

On the Interaction between Content-oriented Traffic Scheduling and Revenue Sharing among Providers

Hyojung Lee[†], Hyeryung Jang[†], Yung Yi[†], and Jeong-woo Cho[‡]

Abstract—The Internet consists of economically selfish players in terms of access/transit connection, content distribution, and users. Such selfish behaviors often lead to techno-economic inefficiencies such as unstable peering and revenue imbalance. Recent research results suggest that cooperation in revenue sharing (thus multi-level ISP settlements) can be a candidate solution for the problem of unfair revenue share. However, it is unclear whether providers are willing to behave cooperatively. In this paper, we study the interaction between how content-oriented traffic scheduling at the edge is and how stable the intended cooperation is. We consider three traffic scheduling policies having various degrees of content-value preference, compare them in terms of implementation complexity, network neutrality, and stability of cooperation, and present interesting trade-offs among them.

I. INTRODUCTION

The Internet is a system where the entities such as EUs (End Users) and content/eyeball/transit ISPs (Internet Service Providers)¹, having different economic perspectives, compete and cooperate in a highly complex manner. Eyeball/transit ISPs connect EUs to the Internet, and content ISPs inject and deliver content into the Internet [1], e.g., videos, web pages, and files. The major interest of such providers, which is often selfish, is to maximize their profits, sometimes incurring techno-economic inefficiencies in the Internet. For example, ISPs' selective peering with other ISPs may have negative impact on Internet's connectivity. It is reported that some providers express economic complaints on revenue imbalance among them, see e.g., [2], [3]. One of the central issues regarding such complaints is how to fairly distribute the revenue from the users to the providers.

There have been recent research efforts on fair and efficient revenue sharing among providers, using the notion of Shapley value (SV) [4] from cooperative game theory. The SV is a fair payoff distribution scheme and presumes that the *grand coalition* (i.e., the coalition containing all players) is agreed by the players. The SV based revenue sharing hypothesizes that the profit distribution is achieved at a multilateral, global level, rather than a bilateral, local level, leading to the nice features in terms of fairness, efficiency, and interconnection incentives, see e.g., [5] and [6].

However, it is questionable that the providers would actually form the grand coalition, referred to as *stability* of the grand coalition. In [5], the grand coalition is conjectured to be formed

for *under-demanded* networks, where the network bandwidths exceed users' traffic demands. However, the Internet edge is often *over-demanded* due to technical advances of edge devices, e.g., smart phones/pads and smart TVs. An over-demanded network significantly changes how we should technically treat the issue of stability. Roughly, the stability of the grand coalition under SV can be studied by checking the existence of sub-coalitions in which all players in that coalition can be better-off under SV. In the over-demanded case, the individual share depends on how the edge networks *sift out* a part of user demands to meet the capacity. This dependence is in sharp contrast to the earlier SV-based research in micro-economics, e.g., [7]–[9]. In particular, this complex inter-play is mainly due to the fact that the actual individual payoffs (distributed by SV) may dramatically change depending on how traffic is sifted out at the edge.

In this paper, we consider various traffic scheduling policies at the edge, each of which presents different degrees of content-value preference, implementation complexity, and network neutrality, and compare them in terms of coalition worth and stability of the grand coalition. To that end, we define a coalition game, called *Revenue Sharing Game (RSG)*, where the players are eyeball, transit, and *virtual* content ISPs. The notion of virtual content ISP, which is a triple of content ISP, content, and region, is introduced with the goal of accurately modeling the cooperation decision of a content ISP.

Following the defined coalition game, we first prove that in under-demanded networks, the coalition worth is maximized at the grand coalition, which is always stable under SV based revenue sharing. This formalizes the result conjectured in [5]. Second, in over-demanded networks, where traffic scheduling is an important factor, we prove that the scheduling absolutely prioritizing higher-value contents (called PP: Priority Policy) maximizes the worth over all possible scheduling policies. Also, we prove that the scheduling which (relatively) assigns higher weights to more profitable contents (called RPP: Revenue Proportional Policy) always generates more worth than the content-agnostic scheduling (called TPP: Traffic Proportional Policy). However, in terms of the stability of the grand coalition, even under PP, which is a worth-maximizing policy, the grand coalition is not stable. We provide sufficient conditions under which a scheduling is better than another scheduling in terms of stability implying that more content-oriented scheduling tends to be “more stable” than content-agnostic one. Interesting tradeoffs are observed here; PP or RPP requires much more complex operations, such as priority scheduling or weighted fair queueing, whereas TPP can be

[†]: Department of Electrical Engineering, KAIST, South Korea, e-mails: {hyojung_lee,hrjang,yi}@lanada.kaist.ac.kr. [‡]: School of Information and Communication Technology at KTH Royal Institute of Technology, Sweden, e-mail:jwcho@kth.se. This research is supported by KCA grant number KCA-2013-12911-05003.

¹ISP is sometimes called just ‘provider’ throughout this paper.

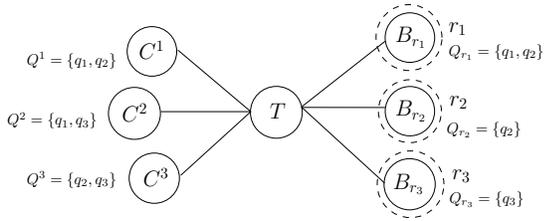


Fig. 1. An example network with content, transit, and eyeball ISPs. The set of contents: $\{q_1, q_2, q_3\}$, which are served by three content ISPs. The set of regions: $\{r_1, r_2, r_3\}$, where for example the users in the region r_1 request the contents $\{q_1, q_2\}$.

realized by a simple FIFO scheduling.

Related work. Related work includes the research on the revenue sharing mechanisms based on proportional fairness and NBS (Nash Bargaining Solution) [10]–[12]. In addition to SV’s application to providers’ settlements, it has also been applied to many other network-economic problems, e.g., peer-assisted services [13], viral marketing [14], and virtual infrastructure sharing [15], which mainly focus on how to apply the notion of SV to revenue sharing under various applications. The stability of the grand coalition under SV has been studied in the micro-economics community, e.g., [7]–[9]. The authors in [13] show that the grand coalition could be unstable in peer-assisted services with multi-providers even if the stability is guaranteed under a single provider. Our paper differs from the above in that our case includes the situation when the worth is a function of some control policy (i.e., traffic scheduling) with distinctive emphasis on the Internet market. In [16], only a small set of examples are provided, hinting that some of traffic scheduling policies can impact the stability based on a couple of simple examples.

II. MODEL

A. Network Model

We consider a network consisting of a transit ISP T , a set \mathcal{C} of content ISPs, and a set \mathcal{B} of eyeball (or access) ISPs, where we denote by $\mathcal{N} = \mathcal{C} \cup \{T\} \cup \mathcal{B}$ the set of all “providers”, i.e., **grand coalition**. The transit ISP offers connectivity between eyeball ISPs and content ISPs. For simplicity and tractable analysis, we assume that there is just a single transit ISP and all eyeball ISPs and content ISPs are connected to the transit ISP, and no direct connection between any content ISP and eyeball ISP exists. Eyeball ISPs connect residential users to the transit ISP. Denote \mathcal{R} as the set of all regions served by the set of eyeball ISPs \mathcal{B} . We also denote by B_r the eyeball ISP which covers the region $r \in \mathcal{R}$, where we assume that there does not exist a region covered by multiple eyeball ISPs. Let n_r be the link capacity between B_r and the transit ISP. The content could be delivered from a content ISP to the region r via the transit ISP to the requesting destination eyeball ISP B_r .

Let \mathcal{Q} be the set of all contents in the network. Note that a content can be served by multiple content ISPs. Each region may have a different set of contents to download, for which we let $X_{r,q}$ be the user population in region r that has demand for content $q \in \mathcal{Q}$. We assume that for a content q , users are

oblivious to a content ISP serving q . We denote \mathcal{C}_q as the set of content ISPs that serve the content q . Let $\mathcal{Q}_r \subseteq \mathcal{Q}$ be the set of contents demanded by the users in region r , and \mathcal{C}_r be the set of content ISPs that serve at least one content in \mathcal{Q}_r . In other words, the set \mathcal{C}_r is the union of the sets of content ISPs that serve the contents in \mathcal{Q}_r , thus $\mathcal{C}_r = \cup_{q \in \mathcal{Q}_r} \mathcal{C}_q$. We let s_q be the traffic volume (in bytes) of the content q , and β_q be the revenue of the content q , i.e., the per-content revenue earned by the content ISPs serving q .

Notation. We use the lower-case i , r , and q to index a content ISP, a region, and a content, respectively. For consistency, we place r and q in subscript and i in superscript. Thus, we often use C^i and \mathcal{Q}^i to refer to the i -th content ISP and the set of contents served by C^i . For any coalition (i.e., a set of ISPs) $S \subset \mathcal{N}$, we denote by $a[S]$ the restriction of a by S , e.g., $\mathcal{R}[S]$ is the set of regions served only in S . See Fig. 1 for an example of our network model.

B. Demand and Traffic Scheduling

Let y_r be the “original” traffic demand in region r , representing the total traffic volume requested by the users in region r , i.e., $y_r = \sum_{q \in \mathcal{Q}_r} s_q X_{r,q}$, where $s_q X_{r,q}$ corresponds to the traffic volume from the region r to access the content q . We say that the region r is *over-demanded* if $y_r > n_r$, i.e., the total traffic demand in region r exceeds the link capacity between the eyeball ISP B_r and the transit ISP, and that a network is over-demanded if there exists at least one over-demanded region in the network, otherwise said to be *under-demanded*.

An eyeball ISP will take some traffic shaping action, called *traffic scheduling*, if its serving region is over-demanded, so that its actually-served traffic volume does not exceed the link capacity. We consider a family of traffic scheduling policies abstracted by a function f subject to the following natural condition:

$$\sum_{q \in \mathcal{Q}_r} s_q X_{r,q} \cdot f(s_q, \beta_q, n_r, X_{r,q}) \leq n_r$$

where $0 \leq f(\cdot) \leq 1$. A scheduling policy can be regarded as a traffic shaper which reduces the original per-content user population $X_{r,q}$, thus the value of $f(\cdot)$ corresponds to the portion of user population that “survives” under a given traffic scheduling policy. We study the following three policies: **TPP** (Traffic Proportional Policy), **RPP** (Revenue Proportional Policy) and **PP** (Priority Policy), as formally stated in Table I². These three policies have diversified degrees of implementation complexity, required information, and content-value preference. For example, PP assigns absolute priority to higher-value contents, whereas TPP is indifferent to the content values, and RPP gives higher weights proportional to the profits generated by contents. We can (arguably) say that TPP is more network-neutral than RPP (similarly, RPP is more network-neutral than PP) due to their different, restricted handling of network traffic, depending on content values.

²To simplify exposition, we assume that the normalized content value β_q/s_q are different across all the contents.

TABLE I
THREE TRAFFIC SCHEDULING POLICIES

Policy	$f(\cdot)$	Operation
TPP	$\min\left(1, \frac{n_r}{\sum_{q' \in \mathcal{Q}_r} s_{q'} X_{r,q'}}\right)$	All traffic is treated neutrally, and $X_{r,q}$ is reduced in proportional to the traffic volume of q , s_q .
RPP	$\min\left(1, \frac{\beta q n_r}{s_q \sum_{q' \in \mathcal{Q}_r} \beta_{q'} X_{r,q'}}\right)$	$X_{r,q}$ is reduced in proportion to the amount of revenue of q , β_q .
PP	$\min\left(1, \left[\frac{n_r - \sum_{q' \in \mathcal{H}_{r,q}} s_{q'} X_{r,q'}}{s_q X_{r,q}}\right]^+\right)$	A content with higher β_q/s_q is absolutely prioritized. ($\mathcal{H}_{r,q}$ is the set of all $q'' \in \mathcal{Q}_r$ s.t. $\beta_{q''}/s_{q''} > \beta_q/s_q$)

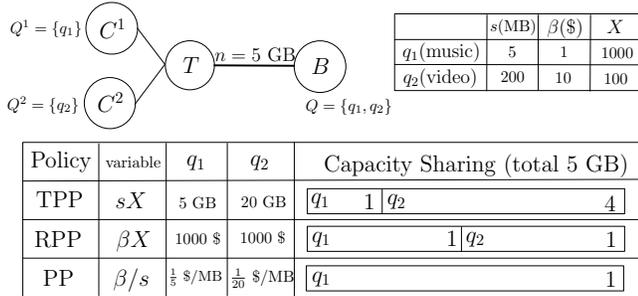


Fig. 2. Example of traffic scheduling policy.

Example. An illustration for three traffic scheduling policies is given in Fig. 2 with two content ISPs, one eyeball and one transit ISP, where two contents (music and video, denoted by q_1 and q_2 , respectively) are served by each content ISP (the region index r is dropped here to simplify exposition). The region is clearly over-demanded, because the total demand ($5\text{MB} \times 1000 + 200\text{MB} \times 100 = 25\text{GB}$) exceeds the capacity (5GB). The music's total traffic volume ($s \cdot X$) and the per-content revenue (β) are smaller than each of those of the movie, but more populations want to download the music. In TPP, the capacity is allocated according to the total traffic volume, i.e., $5 : 20$, whereas in RPP the total revenue is used to split the traffic, i.e., $1\$ \times 1000 : 10\$ \times 100 = 1 : 1$. Finally, in PP, the absolute priority is given to the content with higher normalized value, which is the music. In our example, the music's traffic volume is as same as the capacity n , and thus the video cannot be served at all in PP.

III. REVENUE SHARING GAME (RSG)

In this section, we define a coalition game [17] from cooperative game theory, called Revenue Sharing Game (RSG), followed by necessary preliminaries.

A. Game Formulation

We denote a coalition game with a coalition structure, by $(\mathcal{N}, v, \mathcal{P})$, where \mathcal{N} is a set of players and the game has a transferable utility characterized by a worth function v , which is $v : 2^{\mathcal{N}} \rightarrow \mathbb{R}$ and $v(\emptyset) = 0$. The worth function associates with any coalition $S \subseteq \mathcal{N}$ the value generated by cooperation. A coalition structure \mathcal{P} is a partition of \mathcal{N} . The case when $\mathcal{P} = \{\mathcal{N}\}$ is called *grand coalition*, for which we use just (\mathcal{N}, v) for simplicity, unless confusion arises.

Players in our RSG should be the providers. Transit and eyeball ISPs are naturally included in the player set. For content ISPs, we introduce a notion of *virtual content ISP*, identified

by (i) a content ISP, (ii) a region, and (iii) a content. Denote by $C_{r,q}^i$ the virtual content ISP of i -th content ISP, serving content q to region r . The main objective of introducing virtual content ISPs is to assign finer granularity to the coalition structures for the purpose of reflecting the practice more accurately. Examples include (i) a content ISP such as Google decides to stop servicing some contents to South Korea, or (ii) two content ISPs (excluding the rest of content ISPs in the Internet) form a coalition with just localized contents to serve the population in specific regions.

We define the worth of S as the total revenue earned by the players in the coalition S . As is done in [6], for a given S , we decompose S into *atomic* coalitions $S_{r,q}$, so that a atomic coalition includes an eyeball ISP B_r in region r , a transit ISP T , and the set of virtual content ISPs $C_{r,q}^i$ that serve q to the region r . Then, the coalition worth of S is simply the summation of the worths of the decomposed atomic coalitions, i.e.,

$$v(S) = \sum_{r \in \mathcal{R}[S]} \sum_{q \in \mathcal{Q}_r[S]} v(S_{r,q}), \quad (1)$$

where the worth of each atomic coalition $S_{r,q}$ is defined as the total fee for accessing content q , paid by the users in region r , $v(S_{r,q}) = \beta_q X_{r,q} f(\cdot)$. We also present the concept of *super-additivity* of a coalition game (\mathcal{N}, v) , which means that a coalition achieves larger coalition worth than what is achieved by its arbitrary partition, described as: for any coalition $S, T \subseteq \mathcal{N}$ such that $S \cap T = \emptyset$, $v(S \cup T) \geq v(S) + v(T)$.

B. Shapley Value and Stability

Associated with a coalition game $(\mathcal{N}, v, \mathcal{P})$ is the coalition structure value which is an operator φ which assigns values (or payoffs) to every player in game $(\mathcal{N}, v, \mathcal{P})$. We denote by $\varphi^i(\mathcal{N}, v, \mathcal{P})$ a coalition structure value for player i , and in this paper our focus is on the famous *Shapley value*. Shapley provides an axiomatic approach to determine a coalition structure value φ , which reflects the following desirable properties (as axioms): efficiency, symmetry, additivity, and dummy, see [4] for details. It has been proved that the value satisfying the four axioms is *uniquely* determined for every coalitional game in the premise of grand coalition (\mathcal{N}, v) (i.e., $\mathcal{P} = \{\mathcal{N}\}$), characterized as: for any player i ,

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

where Π is the set of $|\mathcal{N}|!$ orderings of \mathcal{N} and $S(\pi, i)$ is the set of players preceding i in the ordering π , and $\Delta_i(v, S)$ is the marginal contribution of player i for a coalition $S \subseteq \mathcal{N} \setminus \{i\}$, i.e., $\Delta_i(v, S) = v(S \cup \{i\}) - v(S)$. Simply speaking, the

Shapley value can be interpreted by the average marginal contribution over all orderings of players. The axiomatic coalition structure value for any coalition structure \mathcal{P} is called the Aumann-Drèze value (A-D value) [17]. Then, A-D value for a player $i \in S \in \mathcal{P}$ is also denoted by $\varphi^i(S)$ in this paper. For simplicity, we use the term ‘‘Shapley value’’ for both Shapley and A-D values, because their axiomatic structures are the same except that A-D values are computed for arbitrary coalitions.

Definition 1 (Stability of Grand Coalition [7], [8]) The grand coalition is said to be *stable* for a game (\mathcal{N}, v) with respect to the Shapley value φ , if for all $S \subset \mathcal{N}$ there is a player $i \in S$ such that $\varphi^i(\mathcal{N}, v, \mathcal{N}) \geq \varphi^i(\mathcal{N}, v, \{S, \mathcal{N} \setminus S\})$.

Intuitively, the grand coalition is stable under Shapley value, if for any coalition S , there exists at least one player $i = i(S)$ (which may depend on the considered coalition S) that becomes happier in the grand coalition than in S , thus there is no reason to stay out of the grand coalition. We call such a player $i(S)$ *Shapley-advocating player* for a given S . We extend Definition 1 to say that an arbitrary coalition S (which is not necessarily the grand coalition) is stable, when the coalition S is treated as the grand coalition.

IV. COALITION WORTH

It is obvious that for under-demanded networks, the game is super-additive, and thus, the coalition worth is maximized by the grand coalition. Thus, we focus on the case when the network is over-demanded. To differentiate the worth among the policies, we use the notation $v_T(S)$, $v_R(S)$, and $v_P(S)$ to refer to the worth functions of TPP, RPP and PP, respectively, for a given coalition S . We start this section by stating our main result on the coalition worth.

Theorem 1 (Coalition Worth: Over-demanded) Consider an over-demanded network and the corresponding RSG.

- (i) The RSG under PP is super-additive. Thus, the worth is maximized at the grand coalition under PP.
- (ii) For any given coalition S , the following inequality holds: $v_P(S) \geq v_R(S) \geq v_T(S)$ for all $S \subseteq \mathcal{N}$. Moreover, PP is an optimal policy that maximizes the worth over all possible traffic scheduling policies.

Due to space limitation, the proof is presented in [18], and we focus on the interpretations. First, in (i), it is shown that there exists a scheduling policy, which is PP, which ensures that the worth increases as the coalition becomes larger (i.e., super-additivity), and thus the grand coalition may be preferred under PP. This result is intuitive because PP always assigns higher priority to the traffic with higher content values. Second, in (ii), this value-oriented scheduling in PP leads to the result that for any given coalition S , PP is an optimal policy in terms of the worth for S among all other policies.

Similar tendency can be seen for RPP and TPP, as stated in Theorem 1(ii). To be more precise, consider two contents q_i, q_j in a given coalition (with region r), where assume that q_i has a larger normalized content value than q_j , i.e., $\beta_i/s_i \geq \beta_j/s_j$.³

³For notational simplicity, we sometimes use s_i rather than s_{q_i} (similarly, β_i instead of β_{q_i}) throughout this paper, unless confusion arises.

Note that due to this value difference, a scheduling policy that assigns more capacities to the contents with higher β/s will eventually generate more worth than other policies which do not. From Table I, we can check that the ratio of the assigned capacities to each content is given by:

$$\text{RPP: } \frac{\beta_i X_{r,i}}{\beta_j X_{r,j}} \quad \text{and} \quad \text{TPP: } \frac{s_i X_{r,i}}{s_j X_{r,j}}.$$

Also, from $\beta_i/s_i \geq \beta_j/s_j$, we have $\beta_i/\beta_j \geq s_i/s_j$. This means that, for example, supposing that a unit capacity is assigned to q_j , then in RPP, the very content prioritized in case of PP, i.e., q_i , is allocated more capacity than that in TPP. This value-based inter-content preference in RPP allows us to have more total worth in RPP than in TPP.

V. STABILITY OF GRAND COALITION

A. Under-demanded Network

In under-demanded networks, RSG was super-additive, and thus the worth is maximized at the grand coalition for any scheduling policy, and the grand coalition tends to be preferred. However, it is not straight-forward whether the grand coalition is stable in the sense of Definition 1, because there may be a smaller coalition in which the players in that coalition can obtain larger individual shares than in the grand coalition. Examples that a (coalition) game is super-additive but unstable include the famous n -person symmetric majority game [9]. However, in RSG, as stated in Theorem 2, the grand coalition is provably stable under SV, which is conjectured in [5], [19]. To explain why, recall that the Shapley value in (2) quantifies the average marginal contribution, meaning that a player with higher contribution will be given higher share. In n -person symmetric majority game, all players’ contributions on the coalition worth are the same, thus they equally share the coalition worth under SV. The majority coalition’s worth is always one, thus an arbitrary player’s Shapley value under the grand coalition is always smaller than one under the other majority coalition.

However, in RSG, since our model assumes only a single transit ISP, the transit ISP is always included all the positive worth coalition. It implies that the transit ISP’s contribution on the grand coalition’s worth is larger than or equal to that on other coalitions’, i.e. $\varphi^T(\mathcal{N}) \geq \varphi^T(S)$, for any $S \subset \mathcal{N}$. Then, the transit ISP becomes a Shapley-advocating player for any coalition S . Hence, by Definition 1, the grand coalition is stable in under-demanded network in RSG.

Theorem 2 (Stability of GC: Under-demanded) In RSG, the grand coalition is stable under SV for under-demanded networks.

The proof is presented in [18].

B. Stability and Content-value Preference

The challenge of stability analysis for over-demanded networks lies in the fact that a player’s actual SV is computed with traffic scheduling policy. This complex coupling makes the analysis almost impossible in general network with multiple transit ISPs and eyeball ISPs. Thus, for tractable analysis,

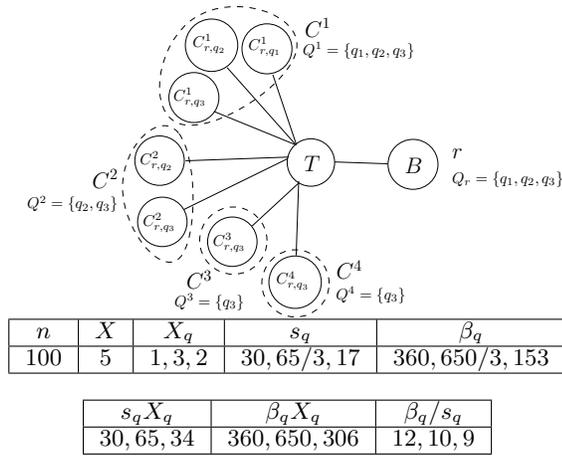


Fig. 3. Example topology: The grand coalition is stable under some policies, but unstable under other policies. Three values in the table refers to ones for three contents q_1 , q_2 , and q_3 .

we consider a simplified model that has one transit and one eyeball ISP under *heavy content regime* as shown in Fig. 3. By heavy content regime, we mean that for a given over-demanded network, we assume that removing one content in the network lets the network be under-demanded, i.e., only the grand coalition is over-demanded and any other smaller coalitions are all under-demanded. This restriction is not just for tractable analysis, but also reflects that “light contents”, whose total traffic volume is not significant, is unlikely to significantly impact the stability. If needed, we use the subscript T , R , or P , to explicitly express SV’s dependence on each scheduling policy, e.g., $\varphi_P^i(\mathcal{N})$ or $\varphi_R^i(\mathcal{N})$.

We define: A policy Π is said to *stability-dominate* another policy Π' (for simplicity, we denote as $\Pi \geq_S \Pi'$), if for a given (over-demanded) network, the stability of the grand coalition under Π' implies that under Π . Lemma 1 is useful in the sense that it suffices to check the Shapley value of either transit or eyeball ISP to check the stability of the grand coalition.

Lemma 1 A policy Π stability-dominates another policy Π' , if the player i , which is either the transit ISP T or the eyeball ISP B , satisfies

$$\varphi_{\Pi}^i(\mathcal{N}) \geq \varphi_{\Pi'}^i(\mathcal{N}). \quad (3)$$

The proof is presented in [18], whose sketch is as follows: Since we assume heavy-content regime, we have $\varphi_{\Pi}^i(S) = \varphi_{\Pi'}^i(S)$ for any S which is strictly smaller than the grand coalition. Hence, it is clear that if $\varphi_{\Pi}^i(\mathcal{N}) \geq \varphi_{\Pi'}^i(\mathcal{N})$ and the grand coalition is stable under Π , then $\varphi_{\Pi'}^i(\mathcal{N}) \geq \varphi_{\Pi'}^i(S)$ for all $S \subseteq \mathcal{N}$, thus the grand coalition is stable under Π' .

In the following section, we will compare the traffic scheduling policies in the context of the stability-dominance for a given network. In Lemma 1, we have proved that the stability dominance among traffic scheduling policies is determined by SV of either transit or eyeball ISP at the grand coalition.

C. Impact of Traffic Scheduling Policies on Stability

We first introduce the notion of Shapley portion.

Shapley Portion. Note that each decomposed coalition $S_{r,q}$ is *simple* [20] in the sense that each player’s marginal contribution is 0 or the entire worth of $S_{r,q}$. Then, it is easy to see that: for any coalition S ,

$$\varphi^i(S) = \sum_{r \in \mathcal{R}[S]} \sum_{q \in \mathcal{Q}_r[S]} \phi^i(S_{r,q}) \cdot v(S_{r,q}) \quad (4)$$

where $\phi^i(M)$ is the *Shapley portion* of the player i for coalition M , defined by:

$$\phi^i(M) \triangleq \frac{1}{|M|!} \sum_{\pi \in \Pi_M} \mathbf{1}_{\{\Delta_i(v, S(\pi, i)) > 0\}}, \quad (5)$$

where $\mathbf{1}_{\{\cdot\}}$ is the indicator function. This observation has also been used in the prior work [5], which facilitates the analysis of this paper. For simplicity, henceforth omit the subscript r , because of our model of the single region network.

GC is not always stable under PP. It might be very tempting to conjecture that the grand coalition is stable under PP, because of the nice properties of PP such as worth maximization at the grand coalition and worth optimality across all possible scheduling policies, from Theorem 1. Counter intuitively, it is not always true, which is exemplified in Fig. 3. In this example, the grand coalition is not stable because the Shapley value of the transit ISP in a smaller coalition can be larger than that of the grand coalition, i.e., $\varphi^T(\mathcal{N}) = 411.83$, and $\varphi^T(\mathcal{N} \setminus \{C_{r,q_1}^1\}) = 413.63$. This instability even in PP can arise from the characteristic of the Shapley value (or the Shapley portion): Note that $\beta_1/s_1 > \beta_2/s_2 > \beta_3/s_3$ from Fig. 3, but $\phi^T(\mathcal{N}_{q_1}) < \phi^T(\mathcal{N}_{q_2}) < \phi^T(\mathcal{N}_{q_3})$. It implies that the grand coalition is not stable since the content q_1 has too small Shapley portion, i.e., q_1 is served by a small number of content ISPs, even though it is a valuable content.

We now present our main result which states which policies are better in terms of stability. Theorem 3 states that a scheduling policy with value-preference is better than other policies that are more agnostic to content values, under some mild condition.

Theorem 3 (Stability of GC: Over-demanded) For an over-demanded network with one transit and one eyeball ISP under the heavy content regime,

$$\text{PP} \geq_S \text{RPP} \geq_S \text{TPP}, \quad (6)$$

if the following condition is met:

$$\mathbf{C1}: \sum_{i,j \in \mathcal{Q}} (s_i s_j X_i X_j) \times \left(\frac{\phi^T(\mathcal{N}_{q_i}) \beta_i}{s_i} - \frac{\phi^T(\mathcal{N}_{q_j}) \beta_j}{s_j} \right) \left(\frac{\beta_i}{s_i} - \frac{\beta_j}{s_j} \right) \geq 0. \quad (7)$$

The proof is presented in [18].

We cannot theoretically guarantee that **C1** is always true, but the following numerical example, as shown in Fig. 4, tells us that **C1** tends true. We consider a network consisting of 50 content ISPs, one transit, and one eyeball ISP. We assign the size and the value of each content following what is seen as Fig. 4(a), the assignment rule of Fig. 4(a) is such that $\beta < \sqrt{s}$. This choice is somewhat artificial, but comes from the trend

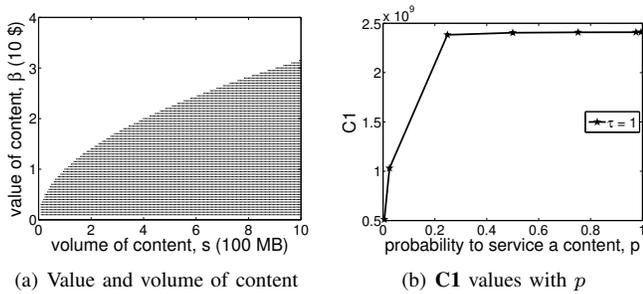


Fig. 4. A numerical test that shows **C1** is true

that the content value does grow in less proportion to the size [21]. Moreover, we assume that the total network population is 100,000 and the distribution of content popularity follows Zipf-like distribution with parameter $\tau = 1$. We vary content availability (i.e., the number of content ISPs serving the same content), by introducing a content selection probability p that corresponds to the probability that a content ISP selects any content (thus, the mean number of contents owned by one content ISP is $8150 \times p$). Then, $|C_{q_i}| \sim \text{Binomial}(8150, p)$, for all contents q_i . Fig. 4(b) shows how **C1** varies for $p = 0.005, 0.025, 0.25, 0.5, 0.75, 0.9975, 0.995$. Finally, we observe that all cases show positive **C1**.

To understand intuitively, let $c_{ij} = (s_i s_j X_i X_j) \left(\frac{\phi^T(\mathcal{N}_{q_i}) \beta_i}{s_i} - \frac{\phi^T(\mathcal{N}_{q_j}) \beta_j}{s_j} \right) \left(\frac{\beta_i}{s_i} - \frac{\beta_j}{s_j} \right)$. The main reason why the condition **C1** is mild is summarized as: (i) c_{ij} is positive in most of content pairs i, j , and (ii) it may not have a large absolute value even when c_{ij} is negative. Note that c_{ij} is a multiplication of two terms, where the first term is $s_i s_j X_i X_j$ and the second term is $\left(\frac{\phi^T(\mathcal{N}_{q_i}) \beta_i}{s_i} - \frac{\phi^T(\mathcal{N}_{q_j}) \beta_j}{s_j} \right) \left(\frac{\beta_i}{s_i} - \frac{\beta_j}{s_j} \right)$. For (i), without loss of generality, we assume that $\beta_i/s_i > \beta_j/s_j$. Then, we will show that it is highly likely true that $\frac{\beta_i/s_i}{\beta_j/s_j} > \frac{\phi^T(\mathcal{N}_{q_j})}{\phi^T(\mathcal{N}_{q_i})}$. The Shapley portion of the transit ISP in the decomposed coalition containing q_i is given by: $\phi^T(\mathcal{N}_{q_i}) = \frac{1}{2} - \frac{1}{(|C_{q_i}|+1)(|C_{q_i}|+2)}$, where $|C_{q_i}|$ is the number of content ISPs serving q_i . It implies that the Shapley portion of content q_i has the value between $1/3$ and $1/2$ ($1/3$ is when $|C_{q_i}|$ is zero and $1/2$ is when $|C_{q_i}|$ goes to ∞ , respectively). However, the difference of normalized content values of two contents, $(\beta_i/s_i)/(\beta_j/s_j)$, is quite significant, e.g., in Fig. 2, the normalized content value of a music content is four times larger than that of a video content. Therefore, the normalized content values of two contents are not similar in most of cases, and thus the second term of c_{ij} tends to be positive. Also, it is hard for us to see the case when the first term is huge, but the second term is negative simultaneously. Even such a case occurs, the same s_i and X_i can also be matched with the case when the second term is positive, contributing to making **C1** positive.

VI. CONCLUSION, LIMITATION, AND FUTURE WORK

In this paper, we have studied the coalition worth and the stability of the grand coalition under Shapley-value based revenue sharing. The main challenges come from the network that is over-demanded, in which case a certain traffic scheduling

algorithm should be applied to sift out the original user demand. This significantly changes how the individual players are assigned their own share, depending on the underlying traffic scheduling policy. Especially, we consider three policies which have different features in terms of content-value preference and implementation complexity. The scheduling with higher content-value preference is shown to have larger coalition worth and tends to form the grand coalition more strongly, but even Priority scheduling which gives absolute priority to higher-value traffic is shown to be unstable for over-demanded networks. Our work is limited in the sense that we only consider a single transit ISP, and sometimes a single region network (as done in the stability analysis). Our future work is the extension to more general networks.

REFERENCES

- [1] A. Dhamadhare and C. Dovrolis, "The Internet is flat: Modeling the transition from a transit hierarchy to a peering mesh," in *Proc. ACM CoNEXT*, Dec. 2010.
- [2] J. Musacchio, G. Schwartz, and J. Walrand, "A two-sided market analysis of provider investment incentives with an application to the net-neutrality issue," *Review of Network Economics*, vol. 8, no. 1, pp. 22–39, 2009.
- [3] A. Odlyzko, "Network neutrality, search neutrality, and the never-ending conflict between efficiency and fairness in markets," *Review of Network Economics*, vol. 8, pp. 40–60, 2009.
- [4] 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.
- [5] R. Ma, D. Chiu, J. Lui, V. Misra, and D. Rubenstein, "On cooperative settlement between content, transit and eyeball Internet service providers," in *Proc. ACM CoNEXT*, Dec. 2008.
- [6] R. Ma, D. Chiu, S. Lui, V. Misra, and D. Rubenstein, "Internet economics: The use of Shapley value for ISP settlement," *ACM/IEEE Trans. Networking*, vol. 18, no. 3, pp. 775–787, 2010.
- [7] S. Hart and M. Kurz, "Endogenous formation of coalitions," *Econometrica*, vol. 51, pp. 1047–1064, 1983.
- [8] A. Tadic, "The Aumann-Drèze value, the Wiese value, and stability: A note," *International Game Theory Review*, vol. 12, no. 2, pp. 189–195, 2010.
- [9] S. Hart and M. Kurz, *Stable Coalition Structures*. Physica-Verlag, 1984.
- [10] L. He and J. Walrand, "Pricing and revenue sharing strategies for Internet service providers," in *Proc. IEEE INFOCOM*, 2005.
- [11] P. Hande, M. Chiang, R. Calderbank, and J. Zhang, "Pricing under constraints in access networks: Revenue maximization and congestion management," in *Proc. IEEE INFOCOM*, 2010.
- [12] H. Jang, H. Lee, and Y. Yi, "On the interaction between ISP revenue sharing and network neutrality," in *Proc. ACM CoNEXT*, Dec. 2010.
- [13] J. Cho and Y. Yi, "On the Shapley-like payoff mechanisms in peer-assisted services with multiple content providers," in *Proc. GameNets*, Apr. 2011.
- [14] Z. Abbassi and V. Misra, "Multi-level revenue sharing for viral marketing," in *Proc. NetEcon*, Jun. 2011.
- [15] P. Antoniadis, S. Fdida, T. Friedman, and V. Misra, "Federation of virtualized infrastructures: Sharing the value of diversity," in *Proc. ACM CoNEXT*, Dec. 2010.
- [16] H. Lee, H. Jang, J. Cho, and Y. Yi, "On the stability of ISPs' coalition structure: Shapley Value based revenue sharing," in *Proc. CISS*, 2012.
- [17] R. Aumann and J. Drèze, "Cooperative games with coalition structures," *International Journal of Game Theory*, vol. 3, pp. 217–237, 1974.
- [18] H. Lee, H. Jang, Y. Yi, and J. Cho, "On the interaction between content-oriented traffic scheduling and inter-provider revenue sharing," Dept. of EE, KAIST, http://lanada.kaist.ac.kr/pub/sv_isp.pdf, Tech. Rep., 2012.
- [19] P. Faratin, D. Clark, P. Gilmore, S. Bauer, A. Berger, and W. Lehr, "Complexity of Internet interconnections: technology, incentives and implications for policy," in *Proc. The 35th Research Conference on Communication*, Sep. 2007.
- [20] P. P. Shenoy, "On coalition formation: a game-theoretical approach," *International Journal of Game Theory*, vol. 8, no. 3, pp. 133–164, 1979.
- [21] A. Odlyzko, "The volume and value of information," *International Journal of Communication*, vol. 6, pp. 920–935, 2012.