Internet flat rates come with a maximum bandwidth, cloud flat rate customers would get a certain quota to spawn VMs.

However, this mechanism is neither sufficient to cap costs that individual customers cause nor to ensure fairness between customers, because customers (i) often deploy different amounts of their quota and (ii) load VMs differently. In particular, (i) means that some customers may only operate a few VMs leaving most of their quota unused, while others fully utilize their quota. (ii) describes that even if customers create exactly the same number and types of VMs, the costs they cause varies depending on how they utilize their VMs. Because VMs on the same host compete for resources, a heavily loaded VM may impair the performance of other VMs on the same host. Therefore, in flat rate or private/commodity clouds, it is desirable to limit VMs of heavy customers in favor of VMs of more moderate customers, such that VMs of moderate customers are not impaired by VMs of heavy customers.

While technical means to enforce fairness in this way exist [3,10,12], cloud fairness is neither sufficiently enforced nor explored [7,8,9,15,28]. What makes the definition of cloud fairness problematic, is that multi-resource fairness has to be defined without access to utility functions. Here, *multi-resource* implies that *bundles*, which are allocated to the consumers, consist of heterogenous resources and the non-accessibility of utility functions implies that it is unknown how much consumers value different bundles (a consumer's *utility function* maps each bundle to a number quantifying the consumer's valuation for the bundle). Utility functions are unknown, because, depending on the workload a VM executes, dependencies between resources differ. Therefore, even for the owner of a VM, the VM's utility function is difficult to determine. Furthermore, no standardized format to describe utility functions exists. Lastly, even if cloud users could determine and express utility functions, they may not want to reveal it to the cloud operator, as it reveals private information about internal processes.

The problem of cloud fairness is often addressed by VM scheduling, *i.e.*, deciding which VM should be started next. This is insufficient to streamline resource utilization in clouds, because it makes static assumptions on utility functions necessary [4,7,9], while VM demands are dynamic. To reach a general applicability (including VM runtime), this paper defines the Greediness Metric (GM) in return to the research question how to quantify the commensurability of VMs' runtime resource utilization, *i.e.*, how fair VMs behave. Because fairness is an intuitive concept, *i.e.*, differs from person to person, the conformance of the GM with an intuitive understanding of fairness is verified by a questionnaire among more than 600 participants.

The remainder of this paper is structured as follows: Section 2 discusses related work leading to the research question stated in Section 2.1. Section 3 outlines the questionnaire, presents its results, and discusses key findings. Based on these outcomes, the new Greediness Metric (GM) is defined in Section 4. Section 5 draws conclusions and outlines future work.

2 Related Work and Problem

For a single resource, the size of bundles can be quantified, and thus, the value of bundles objectively compared. Therefore, fairness can be intuitively defined as Max-Min-Fairness [25] or Proportional Fairness [4] and quantified by metrics such as [14,19,30]. While such allocation problems were extensively studied in computer science [6,25], multi-resource allocation received much less attention.

In data centers, which comprise clusters, grids, and clouds, consumers share resources, such as CPU time, RAM, Disk I/O and space, and network access, wherefore it is necessary to define fairness, when every consumer receives a bundle of heterogenous resources. As noted in [7,9,15,17,23,28], multi-resource allocation in data centers is so far not fully investigated and often reduced to single-resource allocation problems at the cost of efficiency and fairness [9,20]. The assumption of more advanced approaches is that the resource utilization of jobs, which determine those entities consuming resources in clusters and grids, is static or at least follows static ratios [7,9]. This allows to introduce fairness by job scheduling, *i.e.*, which job should be started next. However, VMs, which are the entities consuming resources in clouds, change their resource utilization frequently and dynamically. This not only prohibits reducing cloud fairness to VM scheduling but also prohibits incorporating the concept of envy-freeness [9,11,28]. Envy-freeness, *i.e.*, no consumer prefers to swap his bundle with another, is essential for the definition of fairness in economics. However, this definition is not applicable, when consumer’s utility functions are unknown or highly dynamic, as it is the case for clouds. Subsequently, approaches to multi-resource fairness in data centers (especially their cloud instances) are compared here, while ignoring numerous approaches focusing on a single resource.

Dominant Resource Fairness (DRF) is the most prominent approach to introduce multi-resource fairness in data centers [9]. DRF defines the value of a bundle as the biggest proportion relative to the total supply of any resource in it (cf. Section 3.2). Therefore, to define the value of a bundle only one resource is considered, which is also known as the L_∞ norm. A DRF fair allocation is the allocation that maximizes this value for every consumer. [11] points out that for many other functions (including all other $L_{i \in \mathbb{N}}$ norms) a unique allocation exists that can be found in polynomial time, but that the authors of DRF never argue why their choice is superior. Leontief utility functions model that resources are required in static ratios, *i.e.*, increasing the amount a consumer receives of one

Table 1: Comparison of related approaches.

App.	Fairness	Area	Utility Function
DRF	Max-min for L_∞ norm	Scheduling	Leontief
[4]	Proportional Fairness	Scheduling	Leontief
BBF	Equal share on a bottleneck	Scheduling	Perfectly complementary
[28]	Envy free and Pareto efficient	Micro resources	Cobb-Douglas
[17]	Priority based on metrics	Scheduling	Not needed
GM	Based on questionnaire	Runtime	Not needed

resource does not increase his utility if his share of all other resources is not increased by the same factor. While all proofs for DRF’s desirable properties are based on Leontief utility functions [9,23], the actual DRF scheduling policy, as proposed in [9], allows consumers to request different resource bundles. Therefore, in the scheduling process, consumers can have arbitrary utility functions, but DRF’s properties are only proven for Leontief utility functions. [4] shows that proportional fairness is superior to DRF in terms of efficiency and comes with the same desirable characteristics under realistic assumptions.

Bottleneck-based fairness (BBF) was introduced in [7,8]. An allocation is bottleneck-based fair, if every consumer either is allocated all requested resources or at least the equal share on a congested/bottleneck resource and the other resource in proportion. [29] defines a multi-resource on-line scheduling policy that achieves BBF without knowing consumer’s utility functions in advance.

[28] presents an allocation policy that achieves game-theoretic fairness, *i.e.*, sharing incentive, envy freeness, and pareto-efficiency, when allocating cache capacity and memory bandwidth. Different applications are profiled to convincingly argue that Cobb-douglas utility functions are well suited to model diminishing returns and substitution effects for these resources.

[17] considers the fair sharing of grids between customers and research groups. Contrary to all other approaches discussed in this section, [17] achieves fairness by a penalty/priority function, which is close to our approach. In particular, such function determines the priority of consumers when granting resource requests. Such approach is much more practically oriented, because no assumptions about utility functions need to be made.

Another research field on cloud resource allocation focuses on live-migration [16,26], *i.e.*, rescheduling VMs during runtime. While live-migration is applied during VM runtime it is, just as VM scheduling, orthogonal to this work here, because it does not change priorities of VMs.

Table 1 compares the approaches discussed in terms of the adopted fairness definition, the application area, and the assumed utility function. Most approaches are way too complex to be applied during VM runtime and none of these is based on an intuitive understanding of fairness. This motivates the general research undertaken and specifically this paper’s problem statement.

2.1 Problem Statement

The characteristic of fair cloud resource allocations that is distinct for this context is that customers (between whom fairness is to be defined) utilize resources from different resource pools (hosts) by intermediaries (VMs). While fairness has to be achieved between customers, this has to be done by allocating resources to their VMs. Unfortunately resources can only be moved between VMs, which run on the same host. However, besides this structural dilemma, also a more general problem is faced: Defining and enforcing fairness of multi-resource allocations without knowing consumers’ utility functions. In particular, because bundles consist of multiple heterogenous resources, bundles can contain resources in different amounts. This prohibits an objective comparison of bundles. For example,

some customers may require more CPU for their workloads, while others require more RAM [20]. A third customer may deploy the cloud for backups and, therefore, mostly requires disk-space and bandwidth. Therefore, consumers can have different preferences over the same bundles.

Because utility functions of consumers in clouds are unknown, it is not possible to define fairness via utility functions here and, thus, fairness has to be defined via bundles that VMs serve themselves. In particular, hosts work as “self-serving stores” for VMs, which means that they provide all requested resources to VMs, if possible. Therefore, fairness in clouds has to be defined as constraining VMs of those customers whose VMs overcharge their self-serving stores, *i.e.*, are greedy. While also the concept of greediness has no formal definition, it can be better defined and quantified with the information that is available in clouds.

Therefore, the ***problem statement*** for this paper reads as follows: The greediness of VMs can be defined and quantified based on their multi-resource self-servings. An allocation is fair in such a case, when (i) the aggregates per customer of these quantifications are aligned and (ii) VMs of “greedy” customers are constrained in favor of VMs of “less greedy” customers.

3 Questionnaire

A questionnaire was developed to evaluate the intuitive understanding of fairness and greediness and to justify empirically the design of the new Greediness Metric (GM). The questionnaire can be found in the appendix of [24]. The questionnaire specified real-life resource allocation scenarios in terms of three questions Q1, Q2, and Q3 to not distract participants by technical terms and let them fully concentrate on the question of fairness. While these scenarios were specific, the questions were carefully designed to reach generic insights about intuitive understandings of fairness and greediness. To describe these questions subsequently, resources and consumers are denoted by r_i and c_j , respectively, where $i, j \in \mathbb{N}_{>0}$. Questionnaire participants had to choose between different options of allocations or define rankings of consumers. Additionally, participants were offered to explain their answers in text boxes. The questionnaire did not address any particular target group, to allow evaluating a popular and intuitive understanding of fairness that is not biased by technical notions established by experts. Out of 721 participants, who started the questionnaire, 604 completed it. Q2 addressed the question of how consumer requests should be taken into account, when allocating resources, but did not address the question of how fairness or greediness can be defined. Therefore, Q2 is not directly related to the design of the GM and only discussed in [24].

3.1 Choosing the Most Fair Allocation (Question 1)

DRF is the most prominent approach for fairness in data centers (cf. Section 2). Therefore, DRF’s conformance with an intuitive understanding of fairness was evaluated by Q1. In particular, DRF’s centerpiece, which is using the L_∞ norm

Table 2: The four Allocation Options A11, A12, A13, and A14.

Consumer	A11		A12		A13		A14	
	r_1	r_2	r_1	r_2	r_1	r_2	r_1	r_2
c_1	2	0	3	0	4	0	5	0
c_2	0	4	0	6	0	8	0	10
c_3	4	8	3	6	2	4	1	2

Table 3: The three Scenarios S31, S32, and S33.

Consumer	S31			S32			S33		
	r_1	r_2	r_3	r_1	r_2	r_3	r_1	r_2	r_3
c_1	4	3	3	4	2	4	4	1	4
c_2	2	1	5	1	4	3	1	4	3
c_3	4	2	1	1	6	2	1	6	2
Remainder	2	3	0	6	0	0	6	1	0

to measure the value of a bundle (*c.f.* Section 2 and Section 3.2), was evaluated. The scenario described covers two resources r_1 and r_2 of which six and twelve units were available, respectively. These resources have to be allocated to three consumers c_1 , c_2 , and c_3 . c_1 only requires r_1 , c_2 only r_2 , and c_3 requires for each unit of r_1 two units of r_2 . This results in seven possible allocations to allocate all resources and do not give consumers resources they have no use for. However, most of these allocations are intuitively unfair, *e.g.*, in two of these allocations at least one consumer receives no resources at all.

Because the scenario describes that resources are requested in static ratios, it is transferable to allocation problems in data centers, where these static ratios of resource requests, *i.e.*, Leontief utility functions, are the standard assumption [4,7,9]. Table 2 shows four of the seven allocations and their respective labels. These four allocations were presented to participants numerically and graphically and they had to choose *the* allocation that seemed most fair to them. As expected, A11 and A14 were only chosen by a minority of the 721 participants (0.4% and 1.1%, respectively) and most participants deterred between A12 and A13 (30.0% and 68.5%, respectively). The following arguments in support of A12 are summarized from textual comments received:

- c_1 and c_2 only compete with c_3 for resources, but not with one another. A fair allocation splits resources equally between those who contend for them.
- All receive an equal amount of what they want.
- This is the only allocation where nobody can complain that someone has more of the same resource.

For the following reasons A13 was supported:

- When prices are introduced based on available units, this option gives the same value to all consumers.
- c_1 and c_2 receive $2/3$ of *one* resource and c_3 $1/3$ of *two* resources, which makes $2/3$ for everybody. On a similar note, some participants rejected A12, because c_3 gets as much as c_1 and c_2 combined.
- Because c_1 and c_2 only want one resource, they should get more of it than c_3 , as c_3 wants both resources.
- This option is the result of a simple auction or when all consumers get an equal share of both resources and then trade.

3.2 Estimating Greediness (Question 3)

Q3 was designed to collect information on how the greediness of consumers, who served themselves from a pool of common resources, is perceived. The transferability to clouds is evident, because in clouds VMs serve themselves from their host. Moreover, because in flat rate and commodity clouds no payment on a per VM basis occurs, such clouds are also a common resource shared among customers. In addition, insights were collected on how proportionality and value of resource bundles is perceived, when no information about consumers' utility functions is available. Thus, Q3 provides insights on how resources that different VMs on the same host utilize can be compared.

Q3 is based on three scenarios S31, S32, and S33, where three consumers c_1 , c_2 , and c_3 had served themselves from a pool of three common resources r_1 , r_2 , and r_3 (like VMs on the same host). To split these resources, each consumer had allocated himself a certain bundle as shown in Table 3. The three consumers had to be ranked according to how their greediness was perceived, all being based on the amounts the consumers had allocated themselves.

Metrics: Many participants tackled Q3 by proposing one of the four metrics discussed subsequently.

Price: The price metric is the simplest metric. The value of one unit of resource r_i is defined as p/\overleftarrow{r}_i , where p is a constant and \overleftarrow{r}_i is the number of units available of r_i . The value of a bundle is the sum of values of its resources. For example, for $p = 1$ the value of c_2 's bundle in S31 is $\frac{2}{12} + \frac{1}{9} + \frac{5}{9} = \frac{5}{6}$. This metric is equivalent to the sum-based-penalty function presented in [17].

P×S (Price×Scarcity): The P×S metric is a natural extension of the price metric. The value of one unit of resource r_i is defined as $a(r_i) \cdot p/\overleftarrow{r}_i^2$, where $a(r_i)$ is the amount that is allocated in total of r_i . The value of a bundle is defined as the sum of values of those resources contained. For example, for $p = 1$ the value of c_2 's bundle in S31 is $\frac{2 \cdot 10}{12^2} + \frac{1 \cdot 6}{9^2} + \frac{5 \cdot 9}{9^2} = \frac{29}{54}$.

P∩S (Price∩Scarcity): The P∩S metric is another natural extension of the price metric and defines the value of a resource just as the price metric. However, the value of a bundle is defined only over resources that are depleted, *i.e.*, resources where $a(r_i) = \overleftarrow{r}_i$. For example, for $p = 1$ the value of c_2 's bundle in S31 is $\lfloor \frac{10}{12} \rfloor \cdot \frac{2}{12} + \lfloor \frac{6}{9} \rfloor \cdot \frac{1}{9} + \lfloor \frac{9}{9} \rfloor \cdot \frac{5}{9} = \frac{5}{9}$.

DRF (Dominant Resource Fairness): The DRF metric defines the value of a bundle by the L_∞ norm, *i.e.*, by the biggest share of any resource relative to the overall amount of the resource. For example, the value of c_2 's bundle in S31 is $\max(\frac{2}{12}, \frac{1}{9}, \frac{5}{9}) = \frac{5}{9}$. According to this DRF metric, the bundles of c_1 and c_2 in S33 are equally valuable. This tie is broken by the second biggest share in the bundles, wherefore c_1 's bundle is more valuable.

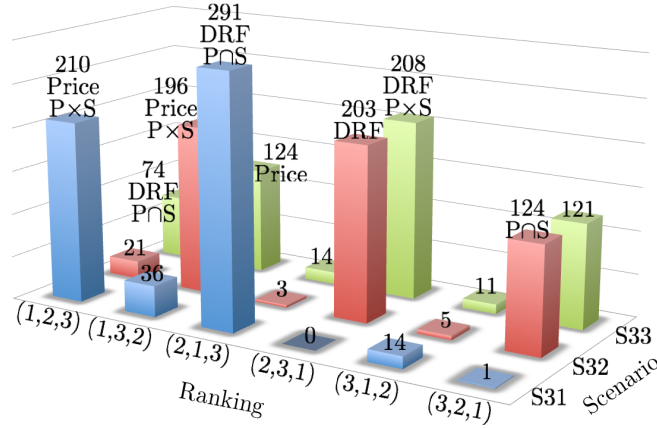


Fig. 1: Ranking Frequencies (represented by triplets) in Q3 scenarios.

Table 4: Percentages of most frequent rankings in Q3.

Metric	S31	S32	S33
DRF	(2,1,3): 52.7%	(2,3,1): 36.8%	(2,3,1): 37.7%
Price	(1,2,3): 38.0%	(1,3,2): 35.5%	(1,3,2): 22.5%
$P \times S$	(1,2,3): 38.0%	(1,3,2): 35.5%	(2,3,1): 37.7%
$P \cap S$	(2,1,3): 52.7%	(3,2,1): 22.5%	(1,2,3): 13.4%

Table 5: A problematic DRF ranking.

	r_1	r_2	r_3	Price	$P \times S$	$P \cap S$	DRF
available	30	30	30				
c_1	18	0	0	1.80	1.68	0.00	0.60
c_2	0	14	17	3.10	3.04	1.40	0.56
c_3	10	16	12	3.80	3.69	1.60	0.53

Frequency Investigations: For some questionnaire rankings the free text indicated that the participant had assumed real life values of those resources and ranked the consumers accordingly. These rankings as well as incomplete rankings were removed, wherefore the presented results are based on 553 answers. Subsequently, consumer rankings are denoted by triplets. For example, the triplet (2,1,3) denotes the first consumer as moderate, the second consumer as most greedy, and the third consumer as least greedy.

Figure 1 illustrates for each scenario how many participants selected each ranking and highlights those rankings that correspond to all metrics discussed in Section 3.2. These numbers in Figure 1 allow to compile Table 4 showing for each metric the respective ranking in the three scenarios and by how many participants this ranking was selected. Figure 1 and Table 4 show that the metrics discussed in Section 3.2 cover the majority of participants' rankings.

[24] analyses the combinations of rankings over the three scenarios, because these rankings given by most participants did not match the same metric over the three scenarios. This reveals that the most frequent combination was selected by 79 participants, conforming to the DRF metric. The second, third, and fourth

most frequent combinations were selected by 55, 43, and 32 participants, conforming to the $P \times S$, Price, and $P \cap S$ metric, respectively.

3.3 Discussion

Only the first two arguments of the questionnaire’s text replies in favor of A12 of the above presented Q1 are correct. The third argument is incorrect, because c_1 receives least of r_2 and c_2 least of r_1 . Therefore, the arguments in favor of A13 are more versatile and sound than those in favor of A12. Because A12 and A13 were chosen by 30.0% and 68.5% of the participants, respectively, it is concluded that A13 determines the intuitively fair allocation.

In Q3, participants who ranked in conformance to the Price and $P \cap S$ metrics, stated that they had applied these metrics to arrive at the ranking. In contrast, participants who ranked conforming to the DRF and $P \times S$ metrics often argued not in conformance to the respective metric. In particular, for participants ranking in conformance to the DRF metric, only one participant argued in conformance to the DRF metric, while the majority argued that those consumers who exceed their equal share are greedy. Thus, in S31 c_2 is the greediest due to the disproportional consumption of r_3 . Analog, in S32 and S33 c_3 is the greediest, because c_3 exceeds the equal share of r_2 by 50% and c_1 is the second greedy, because c_1 exceeds the equal share of r_1 by 33%. Participants who ranked conforming to the $P \times S$ metric mostly argued that the total amount of resources consumed is deceive (the tie of c_1 and c_3 in S33 according to this logic was broken by c_3 ’s 50% overconsumption of r_2). Further details can be found in [24].

Implications for existing metrics: While the metrics discussed in Section 3.2 cover the majority of participants rankings (cf. Table 4 and [24]), none of these metrics captures an intuitive understanding of fairness: The $P \cap S$ metric has a low conformance in S32 and S33 and for S33 results in the inverted ranking of S32. The latter implies that consumers can decrease their score by consuming more. This is not only counter-intuitive, but gives the undesirable incentive to consume more than needed. The Price metric has a low conformance in S33 and identifies c_1 as “greediest” in S31, although c_1 does not cause the bottleneck, but precisely sticks to his equal share (a behavior that is considered humble in S32 and S33). The sum- and the root-based-penalty functions, which [17] identifies as best metrics, result in the same rankings as the Price metric and, therefore, are also not satisfactory. Similar arguments hold for the $P \times S$ metric.

The DRF metric is satisfactory at a first glance: For all three scenarios of Q3 it results in the most frequent ranking and also in the most frequent combination of rankings (cf. Table 4 and [24]). However, only one participant argued according to DRF, while the majority argued that those who exceed their equal share are greedy. Therefore, the high conformance of DRF stems rather from the fact that every consumer exceeds the equal share on at most one resource (because DRF only considers these resources, this allows DRF to produce good results). Table 5 shows an allocation that proves that DRF’s approach to ignore all but one

resource can lead to undesirable results. For the depicted allocation, all other metrics discussed in Section 3.2 give the inverse ranking of DRF. Also, according to the arguments made by the participants, the DRF ranking is unfair: DRF classifies c_1 as the consumer with the most valuable bundle, although c_1 only receives the least loaded resource. c_3 cedes no resource at all and receives most of the only scarce resource, but DRF classifies c_3 as most humble. Also, Q1 of the questionnaire identified A13 as the intuitively fair allocation, while the DRF-fair allocation is A12. Because Leontief utility functions are assumed in Q1, *i.e.*, resources are required in static ratios, and DRF is defined based on this assumption, DRF should result in an intuitively fair allocation. Therefore, while DRF is often applied, when Leontief utility functions do not hold, Q1 shows that already for Leontief utility functions, DRF may result in allocations that are not intuitively fair. Moreover, consumers and resources can be added to this scenario, where consumer c_i requests only resource r_i (and as before one consumer requests all resources evenly). Thereby, the perceived unfairness of DRF can be increased arbitrarily, because the consumer requesting all resources receives as much as all other consumers combined.

Due to these shortcomings of all metrics presented so far, the new Greediness Metric (GM) was developed. This GM is aligned with arguments of this discussion and, thus, (i) classifies A13 in Q1 as the most fair allocation, (ii) results in the most frequent ranking for each of the three scenarios in Q3, (iii) gives the “correct” ranking for the allocation in Table 5, and (iv) captures an intuitive understanding of fairness, which allows defining attractive cloud flat rate models.

4 Greediness Metric

The Greediness Metric (GM) maps each resource bundle in an allocation to a rational number that can be associated to the *greediness* of the consumer, who served himself the bundle. In that sense this GM serves the same purpose as the other metrics presented in Section 3.2, which quantify the value of a bundle. As the most frequent textual comments of the questionnaire suggest, the GM sums up, what exceeds the equal share in each bundle. However, it also deducts what is not consumed of the equal share but handed over to other consumers instead.

Let $R = (r_1, r_2, \dots, r_m)$ be a set of m resources, where resource $r_i \in R$ is available in the amount of \overleftarrow{r}_i . An *allocation* of R to n consumers (c_1, c_2, \dots, c_n) can be denoted by a matrix $A \in \mathbb{R}_{\geq 0}^{m \times n}$ with $\sum_{j=1}^n a_{ij} \leq \overleftarrow{r}_i$, for all $i \in \{1, 2, \dots, m\}$, where c_j receives amount a_{ij} of r_i . The amount of r_i that c_j receives beyond his equal share is then $a_{ij} - \overleftarrow{r}_i/n$ (if the difference is negative, c_j does not utilize its entire equal share of the resource).

If $a_{ij} > \overleftarrow{r}_i/n$, consumers other than c_j have to cede some of their equal share of r_i in order to enable c_j 's share of r_i . Therefore, the amount by which c_j exceeds his equal share is added to the greediness of c_j .

If $a_{ij} = \overleftarrow{r}_i/n$, c_j exactly receives his equal share, wherefore it does not change c_j 's greediness. In particular, if $a_{ij} = \overleftarrow{r}_i/n$ for all $i \in \{1, 2, \dots, m\}$, c_j 's greediness is zero.

Table 6: GM results for Q1 and Q3 of the questionnaire.

γ	Cons.	A12	A13	S31	S32 and S33
	c_1	$0.25-0.5 \cdot x$	$0.5-0.5 \cdot x$	0	$0.1-0.16 \cdot x$
x	c_2	$0.25-0.5 \cdot x$	$0.5-0.5 \cdot x$	0.2	0
	c_3	0.5	0	$-0.2 \cdot x$	$0.16-0.1 \cdot x$
	c_1	-0.25	0	0	-0.05
1	c_2	-0.25	0	0.2	0
	c_3	0.5	0	-0.2	0.05
	c_1	0	0.25	0	0.027
0.5	c_2	0	0.25	0.2	0
	c_3	0.5	0	-0.1	0.1

If $a_{ij} < \overleftarrow{r}_i/n$, c_j 's cession of r_i is credited to c_j , *i.e.*, subtracted from c_j 's greediness, to the extent that other consumers profit from this cession, which is the case, when they utilize r_i beyond their equal share. This extension not only depends on how much of r_i is utilized beyond the equal share by other consumers, but also on how much of r_i is ceded by other consumers. Therefore, the *credit factor* for the cession of r_i is the ratio of what is ceded of r_i to what is consumed beyond the equal share of r_i . To capture this notion formally, the sum of what consumers receive beyond and cede of their equal share of r_i is defined by $\alpha(r_i)$ and $\beta(r_i)$, respectively. Therefore,

$$\alpha(r_i) := \sum_{j=1}^n \max(0, a_{ij} - \overleftarrow{r}_i/n) \quad \text{and} \quad \beta(r_i) := \sum_{j=1}^n \max(0, \overleftarrow{r}_i/n - a_{ij}).$$

Multiplying the amount that c_j cedes of r_i with $\alpha(r_i)/\beta(r_i)$ implements the considerations above. Therefore, the *greediness of c_j* is defined as

$$g(c_j) := \sum_{i=1}^m o(i, j) \cdot n / (m \cdot \overleftarrow{r}_i), \quad (1)$$

where $o(i, j)$ is the *offset* for c_j 's consumption of r_i and defined as

$$o(i, j) := \begin{cases} a_{ij} - \overleftarrow{r}_i/n & \text{if } a_{ij} \geq \overleftarrow{r}_i/n, \\ \gamma \cdot \frac{\alpha(r_i)}{\beta(r_i)} \cdot (a_{ij} - \overleftarrow{r}_i/n) & \text{else.} \end{cases} \quad (2)$$

Note that, if $\beta(r_i) = 0$, no consumer cedes r_i and, therefore, the else-part of Equation 2 is never reached (thus, no division by zero occurs). The factor $n / (m \cdot \overleftarrow{r}_i)$ in Equation 1 normalizes resource units. [24] discusses (i) the choice of this normalization factor in detail, (ii) detailed examples for the calculation of the GM, and (iii) how the GM is applicable to scenarios, where consumers have different endowments (instead of every consumer having an endowment of \overleftarrow{r}_i/n).

The parameter γ defines how strongly the ceding of resources is credited and, thereby, fine-tunes the greediness metric to best comply with the questionnaire results. Table 6 shows the results of the GM dependent on γ for the questionnaire's scenarios. As the table shows, already without the parameter γ

(or, equivalently, $\gamma = 1$), the GM complies with results of Q1 and S31. However, for S32 and S33, the GM results in the ranking (3,2,1) (cf. Table 6), while the ranking most frequently selected by the participants is (2,3,1) (cf. Table 4). This mismatch for $\gamma = 1$ can be explained as follows: In S32 and S33, c_1 exceeds the equal share of r_3 by 33% but also cedes 50% of r_2 to c_3 , while c_2 is “neutral”. Thus, while c_2 ’s greediness is zero independent of γ , c_1 has a negative greediness for $\gamma = 1$, *i.e.*, for $\gamma = 1$, the GM ranks c_2 greedier than c_1 . However, most participants rated c_1 greedier than c_2 , because c_1 over-consumes r_3 while c_2 never exceeds the equal share. Accordingly, $\gamma \in [0, 1]$ regulates how strongly the ceding of resources is credited (in addition to the dynamic regulation by $\frac{\alpha(r_i)}{\beta(r_i)}$), where, the smaller γ is chosen, the harder it gets to compensate for exceeding the equal share. [24] shows that for $\gamma \in]\frac{1}{4}, \frac{2}{3}[$ the GM perfectly complies with the questionnaire results and Table 6 lists that $\gamma = 0.5$ is an appropriate choice.

5 Conclusions and Future Work

Multi-resource fairness for clouds was so far not defined satisfactorily. Therefore, the Greediness Metric (GM) was defined and verified via a questionnaire among more than 600 participants. The GM defines an intuitive understanding of multi-resource fairness without access to utility functions. This questionnaire also revealed that DRF, which is the state-of-the-art in data center fairness, not always conforms with an intuitive understanding of fairness. Because the GM’s definition is based on arguments of non-technical participants, it is intuitively comprehensible. This intuitive definition of cloud applicable fairness, allows for the design of attractive cloud flat rates. In particular, the telecommunications sector has shown that customers often prefer flat rates. Also private/commodity clouds are a perfect use case for the GM, as here no service level agreements guide the resource allocation making fairness an important allocation goal.

In addition to its intuitiveness, the GM is also well suited to be integrated into cloud sharing schemes, because the GM provides the right incentives to chose the configuration of a VM, such that it matches the VM’s subsequent load: When a VM tries to exceed the resources it is configured with, it will receive them, if available. However, the customer’s greediness increases with a potentially negative effect on other VMs of this customer. Because a customer’s greediness only decreases, when resources that his VMs do not utilize are utilized by other VMs, it is not guaranteed that idle resources are credited to the consumer. Therefore, under- as well as over-provisioning VMs is costly. No metric known today (including those in Section 3.2) provides this incentive mechanism.

The GM is currently being evaluated within two different settings: the first setting allows for an evaluation of the GM under idealistic conditions to conclude on the effect of different design aspects, *e.g.*, the function to aggregate the VM greediness to customer greediness. The second setting defines a CloudSim [5] extension, which allows for evaluating the GM under realistic conditions in operation. In turn, the implementation to integrate the GM fairness mechanism into OpenStack [13] is ongoing.

References

1. Jörn Altmann, Björn Rupp, and Pravin Varaiya. Effects of Pricing on Internet User Behavior. *Netnomics*, 3(1):67–84, June 2001.
2. Amazon Web Services, Inc. Amazon EC2 Pricing. <https://aws.amazon.com/ec2/pricing/>, 2016. Online; accessed January 26th, 2016.
3. Andrea Arcangeli, Izik Eidus, and Chris Wright. Increasing Memory Density by Using KSM. *2009 Linux Symposium*, Vol. 1, pp 19–28, Montreal, QC, Canada, July 2009.
4. Thomas Bonald and James Roberts. Enhanced Cluster Computing Performance through Proportional Fairness. *Performance Evaluation*, 79:134–145, April 2014.
5. Rodrigo N. Calheiros, Rajiv Ranjan, Anton Beloglazov, César A. F. De Rose, and Rajkumar Buyya. CloudSim: A Toolkit for Modeling and Simulation of Cloud Computing Environments and Evaluation of Resource Provisioning Algorithms. *Software–Practice & Experience*, 41(1):23–50, January 2011.
6. Abhishek Chandra, Micah Adler, Pawan Goyal, and Prashant Shenoy. Surplus Fair Scheduling: A Proportional-share CPU Scheduling Algorithm for Symmetric Multiprocessors. *4th Conference on Symposium on Operating System Design & Implementation*, OSDI’00, pp 4–4, San Diego, CA, USA, October 2000.
7. Danny Dolev, Dror G. Feitelson, Joseph Y. Halpern, Raz Kupferman, and Nathan Linial. No Justified Complaints: On Fair Sharing of Multiple Resources. *3rd Innovations in Theoretical Computer Science Conference*, ITCS’12, pp 68–75, Cambridge, MA, USA, January 2012.
8. Yoav Etsion, Tal Ben-Nun, and Dror G. Feitelson. A Global Scheduling Framework for Virtualization Environments. *2009 IEEE International Symposium on Parallel Distributed Processing*, IPDPS’09, pp 1–8, Rome, Italy, May 2009.
9. Ali Ghodsi, Matei Zaharia, Benjamin Hindman, Andy Konwinski, Scott Shenker, and Ion Stoica. Dominant Resource Fairness: Fair Allocation of Heterogeneous Resources in Datacenters. Technical Report UCB/EECS-2010-55, EECS Department, University of California, Berkeley, CA, USA, May 2010.
10. Fei Guo. Understanding Memory Resource Management in VMware vSphere 5.0. Performance study, VMware, http://www.vmware.com/files/pdf/mem_mgmt_perf_vsphere5.pdf, Palo Alto, CA, USA, 2011.
11. Avital Gutman and Noam Nisan. Fair Allocation without Trade. *11th International Conference on Autonomous Agents and Multiagent Systems*, Vol. 2 of *AAMAS’12*, pp 719–728, Valencia, Spain, June 2012.
12. IBM. Best Practices for KVM. Technical report, http://www.tdeig.ch/linux/pasche/12_BestPractices_IBM.pdf, Austin, TX, USA, November 2010.
13. Kevin Jackson. *OpenStack Cloud Computing Cookbook*. Packt Publishing, Birmingham, UK, 2012.
14. Rajendra K. Jain, Dah-Ming W. Chiu, and William R. Hawe. A Quantitative Measure of Fairness and Discrimination for Resource Allocation in Shared Computer Systems. Technical Report TR-301, Digital Equipment Corp, Hudson, MA, USA, September 1984.
15. Carlee Joe-Wong, Soumya Sen, Tian Lan, and Mung Chiang. Multi-resource Allocation: Fairness-efficiency Tradeoffs in a Unifying Framework. *31st Annual IEEE International Conference on Computer Communications*, INFOCOM 2012, pp 1206–1214, Orlando, FL, USA, March 2012.
16. Divya Kapil, Emmanuel S. Pilli, and Ramesh C. Joshi. Live Virtual Machine Migration Techniques: Survey and Research Challenges. *IEEE 3rd International*

- Advance Computing Conference*, IACC 2013, pp 963–969, Ghaziabad, UP, India, February 2013.
17. Dalibor Klusáček, Hana Rudová, and Michal Jaroš. Multi Resource Fairness: Problems and Challenges. In: Narayan Desai and Walfredo Cirne (Editors), *Job Scheduling Strategies for Parallel Processing*, Vol. 8429 of *Lecture Notes in Computer Science*, pp 81–95. Springer, Berlin/Heidelberg, Germany, 2014.
 18. Anja Lambrecht and Bernd Skiera. Paying Too Much and Being Happy about It: Existence, Causes, and Consequences of Tariff-Choice Biases. *Journal of Marketing Research*, 43(2):212–223, May 2006.
 19. Tian Lan, David Kao, Mung Chiang, and Ashutosh Sabharwal. An Axiomatic Theory of Fairness in Network Resource Allocation. *29th Annual IEEE International Conference on Computer Communications*, INFOCOM 2010, pp 1–9, San Diego, CA, USA, March 2010.
 20. Gunho Lee, Byung-Gon Chun, and Randy H. Katz. Heterogeneity-aware Resource Allocation and Scheduling in the Cloud. *3rd USENIX Conference on Hot Topics in Cloud Computing*, HotCloud 2011, pp 4–4, Portland, OR, USA, June 2011.
 21. David Levinson and Andrew Odlyzko. Too Expensive to Meter: the Influence of Transaction Costs in Transportation and Communication. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 366(1872):2033–2046, June 2008.
 22. Andrew Odlyzko. Internet Pricing and the History of Communications. *Computer Networks*, 36(5-6):493–517, June 2001.
 23. David C. Parkes, Ariel D. Procaccia, and Nisarg Shah. Beyond Dominant Resource Fairness: Extensions, Limitations, and Indivisibilities. *13th ACM Conference on Electronic Commerce*, EC’12, pp 808–825, Valencia, Spain, June 2012.
 24. Patrick Poullie and Burkhard Stiller. Cloud Flat Rates Enabled via Fair Multi-resource Consumption. Technical Report IFI-2015.03, <https://files.ifi.uzh.ch/CSG/staff/poullie/extern/publications/IFI-2015.03.pdf>, Universität Zürich, Zurich, Switzerland, October 2015.
 25. Sally Floyd, Ed. Metrics for the Evaluation of Congestion Control Mechanisms. RFC 5166, IETF, Berkeley, CA, USA, March 2008.
 26. Anja Strunk. Costs of Virtual Machine Live Migration: A Survey. *2012 IEEE 8th World Congress on Services*, SERVICES 2012, pp 323–329, Honolulu, HI, USA, June 2012.
 27. G. Wei, A. Vasilakos, Y. Zheng, and N. Xiong. A Game-theoretic Method of Fair Resource Allocation for Cloud Computing Services. *The Journal of Supercomputing*, 54(2):252–269, November 2010.
 28. Seyed Majid Zahedi and Benjamin C. Lee. REF: Resource Elasticity Fairness with Sharing Incentives for Multiprocessors. *19th International Conference on Architectural Support for Programming Languages and Operating Systems*, ASPLOS 2014, pp 145–160, Salt Lake City, UT, USA, March 2014.
 29. Yoel Zeldes and Dror G. Feitelson. On-line Fair Allocations Based on Bottlenecks and Global Priorities. *4th ACM/SPEC International Conference on Performance Engineering*, ICPE’13, pp 229–240, Prague, Czech Republic, April 2013.
 30. Moshe Zukerman, Liansheng Tan, Hanwu Wang, and Iradj Ouveysi. Efficiency-fairness Tradeoff in Telecommunications Networks. *IEEE Communications Letters*, 9(7):643–645, July 2005.