Network Bandwidth allocation with end-to-end QoS constraints and Revenue Sharing in Multidomain Federations

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

[1] Facultad de Ingeniería, Universidad de la República, Montevideo, Uruguay.
[2] Telecom Bretagne, Brest, France.

Abstract. Internet is evolving, traffic continues to grow, new revenue sources are sought by Network and Service Providers. Value added services with real time characteristics are likely to be common currency in the near future. Quality of Service (QoS) could allow Application/Service Providers (APs) to offer better services to the end users. At the same time, all actors claim for a fair distribution of revenues. Inspired by this scenario, we propose a complete framework for selling interdomain quality assured services, and subsequently distributing revenues, in an Autonomous System (AS) association context. We state the problem as a network utility maximization problem with QoS constraints and show that a distributed solution can be carried out. In order to fairly share the resulting revenue we study concepts from coalitional game theory and propose a solution based on the Shapley value and statistics on the revenues. Simulations of the whole proposal are shown.

Keywords: Auctions, QoS, Shapley Value.

1 Introduction

Internet traffic is likely to continue increasing in a non-stop fashion. Recent studies [2] show that not only the tendency is to increase in amount but in quality requirements as well, since the applications with real time characteristics, such as Gaming and Video on Demand, are envisioned to have great increase.
Nowadays, the focus of telecommunication market is on best effort traffic and in order to meet customer expectations telecommunications companies are forced to invest in capacity, without getting sufficient return on these investments to have sustainable businesses. The ever evolving features provided by the handset terminals, and the growing number of connection capable equipments, constitute more evidence in favor of the forecast of Internet traffic increase.
Moreover, emerging technologies such as telepresence or cloud computing not only generate large volumes of traffic with real time requirements, but are also used to interconnect sites around the globe. As a consequence, in addition to a QoS capable network, this kind of services require an end-to-end QoS enabled chain crossing heterogeneous carrier networks [15].

In this scenario, nowadays Internet business rules for domain interconnection may not be able to provide a sustainable economy for all actors in the value chain (Application Providers, Network Service Providers, etc.). Indeed, these rules (peering agreements) are not aware of the QoS capabilities of the domains and most of them are based on a traffic-symmetry premise that may no longer be valid in evolving services (for instance HD video on demand). Moreover, a common way of pricing for Internet connection is a monthly flat rate, while other actors, e.g. APs or the so-called Over the Top Providers (OTTs) receive revenues on a per bandwidth-consumed basis, relying their services on the existent network infrastructure but not remunerating Network Providers adequately [11]. Taking into account the previous considerations many companies and academic groups are analyzing future scenarios so as to meet the end-to-end requirements and business models. As a possible architecture to provide these services, the ASs alliances or federations have emerged (see for instance [1]). In this kind of interconnection market there exists a cooperation on infrastructure, policies and incentives for rational usage of resources and agreements for providing end-to-end QoS. At the same time, interesting issues arise, such as priorities and revenue sharing. In this work, we aim at providing a framework in that sense.
We shall focus on a scenario in which ASs work together in a collaborative way in order to sell end-to-end quality assured bit pipes. The pipes are not necessarily sold to the final user but are rather sold to intermediate actors like brokers or OTT which will in turn resell them to the final user, by providing their own services through a quality assured path.
In this context, our contribution is actually twofold. On the one hand, we address the bandwidth allocation problem providing a solution through which the end-to-end quality parameters are assured and the revenue of the whole alliance is maximized. In addition, we prove that this mechanism can be carried out on a distributed fashion. On the other hand, we cover a subsequent problem that is how to distribute the revenues among all the members of the alliance. In this regard, we provide a mechanism that has fairness properties and provides incentives to the ASs to increase their features towards the federation. Beyond the specific contributions, the proposed framework links the revenue income mechanism with the revenue sharing one, which to the best of our knowledge has not been proposed in this context before.

2 Bandwidth allocation with end-to-end QoS Constraints

We are interested, as aforementioned, on a scenario where several ASs work together to sell capacity on a multidomain quality assured path. We shall refer to the quality assured path as *QoS pipe* or as *path*.
In this scenario, the capacity dedicated by each AS to sell by this means is a portion of their already deployed capacity. That is to say, ASs have their infrastructure through which traditional services are sold following the best-effort paradigm and they decide to dedicate some portion of their capacity to the federation.

For each QoS pipe there is a group of users or buyers interested on getting a portion of bandwidth on that pipe. The amount of money this group is willing to pay for each value of bandwidth is the so-called utility function. The objective is to sell the available resources in such a way that the revenue of the whole alliance is maximized while the end-to-end constraints are accomplished. We shall work with the end-to-end delay.
Let us introduce some notations so as to formally represent the scenario described above. Each AS in the alliance is abstracted to a node indexed by n with an equivalent capacity of c_n. The complete set of nodes is denoted by N. The available pipes are the ones in the set S and are indexed by s. The constraint on the delay on path s (i.e. the maximum admissible delay) is denoted by D_s. We assume that the routes within the alliance are fixed and single-path. We represent these routes with the $|N| \times |S|$ matrix R, where the notation $|\cdot|$ refers to the cardinal of the set. The entry $R_{i,j}$ is equal to 1 if the route of the pipe j traverses the node i and is equal to zero otherwise. We denote pipe's s route as $r(s)$. The bandwidth allocated to pipe s (i.e. the amount of traffic sold to the buyers associated to path s) is denoted by a_s. The utility function associated to each path s is called $U_s(a_s)$. We assume that $U_s(a_s)$ is known and, as usual in this context, it is a strictly concave function of the bandwidth.
Please note that the QoS pipes are defined by an ingress and egress point along with a maximum delay. This implies that two QoS pipes are considered different even if they share exactly the same physical path but provide different delay bounds.
Let us now state some additional assumptions. The delay introduced by each node in a path is an increasing convex function of the bandwidth carried by all the paths traversing the node. We assume that this function can be somehow learned or estimated by the domain, and we leave out of the scope of this paper the means for computing it. The delay function of node n is denoted as $f_n(a_n)$ where $a_n = \sum_{s\in S:n\in r(s)} a_s$.
The amount of traffic sold to all paths must be such that the revenue perceived by the alliance is maximized while the QoS constraints are fulfilled. This is formalized in the following bandwidth allocation problem:

Problem 1.

$$\begin{aligned} &\max_{a_s} \sum_{s:s\in S} U_s(a_s) \\ &\text{s.t.} \sum_{n:n\in r(s)} f_n(a_n) \le D_s, \; \forall s \in S. \end{aligned}$$

Remark 1. In Prob. 1 we have not included a capacity constraint which is assumed to be taken into account in f_n. Indeed, if f_n is a barrier function (i.e. it approaches infinity as the bandwidth approaches the capacity) we can safely ignore any capacity constraint.

Remark 2. The fact that the association may not want to sell bandwidth on a certain path if the incomes perceived by doing so are lower than a certain bound is not considered either. However, we can model this situation by defining a cost function of the allocated bandwidth $\kappa_s(a_s)$

for each service $s \in S$ and modifying the objective function in Prob. 1 by $\sum_{s\in S}[U_s(a_s) - \kappa_s(a_s)]$. Provided the cost function is convex, the new problem would be analogous to Prob. 1. For the sake of notations simplicity we shall not consider the cost function hereafter.

We aim at solving Prob. 1 in a distributed way. Hence, we shall explore a primal-dual approach for Prob. 1, whose associated Lagrangian is:

$$L(a, \lambda) = \sum_{s:s\in S} \left[U_s(a_s) + \lambda_s \cdot \left(D_s - \sum_{n:n\in r(s)} f_n(a_n) \right) \right], \qquad (1)$$

where $\lambda = \{\lambda_s\}_{s\in S}$ is the vector of Lagrange multipliers.
To find a saddle point of (1) (i.e. the optimum of Prob. 1) we use the gradient-projection algorithm updating the primal and dual variables as follows:

$$a_s^{t+1} = \left[a_s + \gamma_s \left(U_s^{'}(a_s) - \sum_{n:n\in r(s)} \sum_{v:n\in r(v)} \lambda_v f_n^{'}(a) \right) \right]^+ \qquad (2)$$

$$\lambda_s^{t+1} = \left[\lambda_s - \alpha_s \left(D_s - \sum_{n:n\in r(s)} f_n(a_n) \right) \right]^+ \qquad (3)$$

where $[\cdot]^+ = \max\{0, \cdot\}$ and α_s, γ_s are step sizes.
The updates (2,3) are performed iteratively on each edge router of a pipe, which we call the *source*. Every source sends an initial value for λ_s and a_s through route $r(s)$. Each node receives all the values and computes the delay, the derivative of the delay times the sum of the lambdas it has received and sends them to the source. All these values can be accumulated in two sums in the way back to the source, thus only two values are needed to be sent back to the source on each iteration. Once the source receives such values it proceeds to update the value in λ_s and in a_s. This is repeated iteratively in the control plane and it is run prior to any resource allocation.
The following theorem proves the convergence of the algorithm.

Theorem 1. *Convergence of the primal-dual algorithm. Given Prob. 1 let* $\sum_s U_s(a_s)$ *be a strictly concave function and* $f_n(a_n)\,\forall n \in N$ *convex functions. Then the iterations* $a_s^t\,\forall s \in S$ *as defined in* (2) *and* (3) *converge asymptotically to the solution of Prob. 1.*

The proof is not provided here for lack of space reasons. The main idea of it is proving that the problem can be reduced to the one in [8].

2.1 Application: Multidomain Network Auctions

We now discuss an example that fits to the model proposed before. We associate to each pipe a service to be sold which has a certain bandwidth σ_s and an assured delay D_s (for instance, this service can be a VoD movie). Several instances of a service are sold through the same pipe.

These services are sold by means of network bandwidth auctions. In particular, we shall follow the first price auctions model where the winner user is charged with the amount he/she bids. This bidding mechanism is the most suitable to our problem as explained in Sect. 5.
We shall first consider the case of one-shot bandwidth auctions. That is to say, that the whole capacity available for providing the services is going to be auctioned at a certain moment.
Let us introduce some new notations. For each service s there are N_s buyers or users, which participate in the auction for obtaining an instance of the service. Each of the N_s users bids $b_s^{(i)}$ which we order as

$$b_s^{(1)} \geq b_s^{(2)} \geq \cdots \geq b_s^{(N_s)}. \tag{4}$$

The resource allocation decision is to find which of these bids to accept, so as to maximize the profit of the whole alliance while the per-route delay remains smaller than a given bound, under a first-price auction. Since for each s all bids are for the same bandwidth and delay constraint, the optimal solution is accepting the highest bids per service. We define the variable $\psi_{s,i}$ which is equal to 1 if bid i for service s is accepted, and zero otherwise. Then, defining the variable m_s as the number of bids accepted for service s we have the following equality:

$$\sum_{i=1}^{N_s} b_s^{(i)} \psi_{s,i} = \sum_{i=1}^{m_s} b_s^{(i)}. \tag{5}$$

Accepting m_s bids would render a total accepted rate of a_s where $a_s = \sigma_s m_s$. Thus, the utility per service can be defined as a function of a_s as

$$U_s(a_s) = \sum_{i=1}^{a_s/\sigma_s} b_s^{(i)}. \tag{6}$$

Equation (6) is defined for discrete values of a_s (the multiples of σ_s). We extend it to a piecewise linear concave function of a_s by linear interpolation.
Altogether, we can write the optimization problem as follows:

Problem 2.

$$\max_{a_s} \sum_{s \in S} U_s(a_s)$$
$$\text{s.t.} \sum_{n: n \in r(s)} f_n(a_n) \leq D_s,\ \forall s \in S,\ a_s/\sigma_s \ \in \mathbb{Z}.$$

In Prob. 2 the objective function is concave but not strictly concave (as in Prob. 1) and an integer restriction has been added. Since integer programming is NP hard, we have strong indication of the difficulty of this problem, not easy to overcome even allowing for centralized computation. We will thus accept a sub-optimal allocation which involves solving the convex relaxation, and rounding off to satisfy the integer constraints. The not strictly concaveness of the utility function may compromise the convergence of the algorithm by producing, in some cases, a hopping

result between two consecutive integer values. In order to avoid oscillations we shall use, as proposed in [10], the so-called proximal optimization method which implies modifying the optimization problem by an equivalent one so as to have a strictly concave function as objective without changing the point at which the solution is attained. For lack of space reasons we do not provide further detail on such method.
For selling the services we repeat the process described above in a periodic fashion. Every period of time T, bids are collected and bandwidth is allocated. Most previous work on multi-period auctions (e.g. [7]) allow future bidders to compete with incumbent ones, albeit given the latter some advantage [17]. A different approach (e.g. [3]) is to impose the condition that once bandwidth has been allocated in an auction, the successful bidder has a reservation for the duration of his/her connection. The specific solution for multi-period auctions problem is out of the scope of this paper and any of the previous proposals can be adopted.

2.2 Simulations

We present an illustrative example of the one-shot allocation mechanism. Consider Fig. 1(a), where four ASs associate to provide two services. The equivalent capacities of all the ASs are equal to 40. Service 1 (plain path) has a delay bound $D_1 = 2$ while service 2 (dashed path) provides a delay bound $D_2 = 0.5$. Both services offer an amount of bandwidth of 8. All values are expressed in a certain coherent unit. For both services 10 buyers offer their bids. In Fig. 1(b) the resulting utility function for each service is shown. In this case the service with the most constrictive delay bound has received higher bids. Figure 1(c) shows the evolution

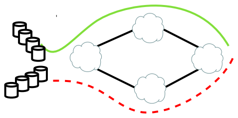

(a) Topology.

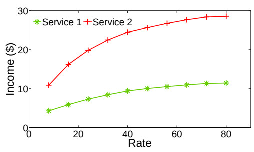

(b) Utility function.

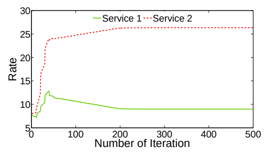

(c) Evolution of the rate (i.e. a_s) for each service.

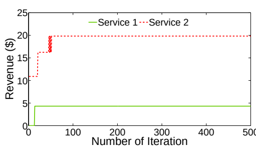

(d) Evolution of the revenue.

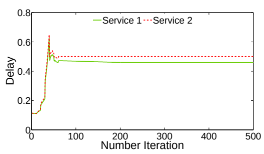

(e) Evolution of the delay for each service.

Fig. 1. Bandwidth auctions with QoS constraints, one-shot allocation. Simulations.

of the rate for each service throughout the iterations needed for the convergence of the distributed algorithm. Results show that the service which implies more incomes is the one that gets more rate accepted. Figure 1(e) shows the evolution of the delays and that both constraints are accomplished. Finally, Fig. 1(d) shows the evolution of the revenue perceived per service.

3 Revenue Sharing

As we have claimed in Sect. 1, traditional peer paying in the Internet may not be suitable for these kinds of assured quality services. We aim at performing the revenue sharing based on some fairness principles. Idelay, the revenue perceived by each AS should be proportional to the profit it provides to the federation. Moreover, the AS who is responsible for the end-to-end QoS degradation or bottleneck, or somehow limits incomes, should be encouraged to increase the resources dedicated to the alliance. In the following subsection we shall explore the concepts of coalitional game theory as a means of achieving the objectives mentioned above.

3.1 The Shapley Value

The Shapley value, proposed by Lloyd Shapley in 1952 [19], provides a means for performing the revenue sharing of an association or coalition. It has been widely used in the literature for its good properties. We now briefly recall some related concepts. The interested reader is referred to [20] for a complete review on Coalitional Game Theory.

A Coalitional Game with Transferable Utility is a pair (M, v) where M is a finite set of players and $v : 2^M \rightarrow \mathbb{R}$, the worth function which associates with each coalition $Q \subseteq M$ a real-valued payoff $v(Q)$ that the members can distribute among them.

Given a game $G = (M, v)$, we shall call $x = \{x_i\}_{\forall i \in M}$ the payoff vector, where x_i represents the share of the grand coalition's (i.e. M) payoff that player $i \in M$ receives. A Pre-imputation is the set of payoff vectors such that the sum of all x_i is equal to $v(M)$. A Dummy player is a player whose contribution to the coalition is the same as the one he/she would achieve on his/her own. With these definitions the axioms of Symmetry (for any v if i and j contribute the same to any coalition then $x_i = x_j$), Dummy player (for any v if i is a dummy player then $x_i = v(\{i\})$) and Additivity are introduced, and the Shapley value is defined as follows.

Theorem 2. *Shapley Value [19]. Given a coalitional game (M, v) there is a unique pre-imputation $\phi(M, v)$ that satisfies the symmetry, dummy player and additivity axioms and it is called the Shapley Value. It is defined, for player i as:*

$$\phi_i(M, v) = \frac{1}{|M|!} \sum_{Q \subseteq M \setminus \{i\}} |Q|!(|M| - |Q| - 1)! \left[v(Q \cup \{i\}) - v(Q) \right].$$

In addition to the properties stated on Theorem 2, the Shapley Value is efficient (it shares the total revenue) and fair. Fairness is defined in terms that for any two players $i, j \in M$, i's contribution to j is equal to j's contribution to i, that is $\phi_i(M, v) - \phi_i(M \setminus \{j\}, v) = \phi_j(M, v) - \phi_j(M \setminus \{i\}, v)$. We shall explore in the following subsection if it incentives the AS to provide better resources towards the association.

3.2 Combining the Shapley Value and the Mean Utility

In order to share the incomes perceived by means of the mechanism introduced in Sect. 2 we propose to manage two time scales. One timescale, say hourly, in which the bandwidth allocation is performed and revenue is collected. A long one, say monthly, in which the collected revenue is shared among all the ASs of the alliance. This allows for adapting the mechanism to a dynamic approach in which allocations are performed online and decentralized, and a centralized stage in which the revenue sharing is computed offline.

We define a game where the players are the set N of ASs in the association and the worth function is defined as follows. We introduce the assumption that the bids are drawn independently from a continuous probability distribution for each service. Provided this, we can safely represent the utility function of several auctions occurred during a certain period of time by the mean of all the utilities of that period. Thus, we define

$$\overline{U}_s(a_s) = E[U_s(a_s)], \tag{7}$$

which is still a strictly concave function of a_s in the general case, or a piecewise linear concave function in the case introduced in Subsect. 2.1. In addition, we assume that the delay function of every AS (i.e. f_n) remains unchanged during the considered time period.

Finally, the worth function v is defined for each sub-coalition $Q \subseteq N$ as the solution to Prob. 3, defined as:

Problem 3.

$$\max_{a_s} \sum_{s \in S^Q} \overline{U}_s(a_s)$$
$$\text{s.t.} \sum_{n : n \in r(s)} f_n(a_n) \le D_s, \ \forall s \in S^Q,$$

where $S^Q \subseteq S$ is the set of services that can be provided by Q.

Once the revenue is collected, during several phases of bandwidth allocation, it is shared among all ASs proportional to the Shapley value. That is to say, we compute $v(Q)\, \forall Q \subseteq N$ according to Prob. 3 and with these values we compute the Shapley value $\phi_n \, \forall n \in N$. Finally, node's i revenue is computed as $\Phi_i = \phi_i \times V / \sum_{j \in n}(\phi_j)$, where V is the total revenue perceived by the coalition on the considered period

We claim that the proposed mechanism provides incentives for the ASs in the association to improve their features towards it. The features we are interested in are the ones that constitute constraints to the incomes (i.e. to Prob. 1). These features are thus captured in the node's delay function (i.e. f_n) and we refer to them as an equivalent capacity for each AS. In the remainder of this section we formalize this property.

Theorem 3. *Incentive for improving capacities. Let* (N, v, c) *be a coalitional game where the set of nodes* N *are the players,* c *represents the equivalent capacities of the nodes in* N *and* v *is the worth function defined by Prob. 3. If* $i \in N$ *increases its capacity then its sharing coefficient (i.e.* ϕ_i*) will be not decreased. That is, letting* c^* *represent the capacities of the nodes where* i*'s capacity is increased,* $\phi_i(N, v, c*) \geq \phi_i(N, v, c)$*, where* $\phi_i(N, v, c)$ *is the Shapley value of node* i *given the game* (N, v) *and the capacities* c*.*

Proof. By definition of Shapley value $\phi_i(N, v, c*) = \frac{K}{N!} \sum_{Q \subseteq N \setminus \{i\}} [v(Q \cup \{i\}, c^*) - v(Q, c^*)]$, where K is a constant and $v(Q, c^*)$ represents the worth function for subcoalition Q when the capacities are given by c^*.

$$\phi_i(N, v, c*) = \frac{K}{N!} \sum_{Q \subseteq N \setminus \{i\}} [v(Q \cup \{i\}, c^*) - v(Q, c)],$$

holds since the worth function of any coalition without i is the same, regardless the capacity of i. By subtracting i's share coefficient with and without increasing its capacity we have:

$$\phi_i(N, v, c*) - \phi_i(N, v, c) = \frac{K}{N!} \sum_{Q \subseteq N \setminus \{i\}} [v(Q \cup \{i\}, c^*) - v(Q \cup \{i\}, c)].$$

We now determine if the inequality $v(Q \cup \{i\}, c^*) \geq v(Q \cup \{i\}, c)\ \forall Q \subseteq N$ holds. Indeed, v is the solution to Prob. 3 which is the maximization of a concave function with convex constraints. By increasing the capacity we relax such problem, thus doing so yields to greater or equal solutions. □

Theorem 3 proves that if node i increases its capacity its sharing coefficient increases as well or remains the same. It is now left to be proved that the total revenue perceived by the federation in the considered period (i.e. V) does not decrease either (recall $\Phi_i = \phi_i \times V / \sum_{j \in n}(\phi_j)$). Indeed, if node's i capacity is increased either the association can allocate more bandwidth (and revenue increases) either it can allocate the same amount of bandwidth (and revenue remains the same). An argument similar to the one used before can be used to formalize this reasoning, but now considering Prob. 1 instead of Prob. 3.

3.3 Simulations

We illustrate the proposed method via a simulation with a simple example. Consider the topology shown in Fig. 2(a), where the capacities of the three ASs and their delay functions are the same. Buyers' bids are random. Results of the accumulated revenue for each AS can be seen in Fig. 2(b) represented with thin lines.

In order to explore the influence of the available capacity on the revenue sharing, we consider the topology in Fig. 2(a) but now the equivalent capacity of the shaded node is increased. The cumulative revenue sharing is shown for each AS in Fig. 2(b) in thick lines, we can see that the revenue of the AS that increases its capacity perceives an improvement.

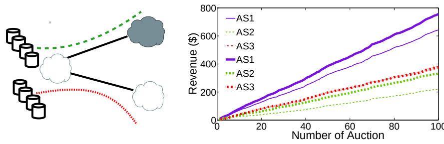

(a) Topology. (b) Cumulative revenue sharing.

Fig. 2. Incentive for increasing capacity towards the federation. Simulation.

4 Implementation Considerations

The multidomain scenario poses new problems that are not experienced in the context of the intradomain one. For instance, political aspects (confidentiality, trust), technical aspects (interoperability, scalability) and economical ones (revenue sharing). We now briefly comment on them.
In the AS Federation context, it is usually considered that the ASs tell the truth and fulfill their common interests. Nevertheless, the ASs in the federations may ask for confidentiality, privacy on committed agreements and freedom on pricing [16].
In the distributed stage, the delay of traversing the AS and its derivative are passed from one AS to another. In the centralized stage, all the ASs in the federation send their delay function and the mean utility function to a centralized trusted entity. Thus, this framework preserves confidentiality. Pricing can be freely defined at the per service level for the premium services, and at a per AS level for best-effort traffic.
Finally, the proposed solution appears to scale well. For the rate allocation, a few bytes in the forward and backward direction are needed during a preallocation iteration phase. For the revenue sharing, the ASs need to send reduced information to the centralized entity. The computation of the Shapley value is often #P-complete [5]. However, in our working context, the associations would rarely consist of more than ten ASs. For instance, the average AS path in the Internet is of four ASs [6]. In addition, the computation is proposed to be performed offline.

5 Related Work

The topics discussed on this paper are covered in several articles, of which we shall mention only a few of them.
Several works in the literature have proposed bandwidth network auctions for solving the bandwidth allocation problem. Most of them seek bids' truth reveling mechanisms. For instance, the ones based on Vickrey's second price auctions (e.g. [4]) where the winning user is charged the second highest bid, or the ones based on Vickrey-Clark-Groves (VCG) mechanisms (e.g. [4]) as in [7, 9, 12, 17]. Most of these mechanisms need for centralized computation, some of them assume certain network topology while others assume the buyer knows the network topology. In these cases the objective is welfare maximization. Other proposals (e.g. [3])

work with first price auctions. In this kind of auction, revenue maximization is sought and the implementation complexity is much lower than the one present in second price auctions. Moreover, in [13] it is shown that VCG mechanisms can hardly be applied on multidomain networks.
For the reasons exposed above, our auctions proposal is aligned with the one in [3]. However, we consider a multidomain federation scenario rather than a single domain and we incorporate an end-to-end QoS constraint rather than only considering capacity constraints. With respect to this last aspect [18] states a similar problem, but its context and the way it is solved differ significantly from ours.
Regarding revenue sharing, for instance, in [11] the proposal is to change the Internet economics by business contracts whose payment is determined by the Shapley value. In [21, 14] the aim is to optimize routing within an alliance of ASs and revenue is shared by means of Shapley value. We share with them the choice of using the Shapley value. However, our proposal incorporates the sell of premium services which are the sources of the revenue, and links the Shapley value with it. In addition, our approach also takes into account the features the ASs provide to the alliance rather that only considering the routing.

6 Conclusion and Future Work

We have proposed a framework for covering the complete cycle for selling end-to-end quality assured services in the context of AS federations. We have stated the problem of network bandwidth allocation with QoS constraints and showed a distributed solution. An application based on network bandwidth auctions for using such problem as the means for selling quality assured paths was shown. A mechanism for performing the revenue sharing of the federation, based on the Shapley value and the mean utility function was proposed. Such mechanism has fairness properties and was proven to incentivize ASs to increase its capacities. The behavior of the whole solution was studied through simulations.
In future work we shall enhance the interdomain network model and deepen on the delay function. In addition, we shall continue the research on revenue sharing, seeking for more properties such as the ones involving the stability of the federations and incentives to collaborate.

Acknowledgment This work was funded by the ETICS project (EC FP7 248567), cf. `www.ict-etics.eu` and the Uruguayan Agency for Research and Innovation (ANII) (PR-POS-2008-003 and FCE 2158).

References

1. ETICS: Economics and Technologies for Inter-carrier Services. European research project, supported by the 7th Framework Programme of the European Union., `http://www.ict-etics.eu`
2. Cisco Systems: Hyperconnectivity and the Approaching Zettabyte Era. Tech. rep. (Jun 2010)

3. Belzarena, P., Ferragut, A., Paganini, F.: Bandwidth Allocation via Distributed Auctions with Time Reservations. In: IEEE INFOCOM, Procedings. Rio de Janeiro, Brazil, 2009 (2009)
4. Courcoubetis, C., Weber, R.: Pricing and Communications Networks. John Wiley & Sons, Ltd (2003)
5. Deng, X., Papadimitriou, C.H.: On the complexity of cooperative solution concepts. Math. Oper. Res. 19, 257–266 (May 1994)
6. Dhamdhere, A., Dovrolis, C.: Ten years in the evolution of the internet ecosystem. In: Proceedings of the 8th ACM SIGCOMM. pp. 183–196. ACM (2008)
7. Dramitinos, M., Stamoulis, G.D., Courcoubetis, C.: An auction mechanism for allocating the bandwidth of networks to their users. Comput. Netw. 51, 4979–4996 (Dec 2007)
8. Feijer, D., Paganini, F.: Stability of primal-dual gradient dynamics and applications to network optimization. Automatica 46, 1974–1981 (Dec 2010)
9. Lazar, A.A., Semret, N.: Design and Analysis of the Progressive Second Price Auction for Network Bandwidth Sharing (1999)
10. Lin, X., Shroff, N.B.: Utility Maximization for Communication Networks With Multipath Routing. IEEE Transactions on Automatic Control 51(5), 766–781 (May 2006)
11. 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)
12. Maillé, P., Tuffin, B.: Pricing the internet with multibid auctions. IEEE/ACM Trans. Netw. 14, 992–1004 (Oct 2006)
13. Maillé, P., Tuffin, B.: Why VCG auctions can hardly be applied to the pricing of inter-domain and ad hoc networks. In: 3rd EuroNGI conference, Trondheim, Norway. pp. 36 – 39 (2007)
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. Pouyllau, H., Douville, R.: End-to-end QoS negotiation in network federations. In: IEEE/IFIP Network Operations and Management Symposium Workshops (NOMS Wksps). pp. 173–176 (2010)
17. Reichl, P., Bessler, S., Stiller, B.: Second-chance auctions for multimedia session pricing. In: PROC. MIPS03 (2003)
18. Saad, M., Leon-garcia, A., Yu, W.: Rate Allocation under Network End-to-End Quality-of-Service Requirements. In: GLOBECOM 2006
19. Shapley, L.: A value for n-person games. In: Kuhn, H., Tucker, A. (eds.) Contributions to the Theory of Games
20. Shoham, Y., Leyton-Brown, K.: Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations. Cambridge University Press, Cambridge, UK (2009)
21. Stefano Secci, Jean-louis Rougier et al.: Connection-oriented Service Management in Provider Alliances: a Shapley Value Perspective. In: EuroNF 5th Int. Workshop on Traffic Management and Engineering for the Future Internet (2009)