Demand management for telecommunications services

Grigorios Zachariadis, Javier A. Barria *

Intelligent Systems and Networks Group, Department of Electrical and Electronic Engineering, Imperial College London, London SW7 2AZ, United Kingdom

Abstract

In this paper we develop a novel method of controlling the demand in a multi-class, QoS-enabled network, using pricing and resource allocation for income maximisation. We first present a solution to the problem of calculating the optimal prices and QoS for a single link using a limiting regime approximation, which reduces the associated computational burden. A heuristic algorithm is then proposed that improves the limiting regime solution, achieving better results for links with small capacity. We further extend this approach to a multi-link network, where a distributed iterative algorithm is developed based on the solution of the single link model. Results from small and medium size networks show that, even when the assumptions we used do not hold, our approach yields results very close to the optimal ones (0.17–2.95% difference), which are computed by exhaustively searching in the decision space. Moreover, the calculation time using the proposed approach is approximately 1.5 min for problems which took more than 240 min to solve using exhaustive search.
© 2007 Elsevier B.V. All rights reserved.

Keywords: Revenue management; Pricing; Utility; Asymptotic approximation; Resource allocation; Quality of service

1. Introduction

In this paper we propose an approach to managing a telecommunication network by controlling the demand for services. In contrast to most existing approaches, which allow demand to depend either on prices or on quality, we examine the case where demand for a particular service depends on both these characteristics. We use this approach to solve the problem of maximising the income of a multi-class network provider using pricing and resource allocation.

In the scenario here investigated the network provider is allowed to offer a limited selection of alternative classes of service to users. Each potential user, according to his/her price and quality requirements, will select the most appropriate and request a connection. Calls belonging to a specific class of service have identical characteristics. Our model is similar to the differentiated services model [30] in that: (i) there is a limited number of classes, (ii) packets or flows of packets of the same class are treated in the same way, and (iii) state information is proportional to the

* Corresponding author. Tel.: +44 2075946275; fax: +44 2075946302.

E-mail addresses: grizac@imperial.ac.uk (G. Zachariadis), j.barria@imperial.ac.uk (J.A. Barria).

number of classes rather than the number of flows. In diffserv the classes' quality of service (QoS) differs in terms of delay, jitter, reliability and other parameters. Using a value similar to effective bandwidth, our model may be a good approximation of the diffserv model. The exact mapping of other QoS requirements to bandwidth can be the subject of further research. We also note that, by making decisions for sets of calls (classes) rather than individual calls, the number of decisions to be made is greatly reduced, improving the scalability of our solutions. The tradeoff is that the decisions are not necessarily optimal for each individual user, since not all the users belonging to a class are characterised by identical requirements.

We assume that, for a given price, an increase in the resources allocated for each call of a specific class leads to an increase of the QoS for this class. If this is the case, the number of potential users willing to subscribe to this class increases, since more users will find the offered quality acceptable. However, the extra resources which will now be used by the users which were already satisfied with the lower quality could have been used for other potential users which, even after the increase of QoS, are not satisfied with it. The unsatisfied users do not request connection and hence are lost from the system. For this reason, a proportion of the resources can be considered not optimally used. The exact amount of resources which are not optimally used in this sense, is the difference between the minimum resources needed to satisfy the total active users, and the actual resources used by them. It could be argued that these resources improve the utility of the users, and therefore may be considered optimally used, but from an income maximisation point of view, these resources do not contribute directly to the income of the network provider.

In this paper we develop methods for minimising the amount of resources not used optimally, as defined earlier. These methods depend on pricing, since the willingness to pay is here considered as being related to the quality expected by a user. Setting different prices for each service is a method for ensuring that users do not request admission for a service which offers much better QoS than they need, assuming that users of low requirements are willing to pay a lower price, and in this way it partially alleviates the problem of misused resources mentioned earlier. Therefore, the way price and QoS are determined for each class greatly affects the efficiency of a network.

We firstly explore the problem of finding optimal prices and qualities for the case of a single link, with the aim of maximising the income. We present a solution which is based on limiting regime approximation and then propose methods to improve this solution. We then develop a distributed iterative solution for the problem of a multi-link network, a solution which is based on the single link solution.

1.1. Related work

Most published research uses pricing to control or regulate one of two things: either (i) the incoming rate of new connection requests or (ii) the data rate sent by existing connections or users. In this paper we use pricing to control the incoming rate of new connections.

A method considering pricing for multiservice networks is presented in [10]. The QoS of each service is given and guaranteed, and the problem of pricing is formulated as an expected income optimisation problem. In contrast, in this paper, we take into account QoS as a decision variable, which affects demand when determining the optimal prices.

Dynamic pricing policies for network services have also been widely investigated, and there is a great variety of propositions. Agent based [11], time-of-day [6], usage-based [29], threshold-based [20,7] and congestion-dependent pricing [19,12] are some of them. The work of Paschalidis et al. [16,17] on dynamic, congestion-dependent pricing suggests that static (congestion-independent) pricing is asymptotically optimal, and therefore the benefits from dynamic pricing may not worth the extra complexity especially for large networks with many calls. A distributed mechanism for resource allocation, based on the pricing results of [16,17] is given in [18]. In contrast to the above approaches, our model allows the provider to make only static decisions.

Regarding static pricing, in the "Paris Metro Pricing" (PMP) scheme [15,21] two classes of services exist in the network, and each one has its own queue. Calls of the high class cost more than calls of the low class. In this way, fewer users request calls of the high class, and each call of the high class can thus use more resources, therefore receiving higher QoS. The PMP scheme exploits also the dependence of demand on prices. In this paper we extend the scheme proposed in PMP, as we also control and guarantee QoS.