

November 2004

Online bandwidth auctions for efficient allocation of bandwidth

Ying Peng Que

University of Wollongong, ypq01@uow.edu.au

D. Platt

University of Wollongong, dplatt@uow.edu.au

Follow this and additional works at: <https://ro.uow.edu.au/infopapers>



Part of the [Physical Sciences and Mathematics Commons](#)

Recommended Citation

Que, Ying Peng and Platt, D.: Online bandwidth auctions for efficient allocation of bandwidth 2004.
<https://ro.uow.edu.au/infopapers/160>

Online bandwidth auctions for efficient allocation of bandwidth

Abstract

This paper describes an on-line auction process for dynamic bandwidth allocation in communication networks. The auction is intended to be repeated at short (perhaps ten minute) intervals. The outcome of the auction gives both the bandwidth allocation, and the price to be paid for it. The network is assumed to be MPLS capable, and two variations are provided, one in which a single LSP is available between each source/destination pair, and another in which multiple LSPs are available. We show that the present algorithms are scalable up to about 150 nodes.

Keywords

bandwidth allocation, electronic commerce, multiprotocol label switching, telecommunication networks

Disciplines

Physical Sciences and Mathematics

Publication Details

This paper originally appeared as: Que, YP and Platt, D, Online bandwidth auctions for efficient allocation of bandwidth, Proceedings. 12th IEEE International Conference on Networks, 16-19 November 2004, vol 2, 630-634. Copyright IEEE 2004.

ONLINE BANDWIDTH AUCTIONS FOR EFFICIENT ALLOCATION OF BANDWIDTH

Y. P. Que and D. Platt
University of Wollongong
Wollongong, NSW, 2522, Australia

Abstract - This paper describes an on-line auction process for dynamic bandwidth allocation in communication networks. The auction is intended to be repeated at short (perhaps ten minute) intervals. The outcome of the auction gives both the bandwidth allocation, and the price to be paid for it. The network is assumed to be MPLS capable, and two variations are provided, one in which a single LSP is available between each source/destination pair, and another in which multiple LSPs are available. We show that the present algorithms are scalable up to about 150 nodes.

1. INTRODUCTION

Network providers sell bandwidth to ISPs or large companies between defined points within their network. When the demand for bandwidth is less than what is available, the ISPs simply pay the standard price. However, when the bandwidth demand is greater than what is available, the problem of bandwidth allocation arises. This problem, in one form or another, has been known for a number of years.

The attraction of using a pricing mechanism has been recognized by a number of researchers. In reference [1], the authors introduce a pricing scheme which rewards users who take less bandwidth than what has been allocated to their access pipe. In [2], the authors use a game theoretic approach to the allocation of bandwidth for elastic services in high-speed networks.

A system of dynamic bandwidth allocation is introduced in [3], however, the paper does not address the issue of the allocation of bandwidth when the users require more bandwidth than is available. Reference [4] does not actually price the bandwidth, however it does address the issue of dynamic bandwidth allocation.

Takahashi and Tanaka [5, 6] are the closest to an on-line dynamic bandwidth allocation using an

auction mechanism. However, their method appears to be overly computationally intensive, and they offer results only for an auction over a single link

What we present here is a bandwidth allocation scheme, which uses the actual rates paid for each route. We will offer results over networks up to hundreds of nodes. An ISP, which is prepared to pay more for a block of bandwidth, will get that bandwidth. The network provider will get the maximum income available from selling the resources.

In a public auction, the participants make their bids for an item, and they are also able to see what others are willing to pay for the same item. It is the intention of this paper to offer a similar method. The auction would be carried out automatically, and would be repeated at short intervals, say every ten minutes. After the auction has been completed, the bidders are guaranteed the bandwidth they are awarded, at the finishing price of the auction over the following ten minutes. We suggest that this is a short enough period for bandwidth demand to be tracked closely. They are also committed to paying for that bandwidth at the price which finished the auction.

2. MODEL OF SIMPLE AUCTION

Let us suppose that we are to auction the bandwidth across a single link. The total bandwidth available is fixed, but the bidders (ISPs) together ask for more than this. The auction would proceed with the auctioneer suggesting a price, and each bidder stating how much bandwidth they are willing to buy at that price. Presumably, as the price increases, the amount of bandwidth that each bidder will buy, must decrease. Thus, as the auction proceeds, the price increases, and the total amount of

bandwidth required decreases. Eventually, it will decrease to an amount just equal to the available bandwidth, and the auction stops at that point. Every bidder is awarded the amount of bandwidth for which they bid, at the final price.

As it has been described above, the auction is an iterated process, but it could be condensed to a single step if the bidders entered a graph of the bandwidth they would buy at the price they would pay. Figure 1 shows such a graph.




Figure 1. Bandwidth vs Price for One Bidder

The graph must have a zero or negative slope at all points. Other bidders would have similar graphs, and we show in Figure 2 the graphs for three bidders. The auction process reaches the point "A" on the graph, where the bandwidths all add to exactly equal the available bandwidth.




Figure 2. Auction for a Single Link

This requires each bidder to disclose, at least to the auctioneer, the bandwidth it would have bid for prices higher than the finishing price (ruling price). The auctioneer would only report to the bidders the finishing price, and the bandwidth they have been allocated. There is an issue around the confidentiality of each bidder's graph. It may be that the bidders are content to simply learn their own result, or they may prefer that the bandwidth allocated to all bidders be made

known to them all. The first option would make collusion among the ISPs difficult or impossible. The second option would allow them to check that all were treated equally.

2.1 MPLS Enabled Network

We will deal with an IP network, enabled for MPLS. Each ISP will require bandwidth between a number of nodes within the network. Congestion in the network must be controlled by the use of appropriate label switched paths (LSPs), and we will consider two ways of doing this in this paper. The first method consists in establishing a single LSP between every source/destination (s/d) pair of nodes. If parts of the network become congested, we can ease the congestion by setting up a different LSP for selected s/d pairs. We leave it to other functions of network management to achieve this. For this paper, it means that there is only one LSP between s/d pairs, and therefore only one LSP to be the subject of an auction. Of course, the auction of one LSP affects the bandwidth available to other LSPs.

Another method to deal with congestion in parts of a network is to allocate several LSPs permanently to each s/d pair. Then a LSP which passes through a congested part of the network is likely to receive relatively little of the total bandwidth allocated to the s/d pair. Most of the bandwidth would go to the LSPs which do not pass through heavily congested areas. This method makes the management of LSPs less of a problem, but increases the complexity of the auction algorithm.

2.2 Method 1 – Single LSP

An algorithm must take as its inputs the topology of the network, the available bandwidth in each link, and the bids of each ISP for each s/d pair. Each bid must be of the form shown in Figure 1, that is, non-increasing, and touching each axis. The first step is to aggregate all the bids from the ISPs for each LSP, as shown in the example of Figure 2. This gives us a combined bid for the LSP concerned.

These combined bids must then compete with the bids for other LSPs for the bandwidth in the links they use. The next step is to look at each link individually, and combine the bids of all LSPs using that link. We then find the ruling price for each link, in the same way as is

illustrated in Figure 2. The algorithm proceeds as below.

1. For each LSP, aggregate bandwidth from all bids. This will produce a curve such as the one in Figure 2. Create a list of all LSPs available for bidding.
2. Aggregate the bandwidth for each link, from the curves for each LSP using the link (similar to Figure 2). Find the ruling price for each link. This is the price where the aggregate curve for the link crosses the available bandwidth in that link.
3. Select the link with the highest ruling price.
4. Set this price for all LSPs using this link, and for all ISP bids for these LSPs. Allocate the corresponding bandwidth to these LSPs and ISPs.
5. Reduce the available bandwidth on all links used by these LSPs by the amounts determined in step 4, above, and remove these LSPs from the list.
6. If all LSPs have been removed from the list, stop. Otherwise, return to step 2.

2.3 Method 2 – Multiple LSPs

When multiple LSPs are used, the algorithm is more complicated. It turns out to be easier if we approach the problem starting from the highest prices, and work down, until all possible links are fully utilized. We also allocate the bandwidth to each LSP in such a way that the bandwidth still left available is the same in each of those LSPs. This mechanism ensures that traffic is directed away from those parts of the network where the traffic would be heaviest.

3 Test Procedures

A series of networks was generated, in software, using a random network generator. The multi-LSP system was run with three parallel paths. In order to give the multi-LSP system a chance to show its potential, the average number of links connecting each node was about five. The bandwidth of each link was set at 10 Gbps. LSPs were generated using Dijkstra's algorithm. As far as possible, the LSPs between a s/d pair were constrained not to use the same links. The collective action of the ISPs was simulated by producing a bandwidth/price curve for each s/d pair. This would represent the total bandwidth from all the ISPs at each price. That is, an aggregate curve is used for each s/d pair. The tests were carried out for both algorithms, and

for varying numbers of nodes, N . The complete test of network generation and auction algorithm was run five times for each value of N .

Two different scenarios were tested, neither of which could be satisfied by simply allocating the floor price. In one scenario, the demand was high for all s/d pairs, but uniform. In the other, the demand for 30% of the s/d pairs was significantly higher than for the rest.

To compare the effectiveness of each method, the total amount of revenue accepted was computed for similar network situations.

The average run time of the algorithm was recorded.

4 Results

Figures 3 and 4 show the ruling prices at the end of the auction process for a network with 20 nodes, using the single LSP system. The s/d pairs are graphed in the order in which the ruling prices were determined, which is why the prices decrease uniformly. In the case of the bids being even, shown in Figure 3, the bulk of the prices are close to the highest price. The last few s/d pairs determined did not need to go through the auction process, and the floor price was used. On the other hand, Figure 4 shows the results for uneven bids, and the first prices determined are considerably above the prices achieved for the bulk of s/d pairs. These results are as one would expect, and the result is similar for the multiple LSP system.




Figure 3. Ruling price for even bids, single LSP




Figure 4. Ruling prices for uneven bids, multiple LSPs

Figure 5 shows the revenue generated by both systems on different sized networks. At all points except $N=10$, the single LSP system generates more income. This is perhaps a surprising result, since the multi-LSP system seems to offer significantly more flexibility. When the bidding pattern was uneven, it would be the most likely that the multi-LSP system would have the best chance of producing the highest revenue. The results, shown in Figure 6, do not support this for networks above about 26 nodes.




Figure 5. Revenue generated with even bids




Figure 6. Revenue generated with uneven bids

This can be explained by the average number of links used in the LSPs when parallel paths have to be used.




Figure 7. Average number of links per LSP

The single LSP always uses the least number of links. Mostly, parallel paths need more links, so the multi-LSP system fills up the network capacity more quickly as the auction progresses.

The single LSP system, being simpler, would be expected to take less time to execute. The comparison is shown in Figure 8, in both cases, the program being run on a 2.7 GHz Pentium 4 processor. The graph for the single LSP case is extended beyond a 50 node network in Figure 9, which shows that for $N = 316$, the average execution time increases to 3×10^4 seconds, or 500 minutes. When $N = 150$, the execution time is about ten minutes.




Figure 8. Execution times




Figure 9. Execution times for single LSP

5 Comments and Conclusions

This paper has proposed a novel system to distribute bandwidth in a communications network among organizations such as ISPs who then on-sell it. The mechanism is based on the traditional idea of an auction. We believe that it offers the best possible bandwidth allocation in the sense that the network provider receives the best possible price for their resources, and bandwidth is allocated to the bidders on the basis of the money they are prepared to pay for it.

Two methods, involving a single LSP between two points, and multiple LSPs were investigated. The single LSP system performed somewhat better than the multi-LSP method on several measures, not least of which was program execution time.

If bandwidth is auctioned between all pairs of nodes, the execution time for the single LSP system takes about ten minutes when the network has 150 nodes. If this is to be used as a real time system, this is about as far as the present system can be taken. For networks up to this size, this is a system which is responsive to commercial reality, and guarantees the highest revenue to the network provider.

6 References

- [1] Costas Courcoubetis† and Vasilios A. Siris, "An Evaluation of Pricing Schemes that are based on Effective Usage", *1998 IEEE International Conference on Communications (ICC 1998)*, Vol.1, pp. 265-269, 7-11 June 1998.
- [2] Haikel Yaiche, Tavi R. Mazumdar, Catherine Rosenberg, "A Game Theoretic Framework for Bandwidth Allocation and Pricing in Broadband Networks", *IEEE/ACM TRANSACTIONS ON NETWORKING*, Vol.8, No.5, 667-678, Oct. 2000.
- [3] Adel S. Elmaghraby, Anup Kumar, Mehmed M Kantardzic, Mostafa Gamal Mostafa, "Bandwidth Allocation in a Dynamic Environment Using a Variable Pricing Policy", *In Proceedings of the Seventh International Symposium on Computers and Communications 2002 (ISCC 2002)*, pp. 589-594, 1-4 July 2002.
- [4] Ana Luiza B.P.B.DINIZ, Carlos C. GOULART, Jose Marcos S. NOGUEIRA, "Dynamic Bandwidth Allocation Service for ATM-based Networks supporting Multimedia Applications", *1998 first IEEE International Conference on ATM (ICATM 1998)*, pp. 292 -299, 22-24 June 1998.
- [5] Eiji Takahashi, Yoshiaki Tanaka, "Auction Method and its Performance in a Dynamic Bandwidth Allocation Service", *1st European Conference on Universal Multiservice Network (ECUMN 2000)*, pp. 328 -334, 2-4 Oct. 2000.
- [6] Eiji Takahashi, Yoshiaki Tanaka, "Auction-based Effective Bandwidth Allocation Mechanism", *10th International Conference on Telecommunications 2003 (ICT 2003)*, Vol. 2, pp.1046 - 1050, 23 Feb. - 1 Mar. 2003.