Multi-bid versus Progressive Second Price Auctions in a Stochastic Environment

Patrick Maillé[1] and Bruno Tuffin[2]

[1] ENST Bretagne, BP 78, 2 rue de la Châtaigneraie
35512 Cesson Sévigné Cedex, France
pmaille@rennes.enst-bretagne.fr
[2] IRISA-INRIA, Campus universitaire de Beaulieu
35042 Rennes Cedex, France
btuffin@irisa.fr

Abstract. Pricing is considered a relevant way to control congestion and differentiate services in communication networks. Among all pricing schemes, auctioning for bandwidth has received a lot of attention. We aim in this paper at comparing a recently designed auction scheme called multi-bid auction with the often referenced progressive second price auction. We especially focus on the case of a stochastic environment, with players/users entering and leaving the game. We illustrate the gain that can be obtained with multi-bids, in terms of complexity, revenue and social welfare in both transient and steady-state regime.

1 Introduction

To cope with congestion in communication networks, it has been proposed to switch from current flat-rate pricing to usage-based or congestion-based pricing schemes (see for instance [3, 4] for surveys on pricing in telecommunication networks, describing the range of possibilities; for the sake of conciseness, we do not describe all the schemes here). Among those pricing schemes, auctioning has appeared as a possibility to share bandwidth. The first time auctioning was proposed was in the seminal smart-market scheme of MacKie-Mason and Varian [6], where each packet contains a bid and, if served, pays the highest bid of the packets which are denied service. This scheme requires a high engineering cost, but has pioneered the auction-based pricing activity in the networking community.

Progressive Second Price (PSP) Auction [5, 11, 12] has recently been proposed as a trade-off between engineering feasability and economic efficiency. In PSP, players submit bids at different epochs, each bid consisting of the required amount of bandwidth and the associated unit-price, until a (Nash) equilibrium is reached. The scheme has been proved to be incentive compatible and efficient. Variants of PSP have been designed in [8, 14] in order to fix some of its drawbacks.

In [9], multi-bid auction (a one-shot version of PSP) has been proposed. It consists for each player in submitting multiple bids once only, providing therefore an approximation of her own valuation function. Market clearing price and allocation can be subsequently computed. Here again, incentive compatibility and efficiency are proved (up to a given constant). The scheme presents the advantage, with respect to PSP, that no bid profile diffusion is necessary along the network, and that there is no convergence phase up to equilibrium, then yielding a gain in engineering and economic efficiency, especially when players enter and leave the game randomly.

The goal of this paper is to numerically highlight and illustrate the gain that can be obtained by multi-bids over PSP. We place ourselves in a stochastic environment, with users of different types entering and leaving the game at random times, and investigate the transient (for a given trajectory) behavior and steady-state performance of both multi-bids and PSP. We especially focus on three criteria: network revenue, social welfare and computational complexity.

Note finally that there exist other auction schemes in the literature [2, 10, 13], but due to space limitation, and since our main purpose was to emphasize the degree of improvement when using the multi-bids instead of PSP, we do not include them here.

The layout of the paper is as follows. In Section 2, we present the stochastic model that will be used to describe the system behavior. In Section 3 we present the PSP mechanism and its properties; the same is done for the multi-bid scheme in Section 4. Section 5 illustrates the gain that can be obtained by using the multi-bid scheme in a stochastic environment; both transient and steady-state results are provided. Finally, we conclude in Section 6.

2 General model

In order to look at the auction schemes' behavior in a stochastic environment, we model for convenience the system by a Markov process.

Consider a single communication link of capacity Q. We assume that there exists a finite number T of different valuation (or willingness-to-pay) functions, corresponding for instance to different sets of applications. A player/user i is then characterized by her *type* $t_i \in \{1, ..., T\}$.

Players compete for bandwidth. To model their behavior, we represent their perception/valuation of the service they can get by a quasi-linear utility function of the form

$$U_i(s) = \theta_{t_i}(a_i(s)) - c_i(s), \tag{1}$$

where θ_{t_i} is the *valuation function* of a type-t_i player, which depends on the quantity of resource received a_i. Quantity c_i is for the total cost charged to player i. Both c_i and a_i will depend on the auction scheme used and on the whole set of bids s (where the term "bid" will depend on the auction scheme).

We assume that new players enter the game according to a Poisson process with rate λ, and that the type of a new player is chosen according to a discrete probability distribution $\mathbb{P}_t$, so that the arrival rate of type u is $\lambda_u = \lambda \mathbb{P}_t(u)$. We also assume that each type-u player sojourn time is exponentially distributed with rate μ_u (independent here of the obtained accumulated bandwidth, like for real-time applications). Let $\mathcal{I}(\tau)$ be the set of active players at time τ (i.e. the set of players present in the game at this time) and $I(\tau)$ be the total number of players at time τ.

To ensure that the bandwidth is not sold at a too low level, the seller can thus be seen as a (permanent) player, noted by 0, with valuation function $\theta_0(q) = p_0 q$. p_0, the reserve price, guarantees that no bandwidth will be sold at a unit price under p_0.

Our goal is to compare the behavior of PSP and multi-bids. Let us now recall the basic concepts of both schemes.

3 Progressive Second Price Auction [5, 11]

In PSP, a player i submits a 2-dimensional bid $s_i = (q_i, p_i) \in \mathcal{S}_i = [0, Q] \times [0, +\infty)$, where q_i is the desired quantity of resource and p_i the *unit* price player i is willing to pay for that resource. $s = (s_1, ..., s_I)$ will denote the bid profile, and $s_{-i} = (s_1, ..., s_{i-1}, s_{i+1}, ..., s_I)$ will be the bid profile that player i faces, so that $s = (s_i; s_{-i})$ (the dependence on time τ is omitted to simplify the notations).

PSP allocation and charge to player i are

$$a_i(s) = q_i \wedge \left[Q - \sum_{p_k \geq p_i, k \neq i} q_k \right]^+ \tag{2}$$

$$c_i(s) = \sum_{j \in \mathcal{I}(\tau),\ j \neq i} p_j \left[a_j(s_{-i}) - a_j(s_i; s_{-i}) \right], \tag{3}$$

so that players bidding the highest get the bandwidth they request and total charge corresponds to declared willingness to pay of players who are excluded by player i's bid.

Each time a player submits a bid, she tries to maximize her utility, and a bid fee ε is charged to her. Under some concavity and regularity assumptions over functions θ_u, when the number of players is fixed and players bid sequentially, the game is proved to converge to a so-called ε-Nash equilibrium, so that no player can improve unilaterally her utility by more than ε. The scheme is also proved to be incentive compatible (meaning that users' best interest is to truly reveal their willingness to pay), and efficient in the sense that the social welfare $\sum_{i \in \mathcal{I}(\tau) \cup \{0\}} \theta_i(a_i)$ is asymptotically maximized (when the algorithm has converged).

Based on the assumption that users enter or leave the game, efficiency might become an issue. We suppose here that each type-u player in the game has the opportunity to submit a new bid at different times. Inter-bid times are assumed to follow an exponential distribution with parameter ν_u, independent of all other random variables. When a new player arrives, she is assumed to submit an optimal bid (meaning that she knows the bid profile).

4 Multi-Bid Auction [9]

In the multi-bid scheme, users, when they enter the game, submit a set of M 2-dimensional bids $s_i = \{s_i^1, ..., s_i^M\}$, where for all $m, 1 \leq m \leq M, s_i^m = (q_i^m, p_i^m)$

is as in PSP (the seller just submits one 2-dimensional bid $s_0 = (q_0, p_0)$ with $q_0 > Q$ and p_0 the reserve price). We assume without loss of generality that bids are sorted such that $p_i^1 \leq p_i^2 \leq ... \leq p_i^M$. With respect to PSP, the bids are submitted just once, so that users do not submit new bids at given epochs. This reduces the signaling overhead.

From the multi-bids of all competing players at time τ, the so-called pseudo-demand function of user i can be computed as the function $\bar{d}_i : \mathbb{R}^+ \to \mathbb{R}^+$, defined by

$$\bar{d}_i(p) = \begin{cases} 0 & \text{if } p_i^M < p \\ \max\limits_{1 \leq m \leq M} \{q_i^m : p_i^m \geq p\} & \text{otherwise.} \end{cases} \tag{4}$$

The pseudo-aggregated demand function is the function $\bar{d} : \mathbb{R}^+ \to \mathbb{R}^+$ defined by $\bar{d}(p) = \sum_{i \in \mathcal{I}(\tau) \cup \{0\}} \bar{d}_i(p)$, where $\bar{d}_0(p) = q_0 \mathbb{1}_{p \leq p_0}$ (apply (4) for $M = 1$) .

From the pseudo-aggregated demand function, we define the pseudo-market clearing price $\bar{u}$ by

$$\bar{u} = \sup \left\{ p : \bar{d}(p) > Q \right\}. \tag{5}$$

Such a $\bar{u}$ always exists since $\bar{d}(0) \geq \bar{d}_0(0) = q_0 > Q$. Moreover for $p > \max_{i \in \mathcal{I}(\tau) \cup \{0\}}(p_i^M)$ we have $\bar{d}(p) = 0$, and therefore $\bar{u} < +\infty$.

Describe now the allocation and pricing rules. First define, for every function $f : \mathbb{R} \to \mathbb{R}$ and all $x \in \mathbb{R}$, $f(x^+) = \lim_{z \to x, z > x} f(z)$.

The allocation, recomputed each time a player enters or leaves the game, is

$$a_i(s_i, s_{-i}) = \bar{d}_i(\bar{u}^+) + \frac{\bar{d}_i(\bar{u}) - \bar{d}_i(\bar{u}^+)}{\bar{d}(\bar{u}) - \bar{d}(\bar{u}^+)} (Q - \bar{d}(\bar{u}^+)), \tag{6}$$

meaning that each player receives the quantity she asks at the lowest price $\bar{u}^+$ for which supply excesses pseudo-demand, $\bar{d}_i(\bar{u}^+)$, and the excess of resource is shared among players who submitted a bid with price $\bar{u}$.

The total charge is computed according to the second-price principle [1, 15] (but using the pseudo-demand functions instead of the real ones):

$$c_i(s_i, s_{-i}) = \sum_{j \in \mathcal{I}(\tau) \cup \{0\}, j \neq i} \int_{a_j(s)}^{a_j(s_{-i})} \bar{\theta}'_{t_j}(q) dq, \tag{7}$$

with $\bar{\theta}'_{t_j}$ pseudo-marginal valuation function of j, defined by

$$\bar{\theta}'_{t_j}(q) = \begin{cases} 0 & \text{if } q_j^1 < q \\ \max\limits_{1 \leq m \leq M} \{p_j^m : q_j^m \geq q\} & \text{otherwise.} \end{cases} \tag{8}$$

As for PSP, incentive compatibility (each user i should better reveal its bandwidth valuation, i.e., $p_i^m = \theta'_{t_i}(q_i^m)\ \forall m$), and efficiency are proved, but up to a controlled constant here (see [9] for details).

It is shown in [9] that it is in the players' interest to submit a uniform quantile repartition of their bids, i.e., $(q_i^m, p_i^m = \theta'_{t_i}(q_i^m))\ \forall 1 \leq m \leq M$ such that

$$\int_{d_i(p_i^{m+1})}^{d_i(p_i^m)} (\theta'_{t_i}(q) - p_i^m) dq = C_i \qquad \forall m, \text{ where } \begin{cases} p_i^{M+1} = \theta'_{t_i}(0) \\ p_i^0 = p_0. \end{cases} \tag{9}$$

5 Comparison of performance

Multi-bids present the following advantages with respect to PSP:

– since the bids are submitted exactly once, no convergence phase is required by resubmitting new bids until an equilibrium is reached. It might be argued that the mean number of re-submission up to equilibrium is less that the number M of multi-bids in some cases; this situation is less likely to occur in the situation of customers arriving or leaving the game, meaning that a new re-submission phase is required for each player in PSP, whereas nothing has to be done for already present players when using multi-bids.
– Following the same idea, when submitting a new bid in PSP, each player is assumed to know the bid profile, meaning that it is advertised to all players. This is not required for the multi-bid scheme, saving then a lot of signaling overhead.

We propose to illustrate the above advantages of multi-bids in the following sub-sections. We especially wish to show that this gain in terms of signaling/complexity is not at the expense of efficiency, in terms of seller's revenue or social welfare, both on a trajectory and during the convergence phase of PSP, as well as in steady state, and that it even actually is the converse.

5.1 Transient analysis

Figure 1 displays the behavior of PSP and multi-bids when the number of players is fixed and until equilibrium is reached for PSP, with two types of players, three type-1 and two type-2 players. The upper left-hand side figure displays the valuation and marginal valuation functions for both types of players, that we used in all our simulations. The lower left-hand side represents social welfare $\sum_{i\in\mathcal{I}(\tau)\cup\{0\}}\theta_i(a_i)$. Since the number of players is fixed during the simulation, multi-bid allocations and charges are fixed, and it can be observed that the social welfare is 60.37, very close to the optimal one 60.71. On the other hand, for PSP auctions the social welfare changes at each re-submission from a user (resulting in the discontinuous curve), reaching equilibrium (with value 60.68) around time $\tau = 26$, but showing a lower social welfare than multi-bids before reaching equilibrium. The lower right-hand side of Figure 1 represents the network revenue for both schemes. Again, multi-bid revenue is constant due to the fact that the number of players is fixed. Also, the revenue for PSP is first increasing, overtaking the one with multi-bids after a while, but then dropping under it just before reaching equilibrium. Actually, we proved in [7] that when the total demand at the unit price p_0 exceeds the available capacity Q, the revenue with PSP in equilibrium tends to $p_0 \times Q$ (i.e. all the resource is sold at the reserve price) when the bid fee ε tends to 0.

Figure 2 illustrates the behavior of both schemes on a trajectory, with players entering and leaving the game[3]. Here the number of players of each type varies, as

[3] The parameters we chose are precised in the figure, and were also used for the study of steady-state performance.

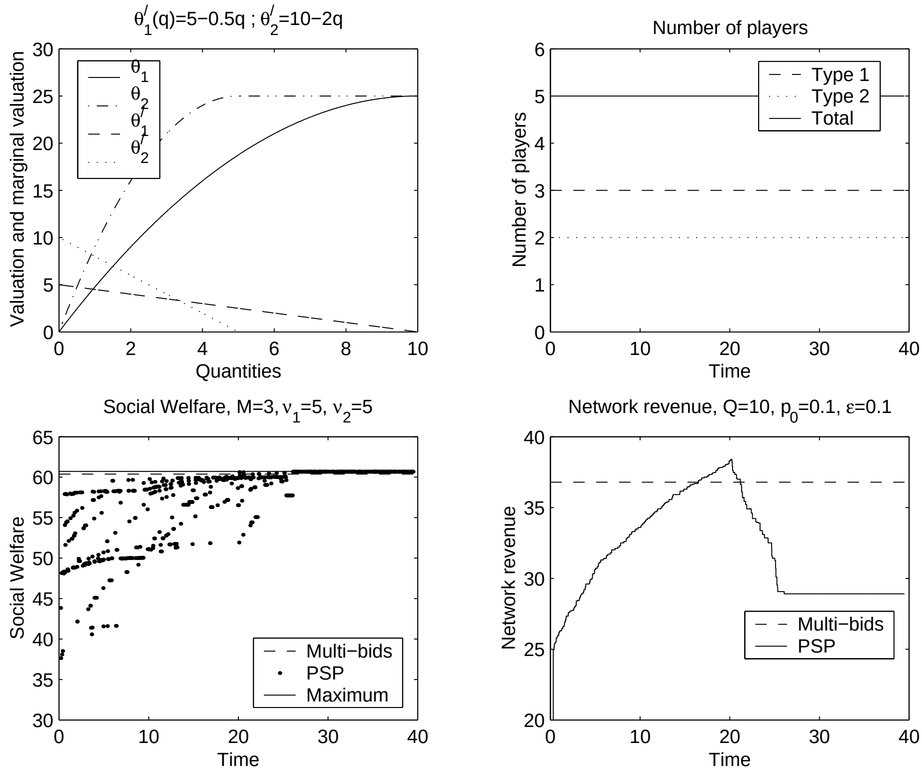

Fig. 1. Comparison of PSP and multi-bids for a fixed number of players, until convergence is reached for PSP

described on the upper right-hand figure. The curves of social welfare show that, when using multi-bids, the resulting social welfare is always very close to the optimal one, whereas, due to the convergence phase, there is a loss of efficiency when using PSP. Similarly, on this trajectory, the network revenue generated by multi-bids is significantly larger than the one generated by PSP.

5.2 Steady-state analysis

Figure 3 illustrates the evolution of the mean efficiency ratio (obtained steady-state social welfare divided by the optimal one), the mean network revenue and the complexity of the algorithm when the number M of multiple bids increases, and compares those performance measures with the ones obtained for PSP. The complexity of computing PSP allocations and prices is of the order $O(I^2)$ [11], and the complexity of multi-bid auction is of the order $O(M \times I^2)$ [9]. We therefore display the mean number of applications of each auction rule by unit of time, multiplying this number by M for multi-bid auction. This curve does not

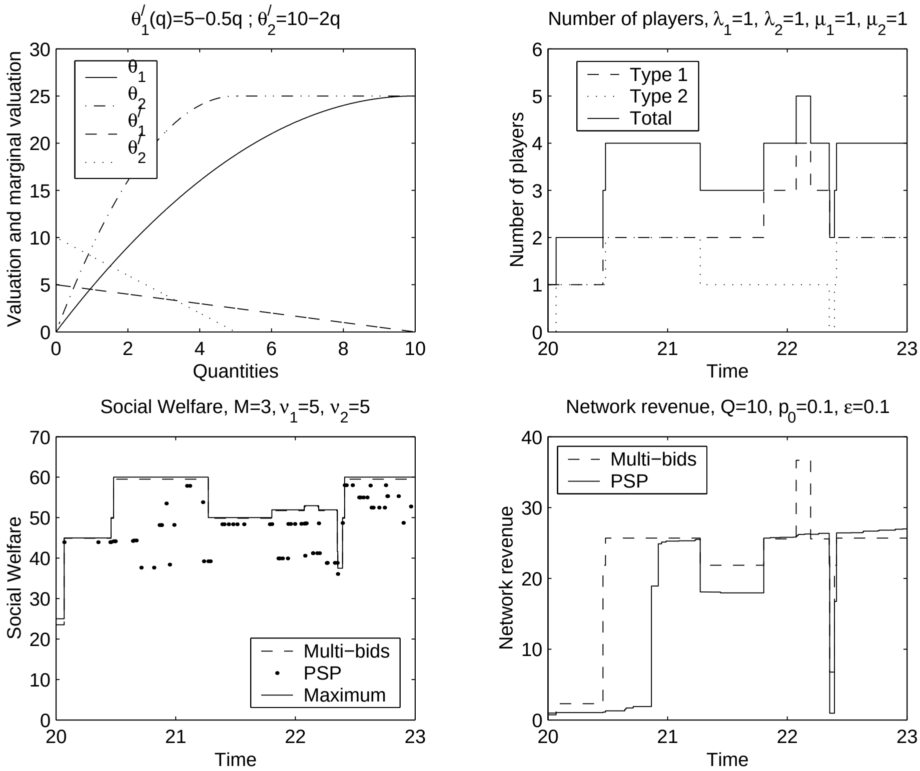

Fig. 2. Behavior of PSP and multi-bids on a trajectory, with players entering and leaving the game

precisely give the number of elementary operations that are conduced, but just gives an idea of how the computational complexity evolves when the parameters vary. Note that this computation of complexity does not include the signaling overhead necessary for PSP. It can be observed then that for small values of M, computational complexity is even smaller also with multi-bids. More important, thanks to the one-shot property of multi-bids (i.e., the fact that no convergence phase is required unlike PSP), steady-state social welfare (for $M \geq 2$) and revenue are larger with multi-bids.

Figure 4 displays the evolution of efficiency ratio, network revenue and complexity when the arrival rate increases (with all other parameters fixed). Again, multi-bids are shown to provide better performance. The difference increases with λ. This is due to the fact that the number of players varies more frequently, so that convergence to optimal values is less likely to occur for PSP, whereas it does not affect multi-bids.

Figure 5 illustrates the three criteria considered in this paper, when the bid-resubmission rate varies for both types of players. Even when this rate increases,

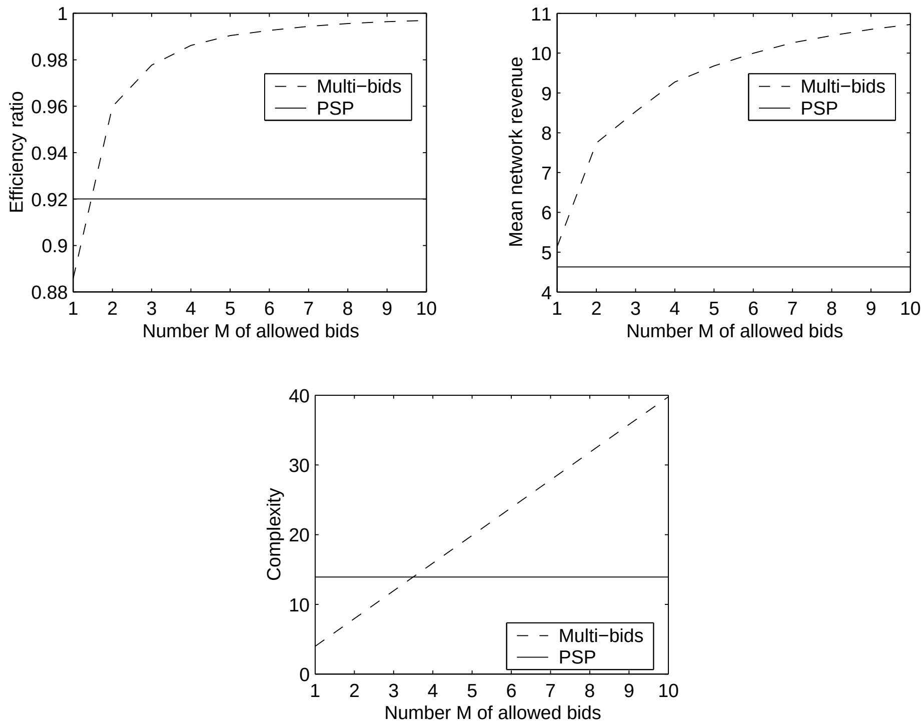

Fig. 3. Steady-state performance of multi-bids for an increasing number M of allowed two-dimensional bids in multi-bid auctions, compared with PSP

leading to a larger computational complexity, we see that the multi-bid auction still outperforms PSP as concerns efficiency and network revenue.

6 Conclusions

The goal of this paper was to compare PSP and multi-bid schemes, two auction mechanisms for bandwidth allocation in telecommunication networks. Based on this purpose, we have considered a model representing a communication link, with players applying for connections at random epochs, and for a random time. Our conclusion is that multi-bid auction scheme significantly reduces the signaling overhead of PSP, but also yields larger social welfare and network revenue (at least for this stochastic context regarding the social welfare).

As future work, we plan to extend the multi-bid auction scheme to a whole network. We already have some results in the case of a tree network, which properly represents the case where the backbone network is overprovisionned and the access networks have a tree structure.

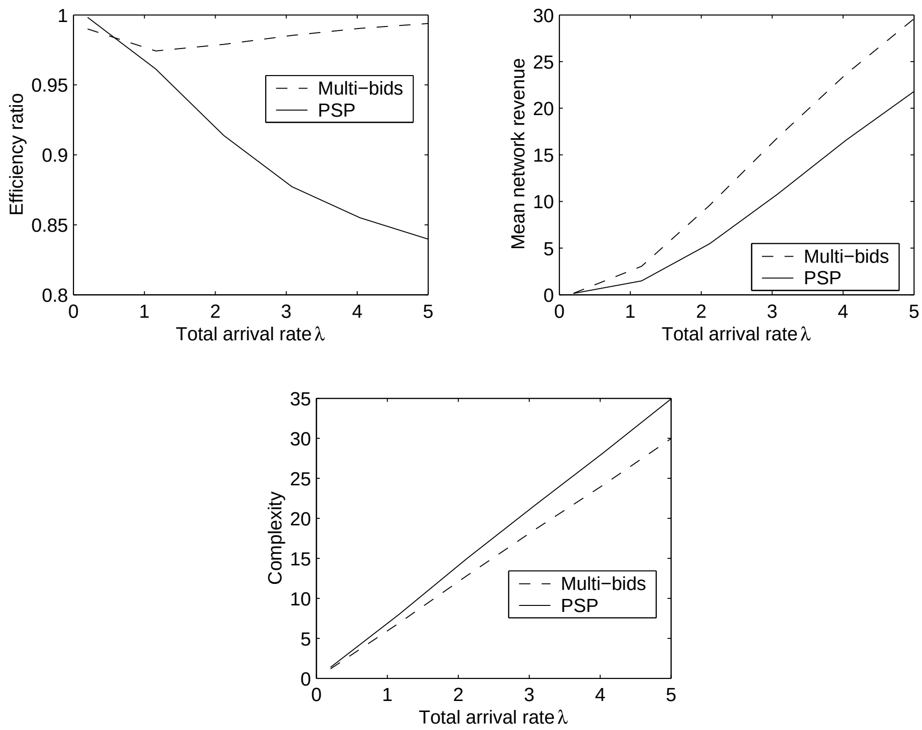

Fig. 4. Steady-state performance comparison when the charge increases

References

1. E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.
2. C. Courcoubetis, M. P. Dramitinos, and G. D. Stamoulis. An auction mechanism for bandwidth allocation over paths. In *Proc. of the 17th International Teletraffic Congress (ITC)*, Dec 2001.
3. L. A. DaSilva. Pricing for QoS-enabled networks: A survey. *IEEE Communications Surveys*, 3(2):2–8, 2000.
4. M. Falkner, M. Devetsikiotis, and I. Lambadaris. An overview of pricing concepts for broadband IP networks. *IEEE Communications Surveys & Tutorials*, 3(2), 2000.
5. A. A. Lazar and N. Semret. Design and analysis of the progressive second price auction for network bandwidth sharing. *Telecommunication Systems - Special issue on Network Economics*, 1999.
6. J. K. MacKie-Mason and H. R. Varian. Pricing the internet. In *Public Access to the Internet*, JFK School of Government, May 1993.
7. P. Maillé. Market clearing price and equilibria of the progressive second price mechanism. Technical Report 1522, IRISA, Mar 2003. submitted.

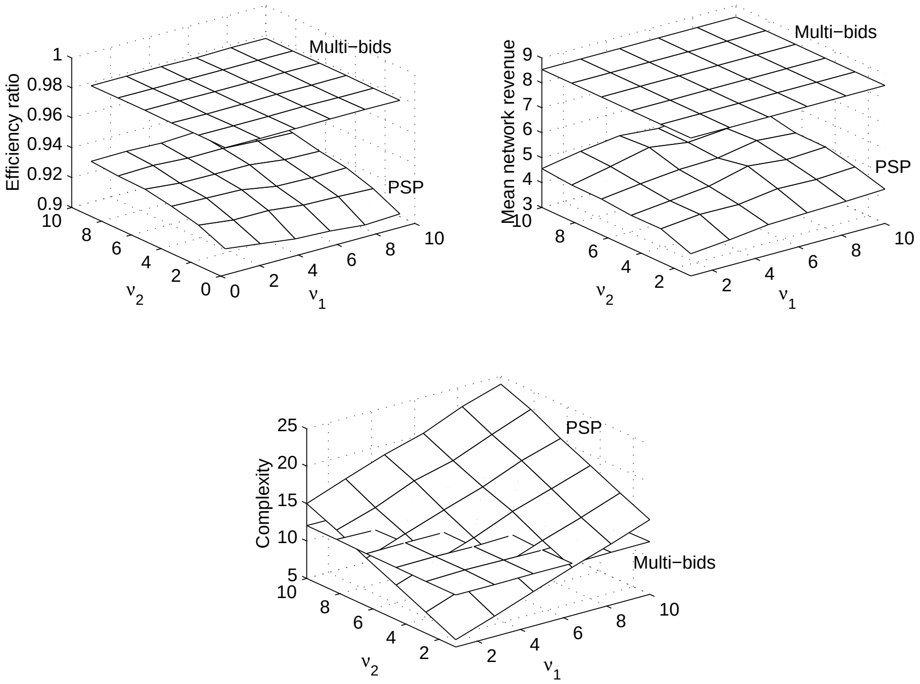

Fig. 5. Steady-state performance comparison when the frequency of re-submission increases in PSP

8. P. Maillé and B. Tuffin. A progressive second price mechanism with a sanction bid from excluded players. In *Proc. of ICQT'03, LNCS 2816*, pages 332–341, Munich, Sept 2003. Springer.
9. P. Maillé and B. Tuffin. Multi-bid auctions for bandwidth allocation in communication networks. In *Proc. of IEEE INFOCOM*, Mar 2004.
10. P. Reichl, G. Fankhauser, and B. Stiller. Auction models for multiprovider internet connections. In *Proc. of Messung, Modelierung und Bewertung (MMB'99)*, Trier, Germany, Sept 1999.
11. N. Semret. *Market Mechanisms for Network Resource Sharing.* PhD thesis, Columbia University, 1999.
12. N. Semret, R. R.-F. Liao, A. T. Campbell, and A. A. Lazar. Pricing, provisionning and peering: Dynamic markets for differentiated internet services and implications for network interconnections. *IEEE Journal on Selected Areas in Communications*, 18(12):2499–2513, Dec 2000.
13. J. Shu and P. Varaiya. Pricing network services. In *Proc. of IEEE INFOCOM*, 2003.
14. B. Tuffin. Revisited progressive second price auctions for charging telecommunication networks. *Telecommunication Systems*, 20(3):255–263, Jul 2002.
15. W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, Mar 1961.