The Shared Assignment Game and Applications to Pricing in Cloud Computing

Gideon Blocq*
Technion
gideon@tx.technion.ac.il

Yoram Bachrach
Microsoft Research
yobach@microsoft.com

Peter Key
Microsoft Research
peter.key@microsoft.com

ABSTRACT

We propose an extension to the Assignment Game [37] in which sellers provide indivisible heterogeneous goods to their buyers. Each good takes up various amounts of resources and each seller has capacity constraints with respect to the total amount of resources it can provide. Hence, the total amount of goods that the seller can provide is dependent on the set of buyers. In this model, we first demonstrate that the core is empty and proceed to suggest a fair allocation of the resulting utility of an optimal match, using the Shapley value. We then examine scenarios where the worth and resource demands of each good are private information of selfish buyers and consider ways in which they can manipulate the system. We show that such Shapley value manipulations are bounded in terms of the gain an agent can achieve by using them. Finally, since this model can be of use when considering elastic resource allocation and utility sharing in cloud computing domains, we provide simulation results which show our approach maximizes welfare and, when used as a pricing scheme, can also increase the revenue of the cloud server providers over what is achieved with the widely-used fixed pricing scheme.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multiagent systems; J.4 [**Social and behavioral sciences**]: Economics

Keywords

Assignment Game; Shapley value; Cloud Pricing

1. INTRODUCTION

The Assignment game (AG) [37] is a widely researched model for a two-sided market in which sellers own an indivisible good, with which they supply buyers, in exchange for money. Originally the model only supported single matchings between sellers and buyers for homogeneous goods, but various generalizations to the model have been proposed, which support many-to-many matchings [40] and heterogeneous goods [31]. For each good, the buyer has a *valuation* or *worth*, i.e., a maximum willingness to pay.

We propose an extension of the AG, called the Shared Assignment Game (SAG). In contrast to the original AG, goods, which are provided by the sellers, require various amounts of resources. Moreover, each seller has a capacity constraint with respect to these resources. Thus, the total amount of goods that a seller provides to buyers is dependent on the set of buyers that it is matched with. Motivation for this setting can be found in a cloud environment where clients (buyers) demand resources in order to execute their tasks (goods) on servers (sellers). As an example, consider the scenario depicted in Figure 1. There are two servers, r_1, r_2, each with two resource types, *CPU* and *Memory*. The capacities of the resources equal (*2 CPUs, 1 GB*) for r_1 and (*5 CPUs, 4 GB*) for r_2. Suppose there are two clients b_1, b_2, where b_1 is interested in executing two different tasks g_1 and g_2, which respectively require (*2 CPUs, 2 GB*), (*1 CPU, 2 GB*) to execute. Moreover, client b_2 is only interested in task g_3 which demands (*3 CPUs, 3 GB*) to execute. The welfare created by executing tasks g_1 and g_2 is \$2 and \$3 for g_3.

Two main questions are: **(1) How should we assign tasks to resources?** In this work we focus on an allocation that maximizes the welfare created through the matchings of resources and clients, hence the optimal assignment delivers goods g_1 and g_2, achieving a total welfare of \$4. **(2) What amount should the clients pay for their allocated resources?** In other words, what share of the welfare achieved would we attribute to the clients and resources? With regards to this question, we propose either the *core* or the *Shapley value* [38] as a fair method to distribute the welfare over all agents. We will show that the core does not necessarily exist and we will continue to focus on the Shapley value. The Shapley value determines a share ϕ_i for agent i (where i can be either a resource or a client), reflecting each agent's importance in achieving the total utility, such that the shares sum up to the total utility obtained by all the agents together.

In addition to the above example, other scenarios exist in which the SAG can be used to match agents and fairly distribute welfare between them. For instance, consider companies (sellers) with resource constraints on the type of workers (buyers) they can hire, where each type of worker takes up different amounts of resources. When more workers are matched with sellers, while abiding the resource constraints, more welfare is created.

*This work was done while the author was an intern at Microsoft Research Cambridge.

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*

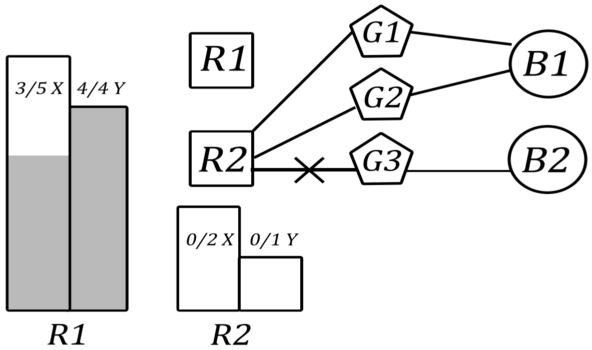

Figure 1: An optimal matching in the SAG.

In a more generalized setting, the above problems represent any two-sided market where welfare is only generated when the resource constraints are satisfied. Each seller can be visualized as solving a 0-1 multi-dimensional knapsack problem, where each item requires a certain amount of capacity of each resource-type.[1] In order to determine the contribution of all agents and fairly distribute the created revenue between the buyers and sellers, we propose a prominent method from *cooperative game theory*, namely the Shapley value, which is the unique utility sharing mechanism fulfilling certain fairness axioms [25, 43].

1.1 Our Contribution

We introduce the Shared Assignment Game, in which sellers have capacity constraints on the different goods that they are able to provide. Moreover, we show that in contrary to the original Assignment Game, the core does not necessarily exist. We continue with a *cooperative game theoretic* approach based on the Shapley value , which brings about a fair distribution of the achieved welfare due to the matching of buyers and sellers. We then discuss the impact of selfish agents holding *private information* regarding the goods' properties, where only the buyer knows the worth and resource demands of its goods, and may lie about them so as to increase its utility under our proposed scheme. We show that our scheme is susceptible to two manipulations: **Splits**, where a buyer pretends its good g is actually two smaller goods g' and g'' in order to reduce the price she pays[2], and **Bluffs**, where a buyer pretends to request a fake good g with the intention of never having it provided, in order to receive a compensation from the sellers. We show that our approach is somewhat resistant to such manipulations and that the utility gain is bounded. The manipulation depends on knowledge regarding other users, as such manipulations can also harm the agent if performed in the wrong environment. Finally, since our model is applicable to pricing in cloud environments with elastic demands, we provide simulation results showing that our approach maximizes social welfare and substantially increases the revenue of the platform over currently used fixed pricing.

[1]In our simulations, we will make use of known approximations to this NP-hard problem.

[2]Indeed we consider indivisible goods, however in cloud settings, programs such as MapReduce, [24], may be used to temporarily "split" a good and merge it back later. Nevertheless, both parts still need to be allocated on a server.

1.2 Related Work

The Assignment Game was proposed in [37] and in [40] is generalized to consider many-to-many matchings. In [31] it is further generalized to consider heterogeneous goods. In contrast, in our model each seller has a capacity constraint for the set of goods it sells and goods vary in their demand for resources. In [19] matchings are considered with externalities and in [4] mechanisms are brought forth for the Shapley-Scarf housing market. Furthermore, the Shapley value and other cooperative solution concepts were used to allocate utility in a stable or fair manner in domains including pollution control [35], network settings [15, 13, 36, 6, 14, 18] and even sharing the costs of joint facilities such as airfields [26, 33]. Alternatively, the Shapley value can be viewed as a *prediction* regarding outcomes reached when bargaining or negotiating in market settings [29, 30, 9]. Previous work studies unweighted coalitional manipulation problems [45] and [3] considers Shapley Value manipulations in Weighted Voting Games. Similar false-name manipulations were also studied in [42, 23, 46, 27]. Fair solutions and mechanisms are combined in [21]. For an overview of the Shapley Value, its axioms and applications, see [44].

The remainder of the paper is organized as follows. In Section 2 we formulate the SAG and show that the core may be empty. Additionally, we consider buyers to have private information on their goods and bring forth bounds on their possible manipulations. An application to the cloud environment is illustrated in Section 3, together with a simulation. Finally, conclusions are represented in Section 4.

2. THE SHARED ASSIGNMENT GAME (SAG)

The SAG is a game over the agent set N, consisting of two disjoint agent types: buyers $B \subseteq N$ and sellers $R \subseteq N$. Each good g, has a worth that buyer b assigns to it, $w_g(b)$.[3] It captures the value of the good in the eyes of the buyer or the maximum price it is willing to pay to receive it. The worth is only split between the seller and buyer if the seller has sufficient resources to provide the good.

Denote the set of goods that a buyer $b \in B$ wishes to receive as G_b. Each good has a tuple $(\mathbf{d_g}, w_g(b))$, where $\mathbf{d_g} = [d_g^1 \dots d_g^L]$ is the demand vector and every d_g^k denotes the demand of g for resource type k.

Seller $j \in R$ has a *capacity vector* $\mathbf{C_j} = [C_j^1 \dots C_j^L]$, where C_j^k is the capacity constraint of seller j for resource type k. An assignment $f : B \to R \cup \{\perp\}$ maps buyers to sellers, where the value $\{\perp\}$ indicates the buyer is not mapped to any seller. We denote the buyers mapped to a seller $r \in R$ in assignment f as $B_r^f = \{b | f(b) = r\}$. We denote the total demand of resource x of seller r under assignment f as $Q_{r,x}^f = \sum_{t \in B_r^f} d_t^x$. An assignment is *feasible* if it respects the capacity constraints of the sellers so for any seller r and any resource type x we have $Q_{r,x}^f \leq C_r^x$ (i.e., the total amount of resource x used in sellers r under the assignment f is at most the available amount of that resource of that seller).

For any coalition $S \subseteq N$ let a_S be a feasible assignment over the agents in S, and let A_S be the set of all feasible assignments over S. Let $G_b(a_S) \subseteq G_b$ be the set of goods that buyer b receives in a_S. Denote the *social welfare*, W_{a_S}, as the total worth of the received goods in a_S, so

[3]We fix the reservation price of all sellers at 0, hence the valuation of the buyer equals the gain of the partnership, i.e., the *worth* of the good.

$W_{a_S} = \sum_{b \in B} \sum_{g \in G_b(a_S)} w_g(b)$. Further, denote by B_{a_S}, the set of buyers that are matched with at least one seller, in a_S. An optimal assignment a^*_S is an assignment that maximizes the social welfare of all feasible assignments over S, i.e., $\arg\max_{a_S \in A_S} W_{a_S}$. We denote the *maximum social welfare* achievable for set S as $W^*_S = \max_{a_S \in A_S} W_{a_S}$. Moreover, we impose a standard assumptions regarding agents and their resource demands: $\forall i, k\ \exists j$ for which $C^k_j \geq d^k_i$, i.e., for any requested good by a buyer, there exists at least one seller with the initial resource capacity to provide it.

TU Cooperative Game: A transferable utility cooperative game $G = (N, v)$ [37], is composed of a set of agents $N = \{1, 2, \ldots, n\}$ and a characteristic function $v : 2^N \to \mathbb{R}$ giving the total value created by various coalitions (agent subsets). By convention, $v(\emptyset) = 0$. By $v(A, B)$ we denote the output of function v applied on the input set $\{A, B\}$. The characteristic function defines the amount of utility any coalition achieves, but selfish agents are only interested in their share of the total utility. An imputation is a division of the grand coalition gains, and is represented by a vector $p = (p_1, \ldots, p_n)$ where $p_i \in \mathbb{R}_+$ is the share of agent i, where the total is the grand coalitions value $\sum_{i=1}^n p_i = v(N)$. An imputation p dominates an imputation $\bar{p}$ if there exists a subset of agents S for which $p_i > \bar{p}_i,\ \forall i \in S$.

The **Shared Assignment Game (SAG)** is the game $H = (N, v)$ over the agents $N = B \cup R$ where the coalition C's value is the total utility achieved in the optimal assignment over the agents in C, i.e., the total social welfare under the optimal assignment for C: $v(C) = W^*_C = \max_{a_C \in A_C} W_{a_C}$. The total worth created due to the optimal match between sellers and buyers, i.e., W^*_N is split between all agents according to a solution concept.

The core: The *core* is the set of imputations that are not dominated. As it is stable under coalitional deviation, the *core* has been used [37] as a fair solution concept. However, the following lemma shows that in contrary to the original assignment game, it does not necessarily exist in our case.

LEMMA 1. *The core in the SAG may not exist.*

PROOF. Consider two sellers, r_1, r_2, with one resource type and capacities $c_{r_1} = 1$ and $c_{r_2} = 2$. Consider three buyers, each interested in one good: g_1, g_2 and g_3 with tuples respectively of $(d_{g_1} = 1, w_{g_1}(b_1) = 2)$, $(d_{g_2} = 1, w_{g_2}(b_2) = 2)$, $(d_{g_3} = 2, w_{g_3}(b_3) = 3)$. It is easy to see that the optimal matching brings about: $W^*_N = 5$. Thus, this worth needs to be distributed between the agents. Assume by contradiction that there exists a core allocation $\mathbf{x}$, so:

$$x_{r_1} + x_{r_2} + x_{b_1} + x_{b_2} + x_{b_3} = 5. \tag{1}$$

Since $\mathbf{x}$ lies in the core it holds that:

$$x_{r_1} + x_{r_2} + x_{b_2} + x_{b_3} \geq v(r_1, r_2, b_2, b_3) = 5, \tag{2}$$

hence $x_{b_1} = 0$. Similarly, it follows that $x_{b_2} = 0$. However

$$x_{r_2} = x_{r_2} + x_{b_1} + x_{b_2} \geq v(r_2, b_1, b_2) = 4, \tag{3}$$

and similarly $x_{r_1} \geq 2$. This is in contradiction to (1). □

Even if the *core* would exist, it would be able to allocate the social welfare in what seems to be unfair to either the sellers or buyers. As an example, The "gloves game" [34], can be represented as a SAG and its *core* may give all the created welfare to either the buyers or the sellers. It is due to this consideration and due to Lemma 1 that we continue to focus on the Shapley value as a *fair* solution concept to allocate the welfare between the agents. As such, we define the utility of each agent $i \in N$ as its Shapley Value: $\phi_i(v)$. The **Shapley value** was proposed as a fair solution for dividing the gains of the grand coalition N of all the agents. It examines the expected contribution an agent adds to its predecessors in an agent order chosen uniformly at random, and is the only solution exhibiting a natural set of fairness axioms [38, 43]. Given a game $H(N, v)$, denote by π a permutation of the N agents ($\pi : \{1, \ldots N\} \to \{1, \ldots N\}$ and π is reversible), by Π_n the set of all agent permutations, by Γ^π_i the predecessors of agent i in π so, $\Gamma^\pi_i = \{j \mid \pi(j) < \pi(i)\}$, and by Ω^π_i the successors of i in π: $\Omega^\pi_i = \{j \mid \pi(j) > \pi(i)\}$. The *marginal contribution* of i to permutation π, denoted $m^\pi_{i,v}$, is its contribution to its predecessors in the permutation, so $m^\pi_{i,v} = v(\Gamma^\pi_i \cup \{i\}) - v(\Gamma^\pi_i)$. The Shapley value of i, denoted $\phi_i(v)$, is i's contribution averaged over all permutations:

$$\phi_i(v) = \frac{1}{N!} \sum_{\pi \in \Pi_n} m^\pi_i = \frac{1}{N!} \sum_{\pi \in \Pi_n} v(\Gamma^\pi_i \cup i) - v(\Gamma^\pi_i) \tag{4}$$

An equivalent formula is:

$$\phi_i(v) = \sum_{S \subseteq N \setminus i} \frac{|S|!(|N - S - 1|)!}{|N|!} m^{\{S\}}_{i,v} \tag{5}$$

where $m^{\{S\}}_{i,v} = v(S \cup i) - v(S)$, i.e., i's *marginal contribution* to $S \subseteq N$. For each S, we call $\frac{|S|!(|N-S-1|)!}{|N|!} m^{\{S\}}_{i,v}$ the proportional marginal contribution of $m^{\{S\}}_{i,v}$.

The SAG by itself only defines the total utility any agent subset can achieve. By applying solution concepts we can use its structure to reach normative and predictive conclusions regarding the sharing of the total revenue. In the normative interpretation, applying the Shapley value to the SAG would give us the importance of each component and its "fair" contribution to the total welfare. In the predictive sense, it estimates how much utility each component is expected to achieve when negotiating with other agents. [4]

We propose using the Shapley value as an *instrument* for sharing the revenue in the SAG, rather than letting agents negotiate through a market mechanism. One caveat is that the SAG assumes monetary transfers between participants. A seller who has a non-zero Shapley value would get it from the pool of payments made. As the Shapley values sum up to the grand coalitions utility, we are guaranteed that the sum of the total payments is zero (i.e. no external subsidy is required as we have budget balance). However, if a buyer does not receive its good in the optimal assignment a^*_N, it might still receive a positive utility, since the Shapley Value looks at its contribution to all possible assignments. This may be difficult to do in certain scenarios, e.g., a cloud computing platform, such as Windows Azure or Amazon EC2. One alternative is to offer a discount to the buyer or task owner, to use in future transactions.

We continue with an example in which the Shapley value distributes revenue in a fair manner, by rewarding the agents that are most critical to the system.

[4]Obviously in any round of negotiation a buyer b either strikes a deal with some seller and achieves its good's worth $w_g(b)$, or it does not strike a deal and achieves zero utility. Its Shapley value is likely to be somewhere in between the two, as it is the expected utility it is likely to have in the uncertain negotiation.

EXAMPLE 1. *Consider servers and tasks of Figure 1. The characteristic function, v, is described by: $v(r_2, b_1) = 4$, $v(r_2, b_2) = 3$, $v(r_2, b_1, b_2) = 4$. Moreover, $\forall S \subseteq N\backslash\{r_1\}$, $v(S, r_1) = v(S)$. For any other $S \subseteq N$, $v(S) = 0$. See Table 1 for the agents' Shapley values.*

As shown in Figure 1, only the goods g_1 and g_2 are provided, hence the social welfare is equal to $v(N) = W_N^* = 4$. This amount is distributed between all agents according to (5). Note that seller r_1 does not contribute anything to any sub-coalition, so using (5), $\phi_{r_1}(v) = 0$. However, seller r_2 receives a large part of the aggregated welfare created as it provides both g_1 and g_2. Finally, good g_3 is not provided, however b_2 receives the remainder payment, considering that though it did not receive g_3, it does marginally contribute to certain sub-coalitions, so according to the "fairness" criteria of the Shapley Value it should receive a compensation. Finally, note that $v(N) = \sum_{i\in N} \phi_i(v)$.

Utility Sellers	Utility Buyers	Discount
$\phi_{r_1} = 0$, $\phi_{r_2} = \frac{30}{12}$	$\phi_{b_1} = 1$	$\phi_{b_2} = \frac{6}{12}$

Table 1: Distribution of welfare according to the Shapley value ϕ_i.

2.1 Scenarios with Private Information

Unlike Example 1, in certain domains, such as in cloud environments, the worths and demands of goods are *private information*, supplied by the buyers. Hence, it might be beneficial for a subset of buyers to provide false reports regarding these values, in order to receive a higher Shapley value. Unless stated otherwise, in the remainder of the paper we hold by two assumptions:

- Each buyer is only interested in a single good, i.e., $w_g(b) \equiv w_b$ and $d_g^k \equiv d_b^k$.
- Each buyer only receives its good from a single seller, i.e., we restrict ourselves to many-to-one matching.

For each buyer $b \in B$, we denote the set of sellers that have enough initial capacity to provide its good g, as $R_b \subseteq R$. Thus, $R_b = \{r \mid r \in R,\ C_r^k \geq d_b^k,\ \forall k\}$.

Before we show the robustness of the Shapley value against agent manipulations, we prove two lemmas.

LEMMA 2. *The amount of permutations in which $S \subseteq N$ agents are a subset of the predecessors of b, equals $\frac{N!}{|S|+1}$. Formally: $|\{\pi \mid S \subseteq N, S \subseteq \Gamma_b^\pi\}| = \frac{N!}{|S|+1}$.*

PROOF. Given agent b in a fixed position and $|\Gamma_b^\pi|$ agents as predecessors of b, where $|\Gamma_b^\pi|$ lies between $|S|, \ldots, |N-1|$. When summing over $|\Gamma_b^\pi|$, all permutations of the remaining $N\backslash\{b\}$ agents for which $S \subseteq \Gamma_b^\pi$ equals:

$$\sum_{|\Gamma_b^\pi|=|S|}^{N-1} |\Gamma_b^\pi|!(|N-1| - |\Gamma_b^\pi|)! \cdot \binom{|N-S-1|}{|\Gamma_b^\pi| - |S|}. \qquad (6)$$

The first two factorial elements in (6) represent the cardinality of the set of permutations of agents in Γ_b^π and in Ω_b^π. Except for the $S \subseteq \Gamma_b^\pi$ agents, there remain $|N-S-1|$ agents in $\{\Gamma_b^\pi \cup \Omega_b^\pi\}$ that fit into $|\Gamma_b^\pi| - |S|$ slots. Equation (6) turns into:

$$|N-S-1|!|S|! \sum_{|\Gamma_b^\pi|=|S|}^{N-1} \binom{|\Gamma_b^\pi|}{|S|} = \frac{N!}{|S|+1}. \qquad (7)$$

Equation 7 can be proven by induction. □

We bound the Shapley value that an agent can receive.

LEMMA 3. $\forall b \in B$, $\frac{w_b|R_b|}{|N|\cdot(|N|-|R_b|)} \leq \phi_b(v) \leq \frac{w_b|R_b|}{|R_b|+1}$

PROOF. **Upper Bound:** To produce an upper bound on the Shapley Value, we consider its representation as given in (4) and examine the best-case scenario where $m_{b,v}^\pi = w_b$ if $\exists r \in R_b$ such that $r \in \Gamma_b^\pi$. From Lemma 2 we have: $|\{\pi \mid \forall r \in R_b, r \in \Gamma_b^\pi\}| = \frac{|N|!}{|R_b|+1}$, i.e., the amount of permutations in which all sellers in R_b are predecessors of b.

Hence through the symmetry of permutations it follows that $|\{\pi \mid \forall r \in R_b, r \in \Omega_b^\pi\}| = \frac{|N|!}{|R_b|+1}$. Therefore:

$$|\{\pi \mid \exists r \in R_b, r \in \Gamma_b^\pi\}| = |N!| - \frac{|N|!}{|R_b|+1} = \frac{|N|!|R_b|}{|R_b|+1} \qquad (8)$$

and from (4): $\phi_b(v) \leq \frac{|R_b|w_b}{|R_b|+1}$.

Lower bound: From our standard assumptions it holds that $\forall i, k,\ \exists j$ for which $C_j^k \geq d_i^k$. Thus, from (5), $\forall b \in B$, $m_{g,v}^{\{r\}} = w_b$ if $r \in R_b$. Hence, a buyer will always be able to receive a proportional marginal contribution of:

$$\sum_{S=1}^{|R_b|} \frac{|S|!(|N-S-1|)!}{|N|!} \binom{|R_b|}{S} w_b, \qquad (9)$$

since it is the only buyer and can choose any combination of sellers to provide its good. Equation 9 turns into:

$$\frac{|R_b|!(|N|-|R_b|-1)}{|N|!} \sum_{S=1}^{|R_b|} \binom{|N|-1-S}{|R_b|-S} w_b \qquad (10)$$

$$= \frac{|R_b|!(|N|-|R_b|-1)}{|N|!} \binom{|N|-1}{|R_b|-1} w_b = \frac{|R_b|w_b}{|N|(|N|-|R_b|)}$$

It follows that:

$$\frac{w_b|R_b|}{|N|\cdot(|N|-|R_b|)} \leq \phi_b(v) \leq \frac{w_b|R_b|}{|R_b|+1} \qquad (11)$$

□

We denote by $\hat{w}_b$, the *reported worth* of buyer $b \in B$ for its good g. This represents the worth that the buyer reports to the system or sellers. If a good is not received $\hat{w}_b = w_b = 0$ and when buyer b behaves truthfully it follows that $w_b = \hat{w}_b$. The reported worth does not necessarily equal the actual worth, w_b and agents might resort to such manipulations in order to increase their payoff. Furthermore, we define the ***Manipulation Gain*** of a buyer b as:

$$u_b = w_b - (\hat{w}_b - \phi_b(v)) = \phi_b(v) + w_b - \hat{w}_b. \qquad (12)$$

This captures the real payoff the buyer obtains from receiving its good and when behaving truthfully, the Shapley Value corresponds to this payoff.

We consider two types of agent manipulations:

1. A buyer b who **splits** its good g into g' and g'', with $w_b = \hat{w}_{b'} + \hat{w}_{b''}$ and $\forall k,\ d_b^k = d_{b'}^k + d_{b''}^k$. Since we assume each buyer is only interested in a single good, we also denote this by b splitting itself into b' and b''.

2. A buyer b, who **bluffs** by requesting a fake good from the system, in order to increase its utility.

2.1.1 Splits

Denote the original game, before splitting b into b' and b'', as $H = (N, v)$ and the *new* game after the split, as $H' = (N{+}1, w)$. We show that splitting a good can decrease or increase the Shapley Value of an agent.

EXAMPLE 2. *Consider a seller R with one resource type and capacity $c_R = 2.5$. Furthermore, consider two buyers b_1, b_2 with tuples $(d_{b_1} = 2, w_{b_1} = 2)$ and $(d_{b_2} = 1, w_{g_2} = 1)$. From (5) it follows that $\phi_{b_1}(v) = \frac{2}{3}$. When b_1 decides to split in two, we get a total of three identical "buyers" b_1', b_1'', b_2, with demands $d_{b_1'}, d_{b_1''}, d_{b_2} = 1$ and worths $\hat{w}_{b_1'}, \hat{w}_{b_1''}, w_{b_2} = 1$. It follows from (5) that $\phi_{b_1''}(w) + \phi_{b_1'}(w) = \frac{1}{2}$ and $\frac{\phi_{b_1''}(w)+\phi_{b_1'}(w)}{\phi_{b_1}(v)} = \frac{3}{4}$. This illustrates that a* split *can lower the Shapley Value of a buyer.*

On the other hand, a split might also be profitable, e.g., when evenly splitting b_2 instead of b_1. Then from (5) we get: $\phi_{b_2'}(w) + \phi_{b_2''}(w) = \frac{1}{4}$ and $\frac{\phi_{b_2''}(w)+\phi_{b_2'}(w)}{\phi_{b_2}(v)} = \frac{3}{2}$, thus the Shapley value of b_2 has increased.

To quantify the potential loss in Shapley value do to a split, we first provide a *lower bound* on the ratio of agents' marginal contributions. For any $m_{b,v}^{\{S\}}$ in H, consider four marginal contributions in H': (1) $m_{b',w}^{\{S\}}$, (2) $m_{b'',w}^{\{S\}}$, (3) $m_{b'',w}^{\{S,b'\}}$, (4) $m_{b',w}^{\{S,b''\}}$, where $\forall i, S$, $m_{i,v}^{\{S\}}$ and $m_{i,w}^{\{S\}}$ are defined as in (5). For any $S \subseteq N\backslash\{b\}$ we define the *Marginal Ratio*, $\alpha_{S(b)}$, as the ratio between the *proportional marginal contributions* of the above *marginal contributions* and $m_{b,v}^{\{S\}}$. We seek a lower bound, so $m_{b,v}^{S} > 0$. From (5), $\forall S \subseteq N\backslash\{b\}$, $\alpha_{S(b)}$ is:

$$\frac{\frac{|S|!|N-S|!}{|N+1|!}[m_{b',w}^{\{S\}} + m_{b'',w}^{\{S\}}] + \frac{|S+1|!|N-S-1|!}{|N+1|!}[m_{b'',w}^{\{S,b'\}} + m_{b',w}^{\{S,b''\}}]}{m_{b,v}^{\{S\}}\frac{|S|!|N-S-1|!}{|N|!}}$$

We show a lower bound on the *Marginal Ratio*.

LEMMA 4. *For all $S \subseteq N\backslash\{b\}$, $\alpha_{S(b)} \geq \frac{2}{|N|+1}$.*

PROOF. By splitting a good g into g' and g'', we enlarge the space of possible strategies, since g' and g'' can always be combined together to construct g. Thus, $w(S, b', b'') = w(S, b'', b') \geq v(S, b)$. As $v(S) = w(S)$ we get:

$$\begin{aligned}
\alpha_{S(b)} &= \frac{2[w(S,b',b'') - w(S)](|S+1|)}{|N+1|(v(S,b) - v(S))} \qquad (13)\\
&+ \frac{[w(S,b') + w(S,b'') - 2w(S)](|N-2S-1|)}{|N+1|(v(S,b) - v(S))}\\
&\geq \frac{2}{N+1} + \frac{[2w(S,b',b'') - w(S,b') - w(S,b'')]|S|}{|N+1|(v(S,b) - v(S))}\\
&+ \frac{[w(S,b') + w(S,b'') - 2w(S)](|N-S-1|)}{|N+1|(v(S,b) - v(S))}\\
&\geq \frac{2}{N+1}.
\end{aligned}$$

The inequalities of (13) follow, as $|N-S-1| \geq 0$, $w(S,b',b'') \geq w(S,b') \geq w(S)$ and $w(S,b',b'') \geq w(S,b'') \geq w(S)$. $\square$

As Lemma 4 holds for all S, from (5), this gives a lower bound on the loss of a buyers b's Shapley Value due to a split. We now continue to tighten this bound by using Lemma 4.

THEOREM 1. $\phi_{b'}(w) + \phi_{b''}(w) \geq \left[\frac{1}{|N|} + \frac{|N|-1}{|N|(|N|-|R_b|)}\right]\phi_b(v)$.

PROOF. In H', denote the proportional marginal contributions of $m_{b',w}^{\{r\}}$, $m_{b'',w}^{\{r\}}$, $m_{b',w}^{\{r,b''\}}$ and $m_{b'',w}^{\{r,b'\}}$, where $r \in R_b$, as *initial contributions* of b' and b''. According to (5), their sum equals:

$$\sum_{S=1}^{|R_b|}\left[\frac{|S|!(|N+1-S-1|)!}{|N+1|!} + \frac{|S+1|!(|N+1-(S+1)-1|)!}{|N+1|!}\right]\binom{|R_b|}{S}(\hat{w}_b' + \hat{w}_b'').$$

From (9) and Lemma 3 it follows that,

$$\frac{\phi_{b'}(w) + \phi_{b''}(w)}{\phi_b(v)} = \frac{\frac{|R_b|w_b}{|N|(|N|-|R_b|)} + \bar{X}}{\frac{|R_b|w_b}{|N|(|N|-|R_b|)} + X}, \qquad (14)$$

where, $\bar{X}$ and X equal the sum of all *proportional marginal contributions* of $m_{b',w}^{\{S\}}$ and $m_{b'',w}^{\{S\}}$ for all S, except for the *initial contributions* of b', b'' and b in respectively H' and H. Moreover, as a result from Lemma 4, $\bar{X} \geq \frac{2}{|N|+1}X$. Hence, from Lemma 3, $\phi_{b'}(w) + \phi_{b''}(w)$ is equal to:

$$\begin{aligned}
&\frac{w_b|R_b|}{|N|(|N|-|R_b|)} + \phi_{b'}(w) + \phi_{b''}(w) - \frac{w_b|R_b|}{|N|(|N|-|R_b|)} \qquad (15)\\
&\geq \frac{w_b|R_b|}{|N|(|N|-|R_b|)} + \frac{2}{|N|+1}\left[\phi_b(v) - \frac{w_b|R_b|}{|N|(|N|-|R_b|)}\right].
\end{aligned}$$

By (11) and (15) :

$$\begin{aligned}
\frac{\phi_{b'}(w) + \phi_{b''}(w)}{\phi_b(v)} &\geq \frac{\frac{w_b|R_b|}{|N|(|N|-|R_b|)}}{\frac{w_b|R_b|}{|R_b|+1}}\left[1 - \frac{2}{|N|+1}\right] + \frac{2}{|N|+1} \qquad (16)\\
&= \frac{1}{|N|} + \frac{|N|-1}{|N|(|N|-|R_b|)},
\end{aligned}$$

which equals $\frac{2}{|N|}$ when $|R_b| = 1$. $\square$

Actual Lower Bound: The lower bound of Theorem 1 does not take into account that our SAG deals with indivisible goods. Hence, even if goods are temporarily split by the buyer, e.g., when dealing with a cloud environment through MapReduce, [24], or a similar program, the goods are worthless unless *all* of their parts get provided. So if either g' or g'' is not received in H', $w_{b'} = w_{b''} = 0$. Consider w.l.o.g. that in H', $b' \in B_{a_N^*}$ and $b'' \notin B_{a_N^*}$. For b' and b'', the Manipulation Gain turns into:

$$\begin{aligned}
u_{b'} + u_{b''} &= 0 - (\hat{w}_{b'} - \phi_{b'}(w)) + 0 - (\hat{w}_{b''} - \phi_{b''}(w)) \qquad (17)\\
&= \phi_{b'}(w) + \phi_{b''}(w) - \hat{w}_{b'} - \hat{w}_{b''}\\
&= \phi_{b'}(w) + \phi_{b''}(w) - w_b.
\end{aligned}$$

Applying Theorem 1, and Lemma 3 we obtain:

$$\begin{aligned}
\frac{u_{b'} + u_{b''}}{\phi_b(v)} &\geq \frac{2N - |R_b| - 1}{|N|(|N|-|R_b|)} - \frac{w_b}{\frac{w_b|R_b|}{|N|(|N|-|R_b|)}} \qquad (18)\\
&\geq \frac{2N - |R_b| - 1}{|N|(|N|-|R_b|)} - \frac{|N|(|N|-|R_b|)}{|R_b|},
\end{aligned}$$

which can be strongly negative. The following example illustrates that this bound is rather tight and that an agent stands to lose immensely due to a split.

EXAMPLE 3. *Consider a single seller R with one resource-type and capacity $C_R = 3$. Consider $K-1$ identical buyers with tuples $(d_i = 2, w_i = w)$, for $i = 1, \ldots, K-1$, and a buyer b with tuple $(d_b = 3, w_b = 2w)$, which splits itself into b' and b'' with $d_{b'} = d_{b''} = 1.5$ and $\hat{w}_{b'} = \hat{w}_{b''} = w$. From (5) it follows that $\phi_b(v) = \frac{2w}{(K+1)K}$ and $\phi_{b'}(w) + \phi_{b''}(w) = \frac{2w}{(K+1)K}$. The ratio of manipulation gains equals:*

$$\frac{u_{b'} + u_{b''}}{\phi_b(v)} = \frac{\frac{2w}{(K+1)K} - w}{\frac{2w}{(K+1)K}} = 1 - \frac{K(K+1)}{2}. \tag{19}$$

Thus, when K grows, b's loss is large and when w increases, its manipulation gain in G' decreases linearly.

Example 3 and (18) add to Theorem 1, showing that buyers which split their goods, risk incurring large losses. We show however that a split may also be beneficial for a buyer.

THEOREM 2. *$\forall b \in B$: $\phi_{b'}(w)+\phi_{b''}(w) \leq \frac{|N|(|N|-|R_b|)}{|R_b|+1}\phi_b(v)$. This bound is tight.*

PROOF. From Lemma 3, $\forall b \in B$ we have:

$$\frac{\phi_{b'}(w) + \phi_{b''}(w)}{\phi_b(v)} \leq \frac{|R_b|\frac{\hat{w}_{b'}+\hat{w}_{b''}}{|R_b|+1}}{\frac{w_b|R_b|}{|N|(|N|-|R_b|)}} = \frac{N(|N|-|R_b|)}{|R_b|+1}. \tag{20}$$

□

EXAMPLE 4. *We show tightness through an example: Consider R sellers with a single resource-type and capacities $C_j = 4 - \epsilon$, $\forall j \in R$, where $\epsilon > 0$. Also, consider K identical buyers with tuples $(d_i = 2, w_i = 3)$, $\forall i \in K$. Furthermore, consider a buyer b with tuple $(d_b = 2, w)$. Buyer b splits itself into b', with $(d_{b'} = 2 - \epsilon, \hat{w}_{b'} = w - \epsilon)$ and b'', with $(d_{b''} = \epsilon, \hat{w}_{b''} = \epsilon)$. When taking the limit of ϵ to zero, from (4), it follows that $\phi_{b'}(w) = \hat{w}_{b'}(1 - \frac{1}{|R|+1})$ and $\phi_{b''}(w) = 0$. Finally, $\frac{\phi_{b'}+\phi_{b''}}{\phi_b} \xrightarrow{\epsilon\to 0} \frac{|N|(|N|-|R|)}{|R|+1}$.*

Brittleness: By Theorem 2, a buyer may decide to split its good, especially if it fully knows the other agents' demands and worths. The next example shows that a *small* change in the system or a wrong estimation of the buyer, can be *detrimental* to the buyer's manipulation gain.

EXAMPLE 5. *Consider two identical servers r_1, r_2 with one resource type and capacities $C_{r_1}, C_{r_2} = 3$ and three identical buyers, with tuples $(d_i = 2, w_i \equiv w)$, for $i = 1, 2, 3$. In the game $H = (N, v)$, it holds that $\forall i$, $\phi_{b_i}(v) = \frac{7}{30} \cdot w$. When b_3 splits itself equally into $b_{3'}$ and $b_{3''}$, its Shapley Value rises and equals: $\phi_{b_{3'}} + \phi_{b_{3''}} = \frac{8}{15} \cdot w$. When parameters slightly change and $C_{r_2} = 3 - \epsilon$, where ϵ is small and positive, the Manipulation Gain turns into: $u_{b_{3'}} + u_{b_{3''}} = \left[\frac{19}{60} - 1\right] \cdot \frac{w}{2} < 0$, decreasing linearly with w. Had b_3 chosen not to split itself, its Shapley Value would be $\frac{7}{30} \cdot w$.*

Hence splitting is very risky, as even a small parameter perturbation can cause an loss larger than the potential gain.

2.1.2 Bluff

Another possible manipulation in a game $H = (N, v)$ is *bluffing*: declaring a fake good g. Though the buyer does not intend for g to be delivered, it can request g in order to extract a higher utility from the system. For an upper bound of this utility, we assume that buyer b has complete information on the worth of all the other agents. We define the *lying ratio* as the ratio between the Shapley Value of a buyer $b \notin B_{a^*_N}$ and the maximum social welfare achievable for N: $\frac{\phi_b(v)}{v(N)}$. The following theorem proves that b is limited in the amount of welfare it can receive from the system.

THEOREM 3. *The lying ratio is bounded by $\frac{1}{|R_b|+1} - \frac{1}{|N||R_b|}$.*

PROOF. Just as in Lemma 3, for an upper bound, we consider a best-case scenario, i.e., if there is an $r \in R_b$ such that $r \in \Gamma^\pi_b$ then $m^\pi_{r,v} = w_b$. By (8), $|\{\pi \mid \exists r \in R_b, r \in \Gamma^\pi_b\}| = |N|! - \frac{|N|!}{|R_b|+1} = \frac{|N|!|R_b|}{|R_b|+1}$. Furthermore, if $|\Gamma^\pi_b| = |N-1|$, $m^\pi_{b,v} = 0$, otherwise, $b \in B_{a^*_N}$. Thus, we need to subtract the permutations in which $|\Gamma^\pi_b| = |N-1|$, so that the permutations of interest equal: $\frac{|N|!|R_b|}{|R_b|+1} - |N-1|!$. Our standard modeling assumptions state that $\forall i, k$, $\exists j$ for which $C^k_j \geq d^k_b$. Since $b \notin B_{a^*_N}$, it follows that in the optimal assignment, each resource which can initially supply the good of b, is matched with tasks with aggregated worth higher than w_b, hence $v(N) \geq |R_b| \cdot w_b$. Thus from (4):

$$\frac{\phi_b(v)}{V(N)} \leq \frac{\frac{1}{|N|!}(\frac{|N|!|R_b|}{|R_b|+1} - |N-1|!) \cdot w_b}{|R_b| \cdot w_b} = \frac{1}{|R_b|+1} - \frac{1}{|N||R_b|}. \tag{21}$$

□

The upper bound of Theorem 3 considers a buyer with full information on all other buyers and sellers. Thus the good, g, can be unwillingly provided when a small error is made in estimating other agents' parametersThen, by Lemma 2 and (12), the *manipulation gain* equals:

$$u_b = 0 + \phi_b(v) - \hat{w}_b \leq \hat{w}_b \frac{|R_b|}{|R_b|+1} - \hat{w}_b = -\frac{\hat{w}_b}{|R_b|+1}. \tag{22}$$

Hence, if b's estimation on the values of other agents is erroneous, the best case states that b stands to lose a significant amount in comparison to its potential gain by bluffing. This makes it unlikely for agents to employ *bluff* manipulations.

3. APPLICATIONS TO MARKETS AND CLOUD COMPUTING

One application of our model is allocation and pricing in cloud computing. Public cloud computing services, such as Windows Azure and Amazon EC2, provide clients with scalable, on-demand access to compute resources. Traditionally, a client specifies the type and "size" of the virtual machine (VM) it requires, and pays a *fixed price* per time interval. The price of a machine depends on the virtual machine's "size". Such *fixed price* mechanisms may be problematic as clients only wish to pay for resources that they actually use [20]. For example, a client who needs many CPUs but little memory may feel cheated paying for a large machine, as she is not using most of the memory. Our model can be used in cloud environments, through the "VM" technology, which allows a single server to run multiple tasks, guaranteeing each a specified amount of resources. This is done by equating the sellers in the SAG, to servers with resources in the cloud. The goods they provide correspond to the required resource bundles (or VM's) of tasks that the clients wish to execute on the cloud. Examples of resource types in cloud environments include CPU, memory and bandwidth. We keep our assumption of only allowing many-to-one matches, as each task only needs to be executed by a single server. Finally, the SAG supports buyers with multiple tasks to execute, based on the available resources, incorporating a version of elastic demand.

In fixed pricing, a good or task whose worth is below the fixed price may not be executed at all, even if there are free resources. Further, a fixed-price scheme leaves unspecified how tasks are allocated to sellers or servers, as it only determines where tasks are eligible to run. In the cloud domain, the Shapley value can be seen as a fair way of allocating the gains from optimally executing the available tasks on the available servers, or as a prediction regarding the prices for cloud servers that would emerge under agent negotiation.

In previous work, the revenue of the cloud sever providers under fixed and spot pricing is discussed in [1]. Available pricing methods include truthful auctions, e.g., VCG mechanisms. However, they do not consider pricing from a fairness perspective and when describing the optimal allocation of tasks, their computation can be NP-hard. Also, when relying on approximations, it loses some of its strategy-proofness [2]. Note that our proposed solution is not strategy-proof (see Theorems 1 and 2), nevertheless we do provide approximate fairness guarantees and we use approximations which are computable in polynomial time, [10]. Finally, new methods, [28], allow cloud providers to allocate tasks efficiently and fairly according to their resource demands, which can be combined with our proposed pricing model.

3.1 Simulation

In order to further support our proposed SAG as a pricing mechanism, we use a simulation to show that when applied to a cloud pricing environment, the SAG maximizes social welfare and increases the revenue of the cloud servers over a fixed-price scheme (FPS). We simulated 50 incoming tasks with demands for two resource types (e.g., CPU, MEM) randomly sampled from i.i.d Gaussian distributions $\mathcal{N}(20, 10)$, however similar results hold with other distribution parameters. The worth of every task t, w_t, is computed by taking the *maximum demand* of either resource-type, $\max\{d_t^{CPU}, d_t^{BW}\}$ and adding an offset α (same α for all tasks). First, we compute $B_{a_N^*}$ and W_N^* as solutions to a Multi-Dimensional-Multiple-Knapsack (MDMK) [39]. The optimal value of the MDMK is hard to compute, so we used a greedy approximation [22, 32]. We then approximated the Shapley Value of all the incoming tasks using the algorithm of [10]. The revenue of the SAG is the Shapley Value of all the servers, and the social welfare is $W_{a_N} = \sum_{i \in B_{a_N^*}} w_i$. Figure 2 compares the social welfare and the revenue of SAG and FPS. The x-axis is the *demand ratio*: the ratio of the expected aggregated capacity of the servers, over the expected aggregated demand of the tasks, $\frac{\sum_{i \in T} C_i}{\sum_{j \in R} d_j}$. This ratio is equal for both resource types and we increase it by adding resources. The FPS uses a fixed price $\bar{f}$, where $\bar{f}$ is the revenue optimal fixed price, calculated by summing over the revenue at each demand ratio, for the given input set of tasks. The FPS social welfare equals $\sum_{\{t | w_t \geq \bar{f}\}} w_i$ and the FPS revenue equals $\bar{f} \cdot |\{t \mid w_t \geq \bar{f}\}|$.

Results: The social welfare is approximately maximized under the SAG, as under the MDMK, all tasks in $B_{a_N^*}$ are

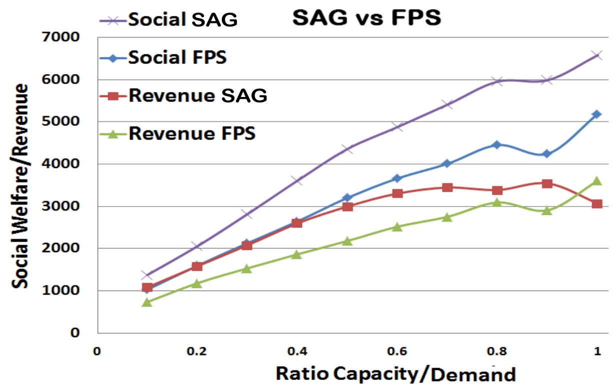

Figure 2: Social Welfare and Revenue Comparison of the FPS and the SAG.

executed. For FPS this is not always the case, as tasks may have a worth that is lower than the fixed price $\bar{f}$. When increasing the demand ratio we get an increase in social welfare and revenue, as more tasks are executed. The revenue under the SAG remains a lot higher than the FPS, as a low ratio corresponds to a high competition over the servers. Consequently the servers receive a higher Shapley Value and the SAG revenue increases. Further, the FPS "filters out" a lot of tasks with a low worth, while the tasks with high worth only pay the threshold $\bar{f}$. When the demand ratio increases, the revenues of the SAG and FPS converge. The FPS has a larger pool of tasks to filter from and the competition for servers decreases, which results in a lower revenue from the SAG. Overall SAG outperforms FPS, maximizing social welfare and increasing revenue, especially when competition over servers is high. As competition decreases, SAG still achieves a higher social welfare, but its advantage in terms of server revenue decreases (and for some parameters and low competition, revenues may even be higher under FPS).

4. CONCLUSION

We introduced the Shared Assignment Game, where sellers are constrained in their resources in a two-sided market, using the Shapley value to fairly share costs. Under *private information*, such as in cloud environments, we showed that agents can manipulate outcome using "splits" and "bluffs". Our bounds on these manipulations show that the SAG is somewhat robust to them. Finally, our simulation shows that SAG outperforms the fixed-price scheme, maximizing social welfare and increasing revenue, especially under high competition for servers. Future research directions include investigating the performance of our approach on real-world cloud datasets, examining manipulations in games where agents are unreliable (see for example [11, 8, 16]), investigating coalitional manipulations and collusion (see [41, 5, 12, 7]) and examining manipulations when task owners adaptively learn from experience. We also wish to apply the Shapley value to new upcoming allocation methods, e.g., [28] and compare the allocated welfare to all agents, with other existing pricing mechanisms, e.g., VCG and [17].

5. REFERENCES

[1] V. Abhishek, I. A. Kash, and P. Key. Fixed and market pricing for cloud services. In *Proceedings of INFOCOM Workshops*, 2012.

[2] L. M. Ausubel and P. Milgrom. The lovely but lonely vickrey auction. In *Combinatorial Auctions, chapter 1*. MIT Press, 2006.

[3] H. Aziz, Y. Bachrach, E. Elkind, and M. Paterson. False-name manipulations in weighted voting games. *JAIR*, 40(1), 2011.

[4] H. Aziz and B. de Keijzer. Housing markets with indifferences: A tale of two mechanisms. In *AAAI'12*.

[5] Y. Bachrach. Honor among thieves: collusion in multi-unit auctions. In *AAMAS*, 2010.

[6] Y. Bachrach. The least-core of threshold network flow games. In *MFCS*. 2011.

[7] Y. Bachrach, E. Elkind, and P. Faliszewski. Coalitional voting manipulation: A game-theoretic perspective. In *AAAI*, 2011.

[8] Y. Bachrach, I. Kash, and N. Shah. Agent failures in totally balanced games and convex games. In *WINE*. 2012.

[9] Y. Bachrach, P. Kohli, and T. Graepel. Rip-off: playing the cooperative negotiation game. *AAMAS*, 2011.

[10] Y. Bachrach, E. Markakis, E. Resnick, A. D. Procaccia, J. S. Rosenschein, and A. Saberi. Approximating power indices: theoretical and empirical analysis. *JAAMAS*, 2010.

[11] Y. Bachrach, R. Meir, M. Feldman, and M. Tennenholtz. Solving cooperative reliability games. *UAI*, 2011.

[12] Y. Bachrach, M. Polukarov, and N. R. Jennings. The good, the bad and the cautious: safety level cooperative games. In *WINE*. 2010.

[13] Y. Bachrach and E. Porat. Path disruption games. In *AAMAS*, 2010.

[14] Y. Bachrach, E. Porat, and J. S. Rosenschein. Sharing rewards in cooperative connectivity games. *JAIR*, 2012.

[15] Y. Bachrach and J. S. Rosenschein. Power in threshold network flow games. *JAAMAS*, 2009.

[16] Y. Bachrach and N. Shah. Reliability weighted voting games. In *SAGT*. 2013.

[17] H. Ballani, P. Costa, T. Karagiannis, and A. Rowstron. The price is right: towards location independent costs in datacenters. In *HotNets-X*, 2011.

[18] G. Blocq and A. Orda. How good is bargained routing? In *Proceedings of INFOCOM*, 2012.

[19] S. Brânzei, T. Michalak, T. Rahwan, K. Larson, and N. R. Jennings. Matchings with externalities and attitudes. In *AAMAS '13*.

[20] R. Buyya, C. Yeo, S. Venugopal, J. Broberg, and I. Brandic. Cloud computing and emerging IT platforms: Vision, hype, and reality for delivering computing as the 5th utility. *FGCS*, 25(6), 2009.

[21] R. Cavallo. Fairness and welfare through redistribution when utility is transferable. In *AAAI'12*.

[22] C. Chekuri and S. Khanna. A PTAS for the multiple knapsack problem. In *SODA'00*.

[23] V. Conitzer and T. Sandholm. Computing shapley values, manipulating value division schemes, and checking core membership in multi-issue domains. In *AAAI'04*.

[24] J. Dean and S. Ghemawat. Mapreduce: Simplified data processing on large clusters. *Comm. ACM*, 2008.

[25] P. Dubey. On the uniqueness of the shapley value. *IJGT*, 4(3), 1975.

[26] P. Dubey. The shapley value as aircraft landing fees–revisited. *Management Science*, 28(8), 1982.

[27] E. Elkind, L. A. Goldberg, P. Goldberg, and M. Wooldridge. Computational complexity of weighted threshold games. In *AAAI*, 2007.

[28] A. Ghodsi, M. Zaharia, B. Hindman, A. Konwinski, S. Shenker, and I. Stoica. Dominant resource fairness: fair allocation of multiple resource types. In *NSDI'11*.

[29] F. Gul. Bargaining foundations of shapely value. *Econometrica*, 1989.

[30] F. Hubert and S. Ikonnikova. Strategic investment and bargaining power in supply chains: A Shapley value analysis of the Eurasian gas market. 2003.

[31] D. Jaume, J. Masso, and A. Neme. The multiple partners assignment game with heterogeneous sales and multi-unit demands: competitive equilibria. *MMOR*, 2012.

[32] H. Kellerer, U. Pferschy, and D. Pisinger. *Knapsack Problems*. 2004.

[33] R. Meir, Y. Bachrach, and J. Rosenschein. Minimal subsidies in expense sharing games. *SAGT*, 2010.

[34] R. Myerson. *Game Theory, Analysis of Conflict*. Harvard U. Press, 1991.

[35] L. Petrosjan and G. Zaccour. Time-consistent shapley value allocation of pollution cost reduction. *JEDC*, 2003.

[36] A. Rey and J. Rothe. Merging and splitting for power indices in weighted voting games and network flow games on hypergraphs. In *STAIRS*, 2010.

[37] L. Shapley and M. Shubik. The assignment game I: The core. *IJGT*, 1, 1971.

[38] L. S. Shapley. A value for n-person games. *In Contributions to the Theory of Games*, II, 1953.

[39] Y. Song, C. Zhang, and Y. Fang. Multiple multidimensional knapsack problem and its applications in cognitive radio networks. In *MILCOM'08*.

[40] M. Sotomayor. *Majumdar M (ed.) Equilibrium and dynamics: Essays in Honor of David Gale Macmillian*, pages 322–336, 1992.

[41] K. Takamiya. Coalition strategy-proofness and monotonicity in shapley–scarf housing markets. *Mathematical Social Sciences*, 41(2):201–213, 2001.

[42] W. Thomson. The manipulability of the shapley-value. *IJGT*, 17(2), 1988.

[43] R. van den Brink. An axiomatization of the Shapley value using a fairness property. *IJGT*, 2002.

[44] E. Winter. The Shapley value. *Handbook of Game Theory with economic applications*, 3, 2002.

[45] L. Xia, M. Zuckerman, A. D. Procaccia, V. Conitzer, and J. S. Rosenschein. Complexity of unweighted coalitional manipulation under some common voting rules. In *IJCAI'09*.

[46] M. Yokoo, V. Conitzer, T. Sandholm, N. Ohta, and A. Iwasaki. Coalitional games in open anonymous environments. In *AAAI'05 (2)*.