

Network bandwidth allocation via distributed auctions with time reservations

Pablo Belzarena*, Andrés Ferragut*[†], and Fernando Paganini [†]

*Facultad de Ingeniería, Universidad de la República

[†]Facultad de Ingeniería, Universidad ORT
Montevideo, Uruguay.

Abstract—This paper studies the problem of allocating network capacity through periodic auctions. Motivated primarily by a service overlay architecture, we impose the following conditions: fully distributed solutions over an arbitrary network topology, and the requirement that resources allocated in a given auction are reserved for the entire duration of the connection, not subject to future contention. Under these conditions, we study the problem of selling capacity to optimize revenue for the operator.

We first study optimal revenue for a single distributed auction in a general network, writing it as an integer program and studying its convex relaxation. Next, the periodic auctions case is considered for a single link, modeling the optimal revenue problem as a Markov Decision Process (MDP); we develop a sequence of receding horizon approximations to its solution. Combining the two approaches we formulate a receding horizon optimization of revenue over a general network topology, leading to a convex program that yields a distributed implementation. The proposal is demonstrated through simulations.

I. INTRODUCTION

The possibility of auctioning bandwidth in real time has been considered by many authors [12], [8], [14], [18], [6], [19], with a variety of applications: diffserv, access control, 3G cellular access, VPNs, etc. Much of this work has focused on game-theoretic considerations, in particular on providing incentives for bidders to reveal their true utilities. The standard theory of auctions [11] provides these mechanisms for the auctioning of a single resource, but it is far more challenging to extend them to a general network topology. Most proposals in this regard require the user (or a broker entity acting on his/her behalf), to place separate bids for internal resources of the network. In particular, the Progressive Second Price (PSP) mechanism of [12] requires each player to coordinate bids at the different nodes on its route, so that each node may run an auction with the allocation and pricing rules of the single resource case. PSP has a long convergence phase, which is improved by a *multibid* method in [14]; however, the latter mechanism only applies to tree topologies. Another approach to bandwidth auctioning for multicast trees or VPNs is proposed in [6], based on Dutch auctions. The mechanism assumes that users interested in a path would try to reserve bandwidth by placing bids simultaneously for all constituent links.

In this paper we are primarily motivated by the Service

Overlay Network (SON) architecture [9], proposed to facilitate the deployment of value-added Internet services with end-to-end quality of service (QoS). An overlay consists of service gateways located in the domain boundaries, and a set of leased tunnels from the underlying operators. Through this infrastructure, the overlay operator can sell for instance video-on-demand with high QoS from a distributed set of servers to clients in all these domains. We are interested in a decentralized and scalable auction mechanism in which each user interacts with the overlay network in the simplest way, with a single bid for the entire end-to-end service, oblivious to the internal topology. Furthermore, we focus on the objective of revenue maximization for the overlay operator who invests in the overlay infrastructure. This leads us to consider in Section II a first-price auction that can be formulated as an integer program, and studied through its convex relaxation, leading to distributed approximate solutions. In Section II-B we argue why this approach, which does not give incentives to truth-revealing bids, is natural for our objective.

Another important aspect of our problem that has not been satisfactorily addressed in previous work are inter-temporal considerations. Most references cover a one-shot auction where bids for the entire duration are known initially. References for multi-period auctions (e.g. [18]) allow future bidders to compete with incumbent ones, albeit given the latter some advantage. This is not an attractive condition for our intended applications. Consider for example 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. In this paper we 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. This means that the operator must assume the risk of future auctions. Optimizing revenue with this risk becomes then a stochastic dynamic optimization problem, that we formulate in Section III as a Markov decision process (MDP) [1], [17], for the single resource case. We introduce a receding horizon approximation that is able to capture the dynamic aspect of the problem in a tractable way, and validate it by simulation.

In Section IV we study multi-period auctions for the general network case, incorporating the reservation requirement. We find an extension of the optimization in Section II that incorporates a receding horizon term, and can lead to distributed computation. In Section V we discuss implementation issues and study the features of the proposed mechanism by simulation.

Conclusions are given in Section VI. A preliminary, abbreviated version of this work was presented in [2].

II. OPTIMAL BANDWIDTH ALLOCATION OVER A NETWORK

In this section we consider a set of users who bid for end-to-end bandwidth in fixed amounts, and the network must make a one-time decision as to how to allocate capacity among them to maximize its revenue. By considering this one-shot decision we postpone any temporal considerations; for the moment the focus is the network topology, and the requirement for a distributed resource allocation method.

We establish some notation. The network is composed of a set of links indexed by l , and a set of end-to-end routes indexed by r . R denotes the routing matrix, $R_{lr} = 1$ iff route r includes link l , otherwise $R_{lr} = 0$. $c = (c_l)$ is the vector of link capacities. Associated with each route r is a class of service defined by a fixed bandwidth σ_r : users bid for this well-defined rate allocation. There could be different classes of service offered over the same topological path; however in that case we use a different index r for each class, so the above formulation involves no loss of generality. For each r , the network receives a set of N_r bids $b_r^{(i)}$, ordered as

$$b_r^{(1)} \geq b_r^{(2)} \geq \dots \geq b_r^{(N_r)}.$$

The resource allocation decision is to find which of these bids to accept, within the capacity constraints of the network, to maximize revenue. We will assume a first-price auction, users will pay their bid; later on we discuss strategic implications. Defining the variable $\xi_{r,i}$ by $\xi_{r,i} = 1$ if bid $b_r^{(i)}$ is accepted, $\xi_{r,i} = 0$ otherwise, the optimal revenue problem is the integer program

$$\max \sum_r \sum_{i=1}^{N_r} b_r^{(i)} \xi_{r,i} \quad (1a)$$

$$\text{s.t.} \sum_r \sum_{i=1}^{N_r} R_{rl} \sigma_r \xi_{r,i} \leq c_l \quad \forall l, \quad (1b)$$

$$\xi_{r,i} \in \{0, 1\}. \quad (1c)$$

We will also use an alternative form. Noting that for fixed r , all bids $b_r^{(i)}$ are for the same amount of bandwidth, the optimal solution will involve the highest bids per route,

$$\sum_{i=1}^{N_r} b_r^{(i)} \xi_{r,i} = \sum_{i=1}^{m_r} b_r^{(i)},$$

where the integer variable m_r is the number of bids accepted in each route. Also denote by a_r the allocated rate in route r , $a_r = \sigma_r m_r$. Now define

$$U_{b_r}(a_r) := \sum_{i=1}^{a_r/\sigma_r} b_r^{(i)}. \quad (2)$$

This function is defined above for discrete values of a_r (the multiples of σ_r). It is convenient to extend it to a function of $a_r \in \mathbb{R}$, by linear interpolation. This piecewise linear function is increasing and concave in a_r , since bids are decreasing. With this notation, we can alternatively rewrite (1) as follows.

Problem 1 (Optimal instantaneous allocation):

$$\max \sum_r U_{b_r}(a_r) \quad (3a)$$

$$\text{s.t.} \sum_r R_{rl} a_r \leq c_l \quad \forall l, \quad (3b)$$

$$a_r/\sigma_r \in \mathbb{Z}. \quad (3c)$$

A. Convex relaxation and distributed solution

Let us ignore for the moment the integer constraint in (3c); the optimization in (3a-3b) has the same form as the *network utility maximization* problem in the congestion control literature [10], [13], [20], but now the utility represents the network revenue. Through the use of duality one can seek decentralized solutions to this convex relaxation. We summarize the method briefly.

Let $\alpha = (\alpha_l)$ be a vector of Lagrange multipliers (prices) associated with the constraints (3b), and let $q_r = \sum_l R_{rl} \alpha_l$ be the accumulated prices per route. Denote $y_l = \sum_r R_{rl} a_r$, $[\cdot]^+ = \max\{\cdot, 0\}$ and let $\gamma_l > 0$. Then the optimum of (3a-3b) can be found dynamically through the gradient projection algorithm

$$a_r := \arg \max_{a_r} [U_{b_r}(a_r) - q_r a_r], \quad (4a)$$

$$\alpha_l := [\alpha_l + \gamma_l (y_l - c_l)]^+. \quad (4b)$$

Here (4a) uses current route prices to fix a rate allocation with maximum ‘‘surplus’’ (utility minus a linear cost). (4b) compares the proposed allocation to link capacity and updates prices (up or down) accordingly. An equilibrium point of (4a-4b) is a saddle point of the Lagrangian

$$L(a, \alpha) = \sum_r U_{b_r}(a_r) + \alpha^T (c - Ra)$$

which corresponds to an optimizing a . In congestion control, the preceding equations are interpreted as describing the *data plane*, in which elastic sources adapt their packet rate and links generate prices based on their instantaneous congestion. They are known to asymptotically reach optimality for strictly concave utilities [20].

In our situation, we think of the above equations as an iteration in the *control plane*, which is run to settle an auction prior to any allocation of resources. In Section V we describe an implementation mechanism.

One issue is that U_{b_r} is not strictly concave, it is piecewise linear, changing slope at the multiples of σ_r . So (4a) might have multiple optima; we could use this freedom to select a solution that satisfies the integer constraint in (3b), but it is not obvious that the algorithm would converge with this choice. If it always did, we would conclude that the convex relaxation is exact. Unfortunately, this is not the case.

Example 1: Consider 4 links with capacity $c_l = 2$, and 5 paths (each with bandwidth requirement $\sigma_r = 1$), with routing matrix

$$R = \begin{bmatrix} 1 & 1 & 1 & 0 & 1 \\ 1 & 1 & 0 & 1 & 1 \\ 1 & 0 & 1 & 1 & 1 \\ 0 & 1 & 1 & 1 & 1 \end{bmatrix}.$$

Bids for the same route are all equal, with the following distribution among routes: $b_1 = b_2 = b_3 = b_4 = 1$, and $1 < b_5 < \frac{4}{3}$. Then, the relaxed convex program (3a-3b) has solution $a^* = (\frac{2}{3}, \frac{2}{3}, \frac{2}{3}, \frac{2}{3}, 0)^T$, with optimum revenue $U^* = \frac{8}{3}$. To see this, note first that a^* satisfies (3b) with equality. Now consider the prices $\alpha_l^* = \frac{1}{3}$, $l = 1, 2, 3, 4$, with aggregate route prices $q^* = (1, 1, 1, 1, \frac{4}{3})^T$. Since $b_5 < q_5^*$, we must have $a_5 = 0$, but the remaining coordinates are indeterminate in $[0, 1]$. So the proposed point (a^*, α^*) is a saddle, but this would not happen with integer coordinates in a .

In fact, here the integer program can be solved by observing that at most two connections can be active over all routes, so the best solution is to give them to the highest bidders in route five, $\tilde{a} = (0, 0, 0, 0, 2)^T$. This gives an optimal integer revenue $\tilde{U} = 2b_5 < \frac{8}{3}$. So the optimal relaxed solution is better than any integer solution. Moreover, the optimal integer solution is not obtained by roundoff of the relaxed solution, it is a radically different allocation.

The above example shows that optimal revenue is not an easy integer program, its convex relaxation is not exact. Since integer programming is NP hard, we have strong indication of a fundamental difficulty in this problem, not easy to overcome even allowing for centralized computation. In practice, this forces us to accept sub-optimal allocations.

Remark 1: Our problem shares similarities with an optimal resource allocation problem studied in [16], under opposite conditions: fixed input demand, minimization of a convex cost subject to integer constraints. Again, except for special cases this integer program does not allow for a convex relaxation.

B. Strategic and game considerations

A large focus of the auction literature has been strategic bidding, and the design of mechanisms in which bidding true-utilities is a dominant strategy. Vickrey's second-price auction [21], where the winning user is

charged the second highest bid, is of this kind. More generally, VCG mechanisms (for Vickrey-Clarke-Groves, see e.g. [5]) have built-in "incentive compatibility", a condition sought in many auction designs for networks [12], [14], [6]. In contrast, we have proposed a first-price auction with incoming bids, ignoring strategic considerations. Are we missing an important point? In this section we argue that the proposed auction is indeed natural for our intended application, and it would *not* be preferable to replace it by VCG-type allocation mechanism.

A first point is that VCG mechanisms are motivated by welfare economics: allocating resources to the set of users that achieves the highest utility. This does not necessarily produce more revenue for the seller: indeed, truth-revelation is subsidized by the seller, possibly at the expense of profit. A fundamental result of the theory of auctions, the *Revenue Equivalence Theorem*, see [5], states that under certain assumptions (mainly, risk neutrality of participants) all auctions have the same expected revenue for the seller. However, under other conditions (e.g., risk-averse buyers) first-price auctions are known to improve revenue [5].

To illustrate these issues for a one-shot auction as considered in this section, consider a simple example, auctioning a single link of capacity C , and assume in a specific auction there are fewer than C competitors (for a unit each). The generalized Vickrey auction would charge a price equal to the highest bid left out, in this case, zero, hence the network receives zero revenue. If, instead, we charge users what they bid, how would strategic bidders behave? *If they knew* that capacity is not scarce, the rational thing would be to submit a bid close to zero; this would confirm revenue equivalence. However, in practice they would not have this information, so there is no clear answer as to what is rational to bid, between zero and their true utility. If they wish to have a good likelihood of securing a circuit, they will be compelled to bid a non-negligible amount. So the seller is better off with a first-price auction.

A second consideration, already brought up in [15], is the *complexity* of truth revealing mechanisms when implemented over a network. Suppose we are allocating bandwidth as discussed above, and accepted that our objective is not revenue but social welfare. A VCG mechanism would be roughly as follows:

- 1) For the current bids, the optimal (maximum welfare) allocation is computed.
- 2) To compute the charge for a certain user: remove its bid, and re-compute the optimal allocation. The difference in welfare the user imposes on others through its presence determines the charge.

Now, each of the allocation problems is equivalent to the integer program discussed in the previous section, which we have seen is hard. Solving a number of such problems of the order of the number of users is not an attractive proposition by any means.

Furthermore, to put such a complex mechanism in place in response to strategic bidders implies that we believe they *themselves* are capable of such complex evaluations, otherwise they will not be able to “game” this system. Such “unbounded rationality” is highly questionable; and, the complexity grows even more when we consider inter-temporal issues. Also, the information requirements are clearly unreasonable: users would have to know the entire network, etc. So, we have serious doubts about the practical value of game theoretic studies at this scale, and will not pursue them further.

III. PERIODIC AUCTIONS FOR ONE LINK

In this section we include the time dimension in the allocation process. Auctions are held periodically, based on bids collected for a period of length T . Once allocated, resources are reserved for a service duration that typically exceeds T , and *reservations* are in place, so that future bids are not allowed to displace incumbent users. We seek an allocation policy that maximizes revenue over time. In this section we study the auction of a single link of capacity C , with a single class of service of bandwidth $\sigma = 1$.

The discrete time index k defines the auction at time kT , involving the ordered bids $b^{k,(i)}$. If a^k represents the admitted rate (in this case equal to the number of admitted connections m^k), the associated revenue is $U_{b^k}(a^k) := \sum_{i=1}^{a^k} b^{k,(i)}$, as in (2). As before, $U_{b^k}(\cdot)$ can be interpolated to define an increasing, piecewise linear and concave function over the real numbers. This extension can be achieved defining U constant for rates greater than a^k .

A. Optimal allocation as a Markov Decision Process

The long-term optimal revenue problem is posed in terms of a stochastic model for the bidding and duration processes. We assume bids are drawn from a certain continuous probability distribution, and service durations are modeled 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.

Remark 2: We assume in this section that the distribution of bids is known to the auctioneer. Also, the number of bids per auction is tentatively assumed fixed at $N \geq C$. In Section V we will consider learning the bids distribution from past observations, and the possibility of bids arriving as a random process.

Remark 3: The most questionable assumption is that of an exponential service duration. Duration is a characteristic of the service being auctioned, and could be considered deterministic. The only natural randomness is that a user might give up before the allotted time (e.g., the end of the movie). Clearly, an exponential distribution does not capture this well, but has been adopted to allow for a Markovian analysis. The duration parameter

μ is associated with the service being auctioned, thus in this section it is common to all users; see Section IV for generalizations.

Given the distribution of bids b , we define the expected revenue function $\bar{U}(a) = E[U_b(a)]$, where we replace the current bids in (2) by their expectation (the order statistics for the known distribution). This is also increasing, piecewise linear and concave.

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 :

$$P[x^{k+1} = i | x^k, a^k] = \binom{x^k + a^k}{i} p^i (1-p)^{x^k + a^k - i}. \quad (5)$$

Problem 2 (Optimal mean revenue, single link):

$$\text{Maximize } \lim_n \frac{1}{n} \sum_{k=0}^{n-1} E[U_{b^k}(a^k)].$$

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 (5). We can also consider the discounted version:

$$\text{Maximize } \sum_{k=0}^{\infty} \rho^k E[U_{b^k}(a^k)], \quad \text{where } 0 < \rho < 1.$$

Both are Markov Decision Processes (MDPs) [3], [17]. The *state* at time k is given by $s_k = (x^k, b^k)$, i.e. the current occupation and the incoming bids. Based on this state, the *action* $a^k = a(s_k)$ decides on how many bids to accept. A solution to the MDP is a *policy* $a(s)$ that results in a minimum cost. In the discounted case $\rho < 1$, this policy satisfies the Bellman equation

$$V^*(x^0, b) = \max_{a \in \mathcal{A}_s} \{U_b(a) + \rho E[V^*(x^1, b')]\}, \quad (6)$$

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 (6) is no longer the optimal cost, but (6) 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 [3].

B. Receding horizon approximation.

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

$$V_1(x^0, b) = \max_{a \leq C-x^0} U_b(a) = U_b(C-x^0).$$

This first step gives the “myopic” policy $a = C-x^0$, that sells all available capacity without regard to the future. In certain parametric scenarios this may be a good policy. To improve on it, we take a second step in the value iteration:

$$\begin{aligned} V_2(x^0, b) &= \max_{a \leq C-x^0} \{U_b(a) + \rho E[V_1(x^1, b')]\} \\ &= \max_{a \leq C-x^0} \{U_b(a) + \rho E[U_{b'}(C-x^1)]\} \\ &= \max_{a \leq C-x^0} \{U_b(a) + \rho E_{x^1} \bar{U}(C-x^1)\}. \end{aligned} \quad (7)$$

In (7), we have taken expectation with respect to the bid b' , using \bar{U} defined above; what remains is the expectation with respect to $x^1 \sim \text{Bin}(x^0 + a, p)$. The policy that solves (7) can be given a *receding horizon* interpretation: the decision optimizes over 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 applied recursively; thus the future is taken into account, but at a limited level of complexity.

The first term in (7) increases with a . To characterize the second, we rewrite it as follows. Consider the function $W(i) = \bar{U}(C) - \bar{U}(C-i)$, piecewise linear, increasing and *convex* in i . Indeed, the increments

$$w(i) := W(i+1) - W(i) = E[b^{(C-i)}], \quad i = 1, \dots, C$$

are non-negative and increasing in i (since bids are decreasing). We now study the expectation with respect to the binomial distribution.

Proposition 1: Define $\bar{W}(x) = E[W(I_x)]$, where $I_x \sim \text{Bin}(x, p)$ for integer x , and extend by linear interpolation. Then $\bar{W}(x)$ is increasing and convex.

Proof: Given $I_x \sim \text{Bin}(x, p)$, integer x , we can generate a $\text{Bin}(x+1, p)$ random variable of the form $I_x + \xi$, where ξ is Bernoulli(p), independent of I_x . Writing

$$W(I_x + \xi) - W(I_x) = w(I_x + 1)\xi$$

and taking expectations, using independence we obtain increments

$$\bar{w}(x+1) := \bar{W}(x+1) - \bar{W}(x) = pE[w(I_x+1)]. \quad (8)$$

It remains to show the last term is increasing in x . Noting that $w(i)$ is increasing, the inequality $w(I_x + \xi + 1) \geq w(I_x + 1)$ holds almost surely; taking expectations, $E[w(I_{x+1} + 1)] = E[w(I_x + \xi + 1)] \geq E[w(I_x + 1)]$. ■

The one-step ahead optimization (7) can now be rewritten as

$$\max_{a \leq C-x^0} U_b(a) - \rho \bar{W}(x^0 + a) + \rho \bar{U}(C). \quad (9)$$

Implicit in (7) and (9) is that a is an integer. In this case, however, the condition can be relaxed without loss of generality, treating (9) as a convex optimization problem. To solve it amounts to looking for a crossing point between the derivatives of $U_b(a)$ and $\rho \bar{W}(x^0 + a)$ (marginal utilities and costs), as depicted in Fig. 1.




Fig. 1. Marginal utility versus marginal cost

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, and have the form $\rho \bar{w}(i)$, with $\bar{w}(\cdot)$ defined in (8); as noted they are increasing in i . Since the bids b are random, the curves of Figure 1 will almost surely cross at a single, integer point. So the convex relaxation is innocuous.

The optimal acceptance policy is the value a such that $b^{(1)} \geq \dots \geq b^{(a)} \geq \rho \bar{w}(i) > b^{(a+1)}$, for $i = x^0 + a$. The values $\rho \bar{w}(i)$; act as successive *thresholds*: to accept a bids, the *lowest* one must exceed $\rho \bar{w}(x^0 + a)$. To accept one more, we require a *more demanding* threshold $\rho \bar{w}(x^0 + a + 1)$ on this (smaller) bid.

A concrete formula for the thresholds as a function of the bid distribution is given (see the Appendix) by

$$\bar{w}(i) = p \sum_{l=0}^{i-1} E(b^{(C-l)}) \binom{i-1}{l} p^l (1-p)^{i-1-l}. \quad (10)$$

Based on knowledge of ρ , p , and the distribution of bids, this expression could be calculated offline and used for carrying out auctions with the policy (7).

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 [7]. The left graph in Fig. 2 shows the acceptance thresholds for both policies: we see the infinite horizon threshold is more demanding. On the right we show the average utility obtained 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 second plot shows there is a clear loss in utility.

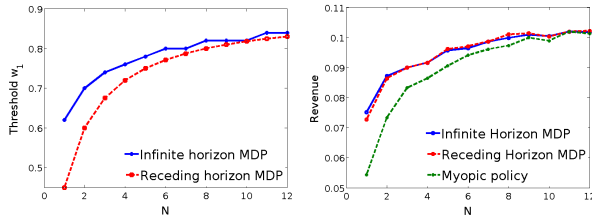


Fig. 2. Comparison between policies, $C = 1$, $p = 0.1$.

C. A fluid approximation

Our ultimate goal is to generalize this allocation policy to a general network topology as studied in Section II. The stochastic calculations involved in (10) appear difficult to generalize, so we adopt a second approximation, replacing the function $\overline{W}(x)$ in (9) by something easier to compute. Namely, define

$$\phi(x) = W(E[I_x]) \text{ for } I_x \sim \text{Bin}(x, p).$$

Since $W(\cdot)$ is convex, this underestimates the one-step cost from before, $\phi(x) \leq \overline{W}(x)$. 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

$$\phi(x) = W(px) = \overline{U}(C) - \overline{U}(C - px). \quad (11)$$

This is still piecewise linear and convex, but easier to compute. The second approximation to the optimal policy is given by

$$\max_{a \leq C - x^0} U_b(a) - \rho\phi(x^0 + a) + \rho\overline{U}(C). \quad (12)$$

Equivalently, by introducing a slack variable z we rewrite the above as the convex program

$$\begin{aligned} & \max U_b(a) + \rho\overline{U}(z), \\ & \text{s.t. } x^0 + a \leq C, \quad p(x^0 + a) + z \leq C. \end{aligned} \quad (13)$$

This is a fluid approximation of (7). At the optimum, the constraint in z is an equality, $z = C - p(x^0 + a)$; note that z (expected future allocation) need not be an integer. On the other hand, a should be an integer; one drawback of the fluid version (12) is that the breakpoints of $\phi(x)$ need not be integers; therefore its solutions might lead to fractional allocations that in practice will have to be approximated; we discuss this in the next section.

IV. PERIODIC AUCTIONS IN THE NETWORK CASE

In this section we consider the full problem of periodic auctions carried out over a general network, with the reservation requirement. In this context, the problem studied in Section II corresponds to the myopic policy of auctioning all bandwidth; we wish to incorporate the consideration of future revenue, generalizing the material of Section III. Given the complexity of this problem

we will only generalize the fluid approximation to the receding horizon policy.

We describe the allocation decision at time $k = 0$, and hence avoid inserting time indices in the bids and other variables. Define column vectors x^0 , a , and z , whose coordinates per route r denote respectively the rate x_r^0 from previous occupation, the rate allocation a_r at the current auction, and the expected rate allocation z_r in the following auction ($t = T$). Recall the definition (2) of the piecewise linear utility $U_{b_r}(a_r)$ based on current bids; analogously define $\overline{U}_r(z_r)$, replacing bids by their expectation. Both are in terms of σ_r , the bandwidth requirement of the class of service associated with r .

Another feature of the class of service is the model for duration: let p_r be the probability that a connection active at $t = 0$ on route r will remain active at $t = T$ ¹. Denoting by $P = \text{diag}(p_r)$ the corresponding diagonal matrix, the expected input rate vector at time $t = T^-$ will be given by $P(a + x^0)$.

Problem 3 (Network receding horizon allocation):

$$\max \sum_r U_{b_r}(a_r) + \rho\overline{U}_r(z_r), \quad (14a)$$

$$\text{s.t. } R(a + x^0) \leq c, \quad RP(a + x^0) + Rz \leq c, \quad (14b)$$

$$a_r / \sigma_r \in \mathbb{Z}. \quad (14c)$$

To get some insight into the above optimization it is useful to rewrite it in a similar manner to the single link case. For a vector x of source rates, define

$$\phi(x) := \max_{Rz \leq c} \sum_r \overline{U}_r(z_r) - \max_{Rz \leq c - RPx} \sum_r \overline{U}_r(z_r). \quad (15)$$

This function plays a similar role as the one defined in (11). For brevity we denote the first term (a constant) by \overline{U}^{\max} , it plays the role of $\overline{U}(C)$ in (11).²

Note that (15) is equivalent to a convex program in z , with x appearing linearly on the right-hand side of the constraints. It follows from ([4], Exc. 5.32) that $\phi(x)$ is convex. Furthermore, since the objective is piecewise linear it can be shown that $\phi(x)$ is piece-wise linear; we omit the details.

Through the function ϕ we can eliminate the variable z from Problem 3, reducing to the following optimization in a , together with the integer constraints (14c):

$$\max_{R(a+x^0) \leq c} \left[\sum_r U_r(a_r) - \rho\phi(a + x^0) + \rho\overline{U}^{\max} \right] \quad (16)$$

Remark 4: One property $\phi(x)$ does *not* have is separability over the components x_r of the vector x . Hence, the above optimization cannot be separated across different routes, it is inherently coupled.

Remark 5: Consider the situation where the constraints in (16) are inactive at the optimum. This means, with the current bids it is advantageous to leave spare

¹For an exponential duration, $p_r = e^{-\mu_r T}$.

²This term is not essential, it merely gives $\phi \geq 0$ and $\phi(0) = 0$.

capacity for the following auction, the predictive term is playing a non-trivial role. In that case, optimum can be found by comparing the marginal utilities (the bids) with the marginal costs, which take a finite number of values due to piecewise linearity of ϕ . Generically, then, these will cross at a single point, as in the scalar case. However, as remarked in that case, here as well the point need not satisfy the integer constraint in a .

The last remark implies that the convex relaxation (14) is not exact; still, we can use it to obtain an approximation. Consider the Lagrangian $L(a, z, \alpha, \beta)$ given by

$$\begin{aligned} L &= \left(\sum_r U_r(a_r) + \rho \bar{U}_r(z_r) \right) + \alpha^T (c - R(a + x^0)) \\ &\quad + \beta^T (c - Rz - RP(a + x^0)) \\ &= \sum_r [U_r(a_r) - (q_r + p_r v_r) a_r] + [\rho \bar{U}_r(z_r) - v_r z_r] \\ &\quad + \alpha^T (c - R x^0) + \beta^T (c - RP x^0) \end{aligned}$$

Here, α and β are the vectors of Lagrange multipliers (prices) for each of the two constraints, and we have defined the aggregate prices per route

$$q = R^T \alpha, \quad v = R^T \beta.$$

We can solve the convex program through a standard dual, gradient projection algorithm which in this case takes the form

$$a_r := \arg \max_{a_r} [U_{b_r}(a_r) - (q_r + p_r v_r) a_r]; \quad (17a)$$

$$z_r := \arg \max_{z_r} [\rho \bar{U}_r(z_r) - v_r z_r]; \quad (17b)$$

$$\alpha := [\alpha + \gamma (R(a + x^0) - c)]^+; \quad (17c)$$

$$\beta := [\beta + \gamma (RP(a + x^0) + Rz - c)]^+. \quad (17d)$$

The above algorithm is very similar to the one in Section II, and suitable for distributed implementation in the control plane. Although there are additional price and rate variables to pass, the message passing is fundamentally the same. (17a) amounts to comparing the bids with the threshold price $q_r + p_r v_r$; (17b) involves the expected bids, and the price v_r / ρ .

The computation still inherits some difficulties of Section II: imposing integer constraints on a_r / σ_r might not yield an equilibrium, but a suboptimal allocation can be found by rounding off a_r in the decreasing direction.

V. IMPLEMENTATION AND SIMULATIONS

Implementing the described allocation algorithm in a real network should be possible with variants of current network protocols. For instance, reservation and price signalling between network elements can be done with the RSVP protocol, as we now briefly describe.

First, user bids are received by the brokers, where each broker is associated with a service and a route from a

network access node to a server. These bids are collected until auction time.

The auction allocation is then performed following the decentralized design of (17), running in the network elements. The rate reservation variables (a_r, z_r) are sent by brokers in RSVP Path messages; prices are accumulated along a path with RSVP Resv messages in the reverse direction. This is iterated until convergence, defined through some tolerance, is reached, or alternatively after a maximum number of iterations. To guarantee feasibility of the final a_r , which might be compromised due to the gap with the integer program, these final values are sent in a last round of RSVP reservations.

An important implementation issue is that the mean user utility function may not be known to the broker. In that case, we use an adaptive method that estimates the function \bar{U} from past bids, through an exponential smoothing of the instantaneous utility function. Namely:

$$\bar{U}^{(k+1)}(z) = (1 - \alpha) \bar{U}^{(k)}(z) + \alpha U_{b^k}(z),$$

where $\bar{U}^{(k)}$ is the current estimate. Note that this requires updating only the values of \bar{U} at multiples of the circuit rate. Furthermore, the iteration applies even if the number of received bids is randomly varying in time.

This procedure allows the allocation mechanism to become independent of the bid distribution, and also of the arrival process. For instance, if bids arrive as a stationary random process (e.g. Poisson), \bar{U} is well defined, but difficult to write explicitly. However, the system can estimate it through smoothing. In Fig. 3 we show an estimation example. In this case, bids arrive as a Poisson process with intensity $\lambda = 10$ bids per auction, each bid having uniform distribution in $[0, 1]$. The averaging is taken over 100 auction periods with $\alpha = 0.05$. The real $\bar{U}(z)$ was calculated numerically.




Fig. 3. Estimation of the $\bar{U}(z)$ for a Poisson process

In order to evaluate the proposed algorithm, we implemented a discrete event simulator in JAVA which runs the allocation algorithm in a configurable network topology, with variable circuit demands, bid distributions and arrival processes. The simulator also implements the myopic policy and the average utility estimation presented above. We present results for three different scenarios.

A. Scenario 1: single link auctions.

We first compare the results of the receding horizon and myopic policies in a single link case, with 30 circuits. Auctions take place each T minutes, and bids arrive periodically with intensity λ bids/min (assumed fixed), totalling $N = \lambda T$ bids per auction.

Bids are assumed independent and uniformly distributed in $[0, 1]$, and rejected bids are discarded after each auction. Accepted jobs are assumed to stay in the system an exponentially distributed time with mean 100 minutes. Hence, T is a critical system parameter: enlarging T will allow more bids to participate in a given auction and circuits to be freed in between, but a very large T will decrease the auction rate, and therefore decrease the revenue per time unit.

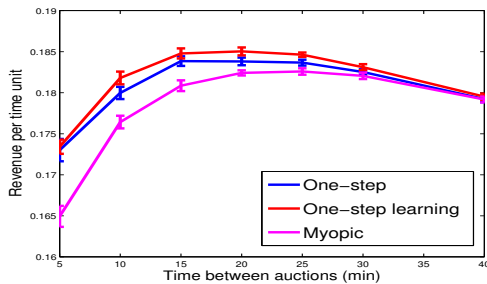


Fig. 4. One link situation: 30 circuits, bid arrival rate $\lambda = 0.5$

In Fig 4 we show the results for $\lambda = 0.5$. In this figure the myopic policy is compared with the one-step ahead policy implemented with the known bids distribution and with the learning version described above. We can see that both one-step ahead policies attain more revenue per time unit than the myopic policy, as expected.

B. Scenario 2: Linear network

We now simulate the linear network topology of Fig. 5. In this case, users in the long route 1 are expected to pay more in order to be allocated resources, since each of its circuits traverses 2 links. In order to emulate a real world situation, the bids arrive as a Poisson process of intensity λ and the learning one-step-ahead policy is used. In the first simulation, we compared the results




Fig. 5. Linear network with varying bids.

of this policy with the myopic one by varying the bid arrival rate λ in every link and keeping the time between auctions $T = 5$ min. We fixed the mean bid of route 1 to be twice of shorter routes. Results are shown in Fig. 6, where the average income per unit time is displayed.

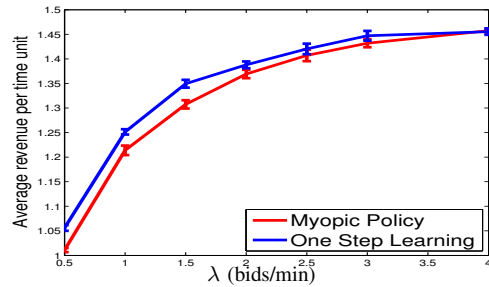


Fig. 6. Linear network with varying bid arrival rate.

As we can see, also in this case the one step ahead policy attains a significant gain over the myopic policy, for a wide range of arrival rates.

Our second experiment deals with varying the mean bid over the long route. In this case, $T = 5$ min. as before and $\lambda = 1$. We assumed independent and uniform bids with mean 1 for the short routes and varying mean for the long route. Results are shown in Table I. As we can see,

TABLE I
EFFECT OF VARYING THE MEAN BID IN THE ALLOCATION.

Avg. bid 1	R_1	R_2	R_3	a_1	a_2	a_3
0.5	0.003	0.643	0.428	0.3	47.5	28.4
1.0	0.051	0.627	0.401	2.8	45.0	25.9
1.5	0.252	0.558	0.312	9.5	38.4	19.2
2.0	0.583	0.469	0.191	17.3	30.7	11.5
2.5	0.918	0.397	0.115	22.1	25.9	6.7

- R_r : revenue per unit time generated by route r .
- a_r : mean allocated rate in route r .

when the mean bid of broker 1 is twice as much as the others, it gets a fair share of connections. Offering more will cause most resources of link 2 to be allocated to broker 1, with broker 2 retaining its share of 10 circuits, and broker 3 will starve.

C. Scenario 3: Overlay network.

In this final scenario, we tested the feasibility of our proposal in the more realistic situation depicted in Figure 7. In this case we have four interconnected servers and several brokers, each one attempting to secure resources of the overlay network. We have two types of demands: each connection in the short routes 1 and 3 consumes 2 circuits representing premium traffic, and the rest consume 1 circuit. The numbers over the links in Figure 7 indicate the number of available circuits. We assume that premium demand is less frequent (20%) but its mean bid is twice the bids of shorter routes.

The results are shown in Table II. We can see that the premium users who only use one link receive a substantial portion of the resources.




Fig. 7. Overlay Network Example.

TABLE II
SIMULATION RESULTS FOR SCENARIO 3

Broker	1	2	3	4	5	6
Links	1	1-2	2	2-3	3-4	4-1
R_r	0.111	0.081	0.115	0.204	0.420	0.211
a_r	30.8	4.5	30.6	11.9	25.3	11.8

- R_r : revenue per unit time generated by route r .
- a_r : mean allocated rate in route r .

VI. CONCLUSIONS

In this work we proposed a mechanism for allocating network capacity through periodic auctions. We formulated the problem of maximizing operator revenue under the following constraints: the solution must be fully distributed, the network has an arbitrary topology, and the resources allocated in a given auction are reserved for the entire duration of the connection. We found near-optimal policies that can be computed via convex optimization. By simulations we see that the distributed algorithm scales well in different network topologies, that the on-line estimation of the bid distribution leads to good approximations, and that the algorithm outperforms the myopic policy of selling all capacity in each auction.

In future work we will study a number of natural extensions: a different receding horizon approximation that might regularize the integer program, multiple-step extensions to the receding horizon, and the multi-path case where each broker competes by placing bids in multiple network paths.

APPENDIX

In this appendix we will get the expression for the threshold $\bar{w}(x)$. From (8) we have that $\bar{w}(x) = E(W(I_x)) - E(W(I_{x-1}))$ with $I_x \sim \text{Bin}(x, p)$. From the definition of W and \bar{U} we can rewrite this as

$$\begin{aligned} \bar{w}(x) &= E_{I_{x-1}} \left(\sum_{l=0}^{C-I_{x-1}} E(b^l) \right) - E_{I_x} \left(\sum_{l=0}^{C-I_x} E(b^l) \right) \\ &= \sum_{j=0}^{x-1} A(x-1, j) \sum_{l=1}^{C-j} E(b^l) - \sum_{j=0}^x A(x, j) \sum_{l=1}^{C-j} E(b^l), \end{aligned}$$

where $A(x, j) := \binom{x}{j} p^j (1-p)^{x-j}$. Operating we have

$$\begin{aligned} \bar{w}(x) &= \sum_{l=C-x+1}^C E(b^l) \left(\sum_{j=0}^{C-l} A(x-1, j) - \sum_{j=0}^{C-l} A(x, j) \right) \\ &= \sum_{l=0}^{x-1} E(b^{C-l}) \left(\sum_{j=0}^l A(x-1, j) - \sum_{j=0}^l A(x, j) \right). \end{aligned} \quad (18)$$

$$\text{Now, } \sum_{j=0}^l A(x-1, j) - \sum_{j=0}^l A(x, j) = pA(x-1, l). \quad (19)$$

can be established by induction. The value of the threshold in (10), follows then from (18) and (19).

Acknowledgment.

This work was partially supported by PDT-Uruguay, project S/C/OP/69/03.

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. Shwartz (Editors) Kluwer, pp. 488-536, 2001.
- [2] P. Belzarena, A. Ferragut, F. Paganini, "Auctions for resource allocation in overlay networks", *Proc. Net-Coop*, Paris, 2008.
- [3] D. Bertsekas, *Dynamic Programming : Deterministic and Stochastic Models*, Prentice-Hall, 1987.
- [4] S. Boyd, L. Vandenberghe, *Convex Optimization*, Cambridge University Press, 2004.
- [5] C. Courcoubetis, R. Weber, *Pricing Communication Networks Economics, Technology and Modelling*, Wiley 2003.
- [6] C. Courcoubetis, M. Dramitinos and G. Stamoulis, "An Auction Mechanism for allocating the bandwidth of networks to their users", *Computer Networks*, September 2007.
- [7] I. Chadès, M-J.Cros, F. Garcia, R. Sabbadin "Markov Decision Process Toolbox for MATLAB", <http://www.inra.fr/internet/Departements/MIA/T/MDPtoolbox/>
- [8] 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.
- [9] 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.
- [10] F. Kelly, A. Maulloo, and D. Tan. Rate control in communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 39:237-252, 1998.
- [11] P. Klemperer, *Auctions: Theory and Practice*, The Toulouse Lectures in Economics, Princeton University Press, 2004.
- [12] A. A. Lazar, N. Semret: Design and Analysis of the Progressive Second Price auction for Network Bandwidth Sharing. *Telecommunication Systems - special issue on Network Economics*, 2000.
- [13] S. H. Low and D. E. Lapsley, "Optimization flow control, I: basic algorithm and convergence", *IEEE/ACM Transactions on Networking*, vol.7, no.6, pp861-874, December 1999.
- [14] P. Maillé, B. Tuffin: Pricing the internet with multibid auctions. *IEEE/ACM Trans. on Networking*. 14(5): 992-1004 (2006)
- [15] P. Maillé, B. Tuffin, "Why VCG auctions can hardly be applied to the pricing of interdomain and ad-hoc networks", *Proc. NGI2007*, Trondheim, Norway.
- [16] A. Ozdaglar and D. Bertsekas, "Routing and Wavelength Assignment in Optical Networks", *IEEE/ACM Transactions on Networking*, Vol 11, No. 2, April 2003.
- [17] M. L. Puterman, *Markov Decision Processes*. Wiley, New Jersey 2005.
- [18] P. Reichl and S. Wrzaczek, "Equilibrium Market Prices for Multi-Period Auctions of Internet Resources". *MMB 2004*: 25-34.
- [19] J. Shu and P. Varaiya, "Pricing network services," in *Proc. IEEE INFOCOM*, 2003, vol. 2, pp. 1221-1230.
- [20] R. Srikant. *The Mathematics of Internet Congestion Control*. Birkhäuser, Boston, MA, 2004.
- [21] W. Vickrey, "Counterspeculation, auctions and competitive sealed tenders", *Journal of Finance* 16, 8-7, 1961.