

# Auctions for resource allocation in overlay networks

Pablo Belzarena<sup>1</sup>, Andrés Ferragut<sup>1,2</sup>, and Fernando Paganini<sup>2</sup>

<sup>1</sup> Universidad de la República

<sup>2</sup> Universidad ORT

Montevideo, Uruguay

**Abstract.** The paper studies the problem of allocating bandwidth resources of a Service Overlay Network, to optimize revenue. Clients bid for network capacity in periodically held auctions, under the condition that resources allocated in an auction are reserved for the entire duration of the connection, not subject to future contention. This makes the optimal allocation coupled over time, which we formulate as a Markov Decision Process (MDP). Studying first the single resource case, we develop a receding horizon approximation to the optimal MDP policy, using current revenue and the expected revenue in the next step to make bandwidth assignments. A second approximation is then found, suitable for generalization to the network case, where bids for different routes compete for shared resources. In that case we develop a distributed implementation of the auction, and demonstrate its performance through simulations.

## 1 Introduction

In recent years, many types of “overlay” networks have been proposed for the Internet. These overlays include content delivery networks, peer-to-peer file sharing, some voice-over-IP services, and testbed networks such as PlanetLab. A particular architecture called Service Overlay Network (SON) [6] has been proposed to deploy value-added Internet services with end-to-end quality of service (QoS). The basic components of the SON architecture are service gateways located in domain boundaries, and a network of tunnels acquired from the underlying domains with guaranteed bandwidth. Through this overlay, a client with good local connectivity in one domain can secure a high QoS connection with a remote domain. The overlay operator invests in the infrastructure and leased bandwidth to offer high-value services, for instance through distributed content servers; a profitable sale of this capacity is thus essential. In this paper we study an access control policy based on auctions for this purpose, where users bid for a service connection and the network gives access to the best bids to maximize its payoff, subject to the QoS constraints for the admitted clients.

Resource sharing policies based on auctions have been analyzed by different previous works [7],[5],[8],[10],[3]. One main difference in our approach is that we impose the condition that once bandwidth has been allocated in an auction, the successful bidder is guaranteed to hold it for the duration of his/her connection.

Previous work on auctions allows future bidders to compete with incumbent ones, albeit given the latter some advantage. Our application scenario does not allow this: consider selling video-on-demand content about 100 minutes long, in auctions every 5 minutes. A consumer will not purchase the service if he/she faces the risk of losing the connection close to the end of the movie.

Reserving bandwidth over multiple auctions means that the operator must assume the risk of future bids. Optimizing revenue with this risk becomes a stochastic dynamic optimization problem, that we formulate as a Markov decision process (MDP) [1, 9]. The optimization involves a tradeoff between the revenue of the current auction, and the expected value of bids the operator will miss in future auctions if it runs out of resources. In Section 2 we analyze the above problem in the case of a single link, and develop a series of approximations to the optimal policy. The aim of our approximations is to allow a distributed implementation of the policy over an arbitrary network topology, where bids are received at the edge for end-to-end services, and the network optimizes the overall revenue. This is described in Section 3, and the method is evaluated by simulation. Conclusions given in Section 4.

## 2 Auctions for one link

In this section we consider auctions for one access link. We make a few simplifying assumptions: all consumers bid for the same amount (unit) of bandwidth, and the link has capacity for  $C$  such connections. Bids are collected for a period of time  $T$ , and an auction is held; we assume for simplicity that the number  $N$  of bids is given, and bids are drawn from a known probability distribution. We denote by  $b_k^{(1)} \geq b_k^{(2)} \geq \dots \geq b_k^{(N)}$  the ordered bids in decreasing order that participate in the auction at time  $kT$ . The bids are for the entire duration of the service, and this is a first-price auction: the admitted users will pay their bid.

Let  $a_k$  represent the number of admitted connections in auction  $k$ . The revenue from this auction is then

$$U_{b_k}(a_k) := \sum_{i=1}^{a_k} b_k^{(i)}. \quad (1)$$

The function  $U_b(a)$  is defined above for integer values of  $a$ ; it is also convenient to extend it to a function of  $a \in \mathbb{R}$ , by linear interpolation. The latter function is increasing and concave in  $a$ . We also define  $\bar{U}(a) = E[U_b(a)]$ , where the expectation is over the distribution of the bids  $b$ ; i.e., we replace the current bids in (1) by their expectation. This is also increasing, piecewise linear and concave.

We will model the connection durations as independent exponential random variables, of mean  $1/\mu$ . Therefore at the end of the period  $T$  each connection has probability  $p := e^{-\mu T}$  of remaining active for the following period. While in practice one would have more deterministic information of the service duration, the above allows for a Markovian treatment of the allocation problem.

Let  $x_k$  denote the number of connections active at  $t = kT^-$ , i.e. before the  $k$ -th auction. The system admits  $a_k$  new connections,  $0 \leq a_k \leq C - x_k$ , taking

the total to  $x_k + a_k$ . By the next auction period,  $t = (k + 1)T^-$ , the number of active connections  $x_{k+1}$  follows then a binomial distribution with parameters  $x_k + a_k$  and  $p$ . Specifically,  $P[x_{k+1} = i | x_k, a_k] = \binom{x_k + a_k}{i} p^i (1 - p)^{x_k + a_k - i}$ . We are ready to state our design objective.

**Optimal revenue problem:** Maximize  $\lim_n \frac{1}{n} \sum_{k=0}^{n-1} E[U_{b_k}(a_k)]$ .

Here the expectation is over two sources of randomness: the vector of bids  $b_k$  and the departure process. The constraints are  $0 \leq a_k \leq C - x_k$  where  $x_k$  follows the binomial transition dynamics defined above. We can also consider the discounted version: Maximize  $\sum_{k=0}^{\infty} \rho^k E[U_{b_k}(a_k)]$ , where  $0 < \rho < 1$ .

Both of these are Markov Decision Processes (MDPs) [2, 9]. The *state* at time  $k$  is  $s_k = (x_k, b_k)$ , i.e. the current occupation and the incoming bids. Based on this state, the *action*  $a_k = a(s_k)$  decides how many bids to accept. Solving the MDP requires finding the *policy*  $a(s)$  that results in a minimum cost. In the discounted case  $\rho < 1$ , this policy satisfies the Bellman equation

$$V^*(x_0, b) = \max_{a \in \mathcal{A}_s} \{U_b(a) + \rho E[V^*(x_1, b')]\}, \quad (2)$$

where  $V^*$  is the value function and the expectation is taken over the binomial distribution of  $x_1 | (x_0, a)$  and the distribution of the next bid  $b'$ . The state-dependent constraints are  $\mathcal{A}_s = \{0 \leq a \leq C - x_0\}$ . For  $\rho = 1$ ,  $V^*$  satisfying (2) is no longer the optimal cost, but (2) still characterizes the optimal action  $a(s)$ .

It is in general difficult to solve the Bellman equation; a commonly used strategy is the *value iteration*  $V_{m+1}(x_0, b) := \max_{a \in \mathcal{A}_s} \{U_b(a) + \rho E[V_m(x_1, b')]\}$ ; starting with an arbitrary  $V_0(s)$ ,  $V_m(s)$  converges to  $V^*(s)$ , and the corresponding maximizing action converges to the optimal action [2].

**Receding horizon approximation.** We will use initial steps of the value iteration to approximate the optimal policy. Starting from  $V_0 \equiv 0$ , we have

$$V_1(x_0, b) = \max_{a \leq C - x_0} U_b(a) = U_b(C - x_0). \quad (3)$$

This first step gives the “myopic” policy  $a = C - x_0$ , that sells all available capacity disregarding the future. To improve on it, we take a second step in the value iteration:

$$\begin{aligned} V_2(x_0, b) &= \max_{a \leq C - x_0} \{U_b(a) + \rho E[V_1(x_1, b')]\} \\ &= \max_{a \leq C - x_0} \{U_b(a) + \rho E[U_{b'}(C - x_1)]\} \\ &= \max_{a \leq C - x_0} \{U_b(a) + \rho E_{x_1} \bar{U}(C - x_1)\}. \end{aligned} \quad (4)$$

In (4), the expectation over the random bid  $b'$  is included in  $\bar{U}$  defined above; the remaining expectation is over  $x_1 \sim \text{Bin}(x_0 + a, p)$ . The policy (4) can be given a *receding horizon* interpretation: the decision considers the current revenue plus the expected revenue of looking one step ahead, assuming all available capacity will be sold off at that time. This decision is then applied recursively; thus the future is taken into account at a limited level of complexity.

The first term in (4) increases with  $a$ . To characterize the second, we rewrite it as follows. Consider the function  $W(i) = \overline{U}(C) - \overline{U}(C - i)$ , piecewise linear and *convex* in  $i$ . Indeed, the increments  $W(i + 1) - W(i) = E[b^{(C-i)}]$ , (expectation of the  $(C - i)$ -th largest offer) are increasing in  $i$ . Now define  $\overline{W}(x) = E[W(x_1)]$ , where  $x_1 \sim \text{Bin}(x, p)$ , again linearly interpolated for non-integer  $x$ . The next statement follows from a stochastic comparison argument, omitted for brevity.


**Proposition 1.**  $\overline{W}(x)$  is increasing and convex in  $x$ .

With this notation, our one-step ahead optimization can be rewritten as

$$\max_{a \leq C - x_0} U_b(a) - \rho \overline{W}(x_0 + a), \quad (5)$$

a convex optimization problem. In it,  $\overline{W}(x_0 + a)$  plays the role of a *cost function* that makes the decision at the current time “internalize” the impact on future decisions.  $\overline{W}(x_0 + a)$  measures the expected loss of revenue in the auction at time  $T$  from having left  $x_0 + a$  occupied circuits at time  $t = 0^+$ .

It is now straightforward to determine the optimal policy  $a(x_0, b)$ . To optimize over  $a$  we plot the derivatives of  $U_b(a)$  and  $\rho \overline{W}(x_0 + a)$  (marginal utilities and costs) and look for a crossing point. This is depicted in Figure 1. The marginal utilities are just the current bids in decreasing order. The marginal costs represent the value of leaving one more free circuit for the next auction.



**Fig. 1.** Marginal utility versus marginal cost

The increasing marginal costs  $\overline{w}_i = \rho(\overline{W}(i) - \overline{W}(i - 1))$  act as successive thresholds for accepting bids. The acceptance policy is the value  $a$  such that for  $i = x_0 + a$  we have  $b^{(1)} \geq b^{(2)} \geq \dots \geq b^{(a)} \geq \overline{w}_i > b^{(a+1)}$ . In words: to accept  $a$  bids, the *lowest* one must exceed  $\overline{w}_i$  for  $i = x_0 + a$ . To accept one more, we require a *more demanding* threshold  $\overline{w}_{i+1}$  on this (smaller) bid<sup>3</sup>.

From the binomial distributions  $j_1 \sim \text{Bin}(i, p)$ ,  $j_2 \sim \text{Bin}(i - 1, p)$  and some combinatorics we obtain for the thresholds the expression

$$\overline{w}_i = \rho[E(W(j_1)) - E(W(j_2))] = \rho p \sum_{l=0}^{i-1} E(b^{(C-l)}) \binom{i-1}{l} p^l (1-p)^{i-1-l}. \quad (6)$$

Based on knowledge of  $\rho$ ,  $p$ , and the distribution of bids, this expression could be calculated offline and used for carrying out auctions with the policy (4).

**Examples:** we evaluate the previous results in a few simple cases (for  $\rho = 1$ ).

<sup>3</sup> Since  $b$ 's are random, the curves of Fig. 1 will generically cross at a single point.

For  $C = 1$ , there is a single link cost  $\bar{w}_1 = pE(b^{(1)})$ , that acts as an admission threshold for bids received when the circuit becomes empty. For instance in the case of  $N$  bids, uniformly distributed in  $[0, b_{max}]$ , we have  $\bar{w}_1 = p\frac{N}{N+1}b_{max}$ .

If the link has capacity  $C = 2$ , there are two marginal costs:  $\bar{w}_1 = pE(b^{(2)})$  for occupying the first connection, and  $\bar{w}_2 = p(E(b^{(2)})(1-p) + E(b^{(1)})p)$  for occupying the second. Again, for uniform in  $[0, b_{max}]$  bids this becomes

$$\bar{w}_1 = p \frac{N-1}{N+1} b_{max}; \quad \bar{w}_2 = p \frac{(N-1)(1-p) + Np}{N+1} b_{max}.$$

We now compare by simulations our receding horizon policy with the optimal infinite-horizon MDP, in the case of one circuit ( $C = 1$ ). In this simple case, the latter is also a threshold policy on the bids, but the optimal threshold does not have a simple formula; we computed it numerically through the value iteration algorithm from [4]. On the left in Figure 2 we show the acceptance thresholds for both policies: we see the infinite horizon threshold is more demanding. On the right we show the average utility obtained empirically by simulation of these two policies. Results are very similar. Therefore, in this case we have managed to extract almost the optimal utility just by looking one-step ahead with the policy. On the other hand, if we apply the myopic policy that always fills the link, the right-hand side plot shows there is a clear loss in utility.




Fig. 2. Optimal MDP, receding horizon and myopic policies.  $C = 1$ ,  $p = 0.1$ .

### A further approximation

Our ultimate goal is to perform an auction over a general network, where bids appear at different routes and are coupled by scarce resources at links. The stochastic calculations involved in (6) appear difficult to generalize, so we adopt a second approximation, replacing the function  $\bar{W}(x)$  in (5) by something easier to compute. Namely, define  $\phi(x) = W(E[x_1])$  for  $x_1 \sim \text{Bin}(x, p)$ . Since  $W(\cdot)$  is convex, this underestimates the one-step cost from before,  $\phi(x) \leq \bar{W}(x)$ . Nevertheless, if  $C$  is large the binomial distribution will be concentrated around its mean and the error is moderate. In return, we have the simple expression  $\phi(x) = W(px) = \bar{U}(C) - \bar{U}(C - px)$ . This is still piecewise linear and convex, but easier to compute. The second approximation to the optimal policy is given by  $\max_{a \leq C-x_0} U_b(a) - \rho\phi(x_0 + a)$ .

Equivalently, we can rewrite the above as the convex program

$$\begin{aligned} & \max U_b(a) + \rho \bar{U}(z), \\ & \text{subject to } x_0 + a \leq C, \quad p(x_0 + a) + z \leq C. \end{aligned} \quad (7)$$

At the optimum, the slack variable  $z$  will satisfy the constraint with equality, i.e.  $z = C - p(x_0 + a)$  is the expected allocation in the next interval.

### 3 The network case

In the following,  $l$  indexes the network links,  $r$  the routes across the network.  $R$  denotes the routing matrix,  $R_{lr} = 1$  iff route  $r$  includes link  $l$ .  $c = (c_l)$  is the vector of link capacities. We describe the allocation decision at time  $t = 0$ .

Define column vectors  $x_0$ ,  $a$ , and  $z$ , whose coordinates per route  $r$  denote respectively the rate  $x_0^r$  from previous occupation, the rate allocation  $a^r$  at the current auction, and the expected allocation  $z^r$  in the following auction ( $t = T$ ). We also define the piecewise linear utility  $U_{br}(a^r)$  based on current bids (1), and the utility  $\bar{U}_r(z^r)$  based on expected bids. Let  $p^r = e^{-\mu^r T}$  be the probability that a connection active at  $t = 0$  will remain active at  $t = T$ ;  $P = \text{diag}(p^r)$  is the corresponding diagonal matrix. Thus,  $P(a + x_0)$  is the expected input rate vector at  $t = T^-$ . The network generalization of (7) is

$$\begin{aligned} & \max \sum_r [U_{br}(a^r) + \rho \bar{U}_r(z^r)], \\ & \text{subject to } R(a + x_0) \leq c; \quad RP(a + x_0) + Rz \leq c. \end{aligned} \quad (8)$$

This optimization can also be rewritten as one for the current allocation  $a$  with a convex cost function  $\phi(x_0 + a)$  that represents the optimization in  $z$ . However, here the cost function would be coupled over the network. A better way to solve (8) is by duality. Consider the Lagrangian  $L(a, z, \alpha, \beta)$  given by

$$\begin{aligned} L &= \sum_r [U_{br}(a^r) + \rho \bar{U}_r(z^r)] + \alpha^T (c - R(a + x_0)) + \beta^T (c - Rz - RP(a + x_0)) \\ &= \sum_r [U_{br}(a^r) - (q^r + p^r v^r) a^r] + [\rho \bar{U}_r(z^r) - v^r z^r] + \alpha^T (c - Rx_0) + \beta^T (c - RPx_0). \end{aligned}$$

Here,  $\alpha$  and  $\beta$  are the vectors of Lagrange multipliers (prices) for each of the two constraints, and we have defined the aggregate prices per route  $q = R^T \alpha$ ,  $v = R^T \beta$ . We can solve the convex program through a dual, gradient projection algorithm similar to those used in the congestion control literature [11], but with additional prices. The algorithm takes the following form (in continuous time):

$$\begin{aligned} a^r &= \arg \max_{a^r} [U_{br}(a^r) - (q^r + p^r v^r) a^r]; & \dot{\alpha} &= [R(a + x_0) - c]_{\alpha}^+; \\ z^r &= \arg \max_{z^r} [\rho \bar{U}_r(z^r) - v^r z^r]; & \dot{\beta} &= [RP(a + x_0) + Rz - c]_{\beta}^+. \end{aligned}$$

An important difference with the congestion control case is that here the algorithm should take place in the *control plane* prior to the allocation of resources.

Despite this difference, we can still obtain a *distributed* computation by message passing between the network links, and a broker entity at the ingress node of route  $r$ . The maximization for  $a^r$  amounts simply to comparing the bids with the threshold price  $q^r + p^r v^r$ . Solving for  $z^r$  involves the expected bids, and the price  $v^r/\rho$ . We must allow time after closing the auction for this distributed algorithm to settle on an allocation vector  $a$ .

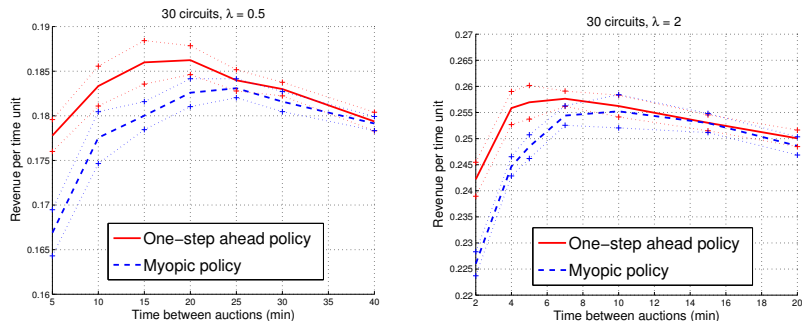
It is useful to compare this algorithm with an implementation of the myopic policy which optimizes the first term in (8); the latter would involve only the variables  $a$ ,  $\alpha$ , but a very similar overhead in terms of message passing.

### Simulations.

To evaluate the proposed algorithm we present some simulation studies, carried out on a flow-level simulator we developed in Java.

First, we compare the one-step ahead policy (7) with the myopic one (3) in a single link case, with  $C = 30$  circuits. Auctions take place every  $T$  minutes, and bids arrive periodically with fixed intensity  $\lambda$ , totalling  $N = \lambda T$  bids per auction. Bids are independent and uniformly distributed in  $[0, 1]$ ; rejected bids are discarded. Job duration is  $\exp(\mu)$  with  $1/\mu = 100$  minutes.

We simulate both policies and compare the respective revenues, for several values of  $T$  which is a critical design parameter: enlarging  $T$  will allow more bids per auction, hence better ones, but decrease the auction rate. Results are shown in Fig. 3 for two different values of  $\lambda$ . In both cases, the one-step ahead policy outperforms the myopic one, as expected; and as remarked before, the gain is achieved with minor extra overhead.



**Fig. 3.** One link situation: 30 circuits, offer arrival rate  $\lambda = 0.5$  (left) and  $\lambda = 2$  (right). Dotted lines are 5% confidence intervals.

Our second study uses the linear network topology of Fig. 4. In this case, users in the long route 1 should pay more to be allocated a circuit, since it occupies two links. We used  $N = 10$  uniform bids with mean 1 for the short routes, and with mean  $\bar{b}$  for route 1. Fig. 4 shows the mean number of connections admitted by the auction algorithm, for several values of  $\bar{b}$ . If  $\bar{b} = 1$ , the users in route 1 starve, but for  $\bar{b} \approx 2$  all routes receive a comparable fraction of resources.




Fig. 4. Linear network, allocation as a function of the mean bid  $\bar{b}$  of route 1.

## 4 Conclusions

In this work we proposed an auction mechanism to assign resources in overlay networks. We formulated the problem of maximizing operator revenue, and found near-optimal policies that can be computed via convex optimization, and allow a distributed implementation over a network. In future work we will study some natural extensions: random numbers of bids per auction, bidding for heterogeneous services which consume different amounts of bandwidth, learning an unknown distribution of bids over time, and multiple-step extensions to the receding horizon policy. We will also study the strategic aspects of the auction.

**Acknowledgment.** This work was supported by a grant from PDT-Uruguay.

## References

1. E. Altman, “Applications of Markov Decision Processes in Communication Networks: a Survey”, in *Markov Decision Processes, Models, Methods, Directions, and Open Problems*, E. Feinberg and A. Schwartz (Editors) Kluwer, pp. 488-536, 2001.
2. Dimitri P. Bertsekas, *Dynamic Programming : Deterministic and Stochastic Models*, Prentice-Hall, 1987.
3. M. Dramitinos, G.D. Stamoulis, C. Courcoubetis, “An auction mechanism for allocating the bandwidth of networks to their users”, *Computer Networks* 51 (2007).
4. I. Chads, M-J.Cros, F. Garcia, R. Sabbadin “Markov Decision Process Toolbox for MATLAB”, <http://www.inra.fr/internet/Departements/MIA/T//MDPtoolbox/>
5. M. Dramitinos, G. Stamoulis and C. Courcoubetis, “Auction-based resource reservation in 2.5/3G Networks” *Mobile Networks and Apps.*, 9, pp. 557-566, 2004.
6. Z. Duan, Z-L.Zhang and Y. T. Hou, Service Overlay Networks: “SLAs, QoS and Bandwidth Provisioning”, *IEEE/ACM Trans. on Networking*, pp 870 - 883, 2003.
7. A. A. Lazar, N. Semret, “Design and Analysis of the Progressive Second Price auction for Network Bandwidth Sharing”. *Telecommunications Systems*, 2000.
8. Patrick Maillé, Bruno Tuffin: Pricing the internet with multibid auctions. *IEEE/ACM Trans. Netw.* 14(5): 992-1004 (2006)
9. Martin L. Puterman, *Markov Decision Processes*. Wiley, New Jersey 2005.
10. Peter Reichl and Stefan Wrzaczek, “Equilibrium Market Prices for Multi-Period Auctions of Internet Resources”. *MMB 2004*: 25-34
11. R. Srikant, *The Mathematics of Internet Congestion Control*, Birkhauser, 2004.