

Incentivizing Peer-Assisted Services: A Fluid Shapley Value Approach

Vishal Misra^{*}
Dept. of Computer Science,
Columbia University,
New York, NY, USA
misra@cs.columbia.edu

Stratis Ioannidis, Augustin Chaintreau,
Laurent Massoulié
Technicolor, Paris, France
firstname.lastname@technicolor.com

ABSTRACT

A new generation of content delivery networks for live streaming, video on demand, and software updates takes advantage of a peer-to-peer architecture to reduce their operating cost. In contrast with previous uncoordinated peer-to-peer schemes, users opt-in to dedicate part of the resources they own to help the content delivery, in exchange for receiving the same service at a reduced price. Such incentive mechanisms are appealing, as they simplify coordination and accounting. However, they also increase a user's expectation that she will receive a fair price for the resources she provides. Addressing this issue carefully is critical in ensuring that all interested parties—including the provider—are willing to participate in such a system, thereby guaranteeing its stability.

In this paper, we take a cooperative game theory approach to identify the ideal incentive structure that follows the axioms formulated by Lloyd Shapley. This ensures that each player, be it the provider or a peer, receives an amount proportional to its contribution and bargaining power when entering the game. In general, the drawback of this ideal incentive structure is its computational complexity. However, we prove that as the number of peers receiving the service becomes large, the Shapley value received by each player approaches a fluid limit. This limit follows a simple closed form expression and can be computed in several scenarios of interest: by applying our technique, we show that several peer-assisted services, deployed on both wired and wireless networks, can benefit from important cost and energy savings with a proper incentive structure that follows simple compensation rules.

^{*}This work was performed while the author was visiting Technicolor, Paris.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGMETRICS'10, June 14–18, 2010, New York, New York, USA.
Copyright 2010 ACM 978-1-4503-0038-4/10/06 ...\$10.00.

Categories and Subject Descriptors

C.2.4 [Distributed Systems]: Distributed Applications;
K.6.2 [Installation Management]: Pricing and resource allocation

General Terms

Theory, Economics

Keywords

Cooperative Game Theory, Incentive Mechanisms

1. INTRODUCTION

Peer to peer networks (P2P) have long been demonstrated to be technically superior to client-server architectures in terms of scalability. With the widespread adoption of broadband access at the residential level, P2P networks promise even larger scalability while maintaining the same quality of service as dedicated data-centers or content delivery networks [9]. Under the radar, the success of BitTorrent makes it probably the world's largest (virtual) content delivery network, transmitting primarily illegal content, albeit scalably and efficiently. Many commercial streaming systems employ P2P networks like Coolstreaming [28], PPLive, UUSEE, PP-Stream¹, *etc.* This trend is not likely to reverse. ISPs and content owners on the other hand are in a constant cat and mouse game with P2P users, trying to disrupt, throttle and shut down P2P traffic. The primary reason for this is loss of revenue for content rights owners. Measures like port blocking, deep packet inspection, application identification, *etc.*, are circumvented by counter measures like random port selection and packet encryption. Our view of this situation is that it requires an *economic rethink* rather than an engineering battle.

Peer-assisted services [5, 12, 24, 25] were introduced recently as an alternative to illegal P2P networks. In such systems, users commit part of their resources (*e.g.*, upload bandwidth and storage capacity) to assist a service provider in the distribution of content. By taking advantage of peer-to-peer scalability to reduce the provider's operating cost, such services can be offered at a significantly reduced price. This ensures that content owners and ISPs can then directly compete with illegal content distributors by enhancing user experience though added features (*e.g.*, ease of use, content authentication, *etc.*). The cost reductions achieved by peer

¹see www.pplive.com, www.uusee.com, www.ppstream.com

assistance, which are reflected on the low price of the service, are significant. For example, a content provider may utilize the set-top boxes (STBs) installed in residential locations as a semi-permanent P2P infrastructure to serve content to users [12, 25]. Studies have shown that this architecture, by replacing monolithic data centers with distributed (nano) data-centers, can result in important savings (60-80%) for the provider in terms of energy costs [25].

However, the key challenge in the deployment of a peer-assisted service is maintaining user participation, *i.e.*, ensuring that users contribute the resources under their control. For example, users can easily throttle their upload bandwidth, disconnect their device whenever they do not actively use it, or simply neglect to maintain it in an operational status. As opposed to current illegal P2P networks, which are based at least partially on altruistic behavior [7], users of a peer-assisted service will require appropriate *incentives* to contribute their resources. Such incentives can be implemented by offering a reduced price only to users that actively contribute their resources. However, such price reductions should be carefully designed: the provider’s respective decrease in revenue should be properly offset by the cost savings achieved by user participation. If not, deploying the peer-assisted service would become unprofitable for the provider. To that end, *the goal of this paper is to develop an economic framework for Peer-Assisted services that creates the right incentives for both users as well as providers to participate.*

Based on the concepts of cooperative games, we design an incentive scheme based on *Shapley values* that achieves the above goal. Our scheme has a wide range of applications, as it allows us to capture the cost reduction achieved by peer assistance at the service provider through a generic cost function. One of the key features of this incentive scheme is that it ensures participation by applying the principle of *balanced contribution*: both the service provider and users share any generated revenues according to the value that they add to the system through their participation.

Our contributions in the paper are the following:

- We show that, under a general condition, both peers and providers have an incentive to participate when the total revenue is shared according to the Shapley value (Section 4, Theorems 4.1 and 4.2).
- Computational complexity has long been the bane of incentive schemes based on Shapley values: in general, the complexity of computing a user’s Shapley value may grow exponentially in the number of participants. We show that the axioms describing the Shapley value of each participant admit fluid a limit as the system gets larger. This allows us to obtain a simple closed form of the Shapley value for a large user population (Section 4, Theorem 4.3).
- The above results are obtained for a generic cost function; we illustrate how the Shapley value depends on this function’s curvature. Moreover, we extend our results to incorporate the existence of several *atomic* players, which capture the special roles played, *e.g.*, by the ISP and/or the content owner (Section 5).
- We apply our model to various scenarios, including file dissemination, live streaming and video on demand. Our analysis shows that a common principle, where

users receive half of their contribution (credited as “upload miles”) to the system efficiency, guarantees user participation in all these cases. We explain the intuition behind this principle, how it extends to the case of a third-party content provider and/or network neutrality, and how it applies to global bandwidth and energy saving in both wired and wireless contexts. (Section 6).

Our results considerably depart from previously proposed incentive schemes, as we adopt a principled approach, building on well-known concepts from the field of economic theory. Though similar revenue sharing principles have been introduced before for IP-routing and cooperative settlement between ISPs [14], to the best of our knowledge, the effect of peer assistance on cost reductions has never been studied before. Moreover, by exploiting the asymptotic behavior of large populations, our paper demonstrates that Shapley values can be incorporated in such systems without suffering from computational complexity restrictions.

2. RELATED WORK

Peer-assisted services and specifically P2P content distribution have attracted a lot of attention recently. Compared to traditional client-server VoD systems [1], a P2P-based VoD solution is less costly and more scalable [9]. In the context of provider managed P2P VoD systems, there are few related papers. [12] proposed a P2P architecture for set-top boxes. [1] empirically evaluated the benefits of using the storage of set-top boxes in a P2P fashion. [24] designed an architecture for pushing content into boxes, and analyzed optimal placement strategies. Empirical evaluation of such an architecture for IPTV is provided in [5]. The NaDa project [25] is postulating enormous energy savings in a peer-assisted architecture for VoD delivery as compared to a traditional monolithic datacenter or CDN approach.

To the best of our knowledge, this is the first paper on designing economic incentive mechanisms for peer-assisted services. The Shapley value concept has been previously studied in the context of network economics in [13–16]. The concept of fluid or continuous Shapley values for non-atomic games was introduced in [2], and was proposed as a basis for instance for sharing telephone call costs at Cornell University in [3] or allocating costs for transportation problems [6].

3. BACKGROUND ON SHAPLEY VALUE

Here, we briefly introduce the concept of Shapley value and its use under our incentive structure context.

3.1 Cooperative games

We consider a set of players denoted as \mathcal{N} . We denote by $N = |\mathcal{N}|$ the number of players in this set. We call any nonempty subset $\mathcal{S} \subseteq \mathcal{N}$ a *coalition* of players. For each coalition \mathcal{S} , we denote by $V(\mathcal{S})$ the *worth function*, which measures the total revenue produced by the service when all players of this coalition \mathcal{S} are active. Let $P_i(\mathcal{S})$ denote the profit (monetary payment received minus operating cost) of player i in the coalition \mathcal{S} , we then have:

$$V(\mathcal{S}) = \sum_{i \in \mathcal{S}} P_i(\mathcal{S}). \quad (1)$$

The contribution of each player to a coalition can be defined as follows.

DEFINITION 3.1. *The marginal contribution of player i to a coalition $\mathcal{S} \subseteq \mathcal{N} \setminus \{i\}$ is defined as $\Delta_i(V, \mathcal{S}) = V(\mathcal{S} \cup \{i\}) - V(\mathcal{S})$.*

Note that the contribution of a player only depends on the worth function V . We introduce two more definitions.

DEFINITION 3.2. *A worth function is superadditive if*

$$V(\mathcal{S} \cup \mathcal{T}) \geq V(\mathcal{S}) + V(\mathcal{T}), \text{ for all } \mathcal{S}, \mathcal{T} \subseteq \mathcal{N} \text{ s.t. } \mathcal{S} \cap \mathcal{T} = \emptyset.$$

I.e., the worth of any two disjoint coalitions is no greater than the worth of their union. In some sense, this indicates that such coalitions are “better off” forming a larger coalition.

DEFINITION 3.3. *A worth function is supermodular if*

$$V(\mathcal{S} \cup \{i\}) - V(\mathcal{S}) \leq V(\mathcal{T} \cup \{i\}) - V(\mathcal{T}), \forall \mathcal{S} \subseteq \mathcal{T} \subseteq \mathcal{N} \setminus \{i\}, \forall i \in \mathcal{N}.$$

A supermodular function exhibits the property of *increasing returns*: a player entering a larger coalition brings “more value” than a player entering a smaller coalition. Note that supermodularity implies superadditivity (though the converse is not true).

3.2 Shapley value: axioms and definition

The Shapley value, originally proposed by Lloyd Shapley [22], serves as an appropriate mechanism for players to share revenues inside a given coalition. In particular, given a cooperative game with worth function V and a coalition \mathcal{S} , the Shapley value determines how the total worth of the coalition, captured by $V(\mathcal{S})$, should be shared among the players in \mathcal{S} .

More specifically, for each player i and coalition \mathcal{S} , the Shapley value of player i is denoted by

$$\varphi_i(\mathcal{S}, V)$$

and is uniquely defined by the following three axioms

Axiom 1 (EFFICIENCY). $\sum_{i \in \mathcal{S}} \varphi_i(\mathcal{S}, V) = V(\mathcal{S})$.

Axiom 2 (SYMMETRY). *If for all $\mathcal{S}' \subseteq \mathcal{S} \setminus \{i, j\}$,*
 $V(\mathcal{S}' \cup \{i\}) = V(\mathcal{S}' \cup \{j\})$

then $\varphi_i(\mathcal{S}, V) = \varphi_j(\mathcal{S}, V)$.

Axiom 3 (FAIRNESS/BALANCED CONTRIBUTION). *For any $i, j \in \mathcal{S}$, j 's contribution to i equals i 's contribution to j , or, in other words*

$$\varphi_i(\mathcal{S}, V) - \varphi_i(\mathcal{S} \setminus \{j\}, V) = \varphi_j(\mathcal{S}, V) - \varphi_j(\mathcal{S} \setminus \{i\}, V).$$

The efficiency axiom states that the total revenue assigned to each player equals the actual profit created by their coalition. In other words, the mechanism does not contribute or receive extra profit. The symmetry axiom requires that if two players contribute the same to every subset of other players, they should receive the same amount of revenue. Finally, the balanced contribution axiom addresses the fairness between any pair of players. It may be illustrated on a two-player system where $\mathcal{N} = \{1, 2\}$. By efficiency we have that, for a coalition of a single player, $\varphi_i(\{i\}, V) = V(\{i\})$. The fairness axiom states that the gain (or loss) of revenue from cooperation, as seen by player 1 and 2, should be the same $\varphi_1(\mathcal{N}, V) - V(\{1\}) = \varphi_2(\mathcal{N}, V) - V(\{2\})$. In that

case it means that the global gain of cooperation, defined as $V(\mathcal{N}) - V(\{1\}) - V(\{2\})$, is split evenly among players. The balanced contribution axiom preserves and generalizes this egalitarian property [17].

Based on the axioms above, one can show that the *Shapley value* φ can be computed as follows [22].

$$\forall i \in \mathcal{S}, \varphi_i(\mathcal{S}, v) = \frac{1}{|\mathcal{S}|!} \sum_{\pi \in \Pi} \Delta_i(v, S(\pi, i)) \quad (2)$$

where Π is the set of all $|\mathcal{S}|!$ orderings of \mathcal{S} and $S(\pi, i)$ is the set of players preceding i in the ordering π . The Shapley value of player i can thus be interpreted as the *expected* marginal contribution $\Delta_i(V, \mathcal{S}')$, where \mathcal{S}' is the set of players in \mathcal{S} preceding i in a uniformly distributed random ordering of \mathcal{S} .

The Shapley value was originally derived axiomatically by Shapley using the axioms of *efficiency*, *symmetry*, *dummy* and *additivity*. In 1977 Myerson [21] replaced the axioms of dummy and additivity by the above axiom of *balanced contribution* (also known as *fairness*) and showed that together with efficiency and symmetry it was enough to uniquely determine the Shapley value, the other two axioms following as properties.

3.3 Pros and cons of the Shapley value

The Shapley value exhibits the following property:

LEMMA 1. *If V is superadditive, then the Shapley value is individually rational, i.e., for all $\mathcal{S} \subset \mathcal{N}$,*

$$\varphi_i(\mathcal{S}, V) \geq V(\{i\}) \quad \forall i \in \mathcal{S}.$$

Intuitively, individual rationality implies that no user has an incentive to abandon a coalition: the return it accrues through the Shapley value exceeds the individual profit that it would gain by abandoning \mathcal{S} . In this sense, individual rationality guarantees the stability of a coalition.

A stronger statement holds if the worth function is supermodular.

LEMMA 2. *If V is supermodular, then, the Shapley value lies in the core of V , i.e.,*

$$\sum_{i \in \mathcal{S}} \varphi_i(\mathcal{N}, V) \geq V(\mathcal{S}), \forall \mathcal{S} \subseteq \mathcal{N}. \quad (3)$$

In other words, if V is supermodular, no given subset of players has an incentive to leave the “grand coalition” \mathcal{N} and form a smaller coalition. This implies that the “most stable” coalition is, in fact, the grand coalition.

In general, the core of a game consists of any payoff scheme that satisfies (3) and the efficiency axiom. An interesting fact is that the Shapley value in fact lies in the *middle* (center of gravity, to be precise [23]) of the core for a supermodular (also known as convex) game. This points to another property of the Shapley value, namely that it is *robust* if the game is convex. Deviations from the Shapley value of any solution, due to imperfect measurements or computations are still likely to remain in the core, whereas other solutions might be marginally stable, and can easily deviate into unstable versions.

Recently, there has been a lot of interest in the application of Shapley values to issues related to Internet Economics. In [14] the authors applied the Shapley value concept to the settlement issue between Content, Eyeball and Transit ISPs.

One of the interesting insights from the analysis was that the current settlement structure of the Internet (customer provider and zero-dollar peering) is potentially unstable, and new settlements (reverse customer provider and paid peering) are needed for stable coalitions. This is a result of the evolution of the Internet from a symmetric entity in terms of traffic to one where there are well defined content providers (Google/Yahoo/EBay/Microsoft etc.) and consumers (end customers), and the authors demonstrated the current settlement structure is far away from the Shapley values, and hence potentially outside the core.

However, the calculation of the Shapley value through (2) involves an exponential time complexity. This has restricted its use mostly as a reference for theoretical interest, and limited its appeal as an implementable revenue distribution mechanism. In this paper we present a method based on fluid approximation for large player population, which drastically reduces the complexity of computing the Shapley value. We show that its computation only involves one or a small number of atomic players and the collective effect of peer-assistance for peers belonging to different classes. This makes it feasible to indeed use the Shapley value directly to incentivize peer-assistance.

4. INCENTIVE FOR PEER-ASSISTANCE

In this section, we first describe a model for peer-assistance as a cooperative game. This model, which focuses on provider cost reduction through peer-assistance, allows to obtain the Shapley value of any user in a simple closed form. As we show in the next sections, this model can be extended and used to understand several scenarios of practical interest.

4.1 A peer-assisted service

Peer-assisted services may be seen as cooperative games.

Players.

We assume here that all the profits and costs in this system are incurred by a single player, called the *provider*, which we denoted by P . Other cases where profit and cost are shared among several players are discussed in Section 5.3. The game contains other players which are the users of the services. Without loss of generality we assume that users belong to m distinct classes, inside which all users are statistically identical. As an example, users may form a flat network and be distinguished only by different upload bandwidths. The set of all users is denoted by \mathcal{U} , with cardinality N , so that the global set containing $(N + 1)$ players is $\mathcal{N} = \{P\} \cup \mathcal{U}$.

Revenue, operational cost.

Assuming the provider decides to enter the game, a service is received by all users of this network, for which they all pay a flat rate R to the provider. The flat rate R can be thought of either as a monthly subscription rate, or as a per content price, say for Video on Demand or a live streaming event. However, the provider incurs an operating cost C^N , which depends on the number of users N and depends on the set of users who decide to join the coalition, as described below.

Coalition, worth function.

As opposed to a regular service, this service allows peer-assistance, which reduces the operational cost C^N . This

works as follows: all users of the service are offered the opportunity to opt-in for a “peer-assistance option”, where their resources can be used by the system to serve other users. This has two consequences: for the provider, these additional resources reduce the value of C^N needed to serve these N users; for the user, a part of this reduction is reflected in a price reduction for the service it receives. Below, we describe how the value of this incentive is set.

Thus, a *coalition* may consist of the provider (if it decides to enter the game), as well as the set of users $\mathcal{U}_a \subseteq \mathcal{U}$ who decide to “opt-in” for this option. Denote by (N_1, \dots, N_m) number of users in each class which sign on for this option, and define the vector containing the respective fractions of users opting in as

$$\bar{X} = \left(\frac{N_1}{N}, \dots, \frac{N_m}{N} \right).$$

Then, we make the following assumption

ASSUMPTION 1. *The operational cost of the provider $C^N : \mathbb{R}^m \rightarrow \mathbb{R}$ is a differentiable function of \bar{X} , i.e., $C^N = C^N(\bar{X})$.*

In other words, the operational cost does not depend on the actual set of users participating in the coalition; instead, it is a function of the *fraction* of participants from each class.

Under this assumption, if the provider enters a coalition $\mathcal{S} = \{P\} \cup \mathcal{U}_a$, the *worth* of the coalition is:

$$V(\{P\} \cup \mathcal{U}_a) = NR - C^N(\bar{X}). \quad (4)$$

On the other hand, the worth of any coalition that *does not include the provider* is 0. This means that the provider is a *veto* player: if it has no incentive to stay in a coalition and departs, the coalition no longer has any value.

DEFINITION 4.1. *We denote by $\varphi_P^N(\bar{X})$ (resp. $\varphi_i^N(\bar{X})$) the Shapley value of the provider (resp. of a user, in class i , that allows peer-assistance), for the coalition $\{P\} \cup \mathcal{U}_a$.*

When no user allows peer-assistance the value of the coalition $\mathcal{S} = \{P\}$ is given by

$$V(\{P\}) = NR - C^N(\bar{0}).$$

In this case, this value is also the Shapley value (or profit) of the provider, since it is the only player in the coalition, and no other player receives any share of the profit.

When a subset of the users allow peer-assistance, the operational cost goes down to $C^N(\bar{X})$, which results in an additional profit for the provider. We propose a simple incentive scheme for the provider and the users to share the total worth of a coalition: all players in a coalition (including, potentially, the service provider) receive their associated Shapley value as revenue. For simple users, this is essentially a subsidy, given to them for the additional profit they bring to the service provider by assisting it. They receive this as an incentive to enter the coalition and assist the provider. Similarly, the Shapley value of the provider is the revenue it accrues, which also incentivizes it to remain in the coalition.

This can be implemented through a “discount price”, where user i pays $R - \varphi_i^N(\bar{X})$ (instead of R) to the provider to receive the same service. Note that if the Shapley value of a user exceeds the value of the flat rate R , this user may even be able to make a net profit; this may be the case if the resources it makes available drastically reduce the cost of the whole system.

The goal of this paper is to show that this incentive mechanism is the right one for both the provider and the users, and to determine *precisely* how the value of the incentive depends on operational costs.

4.2 Main results

Individual rationality and grand coalition stability.

We first prove formally that under a general assumption on the operational cost, namely that it decreases with peer-assistance, our proposed incentive system based on Shapley values is *individually rational*. That is, for any coalition of the form $\mathcal{S} = \{P\} \cup \mathcal{U}_a$, neither the users in \mathcal{U}_a nor the provider P have an incentive to leave the coalition.

THEOREM 4.1. *If the cost function $C^N(\cdot)$ is monotonically decreasing in all of its coordinates, then the incentive structure suggested is individually rational.*

PROOF. It suffices by Lemma 1 to show that V is super-additive, *i.e.*,

$$V(\mathcal{S} \cup \mathcal{T}) \geq V(\mathcal{S}) + V(\mathcal{T}), \text{ for all } \mathcal{S}, \mathcal{T} \subseteq \mathcal{N} \text{ s.t. } \mathcal{S} \cap \mathcal{T} = \emptyset.$$

This is trivially true if neither \mathcal{S} nor \mathcal{T} contain the provider, as then $V(\mathcal{S}) = V(\mathcal{T}) = V(\mathcal{S} \cup \mathcal{T}) = 0$. Otherwise, as $\mathcal{S} \cap \mathcal{T} = \emptyset$, at most one of them will contain the provider, say \mathcal{S} . Then, $V(\mathcal{T}) = 0$ and $V(\mathcal{S} \cup \mathcal{T}) \geq V(\mathcal{S})$ by the monotonicity of C^N , so superadditivity follows. \square

By placing an additional constraint on C^N we can show that the Shapley value is *in the core*, *i.e.*, the “grand coalition” made up by the provider and *all* players is the “most stable” coalition.

THEOREM 4.2. *Assume that the cost function $C^N(\cdot)$ is monotonically decreasing, twice differentiable, and*

$$\frac{\partial^2 C^N}{\partial x_i \partial x_j} \leq 0, \text{ for all classes } i, j. \quad (5)$$

Then the incentive structure suggested lies in the core of the game.

PROOF. By Lemma 2, it suffices to show that V is super-modular, *i.e.*,

$$V(\mathcal{S} \cup \{j\}) - V(\mathcal{S}) \leq V(\mathcal{T} \cup \{j\}) - V(\mathcal{T}), \forall \mathcal{S} \subseteq \mathcal{T} \subseteq \mathcal{M} \setminus \{j\}, \forall i \in \mathcal{N}.$$

The above is trivially true by the monotonicity of C^N if j is the provider P . It is also trivially true, again by the monotonicity of C^N , if $P \in \mathcal{T}$ but $P \notin \mathcal{S}$. We therefore focus in the case where $P \in \mathcal{S}$.

In this case, let $\mathcal{S} = \{P\} \cup \mathcal{U}_a$ and $\mathcal{T} = \{P\} \cup \mathcal{U}'_a$, where $\mathcal{U}_a \subseteq \mathcal{U}'_a$, and denote by \bar{X} and \bar{X}' the fraction vectors corresponding to \mathcal{U}_a and \mathcal{U}'_a , respectively. Then $\bar{X} \leq \bar{X}'$, coordinate-wise.

Assume now that j belongs to the i -th class. Let \bar{e}_i the vector with a 1 at the i -th column and 0 everywhere else. Then,

$$\begin{aligned} V(\mathcal{S} \cup \{j\}) - V(\mathcal{S}) &= -C^N(\bar{X} + \bar{e}_i \frac{1}{N}) + C^N(\bar{X}) \\ &= - \int_{x_i}^{x_i + \frac{1}{N}} \frac{\partial C^N(s, \bar{X}^{-i})}{\partial x_i} ds \end{aligned} \quad (6)$$

where \bar{X}^{-i} is \bar{X} excluding the i -th coordinate. Similarly,

$$\begin{aligned} V(\mathcal{T} \cup \{j\}) - V(\mathcal{T}) &= -C^N(\bar{X}' + \bar{e}_i \frac{1}{N}) + C^N(\bar{X}') \\ &= - \int_{x'_i}^{x'_i + \frac{1}{N}} \frac{\partial C^N(s, \bar{X}'^{-i})}{\partial x_i} ds \end{aligned}$$

Observe that (5) implies that $\frac{\partial C^N(\cdot)}{\partial x_i}$ is decreasing in all its coordinates. Since $[s - (x'_i - x_i), \bar{X}^{-i}] \leq [s, \bar{X}'^{-i}]$ coordinate-wise, we have that

$$\begin{aligned} V(\mathcal{T} \cup \{j\}) - V(\mathcal{T}) &\geq \int_{x'_i}^{x'_i + \frac{1}{N}} \frac{\partial C^N(s - (x'_i - x_i), \bar{X}^{-i})}{\partial x_i} ds \\ &\quad \stackrel{s' = s - x'_i + x_i}{=} \int_{x_i}^{x_i + \frac{1}{N}} \frac{\partial C^N(s', \bar{X}^{-i})}{\partial x_i} ds' \end{aligned}$$

and the theorem follows from (6). \square

These two results establish that, assuming that the cost function C^N is monotone, using the Shapley value provides the right incentives for both the provider and all users to not leave a coalition. With the additional property (5) we can guarantee that the “grand coalition” $\{P\} \cup \mathcal{U}$ is the most stable among all coalitions

In what follows, we prove that if the number of users playing the peer-assistance game becomes large, then the Shapley value approaches a fluid limit. Interestingly, this fluid limit involves two kind of players: the atomic player (*i.e.*, the provider), and infinitesimal players. The impact of the decision of a single infinitesimal player is negligible; nonetheless, the collective effect of their behavior determines the Shapley value of each player.

Limit axioms for a large population of users.

We assume that average cost per user converges to a smooth function of \bar{X} , as N goes to infinity. In other words, we have

$$\lim_{N \rightarrow \infty} \tilde{C}^N(\bar{X}) = C(\bar{X}) \text{ where } \tilde{C}^N(\bar{X}) = \frac{C^N(\bar{X})}{N}.$$

$C(\bar{X})$ can be thought of as the asymptotic operational cost per user served in a very large system where, for each class of users, the fraction of peer-assisting users is fixed. We also introduce the Shapley value of the Provider per user, defined as $\tilde{\varphi}_P^N(\bar{X}) = \frac{1}{N} \varphi_P^N(\bar{X})$. The efficiency axiom can then be rewritten as:

$$N \cdot \tilde{\varphi}_P^N(\bar{X}) + \sum_{i=1}^m N \cdot X_i \varphi_i^N(\bar{X}) = N \cdot R - N \cdot \tilde{C}^N(\bar{X}).$$

while the balanced contribution states, for any i, j

$$\begin{cases} \varphi_i^N(\bar{X}) - \varphi_i^N(\bar{X} - \frac{1}{N} \cdot e_j) = \varphi_j^N(\bar{X}) - \varphi_j^N(\bar{X} - \frac{1}{N} \cdot e_i), \\ \text{and } \varphi_i^N(\bar{X}) = N(\tilde{\varphi}_P^N(\bar{X}) - \tilde{\varphi}_P(\bar{X} - \frac{1}{N} \cdot e_i)), \end{cases}$$

where e_i denotes the vector with 1 at the i -th coordinate and 0 everywhere else.

As the number of users becomes large, the Shapley values φ_i and $\tilde{\varphi}_P$ converge to a smooth function of \bar{X} that should hence satisfy:

$$\begin{cases} \tilde{\varphi}_P(\bar{X}) + \sum_{i=1}^m X_i \varphi_i(\bar{X}) = R - C(\bar{X}), \\ \forall i, j, \frac{\partial \varphi_i}{\partial x_j} = \frac{\partial \varphi_j}{\partial x_i}, \text{ and } \varphi_i(\bar{X}) = \frac{\partial \tilde{\varphi}_P}{\partial x_i}. \end{cases} \quad (7)$$

Solution of the limit axioms.

As we show below, the limit axioms lead to a simple expression of the Shapley value, which greatly simplifies its computation for a given case. We give the general expression for any cost function C . All examples presented later in this paper will use this formula, for a cost function motivated by a specific application.

THEOREM 4.3. *As N tends to infinity, the provider's (per user) Shapley value and the Shapley values of users in all classes converge to the following unique solution of Eq.(7):*

$$\begin{cases} \tilde{\varphi}_P(\bar{X}) = R - \int_0^1 C(s \cdot \bar{X}) ds \\ \forall i, \quad \varphi_i(\bar{X}) = - \int_0^1 s \frac{\partial C}{\partial x_i}(s \bar{X}) ds. \end{cases}$$

PROOF. By the axiom of balanced contribution, we can replace $\varphi_i(\bar{X})$ by $\frac{\partial \tilde{\varphi}_P}{\partial x_i}$ in the first line of (7) which implies

$$\forall \bar{X}, \quad \tilde{\varphi}_P(\bar{X}) + \sum_{i=1}^m X_i \frac{\partial \tilde{\varphi}_P}{\partial x_i}(\bar{X}) = R - C(\bar{X}). \quad (8)$$

The key observation is then that the second term of the LHS is the derivative of a *diagonalized* version of the function $\tilde{\varphi}_P$. More precisely, for any given \bar{X} if we define $\psi_P : s \mapsto \tilde{\varphi}_P(s \cdot \bar{X})$, this function of a real variable is differentiable and we have: $\psi'_P(s) = \sum_{i=1}^m X_i \frac{\partial \tilde{\varphi}_P}{\partial x_i}(s \cdot \bar{X})$. We can then obtain the value of ψ_P and $\tilde{\varphi}_P$ by solving a simple ordinary differential equation of dimension 1.

The following lemma generalizes the above observation:

LEMMA 3. *Let $F, G : \mathbb{R}^m \rightarrow \mathbb{R}$, and $\alpha > 0$*

$$\forall \bar{X}, \quad \alpha \cdot F(\bar{X}) + \sum_{i=1}^m X_i \cdot \frac{\partial F}{\partial x_i}(\bar{X}) = G(\bar{X}).$$

is equivalent to $\forall \bar{X}, F(\bar{X}) = \int_0^1 s^{\alpha-1} G(s \cdot \bar{X}) ds$.

Before proving this lemma, let us observe that it directly implies the Theorem, since we can apply the lemma with $\alpha = 1$ to Eq.(8) and obtain the expression for $\tilde{\varphi}_P$, which is uniquely characterized. A simple derivation w.r.t. x_i then provides the expression for φ_i , which again is uniquely characterized. We can then immediately observe that the balanced contribution among nodes of different class is satisfied, so that this unique solution satisfy all limit axioms (7).

We now prove the lemma: The first condition may be rewritten

$$\forall \bar{X}, \forall s, \quad \alpha \cdot F(s \cdot \bar{X}) + s \cdot \sum_{i=1}^m X_i \cdot \frac{\partial F}{\partial x_i}(s \cdot \bar{X}) = G(s \cdot \bar{X}).$$

Note the additional term in the second term of the LHS, due to the multiplication.

If we introduce, for a given \bar{X} , the function $\psi_F : s \mapsto F(s \cdot \bar{X})$ the equation above may be written as

$$\forall s, \quad \alpha \cdot \psi_F(s) + s \cdot \psi'_F(s) = G(s \cdot \bar{X}),$$

or, equivalently, after multiplying by $s^{\alpha-1}$ each side

$$\forall s, \quad \alpha s^{\alpha-1} \cdot \psi_F(s) + s^\alpha \cdot \psi'_F(s) = s^{\alpha-1} G(s \cdot \bar{X}),$$

The LHS can be recognized as the derivative of the function $s \mapsto s^\alpha \cdot \psi_F(s)$. This function takes value 0 for $s = 0$. The following equality of the derivative is then equivalent to:

$$\forall s, \quad s^\alpha \psi'_F(s) = \int_0^s u^{\alpha-1} G(u \cdot \bar{X}) du.$$

$$\forall X, \forall s, \quad s^\alpha F(s \cdot \bar{X}) = \int_0^s u^{\alpha-1} G(u \cdot \bar{X}) du.$$

which is equivalent to the second condition of the theorem. \square

Remarks

- Our fluid Shapley value results are of similar flavor to the Aumann-Shapley (A-S) prices [2]. The latter are derived for non-atomic games, whereas our problem has a distinguished “atomic” player, the provider. The A-S prices are a special case of our formula for the infinitesimal player in Theorem 4.3, when there is no distinguished atomic player. The only difference is that the term “ s ” multiplying the partial derivative inside the integral disappears.
- The A-S formula is also called the “diagonal” formula, and an intuitive interpretation of the Aumann-Shapley (A-S) value is the following: assume that the vector \bar{X} share in an homogeneous way, starting from 0 and ending at \bar{X} . Suppose also that along the above sharing process each time a “small” proportion (an “infinitesimal” one) of \bar{X} start sharing, the i -th class user produces a marginal cost benefit. Then the average cost benefit per unit of the i th class once \bar{X} are fully sharing will be its A-S value. Hence, the Aumann-Shapley value is computed by calculating the integral of marginal cost benefit for class i along the diagonal of the vector from 0 to \bar{X} . Our formula for the infinitesimal user can be interpreted in the context of the A-S formula in the following way: for every value of s along the diagonal $[0, s\bar{X}]$, the probability that the atomic (veto) player is part of the coalition is simply s . In the absence of the veto player, the value of the coalition and hence the A-S value is 0, in the presence of the veto player the value is the A-S value. Thus, our formula for the infinitesimal user can be interpreted as the expected A-S value, conditioned on the presence/absence of the veto player.
- Note that we are also able to obtain a much simpler and compact derivation of our general formula as compared to the Aumann-Shapley [2] result or Hart [8] that looked at mixed atomic-continuous games. This is because we exploit the axiom of *balanced contribution* extensively to obtain partial derivatives of various Shapley values directly and obtain differential equations that can be easily solved. The axiom of balanced contribution was introduced a few years *after* [21] the results of Aumann-Shapley and Hart were derived.

5. SPECIAL CASES & EXTENSIONS

We now consider some special cases and extensions of the model described above. We start with a study of the qualitative property of the Shapley value in the case of a single class of users. We then consider two possible features that can be incorporated in our analysis: a cost incurred by users for peer-assistance, and a case with multiple atomic players.

5.1 The single-class case

For the case of a single class of users, the Shapley value for users $\phi_i(\bar{X})$ reduces to $\phi(X)$ given by

$$\phi(X) = - \int_0^1 s \frac{dC}{dx}(sX) ds$$

We discuss the effect of the curvature of C on ϕ below; the four cases considered are also illustrated in Figure 1.

Convex cost function. If the cost function is convex, then the user Shapley value decreases monotonically. The interpretation is that increasing sharing brings diminishing returns to the improvement of the cost function. Thus, while the provider continues to benefit, the number of users that share grows faster than the total subsidy sent back. Additional level of sharing brings about “competition” amongst users and reduces their intrinsic value to the provider. Note that even in this case, *every* new sharing user sees an increase in value and so does the provider, so the system still converges to a full sharing mode.

Concave cost function. If the cost function is concave, then the user Shapley value increases monotonically. The implication here is the reverse of the convex case, *i.e.*, increasing levels of sharing brings increasing returns to cost reduction. This super-linear growth in the Shapley value is passed back to the sharing users.

Linear cost function. If the cost function is linear, of the form $A - BX$, then the Shapley value of the users reduces to

$$\phi(X) = - \int_0^1 s \frac{dC}{dx}(Xs) ds = \frac{B}{2}$$

and is thus independent of the level of sharing. This is because of the fact that the marginal contribution of a user stays constant for all levels of sharing. Note however that the provider’s Shapley value per user continues to increase linearly with increased level of sharing. This linear cost function presents a particularly simple and attractive pricing scheme from the users perspective. While the provider can always have access to the precise knowledge of X to compute the exact price for the users, a flat rate pricing scheme is inherently more attractive to users.

Sigmoid cost function. If the cost function is Sigmoid—*i.e.*, the cost function first decreases slowly with small amounts of sharing, then decreases rapidly and finally the decrease saturates—the user Shapley value may have a unique maximum. If such a maximum exists, it will appear after the inflection point of the cost function. Examples of such cost functions would be $A - B/(1 + e^{-x})$ or $A - Bx/(\sqrt{1 + x^2})$. Note that like convex cost functions, new users continue to increase in value and hence everybody still has an incentive to share even after the level of sharing corresponding to the maximum user Shapley value.

5.2 Including peer-assistance cost for users

In our model of cost so far, we have assumed that the operational cost $C^N(\bar{X})$ comes only from the provider, and hence that it decreases with any (component-wise) increase in \bar{X} . Let us now assume that each peer of class i incurs an additional cost δ_i for participating through peer-assistance. The previous expressions of the Shapley value still hold, where the function $C^N(\bar{X})$ is simply replaced by the total sum of cost in the system: $C_P^N(\bar{X}) + \sum_{i=1}^m N \cdot X_i \cdot \delta_i$.

Under the assumption that the cost reduction for the provider is always larger than the cost of peer-assistance to a user, all the results we proved in the previous section still hold. This condition may be written for a finite N , or respectively for large N , as

$$C_P^N(\bar{X}) \leq C_P^N(\bar{X} - \frac{1}{N} \cdot e_i) - \delta_i \quad \text{or, for } N \rightarrow \infty, \quad -\frac{\partial C_P}{\partial x_i} \geq \delta_i.$$

In such case, it is in everyone’s interest, including the provider’s, to join a coalition and receive their associated Shapley value by lowering the overall operational cost. The cost incurred by the user is then compensated by adding this value to the rebate received by the user for enabling peer-assistance. This may be a symbolic cost perceived by users, or a concrete expenditure to pay for energy or bandwidth used. A real world example of such a system is People CDN². This company compensates peers for residential broadband connection and provides peers with residential gateway boxes to be used as content delivery network nodes.

5.3 Multiple atomic players

We now extend the concept to a system with multiple atomic players. An example could be a network provider that has a limited catalog of titles in a video on demand system, and then a third party provider (\mathcal{E}) that extends the catalog. Let us assume that the extension in catalog is priced E dollars per user by the system (*i.e.*, the user pays $R + E$ with the add-on service of the third party provider). We also assume that the operational cost to the provider $C(\cdot)$ increases by an amount $C'(\cdot)$ due to the presence of the extra catalog. Again, this additional cost may change as users opt-in to offer peer-assistance to other users, so that it depends on N and \bar{X} .

We denote by $\varphi_i^N(\bar{X})$ (resp. $\tilde{\varphi}_P^N(\bar{X})$) the Shapley value of a user in class i (resp. the per-user Shapley value of the provider) when \mathcal{E} is not in the coalition. Since the worth function has not changed, all players receive the same Shapley value as seen before. In addition, we denote by $\varphi_i'^N(\bar{X})$ (resp. $\tilde{\varphi}_P'^N(\bar{X})$) the Shapley values when \mathcal{E} decides to join. Let us denote in this case by $\tilde{\varphi}_E'^N(\bar{X})$ the per-user Shapley value for this new player.

As shown in the next theorem, the method we described above can be extended to analyze the Shapley value for this case with multiple players. The proof (omitted due to space constraints, see [18] for details) follows the same type of argument: deducing limit axioms when the population gets large, which leads to simple differential equations.

THEOREM 5.1. *As N tends to infinity, the provider’s (per user) Shapley value and the Shapley values of users in all classes converge to the following smooth functions of \bar{X} :*

$$\begin{cases} \tilde{\varphi}_P'(\bar{X}) = R + \frac{E}{2} - \int_0^1 C(s\bar{X}) ds - \int_0^1 sC'(s\bar{X}) ds \\ \tilde{\varphi}_E'(\bar{X}) = \frac{E}{2} - \int_0^1 sC'(s\bar{X}) ds \\ \forall i, \quad \varphi_i'(\bar{X}) = - \int_0^1 s \frac{\partial C}{\partial x_i}(s\bar{X}) ds - \int_0^1 s^2 \frac{\partial C'}{\partial x_i}(s\bar{X}) ds. \end{cases}$$

An intuitive explanation of the formula is the following: If there are no excess operational costs involved (*i.e.*, $C' = 0$), the third party provider simply gets $E/2$ of the additional per user revenue E . If there are costs involved which may be reduced through peer assistance, users get compensated and half of this compensation is subtracted from the revenue of the third party provider. Revenues and user incentive are equally shared among the network provider and the third party content provider.

²see <http://www.pcdn.info>

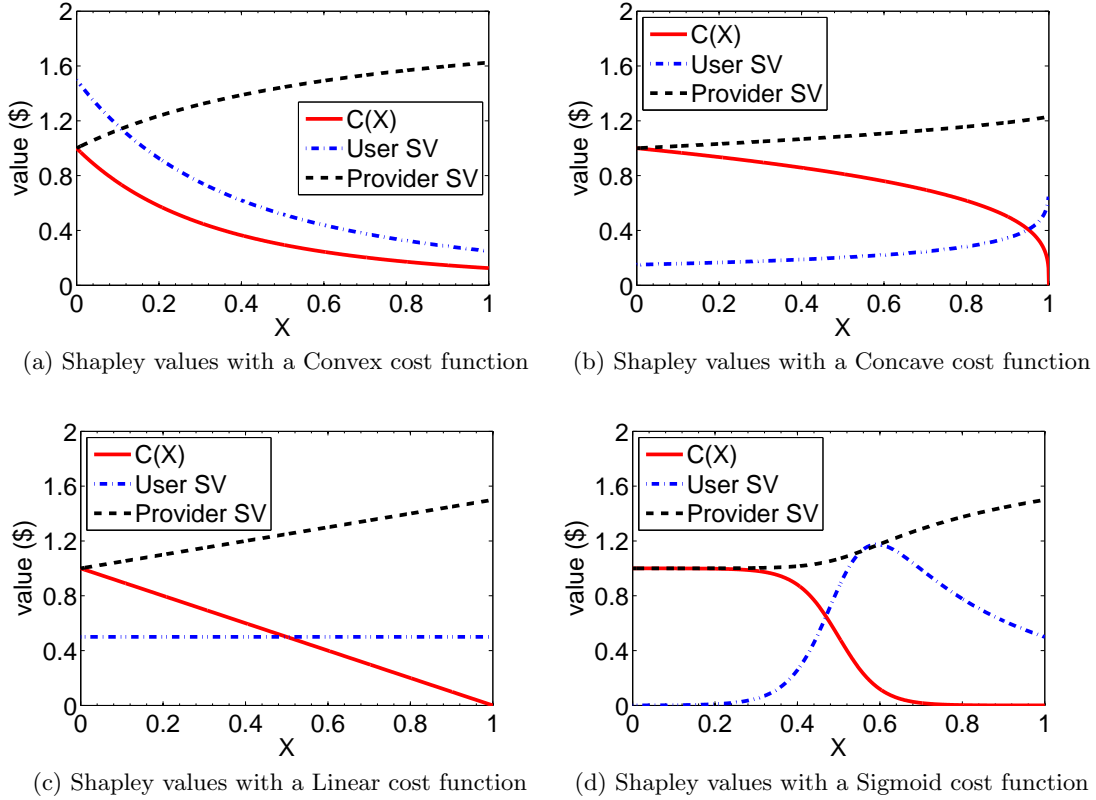


Figure 1: Provider and user Shapley values obtained through Theorem (4.3) in the single class case, for cost functions C with different curvatures. The Shapley values corresponding to a convex cost function are shown in (a): in this case, the user Shapley value decreases monotonically with the fraction X of users sharing, as increasing the latter brings diminishing returns to the cost reduction. The exact opposite behavior is observed in (b), where C is concave. For a linear cost function, shown in (c), the user Shapley value does not depend on X . Finally, when C is sigmoid, there exists a critical fraction of sharing users for which the Shapley value is maximized. Note that, in all four cases, the Shapley value of the provider increases with X ; this is a consequence of the monotonicity of C .

6. APPLICATIONS

We now apply the previous theory to various content access services and related costs. We first treat extensively the case of bandwidth saving in wired networks for three applications: file download, live streaming and video on demand. We exhibit a common simple incentive scheme for all three, which we extend to a more complex value chain containing both content and network providers. We then demonstrate the flexibility of this model by applying this technique to two other applications: incentivizing energy savings and enabling collaborative dissemination of content update in a wireless mobile network.

6.1 Bandwidth cost saving

6.1.1 Three applications with a linear cost model

A recurring theme in the study of P2P systems is that in the fluid limit, feasibility conditions are often expressed as separable, linear functions of the capacities of the P2P users.

File download.

The peer uplink model by Mundinger et al. [19] proposed for analyzing P2P file dissemination, they show that under the fluid limit, the minimal makespan of a provider with peer-assisted content distribution is

$$T^* = \max \left\{ \frac{1}{C_s}, \frac{N}{C_s + \sum_i N_i u_i} \right\},$$

where C_s is the upload capacity of the provider and u_i s are the upload capacities of the N_i participating peers. We can set the desired makespan as a constraint, and then calculate the server capacity required to achieve the makespan. The operational cost can then be computed as the cost of this capacity (which includes bandwidth and server costs, and costs at the user end). The asymptotic cost function $C(\bar{X})$ is then given by

$$C(\bar{X}) = \frac{(1 - \sum_j X_j u_j T^*)}{T^*}.$$

Live streaming.

In [27] the authors studied multichannel P2P live streaming systems. The authors studied both isolated channel (ISO) as well as View-Upload Decoupled (VUD) systems and through queueing theoretic models demonstrated the superiority of the VUD system. They introduce the following threshold parameter α given by

$$\alpha = \frac{r}{\sum_i X_i u_i},$$

Where r is the streaming rate, and X_i is the fraction of users with upload bandwidth u_i . The authors show that if $\alpha < 1$, then the universal streaming probability goes to 1, and all users achieve a desired quality of service. If $\alpha > 1$, then the service provider needs to complement the upload capacity of the users with dedicated servers. Note that P2P streaming provider like UUsee and PPLive maintain dedicated servers to offer quality of service, fitting this analysis [26]. Assuming that the system operates at the critical threshold $\alpha = 1$ the server capacity C_s required per user is

$$C_s = r - \sum_i X_i u_i.$$

Video on demand.

Let us assume a linear interpolation model for Video on Demand. Let γ_s be the cost of transmitting a bit from a datacenter or a CDN, and γ_p be the cost of transmitting the same bit from a P2P user. From a bandwidth perspective, if the P2P user has residential broadband access with no upload limits, then this γ_p can be taken to zero. Otherwise, it can be included as seen in Section 5.2.

The “cost” of serving a movie can be broken down into two components. A fixed overhead A , which includes the costs of licensing fee for the content, storage, maintenance *etc.*, and then a variable component which depends on the cost per bit. Let us assume the size of the file is S bits and that a user in class i uploads S_i bits. If we have N users and N_i in each class share, then

$$NS = \sum_i N_i S_i + NC_s,$$

where C_s is the portion of server bits required to compensate for the capacity shortfall in the system. The bandwidth cost per user is then given by

$$\gamma_s C_s = \gamma_s S - \gamma_i \sum_i X_i S_i.$$

6.1.2 A general incentive scheme

Upload miles.

If the cost is directly proportional to this server capacity (say KC_s), we then have an asymptotically separable linear cost function, with the classes of users characterized by their upload capacities. By Theorem 4.3, the Shapley value for the i^{th} class is then given by $\phi_i = \frac{Ku_i}{2}$ for the live streaming case and file download case, or $\phi_i = \frac{KS_i}{2}$ for the Video on Demand case. Note that there is a subtle difference between u_i and S_i , with former referring to a rate whereas the latter is the total number of uploaded bits—one can go back and forth between them with an appropriate normalization of relevant time periods.

The incentive scheme becomes then a simple one: A peer earns half the cost savings that it provides as “upload miles”. Because of the linear relationship, this translates to an “upload two get one free” scheme as far as distribution costs are concerned. The result is intuitively satisfying, since users providing more resources in terms of upload bandwidth or capacity end up with a proportionately higher Shapley value. In the asymptotic case Shapley value of a specific peer is also independent of what other peers are doing, so for a service provider it is a very easy to implement and explain pricing scheme.

Remarks:

- There is no restriction on the value of u_i or S_i . It can in fact be greater than the per user resource requirement, as a peer may satisfy the demand of multiple users and continue to earn credits. High uploading peers can take over the slack of low uploading/non-sharing users (similar to the design architecture VUD in [27]), as an opportunity to earn more. The peer is compensated directly in proportion to the resources it contributes.
- Our simple analysis assumes a fixed cost of bandwidth to the provider. However, the cost of bandwidth may reduce as the volume purchased gets larger (*i.e.*, “buying in bulk” is sometimes cheaper). The cost function for the provider becomes nonlinear, the rebate given to users for their upload follows then the general form of Theorem 4.3.

Case study: Apple iTunes.

As a case study, we estimate the impact of our framework to Apple iTunes. Apple streams movies, TV shows, and podcasts through its iTunes service. As an example, the movie “Quantum of Solace” is available in High Definition (HD) for renting as well as buying.

We analyze the following renting scenario. Users can rent the HD version for \$3.99, and the file size is 3.6 GB. Note that the HD resolution is 720P, whereas the industrial standard is tending towards 1080P, so called “full” HD. Extrapolating from the ratio of the trailer files that are encoded for 1080P, a 1080P version of the same file would be about 6.8 GB. While CDN prices range anywhere from \$.5 to \$1.00 per GB³, we assume buying in bulk is cheaper and estimate Apple has negotiated an arrangement with its CDN of \$0.10/GB. This results in the raw bandwidth cost for the 1080P version of the movie to be \$.68. Thus, if we look *only* at the bandwidth savings, the users get back \$.34 for a 1080P movie. Given current pricing, this works out to a 5-10% subsidy, which is not insignificant. For the provider, an equal amount comes back as saved costs and thus increased profits. There are fixed overheads like movie licensing rights that the provider has to pay for the content owners like movie studios, but there are other operational savings that are possible, for free content.

For every movie that viewers rent, they typically watch previews of several movies. Assuming a preview size of 150 seconds, this works out to a cost of about \$.01 per viewing, which is provided for free, and is recovered as overhead from the paid content. Then there are millions of free Podcasts downloaded daily by users, again provided by Apple for free.

³see <http://www.cdnpricing.com>

Each one of those bits can be transferred to the P2P infrastructure, and users can earn “miles” by uploading content on behalf of the provider. There are additional operational savings in terms of datacenter costs, storage and maintenance that are possible with this approach, and our mechanism provides an incentive for everyone involved to participate.

6.1.3 Influence of content and network providers

We now assume that a content provider gives access to an extended catalog, as compared to the basic offer from the network provider. Consistent with Section 5.3, we assume that users are charged an additional value E to access this additional catalog. If we assume a single class of users, following previous linear models, the cost incurred by the network provider is $C(x) = C(0) - \gamma_s x$ with the restricted catalog, and an additional cost $C'(x) = C'(0) - \beta_s x$ is incurred due to the increase in bandwidth needed to serve the additional content.

Applying Theorem 5.1 we then have

$$\begin{aligned}\phi_i(x) &= \frac{\gamma_s}{2} + \frac{\beta_s}{3}, \\ \phi_E(x) &= \frac{1}{2} \left[E - C'(x) - \frac{\beta_s x}{3} \right], \\ \phi_P(x) &= R - C(x) + \frac{\gamma x}{2} + \frac{1}{2} \left[E - C'(x) - \frac{\beta_s x}{3} \right].\end{aligned}$$

These equations can be interpreted as follows. When the content provider joins, the additional system benefit, $E - (C'(x))$, is shared equally among the two providers. In addition, any extra cost reduction $\beta_s x$ offered by peer uploads leads to a subsidy of $\beta_s/3$ for each assisting peer or $\beta_s x/3$ in total. This amount is paid for at equal levels by each provider.

In other words, when a peer serves some content, it gets a subsidy equal to the corresponding cost saving, divided by 2 if this belongs to the basic catalog, hence made available by a single veto player (the network provider), and by 3 if it is made available by collaboration of two veto players. This indicates how the simple “miles” scheme previously described in the case of a single operator generalizes in more complex scenarios.

Network neutrality.

Our results with atomic multiple players brings new lights on the issue of network neutrality [11, 20] and revenue for peer-assisted services. With network neutrality in place, the network provider is *no longer a veto player*. It is obliged to provide any and every kind of service that it offers to any third party content provider that uses its network. The network provider gets a flat amount for the bandwidth provision and the subsequent value creation by peer assistance has no bearing on the network providers share. Indeed, the network provider stands to lose if the bandwidth requirements of the content provider are provided by the peers directly, but because of network neutrality the network provider is not in a bargaining position to prevent that from happening. Peer-assistance then results in a cost reduction directly for the third party, and it can be incentivize by the third party provider using the same upload miles scheme built when the network provider was the only atomic player.

6.2 Energy costs for Internet content delivery

In our discussion so far, we assumed that peers did not incur any direct monetary cost for serving others’ requests. This is however not entirely true: peer devices consume extra energy for serving extra jobs, which shows on the end users’ electricity bill.

We shall now apply our theory to identify the incentives required specifically to cover such energy costs, and compensate the savings on the provider’s side. To this end, we will rely on a recent study of energy savings in the so-called Nano-datacenters (NADA) architecture [25]. NADA is a managed system in which content is provided from end users’ set-top boxes or Internet gateways, the latter being the peer devices assisting in service delivery.

The measurements reported in [25] suggest that the energy cost to the provider is linear in the traffic served. We shall denote by γ_s the provider energy cost per Gigabit served. It also follows from [25] that the peer costs are linear as well, and we shall denote by γ_p the corresponding cost per Gigabit.

To estimate the rates γ_s and γ_p , we note from [25] that a typical VoD server consumes 211 Joules per Gigabit (J/Gb). This is further inflated by the so-called Power Usage Efficiency (PUE) of the datacenter hosting the server, which accounts for all the additional costs related to cooling, power transmission and conversion. State-of-the-art datacenters such as Google’s have PUE’s of 1.2, resulting in a total cost of 253.2 J/Gb.

In contrast, the cost of serving content from a standard triple-play gateway, assuming the latter is already active, is of only 100 J/Gb. We may thus take γ_s to equal 253.2 J/Gb and γ_p to equal 100 G/Gb. In fact the NADA architecture brings extra savings due to traffic localization: traffic reduction at routers brings about a reduction of approximately 900 J/Gbit for serving from NADA rather than Datacenters. Thus, if we also take into account router energy costs, we would take instead $\gamma_s = 1153.2$ J/Gb.

For the sake of simplicity, we chose here to express costs in Joules; it is straightforward to translate cost into dollars, given the energy prices for peers and for providers.

Eventually, with the provider cost function of $C_s(x) = \gamma_s V(1 - x)$, where V is the volume (in Gigabits) served per user, and the peer cost function of $C_p(x) = \gamma_p V x$, by applying the general formulas (4.3), we see that the provider would return to an assisting peer an amount of $\gamma_p V$, covering its energy cost, plus an incentive of $(\gamma_s - \gamma_p)V/2$.

If we take the above values, for one joule spent by a peer, it is reimbursed of $1 + (2.53 - 1)/2 = 1.765$ Joules, providing a non-negligible incentive. If we further factor in the energy savings at routers, we would instead take $\gamma_s = 1153$ J/Gb. The reimbursement for one joule spent then becomes $1 + (11.53 - 1)/2 = 6.265$ Joules, which becomes rather compelling. Another interesting outcomes is that, in addition to save their energy expenditure, the players of this game are naturally given incentive to limit their energy as much as possible.

6.3 Update distribution over a mobile network

In [4], a wireless service provider pushes content updates for certain content that changes dynamically (e.g., a news-feed or a blog) to N mobile users. The mobile users can share this content with each other in a peer-to-peer fashion

whenever they meet: a user whose content is most recent pushes it to the one whose content is outdated.

We consider a multiple class case, in which users are partitioned in M classes of equal size. Every user in class i meets other users in its own class uniformly at random, with an aggregate contact rate λ_i . Suppose that only a fraction X_i users in class i form a coalition and share their content. We will assume that users share their content *only* with other users in their class who are part of the coalition, *i.e.*, *also share their content*.

Suppose now that μ_i is the rate with which the service provider pushes updates to users in class i that share their content, while ν_i is the rate with which it pushes content updates to users that do not share their content. Then, the analysis in [4] implies that the expected age of the content at a user in the coalition is

$$Y_i = \frac{1}{X_i \lambda_i} \log \frac{X_i \lambda_i + \mu_i}{\mu_i} \quad (9)$$

where X_i is the fraction of users sharing their content in class i . On the other hand, the expected age of content at a user not sharing content will be

$$Z_i = \frac{1}{\nu_i}. \quad (10)$$

Assuming that the cost to the service provider is proportional to its aggregate downlink rate, used to push updates to the users, the average cost per user is

$$\sum_i R[(1 - X_i)\nu_i + X_i\mu_i] \quad (11)$$

where R is the cost per unit of bandwidth. We set the following requirement on the quality of the content update delivery service at each user: every user must have an expected age below a threshold τ . Under this constraint, the cost at the service provider can be computed by solving the following optimization problem, whose free variables are μ_i and ν_i :

$$\text{Minimize: } \sum_i R[(1 - X_i)\nu_i + X_i\mu_i]. \quad (12a)$$

$$\text{subj. to: } Y_i \leq \tau, Z_i \leq \tau, \text{ for every class } i. \quad (12b)$$

where Y_i , Z_i can be computed from (9) and (10), respectively. This is a convex optimization problem (see [10]), and its solution yields the following cost function

$$C(\bar{X}) = \sum_i R \left[\frac{1 - X_i}{\tau} + \frac{X_i^2 \lambda_i}{e^{X_i \lambda_i \tau} - 1} \right]. \quad (13)$$

Single class case.

When users belong to a single class the cost function becomes

$$C(X) = R \left[\frac{1 - X}{\tau} + \frac{X^2 \lambda}{e^{X \lambda \tau} - 1} \right]. \quad (14)$$

Function $C(X)$ is decreasing and its curvature depends on the product $\lambda\tau$, as illustrated in Fig. 2. For $\lambda\tau \leq \alpha$, where $\alpha \approx 3.4368$, C is concave. For $\lambda\tau > \alpha$, C is sigmoid, with an inflection point at $\alpha/(\lambda\tau)$. The discussion in Section 5.1 implies that the maximal user Shapley value is attained after this inflection point.




Figure 2: The cost function $C(X)$ for three different cases of $\lambda\tau$ (normalized by $1/\tau$). When $\lambda\tau \leq \alpha$, where $\alpha \approx 3.4368$, $C(X)$ is concave. For $\lambda\tau > \alpha$, $C(X)$ is sigmoid —though, past the inflection point, the curvature is almost negligible. In fact, for large $\lambda\tau$ the cost function becomes almost linear.

Nonetheless, if $\lambda\tau \gg \alpha$ the second term in (13) becomes negligible, and the cost function becomes approximately linear. We will exploit this to obtain a simpler formulation in the multiple class case.

Dense network with multiple classes.

We investigate the case where contacts occur very frequently, in which

$$\lambda_i \tau \gg \alpha, \text{ for all } i. \quad (15)$$

In such a case, the contact rate among mobile users is very high compared to the constraint placed on the expected age through τ . Under (15), the second term in (13) becomes negligible for all i : intuitively, the cost for serving the users that share is negligible, as any update inserted in a few users is very quickly (w.r.t. τ) propagated within the class. As a result, the cost is dominated from serving the users that do not share, and thus becomes linear

$$C(\bar{X}) \approx \frac{1}{\tau} - \sum_i \frac{R X_i}{\tau},$$

and the Shapley value of a user in class i becomes

$$\phi_i \approx \frac{R}{2\tau}.$$

In short, whenever a user chooses to share, the service provider can reduce its costs by R/τ —as this user can be served at a negligible cost through sharing. Half of this cost saving is then returned, as compensation, to the mobile user (similar to the “upload miles” scheme in Section 6.1.2).

7. CONCLUSION

In this work, we study incentives in peer-assisted service from a new standpoint. Instead of starting by enforcing a set of rules that should constrain what users will be providing and receive, we answer the following question: “How can we value the fair share of profits of all interested parties in a distributed system and how easy is it to compute?”

If feasible, this approach clearly brings the benefit of not having to worry about free riding and users misbehaving in reaction to an unfair advantage gathered by the provider. It also greatly simplifies pricing of peer-assistance, as provider and users can simply focus on the price of the service they receive, and what the rebate is.

The main contribution of this paper is to demonstrate that such approach is indeed feasible. First, because the large number of users participating is an opportunity for computational efficiency (through fluid limit approximation) rather than an obstacle. Second, because this approach can easily accommodate multiple classes of users, any arbitrary relation between the cost incurred by the provider and the peer-assistance, as well as other players of the value chain (provider of content, network operator).

The results we provide in this paper present a first step for a comprehensive study of the value of peer-assistance in distributed services. The connections between the cost function, which depends on peer-assistance, and the properties of the incentive received seems a promising area to explore, especially in the multi-class case. More applications should be considered, in particular in providing reliability and fault tolerance through peer-assistance. It also seems important to account that, for some services, peers do not only act as relay of a P2P network, but may also provide content themselves which increases the value of the service. We believe that all these research challenges can be considered in a new light due to the insights obtained through this analysis.

8. REFERENCES

- [1] M. S. Allen, B. Y. Zhao, and R. Wolski. Deploying video-on-demand services on cable networks. In *Proc. of IEEE ICDCS*, 2007.
- [2] R. Aumann and L. Shapley. *Values of Non-Atomic Games*. Princeton University Press, 1974.
- [3] L. J. Billera, D. C. Heath, and J. Raanan. Internal Telephone Billing Rates-A Novel Application of Non-Atomic Game Theory. *Operations Research*, 1978.
- [4] A. Chaintreau, J.-Y. L. Boudec, and N. Ristanovic. The age of gossip: Spatial mean-field regime. In *Proc. of ACM SIGMETRICS*, 2009.
- [5] Y. F. Chen, Y. Huang, R. Jana, H. Jiang, M. Rabinovich, J. Rahe, B. Wei, and Z. Xiao. Towards capacity and profit optimization of video-on-demand services in a peer-assisted IPTV platform. *Multimedia Systems*, 15(1):19–32, 2009.
- [6] Y. T. Dov Samet and I. Zang. An Application of the Aumann-Shapley Prices for Cost Allocation in Transportation Problems. *Mathematics of Operations Research*, 9(1):25–42, February 1984.
- [7] L. Guo, S. Chen, Z. Xiao, E. Tan, X. Ding, and X. Zhang. Measurements, analysis, and modeling of bittorrent-like systems. In *Proc. of ACM/USENIX IMC*, 2005.
- [8] S. Hart. Values of Mixed Games. *International Journal of Game Theory*, 2:69–85, 1973.
- [9] C. Huang, J. Li, and K. W. Ross. Can internet video-on-demand be profitable? In *Proc. of ACM SIGCOMM*, 2007.
- [10] S. Ioannidis, A. Chaintreau, and L. Massoulié. Optimal and scalable distribution of content updates over a mobile social network. In *Proc. of IEEE INFOCOM*, 2009.
- [11] M. Jamison and J. Hauge. Getting what you pay for: Analyzing the net neutrality debate. *University of Florida, Department of Economics, PURC Working Paper*, 2007. available at ssrn.com/abstract=1081690.
- [12] V. Janardhan and H. Schulzrinne. Peer assisted VoD for set-top box based IP network. In *Proc. of ACM Workshop P2P-TV*, 2007.
- [13] F. P. Kelly. The Clifford Paterson Lecture, 1995: Modelling Communication Networks, Present and Future. *Philosophical Transactions: Mathematical, Physical and Engineering Sciences*, 354(1707):437–463, March 1996.
- [14] R. T. Ma, D. Chiu, J. C. Lui, V. Misra, and D. Rubenstein. On cooperative settlement between content, transit and eyeball internet service providers. In *Proc. of ACM CoNEXT*, 2008.
- [15] R. T. B. Ma, D. Chiu, J. C. Lui, V. Misra, and D. Rubenstein. Interconnecting Eyeballs to Content: A Shapley Value Perspective on ISP Peering and Settlement. *Proc. of ACM Workshop NetEcon*.
- [16] R. T. B. Ma, D. Chiu, J. C. Lui, V. Misra, and D. Rubenstein. Internet Economics: The use of Shapley value for ISP settlement. *Proc. ACM CoNEXT*, 2007.
- [17] A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic theory*. Oxford University Press, 1995.
- [18] V. Misra, S. Ioannidis, A. Chaintreau, and L. Massoulié. Incentivizing Managed Peer to Peer Systems: A Fluid Shapley Value Approach. Technical report, Columbia University, 2009. Available at <http://www.cs.columbia.edu/~misra/fluidshapley-tech.pdf>.
- [19] J. Munding, R. Weber, and G. Weiss. Analysis of peer-to-peer file dissemination. *SIGMETRICS Perform. Eval. Rev.*, 34(3):12–14, 2006.
- [20] J. Musacchio, G. Schwartz, and J. Walrand. Network neutrality and provider investment incentives. *Proc. of Asilomar Conference*, 2007.
- [21] R. Myerson. Graphs and cooperation in games. *Mathematics of Operations Research*, 2:225–229, 1977.
- [22] L. Shapley. A Value for n-Person Games. In H. W. Kuhn and A. W. Tucker, editors, *Contributions to the Theory of Games II*, volume 28 of *Annals of Mathematics Studies*, pages 307–317. Princeton University Press, 1953.
- [23] L. Shapley. Cores of convex games. *International Journal of Game Theory*, 1:11–26, 1971.
- [24] K. Suh, C. Diot, J. Kurose, L. Massoulié, C. Neumann, D. F. Towsley, and M. Varvello. Push-to-peer video-on-demand system: Design and evaluation. *IEEE JSAC*, 25(9):1706–1716, 2007.
- [25] V. Valancius, N. Laoutaris, L. Massoulié, C. Diot, and P. Rodriguez. Greening the Internet with nano Data Centers. In *Proceedings of ACM CoNEXT*, 2009.
- [26] C. Wu, B. Li, and S. Zhao. Multi-channel Live P2P Streaming: Refocusing on Servers. In *Proc. of IEEE INFOCOM*, April 2008.
- [27] D. Wu, Y. Liu, and K. W. Ross. Queuing Network Models for Multi-Channel Live Streaming Systems. In *Proc. of IEEE INFOCOM*, Rio de Janeiro, April 2009.
- [28] X. Zhang, J. Liu, B. Li, and T. P. Yum. Coolstreaming/donet: A data-driven overlay network for peer-to-peer live media streaming. In *Proc. of IEEE INFOCOM*, 2005.