On the Payoff Mechanism in Peer-Assisted Services with Multiple Content Providers

Jeong-woo Cho and Yung Yi

Dept. of Electrical Engineering, KAIST, South Korea
{ggumdol,yiyung}@kaist.ac.kr

Abstract

This paper studies an incentive structure for cooperation and its stability in peer-assisted services when there exists multiple content providers, using a coalition game theoretic approach. We first consider a generalized coalition structure consisting of multiple providers with many assisting peers, where peers assist providers to reduce the operational cost in content distribution. To distribute the profit from cost reduction to players (*i.e.*, providers and peers), we then establish a generalized formula for individual payoffs when a "Shapley-like" payoff mechanism is adopted. We show that the grand coalition is *unstable*, even when the operational cost functions are concave, which is in sharp contrast to the recently studied case of a single provider where the grand coalition is stable. We also show that irrespective of stability of the grand coalition, there always exist coalition structures which are not convergent to the grand coalition. Our results give us an important insight that a provider does not tend to cooperate with other providers in peer-assisted services, and be separated from them. To further study the case of the separated providers, three examples are presented; *(i)* Each peer is underpaid than his due payoff, *(ii)* a service monopoly is possible, and *(iii)* the peer payoffs based on the Shapley-like mechanism exhibit even oscillatory behaviors. Analytical studies and examples in this paper open many new questions such as realistic and efficient incentive structures and the tradeoffs between fairness and individual providers' competition in peer-assisted services.

I. Introduction

The Internet is becoming more content-oriented, and cost-effective and scalable distribution of contents has been the central role of the Internet. Uncoordinated peer-to-peer (P2P) systems, e.g., BitTorrent, has been successful in distributing contents, but the rights of the content owners are not protected well, and most of the P2P contents are in fact illegal. In its response, a new type of service, called *peer-assisted services*, has received significant attentions these days. In peer-assisted services, users commit a part of their resources to assist content providers in content distribution with objective of enjoying both scalability/efficiency in P2P systems and controllability in client-server systems. Examples of application of peer-assisted services include nano data center [1] and IPTV [2], where high potential of operational cost reduction was observed. However, it is clear that most users will not just "donate" their resources to content providers. Thus, the key factor to the success of peer-assisted services is how to (economically) incentivize users to commit their valuable resources and participate in the service.

One of nice mathematical tools to study incentive-compatibility of peer-assisted services is the coalition game theory which covers how payoffs should be distributed and whether such a payoff scheme can be executed by rational individuals or not. In peer-assisted services, the "symbiosis" between providers and peers are sustained when *(i)* the offered payoff scheme guarantees fair assessment of players' contribution under a provider-peer coalition and *(ii)* each individual has no incentive to exit from the coalition. In the coalition game theory, the notions of Shapley value and the core have been popularly applied to address *(i)* and *(ii)*, respectively, when the entire players cooperate, referred to as the grand coalition. A recent paper by Misra *et al.* [3] demonstrates that the Shapley value approach is a promising payoff mechanism to provide right incentives for cooperation in a *single-provider* peer-assisted service under mild assumptions.

However, in practice, the Internet consists of multiple content providers, even if only giant providers are counted. The focus of our paper is to study the cooperation incentives for *multiple* providers. In the multi-provider case, the model clearly becomes more complex, thus even classical analysis adopted in the single-provider case becomes much more challenging, and moreover the results and their implications may experience drastic changes. To motivate further, see an example in Fig. 1 with two providers (Google TV and iTunes) and consider two cases of cooperation: *(i) separated*, where there exists a fixed partition of peers for each provider, and *(ii) coalescent*, where each peer is possible to assist any provider[1]. In the separated case,

[1]We exclude the case when a peer assists both providers.

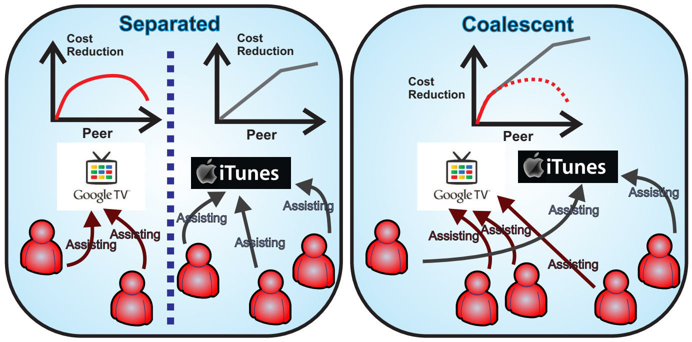

Fig. 1. Coalition Structures for a Dual-Provider Network.

a candidate payoff scheme is based on the Shapley value in each separated coalition. Similarly, in the coalescent case, the Shapley value is also a candidate payoff scheme after the worth function of the grand coalition N (the player set) is defined appropriately. A reasonable definition of the worth function can be the total cost reduction generated by N which is maximized over all combinations of peer partitions to each provider. Then, it is not hard to see that the cost reduction for the coalescent case exceeds that for the separated case, unless the two partitions are equivalent in both cases. This implies that at least one individual in the separated case is *underpaid* than in the coalescent case under the Shapley-value based payoff mechanism. Thus, providers and users are recommended to form the grand coalition and be paid off based on the Shapley value, *i.e.*, the due desert.

However, it is still questionable whether peers will stay in the grand coalition and thus the consequent Shapley-value based payoff mechanism is desirable in the multi-provider setting. In this paper, we anatomize incentive structures in peer-assisted services with multiple content providers and focus on stability issues from two different angles: stability at equilibrium of Shapley value and convergence to the equilibrium.

Our main contributions are summarized as follows:

1) We first provide a closed-form formula of the Shapley value for a general case of multiple providers and peers. To that end, we define a worth function to be a maximum total cost reduction over all possible peer partitions to each provider. Due to the intractability of analytical computation of the Shapley value, we take a fluid-limit approximation that assumes a large number of peers and re-scales the system with the number of peers. This is a non-trivial generalization of the Shapley value for the single-provider case in [3]. In fact, our formula in Theorem 1 establishes the general Shapley value for distinguished *multiple* atomic players and infinitesimal players in the context of the Aumann-Shapley (A-S) prices [4] in coalition game theory.
2) We prove that the Shapley value for the multi-provider case is not in the core under mild conditions, *e.g.*, each provider's cost function is concave. This is in stark contrast to the single-provider case where the concave cost function stabilizes the equilibrium. We also show that, irrespective of stability of the grand coalition, there always exist initial states which are not convergent to the equilibrium. An interesting fact from this part of study is that peers and providers have opposite cooperative preferences, *i.e.*, peers prefer to cooperate with more providers, whereas providers prefer to be separated from other providers.

The insight that our results provide us is the impossibility of cooperation in peer-assisted services with multiple providers. In conjunction with the main contributions mentioned above, we conclude the paper by presenting three examples for non-cooperation among providers: *(i)* the peers are underpaid than the Shapley payoff, *(ii)* a provider with more "advantageous" cost function monopolizes all peers, and *(iii)* Shapley value for each coalition gives rise to an oscillatory behavior of coalition structures. These examples suggest that the system with the separated providers may be unstable as well as unfairness in a peer-assisted service market.

II. Preliminaries

Since this paper deals with multiple content providers and thus a peer can choose any provider to assist, we define a coalition game with a partition (coalition structure), and introduce the payoff mechanisms there.

A. Game with Coalition Structure

A game with coalition structure is a triple $(N, v, \mathcal{P})$ where N is a player set and $v : 2^N \to \mathbb{R}$ (2^N is the set of all subsets of N) is a worth function, $v(\emptyset) = 0$. $v(K)$ is called the worth of a coalition $K \subseteq N$. $\mathcal{P}$ is called a *coalition structure* for (N, v); it is a partition of N where $\mathcal{P}(i) \subset \mathcal{P}$ denotes the coalition containing player i. The *grand coalition* is the partition $\mathcal{P} = \{N\}$. For instance[2], a partition of $N = \{1, 2, 3, 4, 5\}$ is $\mathcal{P} = \{\{1, 2\}, \{3, 4, 5\}\}$, $\mathcal{P}(4) = \{3, 4, 5\}$, and the grand coalition is $\mathcal{P} = \{\{1, 2, 3, 4, 5\}\}$. $\mathcal{P}(N)$ is the set of all partitions of N. For notational simplicity, a game *without* coalition structure $(N, v, \{N\})$ is denoted by (N, v). A value of player i is an operator $\varphi_i(N, v, \mathcal{P})$ that assigns a payoff to player i.

To conduct the equilibrium analysis of coalition games, the notion of *core* has been extensively used to study the stability of the grand coalition $\mathcal{P} = \{N\}$:

Definition 1 (Core) The core is defined as $\{\varphi(N, v) \mid \sum_{i\in N} \varphi_i(N, v) = v(N)$ and $\sum_{i\in K} \varphi_i(N, v) \geq v(K), \forall K \subseteq N\}$.

If a payoff vector $\varphi(N, v)$ lies in the core, no player in N has an incentive to split off to form another coalition K because the worth of the coalition K, $v(K)$, is no more than the payoff sum $\sum_{i\in K} \varphi_i(N, v)$. Note that the definition of the core hypothesizes that the grand coalition is already formed *ex-ante*. We can see the core as an analog of Nash equilibrium from noncooperative games. If a payoff vector $\varphi_N(N, v)$ lies in the core, then the grand coalition is stable with respect to any collusion to break the grand coalition.

B. Shapley Value and Aumann-Drèze Value

On the premise that the player set is not partitioned, *i.e.*, $\mathcal{P} = \{N\}$, the Shapley value is popularly used as a fair distribution of the grand coalition's worth to individual players, defined by:

$$\varphi_i(N, v) = \sum_{S\subseteq N\setminus\{i\}} \frac{|S|!(|N| - |S| - 1)!}{|N|!} \left(v(S \cup \{i\}) - v(S)\right). \tag{1}$$

Shapley [5] gives the following interpretation: "*(i)* Starting with a single member, the coalition adds one player at a time until everybody has been admitted. *(ii)* The order in which players are to join is determined by chance, with all arrangements equally probable. *(iii)* Each player, on his admission, demands and is promised the amount which his adherence contributes to the value of the coalition." The Shapley value quantifies the above that is axiomized (see [5] for the details of the axioms) and has been treated as a worth distribution scheme. The beauty of the Shapley value lies in that the payoff "summarizes" in *one* number all the possibilities of each player's contribution in every coalition structure.

Given a coalition structure $\mathcal{P} \neq \{N\}$, one can obtain the Aumann-Drèze value (A-D value) [6] of player i by taking $\mathcal{P}(i)$, which is the coalition containing player i, to be the player set and by computing the Shapley value of player i of the *reduced* game $(\mathcal{P}(i), v_{\mathcal{P}(i)})$. It is easy to see that the A-D value can be construed as a direct extension of the Shapley value to a game with coalition structure.

III. Coalition Game in Peer-Assisted Services

In this section, we first define a coalition game in a peer-assisted service with multiple content providers by classifying the types of coalition structures as *separated*, where a coalition includes only one provider, and *coalescent*, where a coalition is allowed to include more than one providers (see Fig. 1). Due to the coalition independence of the A-D value, in order to decide the payoffs of a game with a general coalition structure $\mathcal{P}$, it suffices to decide the payoffs of players within each coalition, say $C \in \mathcal{P}$, without considering other coalitions $C \in \mathcal{P}$, $C \neq \mathcal{P}(i)$. We refer the readers to [7] for the details on why it suffices to consider just the two cases. To define the coalition game, we will define a worth function of an arbitrary coalition $S \subseteq N$ for such two cases. The key message of this section is that the rational behavior of the providers makes the

[2] A player i is an *element* of a coalition $C = \mathcal{P}(i)$, which is in turn an *element* of a partition $\mathcal{P}$. Lastly, $\mathcal{P}$ is an element of $\mathcal{P}(N)$ *while* a subset of 2^N.

Shapley value approach *unworkable* because the major premise of the Shapley value, the grand coalition, is not formed in the multi-provider games.

A. Worth Function in Peer-Assisted Services

Assume that players N are divided into two sets, the set of content providers $Z := \{p_1, \cdots, p_\zeta\}$, and the set of peers $H := \{n_1, \cdots, n_\eta\}$., *i.e.*, $N = Z \cup H$. We also assume that the peers are homogeneous, *e.g.*, the same computing powers, disk cache sizes, and upload bandwidths. Later, we discuss that our results can be readily extended to nonhomogeneous peers. The set of peers assisting providers is denoted by $\bar{H} := \{n_1, \cdots, n_{x\cdot\eta}\}$ where $x = |\bar{H}|/\eta$, *i.e.*, the fraction of assisting peers. We define the worth of a coalition S to be the amount of cost reduction due to distribution of the contents with the players in S in both separated and coalescent cases.

Separated case: Denote by $\Omega_p^\eta(x)$ the operational cost of a provider p when the coalition S consists of η peers and x fraction of assisting peers. Since the operational cost cannot be negative, we assume $\Omega_p^\eta(x) > 0$. Note that from the homogeneity assumption of peers, the cost function depends only on the number of peers and the fraction of assisting peers. Then, we define the worth function $\hat{v}(S)$ for the coalition S as:

$$\hat{v}(S) := \Omega_p^\eta(0) - \Omega_p^\eta(x) \tag{2}$$

where $\Omega_p^\eta(0)$ corresponds to the cost when there are no assisting peers.

Coalescent case: In contrast to the separated case, where a coalition includes a single provider, the worth for the coalescent case is not clear yet, since depending on which peers assist which providers the amount of cost reduction may differ. One of reasonable definitions would be the maximum worth out of all peer partitions, *i.e.*, the worth for the coalescent case is defined by:

$$v(S) = \max\left\{\textstyle\sum_{C\in\mathcal{P}} \hat{v}(C) \;\middle|\; \mathcal{P} \in \mathcal{P}(S) \text{ such that } |Z \cap C| = 1 \text{ for all } C \in \mathcal{P}\right\}. \tag{3}$$

The definition above implies that we *view* a coalition containing more than one provider as the most productive coalition whose worth is *maximized* by choosing the optimal partition $\mathcal{P}^*$ among all possible partitions of S. Note that (3) is consistent with the definition (2) for $|Z \cap S| \le 1$, *i.e.*, $v(S) = \hat{v}(S)$ for $|Z \cap S| \le 1$.

Three remarks are in order. First, as opposed to [3] where $\hat{v}(\{p\}) = \eta R - \Omega_p^\eta(0)$ (R is the subscription fee paid by any peer), we simply assume that $\hat{v}(\{p\}) = 0$. Note that, as discussed in [8, Chapter 2.2.1], it is no loss of generality to assume that, initially, each provider has earned no money. In our context, this means that it does not matter how much fraction of peers is subscribing to each provider because each peer has already paid the subscription fee to providers *ex-ante*.

Second, it is also important to note that we cannot always assume that $\Omega_p^\eta(x)$ is monotonically decreasing because providers have to pay the electricity expense of the computers and the maintenance cost of the hard disks of assisting peers. For example, a recent study [9] found that Annualized Failure Rate (AFR) of hard disk drives is over 8.6% for three-year old ones. We discuss in Appendix of [7] that, if we consider a more *intelligent* coalition, the worth function is always non-increasing. However, to simplify the exposition, we assume in this paper that $\Omega_p^\eta(x)$ is non-increasing in x for all $p \in Z$.

Third, the worth function in peer-assisted services can reflect the diversity of peers. It is not difficult to extend our result to the case where peers belong to distinct classes. For example, peers my be distinguished by different upload bandwidths and different hard disk cache sizes. A point at issue for the multiple provider case is whether peers who are *not* subscribing to the content of a provider may be allowed to assist the provider or not. On the assumption that the content is ciphered and not decipherable by the peers who do not know its password which is given only to the subscribers, providers will allow those peers to assist the content distribution. Otherwise, we can easily reflect this issue by dividing the peers into a number of classes where each class is a set of peers subscribing to a certain content.

B. Fluid Aumann-Drèze Value for Multiple-Provider Coalitions

So far we have defined the worth of coalitions. Now let us *distribute* the worth to the players for a given coalition structure $\mathcal{P}$. Recall that the payoffs of players in a coalition are independent from other coalitions by the definition of A-D payoff. Pick a coalition C without loss of generality, and denote the set of providers in C by $\bar{Z} \in Z$. With slight notational abuse, the set of peers assisting $\bar{Z}$ is denoted by $\bar{H}$. Once we find the A-D payoff for a coalition consisting of arbitrary provider set $\bar{Z} \in Z$ and assisting peer set $\bar{H} \in H$, the payoffs for the separated and coalescent cases in Fig. 1 follow from the substitutions $\bar{Z} = Z$ and $\bar{Z} = \{p\}$, respectively. In light of our discussion in Section II-B, it is more reasonable to call a payoff mechanism 'A-D payoff' and 'Shapley payoff' respectively for the partitioned and non-partitioned games $(N, v, \{\bar{Z} \cup \bar{H}, \cdots\})$ and $(N, v, \{Z \cup H\})$[3].

Fluid Limit: We adopt the limit axioms for large population of users to overcome the computational hardness of the A-D payoffs:

$$\lim_{\eta \to \infty} \widetilde{\Omega}_p^\eta(\cdot) = \widetilde{\Omega}_p(\cdot) \text{ where } \widetilde{\Omega}_p^\eta(\cdot) = \tfrac{1}{\eta}\Omega_p^\eta(\cdot) \tag{4}$$

which is the asymptotic operational cost per peer in a very large system. We drop superscript η from notations to denote their limits as $\eta \to \infty$. From the assumption $\Omega_p^\eta(x) > 0$, we have $\widetilde{\Omega}_p(x) \geq 0$. To avoid trivial cases, we also assume $\widetilde{\Omega}_p(x)$ is not constant in the interval $x \in [0, 1]$ for any $p \in Z$. We also introduce the payoff of each provider per user, defined as $\widetilde{\varphi}_p^\eta := \frac{1}{\eta}\varphi_p^\eta$. We now derive the fluid limit equations of the payoffs which can be obtained as $\eta \to \infty$. The proof of the following theorem is given in Appendix of [7].

Theorem 1 (A-D Payoff for Multiple Providers) As η tends to infinity, the A-D payoffs of providers and peers under an arbitrary coalition $C = \bar{Z} \cup \bar{H}$ converge to the following equation:

$$\begin{cases} \widetilde{\varphi}_p^{\bar{Z}}(x) = \widetilde{\Omega}_p(0) - \sum_{S \subseteq \bar{Z} \setminus \{p\}} \int_0^1 u^{|S|}(1-u)^{|\bar{Z}|-1-|S|} \left(M_\Omega^{S \cup \{p\}}(ux) - M_\Omega^S(ux) \right) \mathrm{d}u, & \text{for } p \in \bar{Z} \\ \widetilde{\varphi}_n^{\bar{Z}}(x) = -\sum_{S \subseteq \bar{Z}} \int_0^1 u^{|S|}(1-u)^{|\bar{Z}|-|S|} \frac{\mathrm{d}M_\Omega^S}{\mathrm{d}x}(ux)\mathrm{d}u, & \text{for } n \in \bar{H}. \end{cases} \tag{5}$$

Here $M_\Omega^S(x) := \min \left\{ \sum_{i \in S} \widetilde{\Omega}_i(y_i) \;\middle|\; \sum_{i \in S} y_i \leq x,\; y_i \geq 0 \right\}$ and $M_\Omega^\emptyset(x) := 0$. Note that $M_\Omega^{\{p\}}(x) = \widetilde{\Omega}_p(x)$.

The following corollary is immediate as a special case of Theorem 1, which we will use in Section IV.

Corollary 1 (A-D Payoff for Dual Providers) As η tends to infinity, the A-D payoffs of providers and peers who belong to a dual-provider coalition, *i.e.*, $\bar{Z} = \{p, q\}$, converge to the following equation:

$$\begin{cases} \widetilde{\varphi}_p^{\{p,q\}}(x) = \widetilde{\Omega}_p(0) - \int_0^1 u M_\Omega^{\{p,q\}}(ux)\mathrm{d}u - \int_0^1 (1-u) M_\Omega^{\{p\}}(ux)\mathrm{d}u + \int_0^1 u M_\Omega^{\{q\}}(ux)\mathrm{d}u, & (p,\ q \text{ are interchangeable}) \\ \widetilde{\varphi}_n^{\{p,q\}}(x) = -\int_0^1 u^2 \frac{\mathrm{d}M_\Omega^{\{p,q\}}}{\mathrm{d}x}(ux)\mathrm{d}u - \sum_{i \in \{p,q\}} \int_0^1 u(1-u) \frac{\mathrm{d}M_\Omega^{\{i\}}}{\mathrm{d}x}(ux)\mathrm{d}u, & \text{for } n \in \bar{H}. \end{cases} \tag{6}$$

Note that our A-D payoff formula in Theorem 1 generalizes the formula in Misra *et al.* [3, Theorem 4.3] (*i.e.*, $|Z| = 1$). It also establishes the A-D values for distinguished *multiple* atomic players (the providers) and infinitesimal players (the peers), in the context of the Aumann-Shapley (A-S) prices [4] in coalition game theory.

C. Stability of the Grand Coalition

The stability guarantee of a payoff vector has been an important topic in coalition game theory. For the single provider case, $|Z| = 1$, it was shown in [3, Theorem 4.2] that, if the cost function is decreasing and concave, the Shapley incentive structure lies in the core of the game. What if for $|Z| \geq 2$? Is the grand coalition stable for the multiple provider case? Before answering this question, we need the following definition.

Definition 2 (Noncontributing Provider) A provider $p \in Z$ is called *noncontributing* if $M_\Omega^Z(1) - M_\Omega^{Z \setminus \{p\}}(1) = \widetilde{\Omega}_p(0)$.

To understand this better, note that the above expression is equivalent to the following:

$$\left(\sum_{i \in Z} \widetilde{\Omega}_i(0) - M_\Omega^Z(1) \right) - \left(\sum_{i \in Z \setminus \{p\}} \widetilde{\Omega}_i(0) - M_\Omega^{Z \setminus \{p\}}(1) \right) = 0 \tag{7}$$

which implies that there is no difference in the total cost reductions irrespective of whether the provider p is in the provider set or not. Interestingly, if all cost functions are concave, there exists at least one noncontributing provider.

Lemma 1 Suppose $|Z| \geq 2$. If $\widetilde{\Omega}_p(\cdot)$ is concave for all $p \in Z$, there exists a noncontributing provider.

[3]On the contrary, the term 'Shapley payoff' was used in [3] to refer to the payoff for the game $(N, v, \{\bar{Z} \cup \bar{H}, \cdots\})$ where a proper subset of the peer set assists the content distribution.

To prove this, recall the definition of $M^Z_\Omega(\cdot)$:

$$M^Z_\Omega(x) = \min_{y \in Y(x)} \sum_{i \in \bar{Z}} \widetilde{\Omega}_i(y_i) \quad \text{where } Y(x) := \{(y_1, \cdots, y_{|\bar{Z}|}) \mid \sum_{i \in \bar{Z}} y_i \leq x,\ y_i \geq 0\}.$$

Since the summation of concave functions is concave and the minimum of a concave function over a convex feasible region $Y(x)$ is an *extreme* point of $Y(x)$ as shown in [10, Theorem 3.4.7], we can see that the solutions of the above minimization are the extreme points of $\{(y_1, \cdots, y_{|Z|}) \mid \sum_{i \in \bar{Z}} y_i \leq x,\ y_i \geq 0\}$ which has at least one $p \in Z$ such that $y_p = 0$ if $|Z| \geq 2$.

We are ready to state the following theorem, a direct consequence of Theorem 1. The proof is given in Appendix of [7].

Theorem 2 (Shapley Payoff for Multiple Providers Not in the Core) If there exists a noncontributing provider, the Shapley payoff for the game $(Z \cup H, v)$ does not lie in the core.

It follows from Lemma 1 that, if all operational cost functions are concave and $|Z| \geq 2$, the Shapley payoff does not lie in the core, either. This result appears to be in good agreement with our usual intuition. If there is a provider who does not contribute to the coalition at all in the sense of (7) and is still being paid due to her potential for imaginary contribution assessed by the Shapley formula (1), which is not actually exploited in the current coalition, other players will agree to expel her to improve their payoffs. The condition $|Z| \geq 2$ plays an essential role in the theorem. For $|Z| \geq 2$, the concavity of the cost functions leads to the Shapley value not lying in the core, whereas, for the case $|Z| = 1$, the concavity of the cost function is proven to make the Shapley incentive structure lie in the core [3, Theorem 4.2].

D. Convergence to the Grand Coalition

The notion of the core lends itself to the stability analysis of the grand coalition *on the assumption* that the players are already in the equilibrium, *i.e.*, the grand coalition. Theorem 2 raises a point open to further discussion due to the assumption of concave cost function, *e.g.*, for the cost functions that are not concave, it is possible that the Shapley value is in the core. We present that such cases are unlikely to occur by showing that the grand coalition is not a global attractor under some conditions. To study the convergence of a game with coalition structure to the grand coalition, we define the stability of a game with coalition structure.

Definition 3 (Stable Coalition Structure [11]) We say that a coalition structure $\mathcal{P}'$ blocks $\mathcal{P}$, where $\mathcal{P}'$, $\mathcal{P} \in \mathcal{P}(N)$, with respect to φ if and only if there exists some $C \in \mathcal{P}'$ such that $\varphi_i(N, v, \{C, \cdots\}) > \varphi_i(N, v, \mathcal{P})$ for all $i \in C$. In this case, we also say that C blocks $\mathcal{P}$. If there does not exist any $\mathcal{P}'$ which blocks $\mathcal{P}$, $\mathcal{P}$ is called *stable*.

It is also important to note that, due to the coalition independence of the A-D value, all stability notions defined by Hart and Kurz [11] coincide with the above simplistic definition.

The above definition can be intuitively interpreted that, if there exists any subset of players C who improve their payoffs away from the current coalition structure, they *will* form a new coalition C. In other words, if a coalition structure $\mathcal{P}$ has any blocking coalition C, some rational players will break $\mathcal{P}$ to increase their payoffs. Unsurprisingly, if a payoff vector lies in the core, the grand coalition is stable in the above sense. This reminds us that the core is about the stability of a particular equilibrium, *i.e.*, the grand coalition. The basic premise here is that players are not clairvoyant, *i.e.*, they are interested only in improving their instant payoffs.

Theorem 3 (A-D Payoff for Multiple Providers Does Not Lead to the Grand Coalition) Suppose $|Z| \geq 2$ and $\widetilde{\Omega}_p(y)$ is not constant in the interval $y \in [0, x]$ for any $p \in Z$ where $x = |\bar{H}|/|H|$. The followings hold for all $p \in Z$ and $n \in \bar{H}$.

- The A-D payoff for provider p in coalition $\{p\} \cup \bar{H}$ is larger than that in all coalition $T \cup \bar{H}$ for $\{p\} \subsetneq T \subseteq Z$.
- The A-D payoff of peer n in coalition $\{p\} \cup \bar{H}$ is smaller than that in all coalition $T \cup \bar{H}$ for $\{p\} \subsetneq T \subseteq Z$.

In plain words, a provider, who is in cooperation with a peer set, will receive the highest dividend when she cooperates only with the peers excluding other providers whereas each peer wants to cooperate with as many as possible providers.

It is surprising that, for the multiple provider case, *i.e.*, $|Z| \geq 2$, each provider benefits from forming the single-provider coalition *whether* the cost function is concave *or not*. There is no *positive* incentives for providers to cooperate with each other under the implementation of A-D payoffs. On the contrary, a peer always looses by leaving the grand coalition.

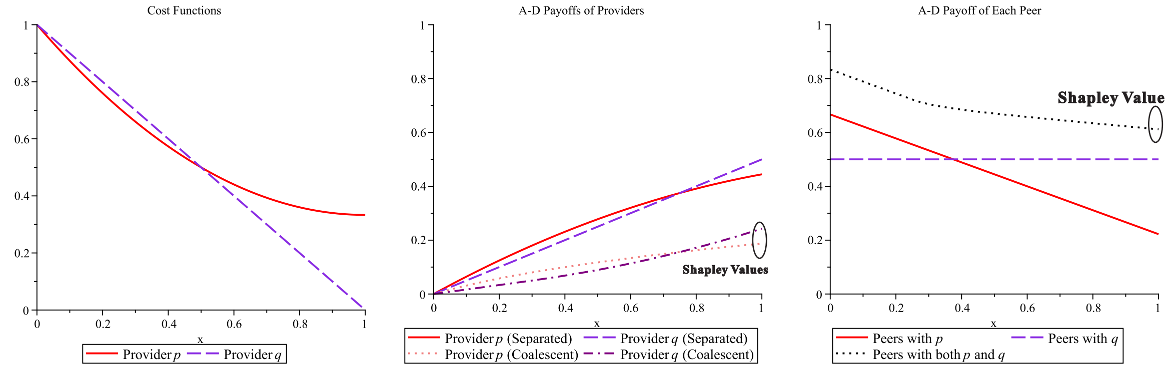

Fig. 2. Example 1: A-D Payoffs of Two Providers and Peers for Convex Cost Functions.

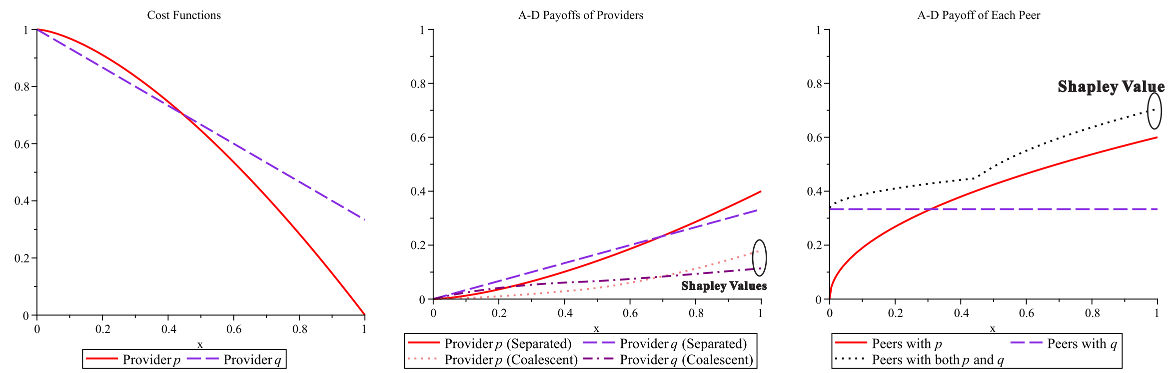

Fig. 3. Example 2: A-D Payoffs of Two Providers and Peers for Concave Cost Functions.

Upon the condition that each provider begins with a single-provider coalition with a large number of peers, one cannot reach the grand coalition because those single-provider coalitions are already *stable* in the sense of the stability in Definition 3. That is, the grand coalition is not the global attractor.

IV. A Critique of the A-D Payoff for Separate Providers

The discussion so far has centered on the stability of the grand coalition. The result in Theorem 2 suggests that if there is a noncontributing (free-riding) provider, the grand coalition will be broken. The situation is aggravated by Theorem 3, stating that single-provider coalitions will persist if providers are rational. In this section, on the major premise that the providers are separated, we illustrate weak points of the A-D payoff with a few representative examples.

Example 1 (Unfairness) Suppose that there are two providers, *i.e.*, $Z = \{p, q\}$, with $\widetilde{\Omega}_p(x) = 2(x-1)^2/3 + 1/3$ and $\widetilde{\Omega}_q(x) = 1 - x$, both of which are decreasing and *convex*. All values are shown in Fig. 2 as functions of x. In line with Theorem 3, providers are paid more than their Shapley values, whereas peers are paid less than theirs. We can see that each peer n will be paid $2/3$ ($\widetilde{\varphi}_n^{\{p\}}(0)$) when he is contained by the coalition $\{p, n\}$ and the payoff decreases with the number of peers in this coalition. On the other hand, provider p wants to be assisted by as many peers as possible because $\widetilde{\varphi}_p^{\{p\}}(x)$ is increasing in x. If it is possible for n to prevent other peers from joining the coalition, he can get $2/3$. However, it is more likely that no peer can kick out other peers. Thus, p will be assisted by $x = 3/8$ fraction of peers, which is the unique solution of $\widetilde{\varphi}_n^{\{p\}}(x) = \widetilde{\varphi}_n^{\{q\}}(x)$ while q assisted by $1 - x = 5/8$ fraction of peers.

Example 2 (Monopoly) Consider a two-provider system $Z = \{p, q\}$ with $\widetilde{\Omega}_p(x) = 1 - x^{3/2}$ and $\widetilde{\Omega}_q(x) = 1 - 2x/3$, both of which are decreasing and *concave*. All values including the Shapley values are shown in Fig. 3. Not to mention unfairness in line with Theorem 3, provider p *monopolizes* the whole peer-assisted services. No provider has an incentive to cooperate with other provider and each peer has to choose between the two providers. It can be seen that all peer will assist provider p because $\widetilde{\varphi}_n^{\{p\}}(x) > \widetilde{\varphi}_n^{\{q\}}(x)$ for $x > 25/81$. Appealing to Definition 3, if the providers are initially separated, the coalition

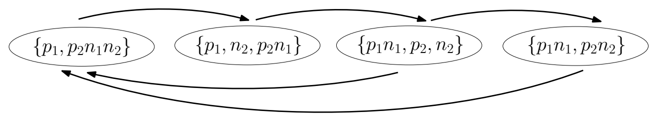

Fig. 4. Example 3: A-D Payoff Mechanism Leads to Oscillatory Coalition Structure.

structure will converge to the service monopoly by p. In line with Lemma 1 and Theorem 2, even if the grand coalition is supposed to be the initial condition, it is not stable in the sense of the core. The noncontributing provider (Definition 2) in this example is q.

Example 3 (Oscillation) Consider a game with two providers and two peers where $N = \{p_1, p_2, n_1, n_2\}$. If $\{n_1\}$, $\{n_2\}$ and $\{n_1, n_2\}$ assist the content distribution of p_1, the reduction of the distribution cost is respectively 10\$, 9\$ and 11\$ per month. However, the hard disk maintenance cost incurred from a peer is 5\$. In the meantime, if $\{n_1\}$, $\{n_2\}$ and $\{n_1, n_2\}$ assist the content distribution of p_2, the reduction of the distribution cost is respectively 6\$, 3\$ and 13\$ per month. In this case, the hard disk maintenance cost incurred from a peer is supposed to be 2\$ due to smaller contents of p_2 as opposed to those of p_1. We can compute the *net* cost reduction for all possible coalitions. For example, if n_1 and n_2 help p_1, the coalition worth becomes $v(\{n_1, n_2, p_1\}) = 11\$ - 5\$ - 5\$ = 1\$$.

Since it is very tedious to compute the A-D payoffs for all coalition structures and to determine their stability, we refer to Appendix of [7] for a detailed analysis. For notational simplicity, we adopt a simplified expression for coalitional structure $\mathcal{P}$: A coalition $\{a, b, c\} \in \mathcal{P}$ is denoted by abc and each singleton set $\{i\}$ is denoted by i. We first observe that the Shapley payoff of this example does not lie in the core. As time tends to infinity, the A-D payoff exhibits an oscillation of the partition $\mathcal{P}$ consisting of the four recurrent coalition structures as shown in Fig. 4. As of now, from the-state-of-the-art in the literature on this behavior, it is not yet clear how this behavior will be developed in large-scale systems.

V. Concluding Remarks

A quote from an interview of BBC iPlayer with CNET UK [12]: "*Some people didn't like their upload bandwidth being used. It was clearly a concern for us, and we want to make sure that everyone is happy, unequivocally, using iPlayer.*"

In this paper, we have studied whether the Shapley incentive structure in peer-assisted services would be in conflict with the pursuit of profits by rational content providers and peers. A lesson from our analysis is summarized as: Even though it is righteous to pay peers more because they become relatively more useful as the number of peer-assisted services increases, the content providers will not admit that peers should receive their due deserts. The providers tend to persist in single-provider coalitions. In the sense of the classical stability notion, called 'core', the cooperation would have been broken even if we had begun with the grand coalition as the initial condition. Lastly, we have illustrated yet another problems when we use the Shapley-like incentive for the exclusive single-provider coalitions. These results suggest that the profit-sharing system, Shapley value, and hence its fairness axioms, are not compatible with the selfishness of the content providers. We believe that a realistic incentive structure in peer-assisted services should reflect a trade-off between fairness and competition among individuals.

References

[1] V. Valancius, N. Laoutaris, L. Massoulié, C. Diot, and P. Rodriguez, "Greening the Internet with Nano Data Centers," in *Proc. ACM CoNEXT*, Dec. 2009.
[2] M. Cha, , P. Rodriguez, S. Moon, and J. Crowcroft, "On next-generation telco-managed P2P TV architectures," in *Proc. USENIX IPTPS*, Feb. 2008.
[3] V. Misra, S. Ioannidis, A. Chaintreau, and L. Massoulié, "Incentivizing peer-assisted services: A fluid Shapley value approach," in *Proc. ACM Sigmetrics*, Jun. 2010.
[4] R. Aumann and L. Shapley, *Values of Non-Atomic Games*. Princeton University Press, 1974.
[5] L. Shapley, *A Value for n-Person Games*. In H. W. Kuhn and A. W. Tucker, editors, Contribution to the Theory of Games II, vol. 28 of Annals of Mathematics Studies, Princeton University Press, 1953.
[6] R. Aumann and J. Drèze, "Cooperative games with coalition structures," *International Journal of Game Theory*, vol. 3, pp. 217–237, 1974.
[7] J. Cho and Y. Yi, "On the payoff mechanism in peer-assisted services with multiple content providers," *Technical Report. Available at http://www.q2s.ntnu.no/~jeongwoo/data/multi_p2p_gamenets.pdf*, Dec. 2010.
[8] B. Peleg and P. Sudhölter, *Introduction to the Theory of Cooperative Games*, 2nd ed. Springer-Verlag, 2007.
[9] E. Pinheiro, W. Weber, and L. A. Barroso, "Failure trends in a large disk drive population," in *Proc. USENIX FAST*, Feb. 2007.
[10] M. S. Bazaraa, H. D. Sherali, and C. M. Shetty, *Nonlinear Programming: Theory and Algorithms*, 2nd ed. John Wiley & Sons Inc., 1993.
[11] S. Hart and M. Kurz, "Endogenous formation of coalitions," *Econometrica*, vol. 51, pp. 1047–1064, 1983.
[12] N. Lanxon, "iPlayer uncovered: What powers the BBC's epic creation?" *CNET UK*, May 2009.