On The Problem of Revenue Sharing in Multi-domain Federations*

Isabel Amigo[1,2], Pablo Belzarena[1], and Sandrine Vaton[2]

[1] Facultad de Ingeniería, Universidad de la República, Montevideo, Uruguay
Julio Herrera y Reissig 565 CP 11300
[2] TELECOM Bretagne, Brest, France
Technopole Brest Iroise CS 83818 29285 Brest Cedex 3
iamigo@fing.edu.uy, belza@fing.edu.uy, sandrine.vaton@telecom-bretagne.eu

Abstract. Autonomous System alliances or federations are envisioned to emerge in the near future as a means of selling end-to-end quality assured services through interdomain networks. This collaborative paradigm mainly responds to the ever increasing Internet traffic volumes that requires assured quality, and constitutes a new business opportunity for Network Service Providers (NSPs). However, current Internet business rules are not likely to satisfy all involved partners in this emerging scenario. How the revenue is shared among NSPs must be agreed in advance, and should enforce economical incentives to join an alliance and remain in it, so that the alliance remains stable. In this paper, we work on the scenario of such federations, where service selling is formulated as a Network Utility Maximization (NUM) problem. In this context, we formally formulate the properties the revenue sharing (RS) method should fulfill and argue why the existing methods are not suitable. Finally, we propose a family of solutions to the RS problem such that the economical stability and efficiency of the alliance in the long term is guaranteed. The proposed method is based on solving a series of Optimization Problems and considering statistics on the incomes.

Keywords: Revenue sharing, Optimization, Autonomous Systems Alliances.

1 Introduction

Internet traffic consumption tendencies are evolving along two main axis. On the one hand, the continuous growth in terms of volume as well as in terms of Quality of Service (QoS) demanding applications, such as telepresence, video or gaming [3]. On the other hand, the need for QoS connectivity from end-to-end across several domains or ASs, which posses political, economical and technical issues [15]. In addition, there is a need for NSPs to find new business cases and technology for fulfilling customer needs and maximizing incomes.

* This work was funded by the ETICS project (EC FP7 248567)[1] and the Uruguayan Agency for Research and Innovation (PR-POS-2008-003)

Nowadays Internet business rules for domain interconnection (peering agreements and customer-provider agreements) may not be able to provide a sustainable economy for all actors in the value chain (i.e. Application Providers, NSPs, etc.). Indeed, non of such rules are aware of the QoS capabilities of the different domains, and peering ones are based on a traffic-symmetry premise that may no longer be valid in evolving services (for instance, on-line video, which is foreseen as one of the services that is going to grow the most [3]).
Further on, a common way of pricing for Internet connection is a monthly flat rate, while other actors like Application Content Providers or the so-called Over-The Top-Providers receive revenues on a per bandwidth-consumed basis, relying their services on the existing network infrastructure but not remunerating network providers adequately [12].
Taking into account the previous considerations many companies and academic groups are analyzing different future scenarios in order to meet the end-to-end requirements by defining inter-domain architectures and business models. As a possible architecture to provide these end-to-end QoS enabled services, the concept of ASs alliances or federations has emerged (see for instance [1]). In this kind of interconnection market scenario there exists a cooperation among ASs in infrastructure, policies and incentives for rational usage of resources and agreements for providing end-to-end QoS. While at the same time, challenging issues arise, such as priorities and revenue sharing.
This work focuses on the previously described context: several ASs that create a federation and offer for sell different services. In particular, we shall work on the framework presented on [4]. In such framework, for each service there is a group of users interested in buying it. The total income is a function of the whole bandwidth demanded for all the instances of all the services. The federation may not be able to sell all the demand of bandwidth for each service, because of capacity constraints in their data networks. The objective of the federation is thus, to allocate bandwidth in such a way that the revenue of the federation is maximized.
Once the services are sold, the income has to be split among all the providers involved in the federation. The problem of how to make this sharing is not an easy one. There are some properties that must be fulfilled by the solution in order to make the sharing fair, and interesting for all the providers, such that the alliance remains stable.
This work aims at shading light into the RS problem. In this sense we provide the following contributions: formal representation of the problem and discussion of the desired properties, evaluation of existing methods which concludes that none of them are suitable for our problem, guidelines for a new method, and a solution proposal. The method is validated through simulation studies.
This paper is organized as follows. Section 2 introduces the notations used throughout the paper and states the desired properties for the RS mechanism. In Section 3 we review the most common sharing rules used in the economics field, and argue on why they are not useful for our problem. This yields to presenting in Section 4 a new method, which provides with a solution that guarantees stability and efficiency in economical terms. Simulations results that demonstrate the correct behavior of the proposed method are shown in Section 5. Finally, concluding remarks and future work are addressed in Section 6.

2 Problem Description

2.1 Definitions and Notations

We first introduce the notation needed to represent the interdomain network of providers. Because of confidentiality and scalability issues we need to abstract each provider's topology to a simpler one. In particular we shall consider each AS as a node, which is a very simple abstraction, though reasonable enough for our study. In addition, more complex topology abstraction approaches could also be applied, without implying any change in the mathematical formulation of the problem (see for instance [2],[10]). The set of ASs or nodes is called N, there are $|N|$ nodes in the network, where the notation $|\cdot|$ refers to the cardinal of the set. Each node $n \in N$ has an equivalent capacity associated to it, which we call c_n, $c = \{c_n\}_{n \in N}$ is the vector of nodes capacities. S is the set of services offered by the network, there are $|S|$ services. A service can be abstracted to a tunnel that carries bandwidth from one ingress to one egress node of the network. The routing matrix R indicates the routes of all services in S, i.e. the nodes traversed by each service $s \in S$. More formally, R is a $|N| \times |S|$ matrix whose entries $\{R_{n,s}\}_{n \in N, s \in S}$ are equal to 1 if and only if route s traverses node n, and are zero otherwise.
The amount of bandwidth traversing each service route is indicated by each component a_s, $s \in S$ of column vector $a \in R^{|S|}$. There is a utility function associated to each service $s \in S$ which is called U_s and it is a function of a_s. The utility U_s is the willingness to pay of the group of users interested in service $s \in S$. We assume that $U_s(a_s)$ is known and, as usual in this context, it is a strictly non-decreasing concave function of the bandwidth. For more details on the allocation of the bandwidth traversing the federation the reader is referred to [4].
We now introduce extra notation in order to represent the RS problem. The grand coalition is N, the set of all nodes in the network. The income is assessed by the revenue function $V : 2^{|N|} \to \mathbb{R}$ which associates to each subcoalition $\mathcal{Q} \subseteq N$ (i.e. subgroups of nodes) a real value $V(\mathcal{Q}, c^{\mathcal{Q}})$, where $c^{\mathcal{Q}}$ is the capacities vector restricted to subcoalition $\mathcal{Q}$, that is:

$$c_n^{\mathcal{Q}} = c_n \text{ if } n \in \mathcal{Q} \text{ ; } c_n^{\mathcal{Q}} = 0 \text{ otherwise.} \tag{1}$$

In our problem, the revenue function V is given by the solution of Problem 1. This problem states that services are sold (i.e. bandwidth is allocated) such that the revenue of the coalition is maximized, while respecting the capacity constraints. It is thus formulated as a NUM problem [11].

Problem 1.

$$\max_a \sum_{s \in S} U_s(a_s)$$
$$\text{s.t.} Ra \leq c^{\mathcal{Q}}, a_s/\sigma_s \in \mathbb{N} \, \forall s \in S,$$

where σ_s is the amount of bandwidth provided by service s. We also accept the notation $V(\mathcal{Q})$ to indicate the total revenue of coalition $\mathcal{Q} \subseteq N$, where the capacities are implicit. This optimal revenue problem is not an easy integer program and its convex relaxation can be not exact.

However, since integer programming is NP hard, we accept sub-optimal solutions by convex relaxation that in many cases leads to tight solutions. We define the contribution v_n of node $n \in N$ to the coalition as $v_n = V(N) - V(N \setminus \{n\})$. The total revenue is shared among all the nodes in N according to the sharing function $\Phi : \mathbb{R} \to \mathbb{R}^{|N|}$, $\Phi = \Phi(N, c) = \{\phi_n(N, c)\}_{n \in N}$, where ϕ_n is the share corresponding to node $n \in N$. For convenience, we sometimes also use the shorter notation x to denote the RS vector, where $x \in \mathbb{R}^{|N|}$ is a column vector containing on each component x_n, $n \in N$, the revenue share of node n, when the values of N and c are implicit by context.

2.2 Desired properties of the RS mechanism

We shall now state the properties that a revenue mechanism for ASs alliances should fulfill. The idea that motivates all of them is that the ASs should be encouraged to remain in the coalition, which will occur if that makes sense from the economical point of view. The properties discussed below, and more, are usually discussed in cost/revenue sharing problems, with slightly different definitions (see for instance [9],[6]). We select from them the ones that we believe are of more relevance to our problem and formally define them.

Efficiency. The mechanism should distribute the whole revenue among all the ASs in the federation, that is

$$\sum_{n:n \in N} x_n = V(N). \tag{2}$$

Stability. In order to assure the sustainability of the federation, the mechanism should not provide incentives to any subcoalition to break the grand coalition. That is to say, no subcoalition should have economic incentives to form a smaller coalition outside the grand coalition, since this would lead to instabilities in the federation. This can be written as:

$$\sum_{n:n \in \mathcal{Q}} x_n \geq V(\mathcal{Q}), \, \forall \mathcal{Q} \subseteq N. \tag{3}$$

Please note that this definition also implies another interesting property usually known as *stand alone*. This means that the revenue perceived by every node $i \in N$ in the coalition is not less than the revenue it could achieve alone, i.e. $x_i \geq V(\{i\})$, $\forall i \in N$. The set of points that verify (3) constitutes the so-called *core set* in the context of Coalitional Game Theory. The reader is referred to [17] for more details on the core concept and coalitional game theory.

Monotonicity. The mechanism should provide the right incentives to the nodes to increase their resources towards the coalition. In our model, these resources are considered in the capacity. We formally define this property as follows. Given c and $\hat{c}$ two vectors of nodes capacities, such that $\hat{c}_n = c_n \; \forall n \in N \setminus \{i\}$ and $\hat{c}_i \geq c_i$ then $\phi_i(N, \hat{c}) \geq \phi_i(N, c)$.
Thus, the Monotonicity property means that if an AS increases its capacity then its revenue will as well increase or remain the same.

Fairness. We want the mechanism to be fair in the sharing. There is not a general consensus in the literature regarding the notion of fairness. However, we propose the following intuitive rules to be fulfilled. If $v_i \geq v_j$ then $x_i \geq x_j$, which is usually known as the *order preserving* property and if $v_i = 0$ then $x_i = 0$. We may also accept a weaker notion of fairness, which only asks for $x_i = x_j$ if $v_i = v_j$.

3 Existing Techniques

In this section we present existing RS techniques, which are widely used in the field of economics. A detailed review can be found in [9]. We also comment on why these techniques are not suitable for our problem.

The Proportional Share: One of the simplest way to perform the RS is the one that is proportional to the contribution of each node. Using the definitions introduced in Section 2 we write the proportional share as:

$$x_i^{prop} = \frac{v_i}{\sum_{n \in N} v_n} V(N). \tag{4}$$

The proportional share a priori seems to be a very attractive distribution rule. It fulfills the properties of Efficiency and Fairness and it is very simple to compute. However, it has the drawback that it does not always guarantee neither Stability nor Monotonicity.

The Shapley Value. The Shapley value, proposed by Lloyd Shapley in 1952 [16], provides a means for performing the RS of an association or coalition. It has been widely used in the literature for its good properties. Given a coalitional game, i.e. a pair (N, V) where N is a finite set of players and $V : 2^{|N|} \rightarrow \mathbb{R}$ any worth or revenue function, the Shapley value for player $i \in N$ is defined as:

$$x_i^{sh} = \frac{1}{|N|!} \sum_{\mathcal{Q} \subseteq N \setminus \{i\}} |\mathcal{Q}|!(|N| - |\mathcal{Q}| - 1)! \left[V(\mathcal{Q} \cup \{i\}) - V(\mathcal{Q}) \right]. \tag{5}$$

Among its properties, the Shapley vale has the Efficieny, Monotonicity and a particular case of the Fairness as defined in Subsection 2.2. It also fulfills its own definition of fairness in terms that for any players $i, j \in N$, i's contribution to j is equal to j's contribution to i.
However, the Shapley value does not always provide stable solutions. That is the reason why it is not suitable for our problem. Nonetheless, its great popularity in previous work is due to the fact that it is proven that it provides with stable solutions when the revenue function is a convex function. For instance, the Shapley value has been used in [12], where the proposal is to change the Internet economics by business contracts whose payment is determined by the Shapley value. And also in [14], where the aim is to optimize the routing within an alliance of ASs and the revenue is shared by means of Shapley value. Recently, it has also been use in [18] for splitting cost savings among several domains.
In our problem, the revenue function V is not a convex one and solutions through Shapley value can lie outside the *core*. A simple example can be found using the topology on Fig. 1b and the example in Section 4. In that case, the Shapley value renders $x = (1/2, 3, 3/2)$ which is outside the *core* while we shall show on that section that a non-empty *core* exists.

The Aumann-Shapley Rule. The Aumann-Shapley Rule for cost sharing [5] was introduced by Shapley and Aumann in 1974, and can be applied analogously for a RS problem. The idea of this rule is to compute the revenue share of node $i \in N$ as its average marginal revenue along a certain path going from capacity equal to 0 to c_i. More precisely, the share for node $i \in N$ according to this rule is defined as:

$$x_i^{as} = \int_0^{c_i} \partial_i V(N, \frac{t}{c_i} c) dt = c_i \int_0^1 \partial_i V(N, tc) dt, \tag{6}$$

where the notation $\partial_i V(N, c)$ means the first order derivative of V at c with respect to c_i. Please note that in Equation (6) we have used the alternative notation for V where it is explicitly mentioned its dependency on the subcoalition and the equivalent capacities.
In first place, it must be noticed that the derivative of V with respect to c_i is not defined for all values of c_i. Indeed, consider a simple topology with only one service crossing several nodes, which all have the same capacity. Let $\hat{c}$ be that capacity. If a given node i increases its capacity, the other nodes will act as bottlenecks and the revenue will not change, while if i reduces its capacity then it will itself become the bottleneck and the revenue will decrease. Hence, the derivative of V takes different values at both sides of $\hat{c}$ and is not defined at $c_i = \hat{c}$. In addition, this rule does not fulfill the Monotonicity property, this is due to the characteristics of our revenue function. Furthermore, this rule applied to our problem could even provide incentives to reduce capacity.

The Friedman-Moulin Rule. This rule was proposed by Friedman and Moulin in 1999 [8]. We introduce the operator $\wedge$, which is defined for two vectors a and $b \in \mathbb{R}^{|N|}$ as $a \wedge b = \min(a_i, b_i)$ $i \in N$ and column vector e, which is of dimension $|N|$ and has all its components equal to one. This rule is similar to the Aumman-Shapley one, in terms that it integrates marginal revenues, but in this case the integration is done through a different path. According to the Friedman-Moulin rule, the share for node $i \in N$ is calculated as:

$$x_i^{fm} = \int_0^{c_i} \partial_i V(N, t \cdot e \wedge c) dt. \tag{7}$$

This rule can not be applied in our context since V is not derivable along the whole path, for the same reasons explained above.

4 The Proposed Method

Having seen that existing techniques are not suitable for our problem, we shall now propose a new method to perform the RS in our specific scenario. We focus our attention on two properties: Stability and Efficiency. Nevertheless, we shall present a flexible method which allows for including further properties on future work. We first study the set of possible solutions. For clarity sake, we consider this set of solutions in a simple scenario, which we call the one-shot scenario. In this scenario services are sold through what we call a service selling (SS) phase and RS is performed right afterwards. We shall latter on move to a multi-period scenario.

4.1 The Feasible Solutions Set

In order to have stability in the coalitions inequality (3) must hold. Let us enumerate all the possible subcoalitions $\mathcal{Q} \in N$ and index them using index $j = 1 \dots 2^{|N|}$. We rewrite inequality (3) as a linear system as:

$$Qx \geq \hat{v}, \tag{8}$$

where $Q = \{Q_{j,i}\}$ is a $2^{|N|} \times |N|$ matrix that indicates which nodes belong to each subcoalition (i.e. $Q_{j,i} = 1$ if node i belongs to subcoalition j and 0 otherwise) and $\hat{v} = \{V(\mathcal{Q}_j)\}_{j=1\dots2^{|N|}}$ is the vector that indicates in the j-th component the revenue of subcoalition j.
We must consider at the same time the Efficiency property, which we write as the vector representation of Equation (2):

$$e^T x = V(N). \tag{9}$$

But do these constraints determine a unique point? Or rather they determine a set of points? Is this set empty? We show through the following examples that actually different cases can occur.

An empty feasible set. Consider the network on Fig. 1a. The capacities of all nodes are equal to 1 unit. The three services illustrated on the mentioned figure are sold, each one of them is defined for 1 unit of bandwidth. The utility functions are such that $V(N) = V(\{1,2\}) = 5$, $V(\{2,3\}) = 4$ and $V(\{1,3\}) = 2$.

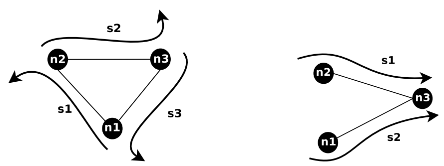

(a) (b)

Fig. 1: The feasible set in different situations.

For achieving stability the total revenue (5 units) must be split in such a way that every route receives at least what they would receive alone. It is not difficult to see that this is not possible at the same time for all routes, since the following inequalities must hold: $x_1+x_2 \geq 5$, $x_1+x_3 \geq 2$, $x_2 + x_3 \geq 4$ and $x_1 + x_2 + x_3 = 5$. Hence, the feasible set is empty.
It is interesting to remark that for different values of the revenues, and same topology, the feasible set could be non-empty.

A feasible region. Consider now the network on Fig. 1b. The capacities are again equal to 1 unit for all nodes and we sell services of 1 unit of bandwidth. Utility functions are such that $V(N) = V(\{2,3\}) = 5$ $V(\{1,3\}) = 2$. A feasible solution must fulfill $x_1 + x_3 \geq 2$, $x_2 + x_3 \geq 5$ and $x_1 + x_2 + x_3 = 5$. The vectors x that satisfy all equations are $\{x = (0, 2+\epsilon, 3-\epsilon) : \epsilon \in \mathbb{R}, 0 \leq \epsilon \leq 3\}$, which corresponds to a segment in $\mathbb{R}^2$.

4.2 One-Shot Scenario

We have seen in the previous subsection that configurations with no solution can exist, in this case we claim that the coalition should not exist as such, since there is no RS method that can make it stable. Therefore, we focus our attention on the case where constraints (8) and (9) determine a region. In order to choose a point from such region we formulate the following Optimization Problem:

Problem 2.

$$\min_{x} \ f(x)$$
$$\text{s.t. } Qx \geq \hat{v},\, e^T x = V(N),$$

where $f(x)$ is a convex function. Please note that we can dispense with the restriction of non negative revenue shares, since it is already considered on the Stability property. Problem (2) constitutes a family of methods which can be tuned to cover additional properties by considering different objective functions. Examples of objective functions are to project the Proportional share or the Shapley value into the feasible set, which means the method inherits their properties when the share is already in the feasible set, and otherwise returns the closest value. Yet another example is the square of the Euclidean norm of the share vector, which would intuitively provide with more even shares among the nodes. Regarding implementation aspects, the proposal is to have a central trusted entity computing the RS. This entity must know the utility functions for each service and the topology of the coalition, at the AS level.

4.3 Multi-period Scenario

We shall now focus our attention on the multi-period scenario, that is to say, when several phases of SS occur. A new period implies a new utility function, thus different values for $U_s, s \in S$. This necessarily leads to a different feasible set. Finding on each period a valid RS vector would involve performing a great number of computations, besides to a great exchange of information among the domains and the central entity solving the RS. In other words, the multi-period case may pose the problem of scalability thus, we face the challenge of providing a scalable approach. One could naively propose as a solution to compute the RS once, and then simply keep the sharing proportion for the subsequent RS phases. However, if we were to use the same RS proportion for a new SS instance, then the new RS vector can lie within the new feasible set or outside of it, which leads us to discard that option.
Altogether, we are motivated to perform the RS on a longer timescale than the SS phase, and work with statistics of the utilities received during the several SS phases considered for a given RS phase. In the following we shall discuss two different approaches for working with such statistics.

Approach 1. In order to model the multi-period situation, let us introduce the assumption that the utility functions depend on a value drawn independently from a continuous probability distribution for each service.

Provided this, we can safely represent the utility functions of several SS phases occurred during a certain period of time by their mean over that period of time. As usually, notation E represents the expectation of a random variable. We define the mean utility function as:

$$\overline{U}_s(a_s) = E[U_s(a_s)], \tag{10}$$

which is still a non-decreasing concave function of $a_s, \forall s \in S$. Finally, we redefine the revenue function V by Problem 3, and call it $\mathcal{V}$.

Problem 3.

$$\max_a \sum_{s \in S} \overline{U}_s(a_s)$$
$$\text{s.t.} Ra \leq c^{\mathcal{Q}}.$$

The procedure then continues solving Problem 2, but considering now $\mathcal{V}$ instead of V for the definition of $\hat{v}$.

The explained mechanism allows us to perform the computation only once in a while (e.g. monthly). In addition, the amount of information exchanged is also kept small, since the only information that has to be transmitted to the central entity on each RS phase is the mean of the utilities over that period. However, can we be sure that the solution provided by this approach is always stable in the long term? The answer to this question is addressed in the following.

Approach 2. Usually, providers' decisions are based on long term behaviors, mainly for keeping network stability. Likewise, the interest of the providers to remain in the alliance would be based on its economical stability in the long term. That is, they would likely need to know if their revenue share is economically attractive in the long term. For considering such situation we compute the long term feasible set, which is obtained after the expectation of the revenues of each subcoaliton, and obtain the RS from such set. This is summarized on Problem 4.

Problem 4.

$$\min_x \ f(x)$$
$$\text{s.t. } Qx \geq E\{\hat{v}\}, \, e^T x = E\{V(N)\}.$$

Please note that $E\{\hat{v}\}$ is obtained by computing the expectation of the output of Problem (1) for each subcoalition $\mathcal{Q} \subseteq N$.

The raised question reduces then to answering if the point chosen by Problem (3) lies within the feasible set of Problem (4) or not. Unfortunately this is not necessary true. Indeed, as shown in [13], where relationships between stochastic non-liner programming problems are demonstrated, the following inequality applies:

$$E\{V(\mathcal{Q})\} \geq \mathcal{V}(\mathcal{Q}), \forall \mathcal{Q} \subseteq N, \tag{11}$$

which means that the feasible set of Approach 2 is contained in the one of Approach 1. However, we have no indication about the tightness of the bound, thus we shall evaluate the impact of using either of both approaches by simulation, in the following section.

5 Simulations

The simulations presented in this section were performed on a regular computer with a i5 processor of 2.67GHz and 3.6 GB of RAM memory. The optimization problems were solved using CPLEX through AMPL.

5.1 One-shot scenario

We shall consider the topologies on Fig. 2, where the arrows indicate the services' paths, $c_n = 10$ for all nodes n and $\sigma_s = 5$ for all services s, all values expressed in some coherent unit. Tabel 1a shows the utility in $ for carrying 5 and 10 units of bandwidth, where the underlined values correspond to the solution of Problem 1 (i.e. the revenue).

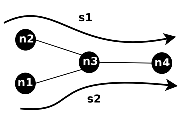

(a) 1.

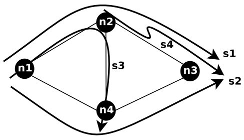

(b) 2.

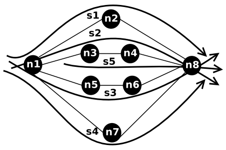

(c) 3.

Fig. 2: Topologies used throughout the simulation studies.

	Utility ($)					
Service	$U_s(5)$	$U_s(10)$	$U_s(5)$	$U_s(10)$	$U_s(5)$	$U_s(10)$
s1	1	2	5	9	<u>7</u>	8
s2	6	<u>9</u>	<u>7</u>	11	2	3
s3	-	-	<u>11</u>	16	5	8
s4	-	-	<u>12</u>	18	5	8
s5	-	-	-	-	<u>6</u>	11
	Topology 1		Topology 2		Topology 3	

(a) Utility values

Topology	Revenue Share (%)					
	x^{norm}	$x_{\perp}^{sv}$	x^{sv}	$x_{\perp}^{prop}$	x^{prop}	v_i
1	0.333	0.284	0.278	0.280	0.280	7
	0	0	0.019	0	0	0
	0.333	0.358	0.352	0.360	0.360	9
	0.333	0.358	0.352	0.360	0.360	9
2	0.200	0.175	0.175	0.200	0.211	12
	0.300	0.353	0.353	0.344	0.333	19
	0.300	0.297	0.297	0.256	0.246	14
	0.200	0.175	0.175	0.200	0.211	12
3	0.077	0.130	0.183	0.077	0.105	2
	0.077	0.023	0.076	0.077	0.105	2
	0.039	0.039	0.106	0.039	0.053	1
	0.039	0.039	0.106	0.039	0.053	1
	0	0	0.014	0	0	0
	0	0	0.014	0	0	0
	0	0	0.036	0	0	0
	0.769	0.769	0.465	0.769	0.684	13

(b) Results using different criteria.

Table 1: Revenue sharing, one-shot scenario.

Revenue shares were computed using the different rules introduced in Section 3 and the criteria introduced in Section 4. Results are shown on Table 1b (where notation $x_{\perp}^{*}$ stands for the projection of * into the feasible set and x^{norm} is the solution when $f(x) = \|x\|^2$), along with the value of v_i for each node. Topologies 2 and 3 constitute examples where the Proportional share does not lie into the feasible region, so do topologies 1 and 3 for the case of the Shapley value. Regarding the different criteria, x^{norm} shows the most even shares. Some fairness notions are also observed for most criteria but the Shapley value; the smaller the v_i, the smaller the x_i, nodes whose contribution is null (i.e. $v_i = 0$) obtain no revenue, while those with same v_i obtain the same RS.

5.2 Multi-period scenario

We now compute the solution according to Approach 1 and 2. In both cases, a number of 50 SS phases were performed before a RS phase and the projection of x^{prop} was used as criteria. Results for topology on Fig. 2b are shown on Fig. 3a. For this topology, on every RS phase the results obtained using both approaches are almost the same. Same thing occurs for all the simulations performed, in particular for the topology on Fig. 2c, whose results for selected nodes are shown on Fig. 3b.

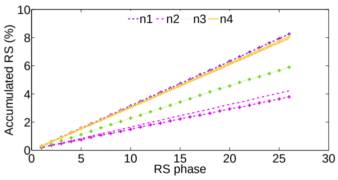

(a) Topology on Fig. 2b.

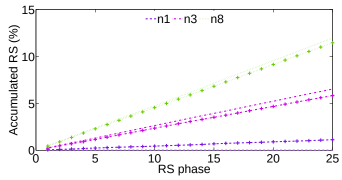

(b) Topology on Fig. 2c.

Fig. 3: RS using Approach 1 (-) and 2 (+).

We now evaluate the time consumed by each approach. We shall consider a simple topology with only one service defined and linearly increase the number of nodes in the service's path. Results show that for both approaches the time consumed by the method increases exponentially with the number of nodes in the network. This is related to the Stability property, since for taking it into account we consider all subcoalitions of nodes (i.e. $2^{|N|}$ cases). For a topology of 8 nodes, Approach 1 consumed on the average 2 ms while Approach 2 consumed 135 ms. However, Approach 2 is still feasible, moreover considering it is proposed to be performed off-line and in a long timescale. In addition, ASs alliances are likely to have no more than 10 nodes, considering for instance, that the average AS path in the Internet is of 4 ASs [7].
All in all, we can claim that Approach 2 provides with a solution that fulfills the sought properties with affordable computation time.

6 Conclusion and Future Work

The present work has addressed the problem of RS in the context of ASs alliances. We have focused on the case where the income of the alliance is determined by the output of a NUM problem. This particular scenario poses new challenges. Previous results for performing RS, which have been reviewed in this work, were found to be inappropriate applied to this case. The desired properties for the RS were formally stated and a new method has been proposed. This method is conceived for providing economical stability and efficiency to the alliance and it is flexible enough to be adapted to fulfill additional properties. The method is based on solving optimization problems and considers statistics on the income. Its proper behavior has been evaluated through simulation studies.
As future work, we shall study the inclusion of further properties into the method, as well as research on the relation between the two approaches

provided for the multi-period scenario. In particular, we are interested in including the consideration of QoS parameters associated to the services and providing incentives through the RS mechanism to guarantee them.

References

1. ETICS: Economics and Technologies for Inter-carrier Services. European research project of the FP7, `http://www.ict-etics.eu`
2. The ATM Forum, Private Network-Network Interface, specification version 1.0, March 1996
3. Cisco Systems: Entering the Zettabyte Era. Tech. rep. (2011)
4. Amigo, I., Belzarena, P., Larroca, F., Vaton, S.: Network bandwidth allocation with end-to-end QoS constraints and revenue sharing in multi-domain federations. In: ICQT 2011. pp. 50–62. Springer-Verlag, Berlin, Heidelberg (2011)
5. Aumann, R.J., Shapley, L.S.: Values of non-atomic games. Princeton University Press (1974)
6. Bogomolnaia, A., Holzman, R., Moulin, H.: Sharing the Cost of a Capacity Network. Mathematics of Operations Research 35(1), 173–192 (Feb 2010)
7. Dhamdhere, A., Dovrolis, C.: Ten years in the evolution of the internet ecosystem. In: ACM SIGCOMM 2008. pp. 183–196. ACM (2008)
8. Friedman, E., Moulin, H.: Three Methods to Share Joint Costs or Surplus. Journal of Economic Theory 87(2), 275–312 (Aug 1999)
9. Hougaard, J.L.: An Introduction to Allocation Rules. Springer (2009)
10. Karaoglu, H.T., Yuksel, M.: Value flows: Inter-domain routing over contract links pp. 327–332 (Dec 2010)
11. Kelly, F.P., Maulloo, A.K., Tan, D.K.H.: Rate Control for Communication Networks: Shadow Prices, Proportional Fairness and Stability. The Journal of Operational Research Society 49(3), 237–252 (1998)
12. Ma, R.T.B., Chiu, D.M., Lui, J.C.S., Misra, V., Rubenstein, D.: Internet Economics: The Use of Shapley Value for ISP Settlement. IEEE/ACM Transactions on Networking 18(3), 775–787 (Jun 2010)
13. Mangasarian, O.L.: Nonlinear Programming Problems with Stochastic Objective Functions. Management Science 10(2) (1964)
14. Mycek, M., Secci, S., Pioro, M., Rougier, J.L., Tomaszewski, A., Pattavina, A.: Cooperative multi-provider routing optimization and income distribution. In: DRCN 2009. pp. 281–288 (Oct 2009)
15. N. Le Sauze et al.: ETICS : QoS-enabled interconnection for Future Internet services. In: Future Network and Mobile Summit (2010)
16. Shapley, L.: A value for n-person games. In: Kuhn, H., Tucker, A. (eds.) Contributions to the Theory of Games
17. Shoham, Y., Leyton-Brown, K.: Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations. Cambridge University Press, Cambridge, UK (2009)
18. Stanojevic, R., Castro, I., Gorinsky, S.: CIPT: using tuangou to reduce IP transit costs. In: CoNEXT 2011. ACM, New York, NY, USA (2011)