

Auction-type framework for selling inter-domain paths

Xavier Misseri, Jean-Louis Rougier
LTCI
TELECOM ParisTech – Paris – France
Email: {misseri,rougier}@telecom-paristech.fr

Stefano Moretti
CNRS/LAMSADE
Université Paris-Dauphine – Paris – France
Email: stefano.moretti@dauphine.fr

Abstract—In the present Internet, inter-domain routing is based on BGP-4 [1] which selects a single path per destination prefix, thereby preventing carriers and end-users to use the vast inherent path diversity [2]. Addition of multi-path capabilities to the Internet have long been advocated for both robustness and traffic engineering purposes. Some works [3], [4] propose inter-domain multipath architectures.

In this paper we consider a new service where carriers offer additional routes to their customers (w.r.t. to BGP default route) as an added-value service. These alternate routes can be used by customers to help them to meet their traffic engineering objectives (better delays etc.) or just for robustness purposes (disjoint alternate routes). Announcing additional paths can lead to scalability issues [5], so one carrier will propagate only the paths that are most interesting for neighboring domains.

We propose an auction-like framework adapted to this specific service, allowing one carrier to select the most interesting paths and determine the prices at which these routes can be sold. We consider the case where routes are sold as infinitely duplicable goods (assuming small demands with regards to route capacities). We design a winner determination mechanism, based on the maximization of the seller's revenue, that enforces fair allocation of goods and is loser collusion proof. We also propose a payment mechanism that is proven to be truthful when each bidder submits one (potentially combinatorial) bid.

I. INTRODUCTION

Internet is composed of Autonomous Systems (i.e., ASes) exchanging routing information. With the current de facto inter-domain routing protocol (i.e., BGP [1]), an AS can only propagate one path to its neighbors. Therefore only one path is put into use and the benefits of path diversity, inherently present in the Internet, are never harvested. At the same time, network service providers (NSPs) are looking for providing new services for customers and path diversity can be the cornerstone of such services [6], [7]. Some proposals [3], [4] highlight ways to propagate and use this diversity. Nevertheless, enabling Internet path diversity faces huge scalability issues as the number of potential available paths can be very important [5]. The insertion of any additional route has an impact on the cost of the network. Network service providers will hardly implement these routes at their expenses and will therefore make their clients pay for this service. Nevertheless clients will pay if the NSP succeed in proposing routes that are interesting, according to the clients' criteria.

Such an approach faces two important problems. First different paths have different characteristics (e.g., length, delay...) and each of these paths may be interesting for a neighbor but

not for one another. Second, network service providers aim at proposing path diversity as a value-added service, in particular to face the scalability issue mentioned above. Therefore the price establishment of a path is something very important which has not yet been studied. **We aim, in this paper, at providing a simple framework that addresses these problems and fits the specific constraints of the inter-domain routing.** Therefore we propose a route allocation framework, inspired by the auction theory, that aims at pricing routes and matching them to interested neighboring domains. The allocation process outputs both route matching and pricing and unifies good properties in order to motivate NSPs to adopt such an approach.

The paper is organized as follows: In Section II, we describe the context and the constraints of the inter-domain routing where the allocation process takes place. In Section III, we detail the notations we use, the wanted properties of the framework and the related work. Then the Section IV presents the allocation framework, including the winner determination, followed in Section V by an example of payment function proven to be truthful when bidders submit one bid. We then underline some evaluations in Section VI before concluding in Section VII by providing some further works extending the framework.

II. BACKGROUND

A. Context

The architectures proposed in [3], [4] allow for the propagation and the use of an important amount of paths. Internet routers currently have in their FIBs 450 000 entries [8]¹, which is already identified as a growing issue [9], [10]. [5] underlines that without any filtering, the amount of propagated paths could reach a non-manageable amount (i.e., order of magnitude of 10^{10}). In order to address such an issue Kwong et al. [10] suggest the use of market mechanisms to limit the increase of the routing table size in a single path Internet. Therefore we adopt the same perspective as [10] to avoid the explosion of the routing table of Internet routers in the multi-path context.

In our context, each domain filters the routes according to the needs of his neighbors and makes neighbors pay according the routes they receive. Filtering according to the

¹450 000 routes or so in november 2012

needs of each neighbor allows to leverage a new commercial approach of route propagation, as each neighboring domain may pay to receive his own specific route. Nevertheless a good knowledge about each neighbor is required in order to propagate the correct sets of routes. This knowledge is very hard to obtain as it depends on the type of neighbors and reflects complex policies of neighboring domains, which may include commercially sensitive information. A matching process is therefore required to associate, for each neighbor, the best set of routes. Therefore a network service provider, which has already received several paths, aims at identifying the paths which deserve to be propagated to some of his **direct** neighbors (cf. Figure 1). The path allocation takes place between the NSP and his neighbors and no further domain is involved in the process.

Therefore the goal of the current work is to answer both following questions: 1/ Path to neighbor matching: to which neighbor do I send which route? 2/ Price computation: what price will each neighbor pay? While the proposed framework differs from pure auctions (cf. Sections II-B and III-C), its design is very close as both problems are well tackled by the auction theory respectively under the name of Winner Determination Problem and Price Establishment [11]. Therefore some comparisons between the current route allocation framework and the auction theory are used all along the paper to highlight important aspects.

The propagation of paths to neighbors, according to the proposed allocation framework, follows these four steps:

- Step 1: **publication of paths**: The provider sends to his neighbors a set of path² information (e.g., AS path, path characteristics, monitoring results [12]...) and the setting of the allocation process (e.g., type of mechanism, constraints...).
- Step 2: **Path bidding**: Each neighbor sends a set of bids, each one containing a set of routes and the corresponding price he is willing to pay for it.
- Step 3: **Path(s) matching and price(s) computation**: after receiving the bids, the provider is able to compute both the path to neighbor matching, which selects, for each neighbor, the set of paths he wins, and the prices neighbors must pay.
- Step 4: **Path configuration and payment**: once matching and pricing are done, the provider allocates the routes by configuring the routing equipments such that each winner is able to use the paths he won, and propagates the full routing information to the winners. These route announcements and configurations can be considered as a grant to use the additional routes.

The framework allows a carrier to offer additional routes to some selected customers in order to help them achieve their respective traffic engineering objectives. However, the end to end paths that are offered may pass through carriers that are not involved in the process. In other words, the

²Only a set of the most “interesting” routes may be sent using eventually coarse filtering functions (e.g. too long paths etc.).

carrier has usually no control on the proposed routes (and their characteristics). Only a privilege to access these alternate routes is provided, with no quality of service (i.e., QoS) guarantees.

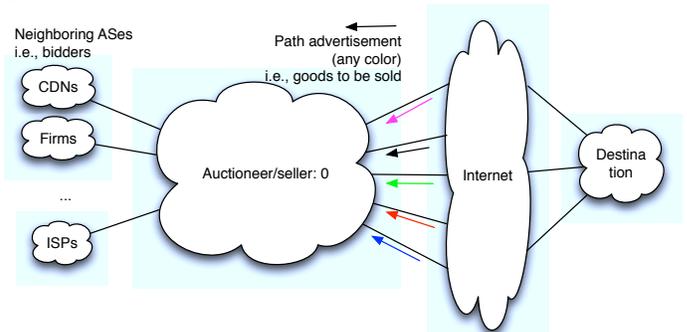


Fig. 1. Goods and actors identification

B. Inter-domain constraints

From an inter-domain point of view, the price of inter-domain peering bandwidth has drastically dropped during the last decade (a two order of magnitude decrease [13]). Therefore we assume that the network service provider that sells paths is able to forward the whole traffic demand of his clients, leading to the assumption that the NSP’s capacities are over-dimensioned compared to the needs of his neighbors.

Moreover, selling an inter-domain path is different from selling a conventional good in the sense that, after the transaction, a path does not belong to the buyer. Indeed instead of really selling a path, the seller provides a grant to use the said path.

As it is currently the case, several buyers can be granted simultaneously to use a common path. Without any bandwidth constraint, all the neighbors could be granted to use the same path. Whereas the number of buyers is limited by the number of neighbors, we consider, for the rest of the chapter and without loss of generality, that an unlimited number of neighbors/bidders can win a single path.

Therefore here are the characteristics of the inter-domain context:

- Neighbors/bidders may want to win several routes to reach the same destination. They must therefore be able to bid on bundle(s) of routes/goods.
- these routes/goods can be **duplicated infinitely** by the seller in the sense that the seller can either propagate the routing information to only one bidder or to all bidders, which can be considered as a duplication of routes.
- the inter-domain allocation process must generate low amount of messages. Therefore a one round allocation process should be preferred³.
- a route is allocated to a bidder at most once. Indeed, once a domain receives the routing information and is able to use the associated path, receiving again the same

³The present work describes one single use of the process, which should be a one round process to minimize the communication overhead. Nevertheless, routes can be re-allocated on a regular basis (e.g., hourly, daily...) by using the same one round process. It could therefore elect different winners with different prices, if the valuations of bidders changed since the previous allocation. The use of repeated allocation processes should be analysed from a repeated game perspective, which is outside the scope of the current work.

routing information does not give him any new grant or any additional capability.

III. FRAMEWORK REQUIRED PROPERTIES

A. Notations

We present in this section the notations, inspired from [14], for our work. A transit domain (i.e., the seller denoted by 0) proposes a set of routes G (G for goods) to all or part of his neighbors B (B for bidders). Figure 1 identifies the seller, the bidders and the routes/goods in the global environment. Each neighbor $i \in B$ is a potential bidder and may therefore answer to the process in order to buy one or several route(s). Each bidder has a utility function $u_i(g)$, which associates each set of routes $g \subseteq G$ with a utility value. The utility function is private and represents the preference of bidder i for each set of routes g . Each bidder i bids a reported valuation $v_i(g)$. These reported valuations are either public or at least known by the seller. After receiving the bids from each bidder the seller elects the winners $W \subseteq B$. The seller also computes the price each winner must pay in order to obtain the set of routes he wins. Bidders that do not win any route are named losers L , with $L \cup W = B$. A bidder i who wins a set of routes $g \subseteq G$ pays a price $p_i(g) \leq v_i(g)$. His payoff is then $\pi_i = u_i(g) - p_i(g)$ whereas the payoff of the seller (i.e., the amount he wins) is $\pi_0 = \sum_{i \in B} p_i$.

B. Properties

The mechanism we design is intended to take place in both market and inter-domain routing contexts. Therefore we consider the following properties as being essential to the successful adoption of this framework:

- **Implementable:** With the current 450 000 IPv4 prefixes, inter-domain routing is already dealing with an important amount of information. The present allocation framework must occur in a context where each one of the said prefixes is associated with several routes (which is not the case with the current one route paradigm) and where several instances of the framework may run in parallel (e.g., one instance per prefix). Therefore the first requirement of the mechanism is to be scalable and implementable.
- **Truthful:** A mechanism is truthful if bidding u_i is a dominant strategy for every bidder i (i.e., $v_i = u_i$). Such a property is interesting for the seller as it allows to assign the paths to those that really evaluates them more. Moreover it helps the seller to design future route allocations by learning the type of routes that are really interesting. Ultimately, it prevents bidders from manipulating bid values in order to influence the price they have to pay.
- **Maximize the income of the seller:** We aim at providing to the seller a substantial income that will make him adopt this mechanism.
- **Deal with unknown valuation functions:** Pricing of alternate routes is a new possibility of the future multi-path Internet. Therefore a seller does not know yet how much each route could be priced. It is therefore impossible to fix prices and let buyers come if the price is interesting for

them. Furthermore, we do not know yet the type of utility functions NSPs have for such a new service. We must be agnostic about this information. Different approaches, which are cited in Section III-C may be used on a long run - i.e., when NSPs have sufficient knowledge about their clients' needs. Nevertheless the unknown valuation function case is more adapted in this emerging market.

The two first properties (i.e., implementability and truthfulness) may prevent the mechanism from closely approaching the maximum revenue of the seller. Nevertheless we consider these two properties essential to the adoption of the framework. Therefore we aim at maximizing the revenue of the seller under the constraint that the two first requirements are respected.

The presented requirements are tackled by the auction theory (cf. related work in Section III-C). This is the reason we chose to set up an auction-like allocation framework that we adapt in order to be compatible with the interdomain constraints presented in Section II-B.

As we underlined in Section II-B, routes are infinitely duplicable. Therefore the seller could potentially allocate a set of routes to every bidder. Taking into account the revenue maximization goal, a common reflex is to sell the routes at the maximum price that bidders are able to pay (i.e., a pay as you bid mechanism: $p_i = v_i$). Nevertheless, in our context, such a mechanism has two consequences. The first is that each bidder is elected as a winner and wins the good he wants. Indeed, for each new bidder that wins a good, the seller increases his revenue (without any good depletion). The maximum revenue is thus reached if all bidders win the auction. As a consequence, competition between bidders totally disappears. Moreover, another consequence is the lack of truth, which encourage bidders to bid lower values. Indeed each bidder knows, before bidding, that he will win the set of goods he wants regardless of his bid. All bidders will thus bid with a valuation of zero, which will lead to a null seller's revenue, which is far from maximizing his revenue. A payment function must therefore be found to maximize the seller's revenue and enforce truthfulness.

C. Related work

Pricing of items is a large and old field that has been well studied to sell physical goods (e.g. [15] for cars, computers...). The emergence of computer science extended the field to the pricing of unlimited supply goods. Some works have been proposed to maximize the revenue of the seller in an unlimited supply context [16], [17], [18], [19], [15], [20]. While the computation of the maximum revenue price vector is NP-hard [16], contributions to the field generally propose price computation algorithms that approximate the maximum revenue of the seller. To the best of our knowledge, no contributions unify the properties presented in Section III-B. [21] deals with the selling of one item while we deal with several items. Some other works do not ensure that players disclose their true utilities [16], [20]. While [22] proposes a way to transform any non-truthful item pricing mechanism into

a truthful mechanism, it degrades the efficiency of the revenue approximation and overall requires a minimum amount of buyers, which is not a reasonable assumption in our context. Some other works [19], [23] take into account historical records (e.g., Bayesian prior) or already know how much each buyer is able to pay [24], which is also not the case in our context.

From a pure auction perspective, our framework gives the opportunity to sell sets of routes to neighbors, which could be compared to a combinatorial auction⁴. Combinatorial auctions have been well studied in the literature [25], [14]. Although they seem to fit our goal, we apply them in a context where **items to be sold can be duplicated** by the seller (i.e., unlimited supply) and allocated simultaneously to different clients/bidders, which is not the case of these works.

Auctioning goods with unlimited supply is not common and some works have been done to address this problem. Goldberg et al. studied in [26] the selling of a single digital good and extended their work to multiple goods in [27]. Nevertheless, their approaches only take into account a single won good per winner and does not deal with combinatorial auctions. Other works also deal with digital auctions, but all focus on multiple identical items [28], [29], [30], [31]. We also aim at setting up our framework to make the bidders give their true valuations. A lot of work has been done in auction theory to set up truthful auctions [32], [33], [11], [14], [34], [35], [36]. Nevertheless all these works focus on a finite number of goods to be sold.

Other works [37], [38], [39] aim at providing frameworks that associate inter-domain routing with auctions or pricing. From a networking point of view, they propose to sell bandwidth along paths that everybody can use and fix transit prices according to pricing mechanisms, whereas we aim at selling the right to use specific paths, regardless of the bandwidth. From a game theory point of view, they address these issues by using a VCG approach, which is not adapted in our infinite supply context [40].

From a pure networking point of view, an architecture has previously been proposed to sell inter-domain routes [41] and advocates for the use of auctions in this context, without providing any auction framework. While this architecture proposal is well adapted to a federation of domains, it is highly centralized and thus not adapted to the Internet (distributed by nature). It requires a strong cooperation between carriers while we adopt an incremental perspective, where each NSP can adopt the framework on his own⁵.

IV. PATH ALLOCATION FRAMEWORK

A. Allocation process

Conventional winner determination computations are based on the maximization of the social welfare (i.e., $X(v) = \operatorname{argmax}_x \sum_i v_i(x)$ [11]). Such maximization selects in general a restricted number of winners as goods are over-

⁴Combinatorial auctions: auctions where bidders can bid simultaneously on several goods - cf. [25]

⁵A NSP currently receives a path diversity thanks to his external BGP peerings [42]. Therefore he can already sell this diversity to his clients without any cooperation with other NSPs.

demand. Nevertheless, in our unlimited supply context, the maximization of the social welfare is equivalent to electing all bidders to be winners, which could lead to a seller's revenue of zero.

It must be noticed that accepting a new winner (or increasing the set of winning bids) may face two contradictory effects:

- the new winner pays, which makes the payoff of the seller π_0 (i.e., his revenue) increase;
- depending on how the prices are computed, the prices paid by winners may decrease, which makes the payoff of the seller π_0 decrease (as illustrated in Figure 2).

For instance, in multiple identical item auctions with limited supply, the VCG⁶ price may reduce if you add another item. In such a case, adding items to make every bidder win makes the price and the seller's revenue drop to zero [11]. Therefore the revenue of the seller is null whether there is no winner (i.e., $|W| = 0$) or every bidder is elected to win (i.e., $W = B$). There must exist at least a set of winners W where $\pi_0 > 0$ (unless the selling process does not make sense). Therefore, despite what can be intuitively thought, maximizing the seller's revenue does not always require to sell routes to all bidders. It depends on the price computation function which is used by the seller.

The revenue of the seller depends on both the set of winning bids and the prices winners pay. We adopt an iterative approach to elect the winners and the prices. We iteratively increase the set of winning bids in order to explore the bid space, compute the corresponding prices and the revenue of the seller at each iteration. We can then deduce the seller's maximum revenue point. This process is described in Algorithm 1.

Algorithm 1 Revenue curve computation

```

Revenue_max = 0;
Winners_max = 0;
ForAll Winning_bids  $\subset$  bids do
  Revenue = compute_revenue(Winning_bids)
  if Revenue > Revenue_max then
    Revenue_max = Revenue
    Winners_max = Winning_bids
  end if
EndFor

```

This algorithm computes the price for different sets of winning bids. A first approach for computing the maximum revenue of the seller is to compute, for all the combinations of winners, the prices and the associated seller's revenue and keep the combination of winners providing the maximum revenue. Nevertheless this approach suffers from the important number of potential winner sets. Therefore we use a ranking approach to add bids to the set of winning bids. A way to rank bidders must be found in order to easily identify, at each computation step of the algorithm, the bid which must be elected as extra winning bid for the next iteration.

⁶Whereas VCG pricing computation is not applicable in our context, it is a good way to illustrate that increasing the set of winners may reduce prices.

B. How to rank bids fairly?

Some rankings have already been proposed in [43], [40], nevertheless their characteristics are proved in a context where a limited number of goods is available and with single-minded bidders (i.e., each bidder is only interested in a single set of goods). In order to set up a fair matching between goods and bidders, a bid can be elected as a winning bid only if all the bidders that proposed more for the same set of routes also win. Let b be a bidder and $v_b(g)$ his bid for the set of routes g (b may also submit other bids). Here are the requirements of a ranking. $v_b(g)$ can be elected as a winning bid:

- requirement 1: if all the higher ranked bids have already been elected as winning bids
- requirement 2: and if no group of lower rank bids can propose more than $v_b(g)$ for the same set g .

These two requirements may be considered as fair conditions. It must be noted that a bidder may have several winning bids and several losing bids (i.e., for different sets of routes).

It is easy to rank bids for the same set of goods. Nevertheless, it is more difficult to rank bids associated to different sets of goods as the sizes of the sets may be different and as the sets may not contain the same routes. In such a case, bid values can not be considered alone and set sizes must also be taken into account. Therefore we rank bids according to the mean reported valuation of the bid - i.e., $\overline{v_b(g)} = \frac{v_b(g)}{|g|}$.

Definition 1: A **mean-bid winner determination** is a determination of the winners such that if a bidder i wins thanks to his bid v_i , each bidder j that has, at least, one higher mean-bid (i.e., $\overline{v_j} > \overline{v_i}$) also wins.

By definition, mean-bid winner determination is compatible with the first requirement of a bid ranking.

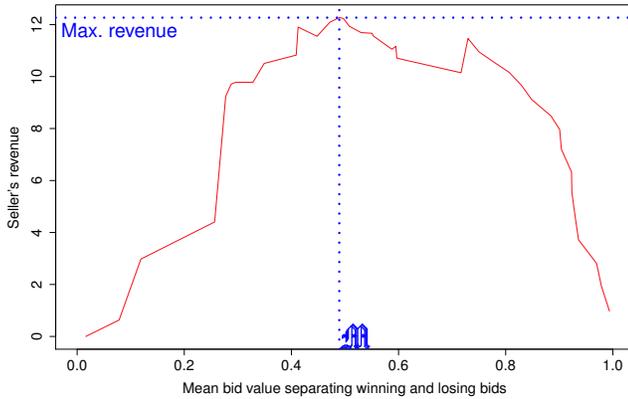


Fig. 2. Example of the evolution of the seller's revenue.

Figure 2 presents what the output of Algorithm 1 could be on a single duplicable item. In this graphic example, 40 bidders in the range $[0, 1]$ and prices are computed according to a price function⁷. The curve traces the revenue that the seller could earn by fixing the mean-bid-value separating winners from losers. **The point providing the highest revenue is the point $\overline{\mathfrak{M}}$.** Each bid whose mean value is lower than $\overline{\mathfrak{M}}$ is a

⁷Here, the first-loser-mean-bid payment function (detailed in Section V) is used but other payment functions may be used.

losing bid. If a bidder has only losing bids, he does not win anything. **In the case he has several bids higher than $\overline{\mathfrak{M}}$, he wins the set of goods which has the highest reported valuation** - i.e., $x_i = \max_g \{v_i(g) | \overline{v_i(g)} > \overline{\mathfrak{M}}\}$.

Theorem 1: When goods are infinitely duplicable by the seller, a mean-bid winner determination is loser collusion proof.

Proof: Theorem 1 advances that no losing bids or coalition of losing bids can propose to the seller more than what winners propose, for the same set of routes/goods. We rank all the bids by their mean values. For any winner bid v_1 , we perform a reductio ad absurdum by imagining that a set of lower rank bids (i.e., $\{v_2, \dots, v_n\}$ with $\forall i \in \{2, \dots, n\}, \overline{v_i} < \overline{v_1}$) can propose a better common bid $v_{\{2, \dots, n\}}$ for exactly the same set of goods $g = g_1 = \cup_{i \in \{2, \dots, n\}} g_i$ (with g_i the set that bidder i bids on). We have to note that $\forall i, j \in \{2, \dots, n\}, g_i \cap g_j = \emptyset$ as a good is only present once in g and that goods are not duplicable by bidders (cf. Section II-B). If the coalition is successful, its valuation should be higher than the one of bidder 1 - i.e., $v_{\{2, \dots, n\}} = \sum_{i \in \{2, \dots, n\}} v_i > v_1$ and therefore $\overline{v_{\{2, \dots, n\}}} = \frac{\sum_{i \in \{2, \dots, n\}} v_i}{|g|} \geq \overline{v_1} = \frac{v_1}{|g|}$. Nevertheless we have the following, which yields a contradiction with the previous relation, and the proof follows:

$$\begin{aligned} \overline{v_{\{2, \dots, n\}}} &= \frac{\sum_{i \in \{2, \dots, n\}} v_i}{|g|} \\ &= \frac{\sum_{i \in \{2, \dots, n\}} \frac{v_i \times |g_i|}{|g_i|}}{|g|} \\ &< \frac{v_1}{|g_1|} \times \frac{\sum_{i \in \{2, \dots, n\}} |g_i|}{|g|} \\ &< \frac{v_1}{|g_1|} = \overline{v_1} \end{aligned}$$

Theorem 1 proves that the mean-bid winner determination is compatible with the second requirement of a bid ranking. ■

V. EXAMPLE OF SIMPLE PAYMENT FUNCTION

We analyse in this section the characteristics of a simple payment function, the first-loser-mean-bid payment. Despite the apparent simplicity of this payment function, it gives to bidders incentive to tell the truth when each of them submit a unique bid (cf. Section V-B).

A. Payment function presentation

Definition 2: A **first-loser-mean-bid payment** is a payment function where the winners pay each of their won goods the highest mean-valuation of the losing bidders:

$$\begin{aligned} \text{for each } g \subseteq G, p(g) &= |g| \times \overline{\mathfrak{M}} \\ \text{and } \overline{\mathfrak{M}} &= \max\{\overline{v_i} | i \in L\} \end{aligned}$$

Here is a small example of the allocation process with 4 bidders (i.e., bidders 1, 2, 3 and 4) competing for 2 goods (i.e., a and b). Table I presents an example of bids, which are then ranked according to the mean in Table II and processed to compute both the seller's revenue and the good allocation in Table III.

Having ranked the different bids according to the mean allows the seller to scan all the possible winner combinations by scanning it from high to low mean values (cf. Table II). At each step, the revenue of the seller is computed (here, using the first-loser-mean-bid payment) in order to select the maximum revenue good allocation.

As a first step, the seller states that only bidders who bid more than 7, as mean bid value, are elected as winners (cf. first column in Table III). Therefore bidder (1) wins b and the price of each won route is 6 (i.e., the mean bid of the first loser), which gives to the seller a total revenue of 6. As a second step, the seller elects the more-than-6-mean-value bidders as winners. The price which is to be paid is 5 per route as only bidders 3 and 4 remain losers and their maximum mean-bid value is 5. The revenue is therefore 10. Every step is computed and the seller selects the one maximizing his revenue (i.e., step 3, in red, with a revenue of 15).

	(1)	(2)	(3)	(4)
a	-	6	5	2
b	7	-	5	-
ab	11	9	6	3

TABLE I
BIDS

Mean	Bidder	Goods
7	(1)	b
6	(2)	a
5.5	(1)	ab
5	(3)	$a \oplus b$
4.5	(2)	ab
3	(3)	ab
2	(4)	a
1.5	(4)	ab

TABLE II
BIDS RANKED ACCORDING TO THE MEAN

	Minimum winning mean bid							
	7	6	5.5	5	4.5	3	2	1.5
(1)	b	b	ab	ab	ab	ab	ab	ab
(2)	\emptyset	a	a	a	ab	ab	ab	ab
(3)	\emptyset	\emptyset	\emptyset	$a \oplus b$	$a \oplus b$	ab	ab	ab
(4)	\emptyset	\emptyset	\emptyset	\emptyset	\emptyset	\emptyset	a	ab
Route price	6	5	5	2	2	2	0	0
Revenue	6	10	15	8	10	12	0	0

TABLE III
REVENUE COMPUTATION (\oplus) FOR THE CONVENTIONAL XOR OPERATOR)

B. Telling the truth

We prove here that if bidders submit only one bid, the proposed process and the first-loser-mean bid payment is truthful. It must be noted that the seller can easily enforce bidders to submit only one bid (potentially containing several routes/goods).

Theorem 2: In the present allocation framework (duplicable goods, mean bid-winner determination based on the maximization of the seller's revenue and first-loser-mean-bid payment), if each bidder submit one bid, truth telling is a dominant strategy.

Proof: Here bidders submit only one bid (potentially containing several goods). A winner may lie in order to reduce the price he is about to pay and a loser may lie in order to become winner.

First, by modifying their bids, winning bidders are not capable to modify the price they pay as the price is computed thanks to bids of losing bidders.

Nevertheless a loser can make his losing bid become a winning bid by increasing the associated bid value. More

formally, a bidder i lies and bids a valuation $v_i(g) > u_i(g)$ such that he is able to win the set of goods $g \subseteq G$ that he would have lost by telling the truth (i.e., when $v_i(g) = u_i(g)$). What price would bidder i pay? It is obvious that i is not willing to pay the price $p(g)^T > u_i(g)$ (i.e., where $p(g)^T$ is the value of $p(g)$ when i tells the truth). Nevertheless $p(g)$ could be reduced by the lie of i and take an interesting value $p(g)^L < u_i(g) < p(g)^T$ (i.e., L when i lies). Therefore the question is: can $p(g)$ decrease and take a value lower than $u_i(g)$ by increasing $v_i(g)$.

Figure 3 shows a zoom around the maximum revenue point and illustrates the elevation of the revenue curve because of the lie of a bidder. The black arrow shows the lie of bidder i (from a mean value $v_i(g)^T = 0.41$ to a mean value $v_i(g)^L = 0.52$ - cf. point 0 in Figure 3). Without the lie of the bidder, the revenue curve is represented by the red curve whereas the new revenue values, influenced by the lie, are represented by the blue diamonds. This bid modification makes some potential revenue points increase (cf. blue arrows that lead to points 1, 2, 3, 4 and 5), whereas some other points are not modified.

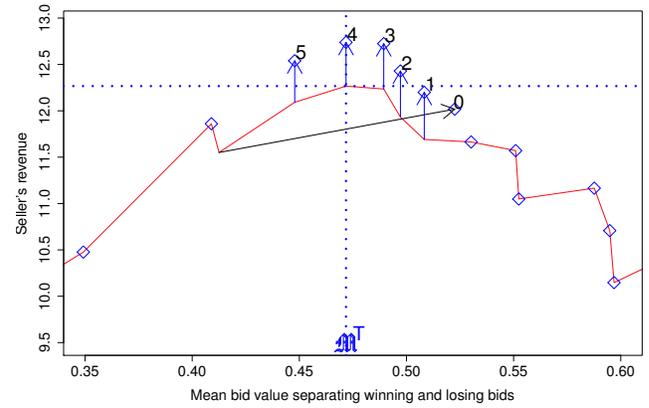


Fig. 3. Revenue curve deformation.

We analyse here the increase of the revenue curve led by the lie. The following applies either if $\mathfrak{M}^T \in (u_i(g), v_i(g))$ or not.

The lie (i.e., $v_i(g) > u_i(g)$) makes the seller's revenue curve change (as illustrated in Figure 3). It is important to notice that the modification of the seller's revenue curve only occurs in the value interval $(u_i(g), v_i(g))$. Indeed $\forall m \notin (u_i(g), v_i(g))$, the number of won goods is not changed and the revenue remains the same.

Then $\forall m \in (u_i(g), v_i(g))$, the value of the revenue is increased by the price paid by i (i.e., $m \times g$). Therefore either the maximum revenue point remains the same (i.e., $\mathfrak{M}^T = \mathfrak{M}^L$) or $\exists m \in (u_i(g), v_i(g))$ such that $r(m) \geq \mathfrak{M}^T$. In both cases the new maximum revenue point \mathfrak{M}^L provides a price that is higher than $u_i(g)$ leading to a negative payoff for i .

As a conclusion, truth telling gives to bidders their maximum payoff and is therefore a dominant strategy. ■

VI. EVALUATIONS

A. Computational complexity

From a computational perspective, the process can be summarized by the ranking of bids and the computation of the revenue curve. The ranking of the bids is an $O(|Bids| \times \log(|Bids|))$ [44] operation (with $|Bids| = |B| \times |G|$). Once the ranking of bids is done, the computation of the revenue curve, according to the first-loser-mean-bid payment, requires the computation of $|B|$ revenue computations. Each revenue computation is linear in the number of bids as it only looks at all the bids to select the ones that are elected as winners and computes the revenue with the price given by the first-loser-mean-bid payment function. The overall process complexity is therefore in $O(|B|^2)$ and in $O(|G|\log(|G|))$. Figure 4 shows the square root of the computation time according to the number of bidders, for different processes. It shows that the allocation process is less than linear regarding the square of the number of bidders. It reinforces the square complexity of the allocation process, regarding the number of bidders.

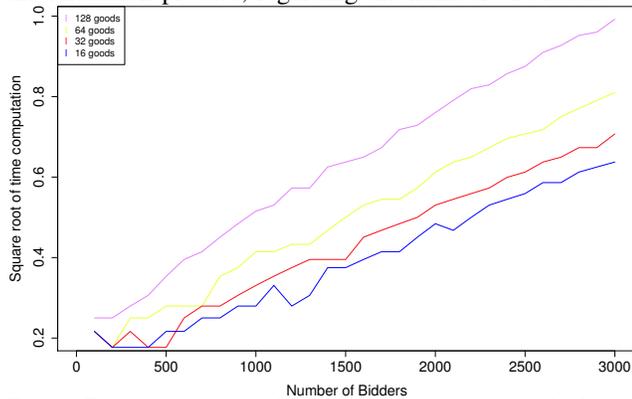


Fig. 4. Time computation evolution (square roots) regarding the Number of bidders.

B. Ratio between revenue and winners' valuation

Figure 5 shows the ratio between the revenue of the seller and the valuation that winners could have paid to win their routes. We compute this ratio for different numbers of goods/routes and for different numbers of bidders.

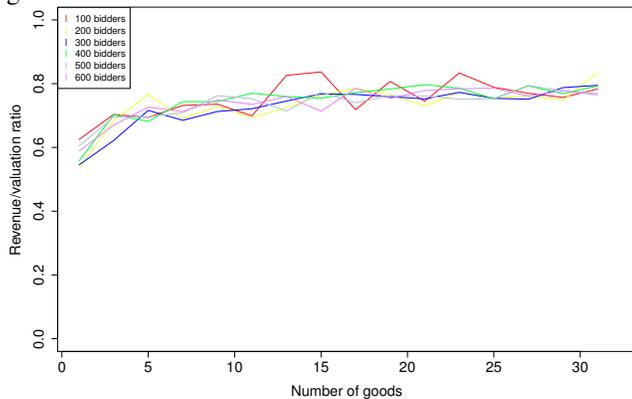


Fig. 5. Revenue over valuation ratio.

We find that the seller receives as a revenue between 55% and 80% of the prices that winners could have paid (i.e., their valuations). The revenue that is not perceived by the seller is the winners' payoff. Nevertheless, Figure 5 is an evaluation of

the allocation when bidders submit several bids, which is not proven to be truthful.

Figure 6 shows the same evaluation (i.e., revenue over valuation) for allocations which accept only one bid per bidder. This evaluation is important as we proved in Section V that truth telling is a dominant strategy for this type of mechanisms. Therefore, whereas the seller is not sure that the valuations provided by the bidders in the evaluation of Figure 5 are true, the ones provided for the allocations simulated in Figure 6 can be assumed to be equal to the real utility values.

We can see that the ratio between revenue and valuation is almost the same (i.e., approximately 60%) for different numbers of bidders and for different numbers of routes. It is interesting to note that this revenue over valuation values are coherent with the ones of the one good sales of Figure 5 as the one good sales, which leads to a one bid per bidder submission, provide a ration between revenue and valuation of around 60%. The evaluations of Figure 6 give a less interesting ratio than Figure 5. It can be considered as the price the seller has to pay in order to enforce truthfulness - i.e., the price of truthfulness.

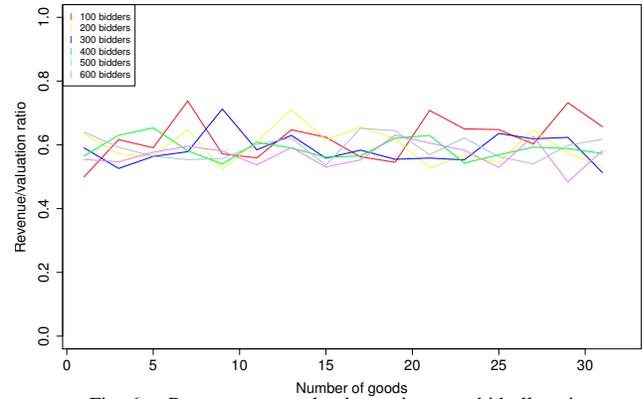


Fig. 6. Revenue over valuation ratio - one bid allocation.

VII. CONCLUSION

Today's Internet displays vast potential path diversity, which is truncated by the inter-domain routing protocol BGP. Some works [3], [4] propose some architectures to enable the propagation and the use of this path diversity. Nevertheless, the enabling of such an amount of paths would lead to import scalability issues [5] and ISPs need a mechanism to provide limited access to these alternate routes.

We propose in this paper a path allocation process based on auctions, organized by a network service provider to locally sell routes to its neighbors. Generally speaking, this auction-like process has for benefits to both propagate to each neighbor only the paths he is interested in and to establish a price according to what he can pay. In addition, the allocation framework we propose has the advantage that the winner determination, which is based on the maximization of the seller's revenue, is loser coalition proof, which means that no loser or collusion of losers can question the seller's selection of winners. Furthermore we propose a payment function that is dominant truthful when bidders submit one bid, meaning that

each bidder has an incentive to bid the true value he evaluates the set of paths he bids on.

As a further work, we aim at studying this allocation framework if other constraints come into play (e.g., bandwidth constraints...) and characterize the performance of the mechanism, in term of revenue.

ACKNOWLEDGMENTS

This work has received funding from the European Unions's Seventh Framework Program (FP7/2007-2013) under grant agreement 248567 for the ETICS project.

REFERENCES

- [1] Y. Rekhter, T. Li, and S. Hares, "A Border Gateway Protocol 4 (BGP-4)," IETF RFC 4271 (Draft Standard), Internet Engineering Task Force, 2006.
- [2] W. Mühlbauer, A. Feldmann, O. Maennel, M. Roughan, and S. Uhlig, "Building an as-topology model that captures route diversity," in *Proceedings of the 2006 conference on Applications, technologies, architectures, and protocols for computer communications*, ser. SIGCOMM '06. New York, NY, USA: ACM, 2006.
- [3] W. Xu and J. Rexford, "MIRO: Multi-path Interdomain ROuting," in *Proc. ACM SIGCOMM '06*, September 2006, pp. 171–182.
- [4] X. Misseri, I. Gojmerac, and J.-I. Rougier, "IDRD: Enabling Inter-Domain Route Diversity," in *IEEE International Conference on Communications (ICC) 2013*, Budapest, Hungary, 2013.
- [5] —, "Internet-wide multipath: a scalability analysis of path identification schemes," in *IEEE Network of the Future*, Tunis, 2012.
- [6] K. Lakshminarayanan, I. Stoica, S. Shenker, and J. Rexford, "Routing as a Service," UC Berkeley Technical Report, Tech. Rep., 2004.
- [7] S. Tao and R. Guérin, "Application-specific path switching: a case study for streaming video," in *Proceedings of the 12th annual ACM international conference on Multimedia*, ser. MULTIMEDIA '04. New York, NY, USA: ACM, 2004, pp. 136–143.
- [8] G. Huston, "BGP Routing Table Analysis Report." [Online]. Available: bgp.potaroo.net
- [9] L. Zhang, D. Meyer, and K. Fall, "Report from the IAB Workshop on Routing and Addressing," IETF RFC 4984, Internet Engineering Task Force, 2007.
- [10] K.-W. Kwong and R. Guérin, "Controlling the growth of internet routing tables through market mechanisms," in *Proceedings of the Re-Architecting the Internet Workshop*, ser. ReARCH '10. New York, NY, USA: ACM, 2010, pp. 2:1–2:6.
- [11] P. Milgrom, *Putting Auction Theory to Work*. Cambridge University Press, 2004.
- [12] J. Seedorf and E. Burger, "Application-Layer Traffic Optimization (ALTO) Problem Statement," RFC 5693 (Informational), Internet Engineering Task Force, Oct. 2009. [Online]. Available: <http://www.ietf.org/rfc/rfc5693.txt>
- [13] "DrPeering International." [Online]. Available: <http://drpeering.net/>
- [14] R. Day and P. Milgrom, "Core-selecting package auctions," *International Journal of Game Theory*, vol. 36, no. 3-4, pp. 393–407, Jul. 2007.
- [15] G. Aggarwal, T. Feder, R. Motwani, and A. Zhu, "Algorithms for multi-product pricing," in *In Proceedings of the International Colloquium on Automata, Languages, and Programming*, 2004, pp. 72–83.
- [16] P. Briest and P. Krysta, "Single-minded unlimited supply pricing on sparse instances," in *In Proceedings of the 17th ACM-SIAM Symposium on Discrete Algorithms*, 2006, pp. 1093–1102.
- [17] M.-F. Balcan and F. Constantin, "Sequential item pricing for unlimited supply," in *Proceedings of the 6th international conference on Internet and network economics*, ser. WINE'10. Berlin, Heidelberg: Springer-Verlag, 2010, pp. 50–62.
- [18] M.-F. Balcan and A. Blum, "Approximation algorithms and online mechanisms for item pricing," in *Proceedings of the 7th ACM conference on Electronic commerce*, ser. EC '06. New York, NY, USA: ACM, 2006, pp. 29–35.
- [19] M.-F. Balcan, A. Blum, and Y. Mansour, "Item pricing for revenue maximization," in *Proceedings of the 9th ACM conference on Electronic commerce*, ser. EC '08. New York, NY, USA: ACM, 2008, pp. 50–59.
- [20] J. D. Hartline and V. Koltun, "Near-optimal pricing in near-linear time," in *In 9th Workshop on Algorithms and Data Structures*, 2005, pp. 422–431.
- [21] J. D. Hartline and R. McGrew, "From optimal limited to unlimited supply auctions," in *Proceedings of the 6th ACM conference on Electronic commerce*, ser. EC '05. New York, NY, USA: ACM, 2005, pp. 175–182.
- [22] M.-F. Balcan and A. Blum, "Mechanism design via machine learning," in *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, ser. FOCS '05. Washington, DC, USA: IEEE Computer Society, 2005, pp. 605–614.
- [23] I. Segal, "Optimal pricing mechanisms with unknown demand," *American Economic Review*, vol. 93, p. 2003, 2003.
- [24] V. Guruswami, J. D. Hartline, A. R. Karlin, D. Kempe, C. Kenyon, and F. McSherry, "On profit-maximizing envy-free pricing," in *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, ser. SODA '05. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics, 2005, pp. 1164–1173.
- [25] P. Cramton, Y. Shoham, and R. Steinberg, *Combinatorial auctions*. MIT Press, 2006.
- [26] A. V. Goldberg, J. D. Hartline, and A. Wright, "Competitive auctions and digital goods," in *Proceedings of the twelfth annual ACM-SIAM symposium on Discrete algorithms*, ser. SODA '01. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics, 2001.
- [27] A. Goldberg, Jason, H. May, A. V. Goldberg, and J. D. Hartline, "Competitive auctions for multiple digital goods," in *In Proc. 9th European Symposium on Algorithms*, 2001.
- [28] Z. Bar-Yossef, K. Hildrum, and F. Wu, "Incentive-compatible on-line auctions for digital goods," in *Proceedings of the thirteenth annual ACM-SIAM symposium on Discrete algorithms*, ser. SODA '02. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics, 2002, pp. 964–970.
- [29] A. V. Goldberg and J. D. Hartline, "Envy-free auctions for digital goods," in *Proceedings of the 4th ACM conference on Electronic commerce*, ser. EC '03, 2003, pp. 29–35.
- [30] A. V. Goldberg, J. D. Hartline, A. R. Karlin, and M. Saks, "A lower bound on the competitive ratio of truthful auctions," in *In Proceedings 21st Symposium on Theoretical Aspects of Computer Science*. Springer, 2004, pp. 644–655.
- [31] A. V. Goldberg, J. D. Hartline, A. R. Karlin, M. Saks, and A. Wright, "Competitive auctions," *Games and Economic Behavior*, vol. 55, no. 2, 2006.
- [32] L. M. Ausubel and P. Milgrom, *The Lovely but Lonely Vickrey Auction*, P. Cramton, Y. Shoham, and R. Steinberg, Eds. MIT Press, 2006, vol. 94305, no. 03.
- [33] B. Day and P. Milgrom, "Optimal Incentives in Core-Selecting Auctions," in *Handbook of Market Design*, 2010.
- [34] A. Erdil and P. Klemperer, "A new payment rule for core-selecting package auctions," *Journal of the European Economic Association*, vol. 8, no. 2-3, pp. 537–547, 2010.
- [35] *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, 2009, vol. 54, no. 1-4.
- [36] R. W. Day and S. Raghavan, "Computing core payments in combinatorial auctions," *SIGecom Exch.*, vol. 7, no. 1, pp. 22–24, Dec. 2007.
- [37] P. Maille and B. Tuffin, "Multi-bid auctions for bandwidth allocation in communication networks," in *In Proc. of IEEE INFOCOM*, 2004.
- [38] M. Dramitinos, G. D. Stamoulis, and C. Courcoubetis, "An auction mechanism for allocating the bandwidth of networks to their users," *Computer Networks*, vol. 51, no. 18, pp. 4979 – 4996, 2007.
- [39] J. Feigenbaum, C. Papadimitriou, R. Sami, and S. Shenker, "A bgp-based mechanism for lowest-cost routing," *Distrib. Comput.*, vol. 18, no. 1, pp. 61–72, Jul. 2005.
- [40] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [41] V. Valancius, N. Feamster, R. Johari, and V. Vazirani, "Mint: a market for internet transit," in *Proceedings of the 2008 ACM CoNEXT Conference*, ser. CoNEXT '08. New York, NY, USA: ACM, 2008.
- [42] X. Misseri, J.-L. Rougier, and D. Saucez, "Internet routing diversity for stub networks with a Map-and-Encap scheme," in *ICC*, 2012.
- [43] D. Lehmann, L. I. O'Callaghan, and Y. Shoham, "Truth revelation in approximately efficient combinatorial auctions," *J. ACM*, vol. 49, no. 5, pp. 577–602, Sep. 2002.
- [44] C. H. Papadimitriou, "Computational complexity," in *Encyclopedia of Computer Science*. Chichester, UK: John Wiley and Sons Ltd., pp. 260–265.