Pricing QoS in Campus Networks

Nazli Mollah, Marcos S. Pinto
Melville University, City Tech, CUNY

Abstract

Campus Quality of Service (QoS) concerns the transformation of current campuses internets into systems that can offer one or more levels of QoS. Choosing the "best" QoS will require making decisions about the relative costs of network capacity and the machinery to manage it. Pricing is not only needed for recovering costs, but is also needed as a method of control. The congestion that has plagued the Internet, where pricing is based largely on a flat rate pricing, highlights the fact that without a more efficient pricing method it is difficult to control congestion and share network resources. This research uses the time-volume charging scheme, which determines that both the service time and transmitted volume are able to reflect the usage of a network. Service contracts for communications services do not specify fully the resources that are required to produce a unit of output, and therefore prices can be used both to regulate resource sharing and to maximize social efficiency.

1. Introduction

Even if we will have much more bandwidth in the future than now, the control of network resource utilization remains essential for the support of applications with special demands and for the prevention of (malicious or accidental) waste of bandwidth. Charging provides a possibility to control utilization and sharing of network resources. The research consists of presenting a refinement to the Time-Volume pricing model through the use of simulation network software, OPNET, where different scenarios with varying values of the parameters that compose the mathematical representation of the model produce results that will be analyzed for optimization (minimization) of the charges incurred in the usage of the network. The network model used in the simulation reflects the switching-based nature of existing campus networks.

2. Cost of Traffic in Campus Networks

Campuses across the country are recognizing the need to centrally plan and coordinate network services, standards and policies. Central planning and coordination are needed to provide greater functionality, security, and cost savings in the implementation of new technology across the network. Networks are increasing in cost to universities as more services and new technology is provided. The network pricing objectives can be listed as: to provide capital funds for adding capacity to the network, to recover transfer costs for external network connectivity, especially if those costs correlate to individual usage, to moderate demand for the enterprise network, or the external connections, on the grounds that free goods tend to be used inefficiently, and to provide budget predictability, which might tend to favor flat-rate quota-based schemes rather than usage-based billing. The best effort service that dominates current communications on IP networks cannot cope with the increasing demands of services with different Quality of Service (QoS) requirements. Hence, the Integrated Service model (IntServ) [1] and the Differentiated Services model (DiffServ) [2] have been proposed to provide different classes of service with different QoS. The IntServ model provides resource reservations and service guarantees. However, its implementation is complicated and is not scalable. In this project, we will focus on the DiffServ model because of its favorable features, such as simplicity and scalability. This research will use time-volume pricing scheme interested on the architectural issues such as where charges are computed and how multicast sessions and receivers can be charged, and not on how to compute usage charges.

3. Time-Volume Pricing Model

In this research the time-volume model is applied to traffic with a specific class type, i.e., traffic with specific service requirement that is to be satisfied by the service provider via a contract, a service level agreement (SLA), in which the customer must send at no more than a maximum rate *h* (the a priori information). Assume that a connection uses this contract and sends data at a mean rate *m* (the a posteriori information). It can be shown that given *h*, *m*, then the traffic with the greatest effective bandwidth is one that is periodically on and off, and has on and off phases of long duration. The effective bandwidth typically lies between the average rate and the peak rate and is given by

$$\alpha_{on-off}(s,t) = \frac{1}{st}\log\left[1+\frac{m}{h}(s^{sth}-1)\right] \qquad (1)$$

s and *t* are defined by the operating point of the multiplexer. We can see that $\alpha_{\text{on-off}}$ is a function of *m* which is unknown to the network when the contract is established.
We define a family of tariff lines $f_m(M)$, parameterized by the parameter *m*, each of which takes the form:

$$f_m(M) = a(m) + b(m)M \qquad (2)$$

which as a function of *M* lies above the curve $\alpha_{\text{on-off}}$ (M) and is tangent to it at m = M. The user chooses a tariff, or equivalently *states a value of m*. The final charge *f(m)* is given by:

$$f(m) = T\,[a(m) + b(m)M] \qquad (3)$$

where *M* is the measured mean rate of the user's traffic during *T* - the connection duration.
Equivalently, the charge c(m) is:

$$c(m) = a(m)T + b(m)V \qquad (4)$$

where

$$V = TM \qquad (5)$$

which is the volume of traffic carried (measured in cells or bytes).

Figure 1 shows the graphical relationship between effective bandwidth $\alpha_{on\text{-}off}$ and the mean rate M. The tangents fm(M) and fm'(M) to the curve are the tariffs. The effective bandwidth is plotted against the mean rate *M* for a fixed peak rate *h*. The user is free to choose any tangent to this curve, and is then charged *a(m)* per unit time and *b(m)* per unit volume. He minimizes his average charge rate to 2 unit of money by selecting his sending data rate equal to 1 and being equal to its expected value, i.e. $m = E[M] = 1$. If he chooses the tariff indexed by some other value, say *m'*, the average charge will be greater, here $E[f_{m'}(M)] = 2.4$

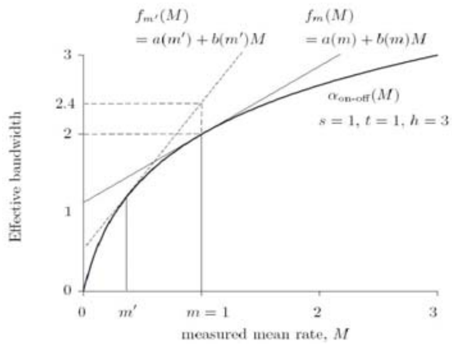

Fig 1. Implicit pricing of an effective

Equations (4) and (5), although represent the simplest usage-based scheme, are firmly based in the theory of effective bandwidths. This scheme is based only on time and volume. To optimize (4) we assumed that each user utilizes a set of applications for each of which: (i) there is a set of permissible QoS classes, (ii) an estimate for the anticipated mean rate per QoS class, and (iii) the particular QoS class to serve the application each time is specified in the SLA. The tariff coefficients *a*, *b* depend on the parameters *h*, *s*, *t* of the traffic contract, SLA, and are computed using sophisticated techniques such as effective bandwidths and the link operating point parameters *s*, *t*. The coefficient are, though, hidden from the user. We, however, will restrict the analysis to simple contracts that specify only the peak rate h of the source, and thus we will be able to calculate values for *a* and *b*. In practice, a constant coefficient d is added to the tariff. The tariff takes the form aT + bV + d, where d is chosen to discourage traffic splitting. We could look at the issue of providing incentives to avoid this illegal activity. However, we will limit our research to the cases where "shadow price" is not included in the mathematical formula in order to enforce traffic contract. Such approach can be viewed as a way to force discipline into the market of service providing of telecommunications.

4. The Simulation

In our simulation scenario, the first congestion point a packet leaving a desktop system is likely to encounter would be the subnet router interface. This scenario is attractive because it suggests that most of the QoS complexity in an enterprise network can safely be relegated to routers, and the edge switches don't need to be very smart, meaning that they can be cheap (i.e. cheaper than they otherwise would be.) The figure 2 below shows the configuration that will be used for the simulation.

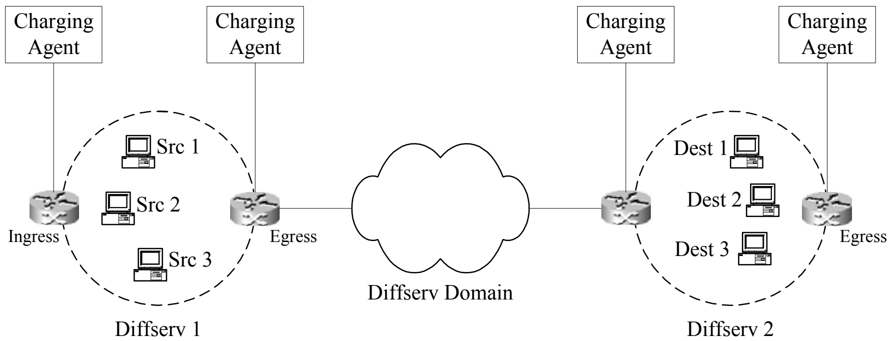

Fig 2. Simulation Configuration

The project does not have mixed domains of QoS, i.e., DS-enabled domains and non-DS-enabled domains, in the internetworks. Thus, the simulation will only contain DS-enabled networks. The simulation consists of several 5 minutes runs of typical campus network traffic originating from web browsing, ftp, email, and database queries.

We compiled values for the pricing after each run by varying the traffic coefficients and the mean rate. The goal is to list the combinations of (*h, a, b*) that would result in the lowest values of the charges. We simulated for different rates (h) and links with different bandwidths. The routers were the agents to collect traffic time and volume. Each source's application generated a specified amount (volume) of traffic and its departure time was recorded. At the receiving end, the time the traffic arrived at the router and the amount of traffic were also recorded; then, the charge was computed using these values of time and volume. We assume that the user has selected the traffic he will send, hence his mean rate. We also applied leaky bucket algorithm to the wide area network that interconnect the two DS domains. The ingress and egress interfaces (routers) carry leaky buckets algorithm; the simulator software provides such facility. Fig. 3 shows the charges for high and medium classes of traffic for the simulation run. As it can be observed, the charges are lower when leaky bucket is included since without it the dropped packets are not accounted for. This provides a better fairness to the users in this simulation case.

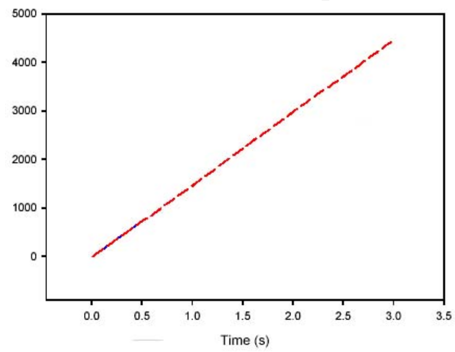

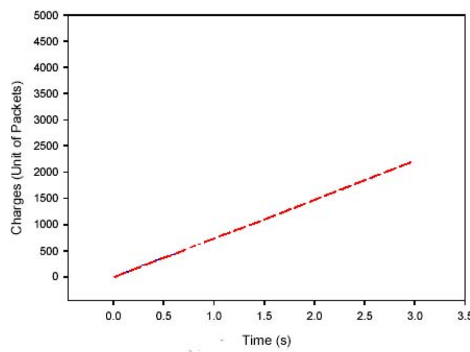

Figure 3. (a) High Priority Traffic (b) Medium Priority Traffic

5. Conclusion

Campus networks of the future are going to provide a diverse group of services to a diverse group of users with different service requirements. An array of new technologies such as global video-conferences, video/audio on-demand, virtual classrooms, priority messaging, maintenance of knowledge bases, and so on, will require specialized networks. The only possible way to realize such a network is by allowing a multi-service class network with possibly diverse data transmission protocols to cater to different service requirements and/or specifications. Networks are increasing in cost to universities as more services and new technology is provided. Pricing becomes an important issue and campus administrators request efficient and fair methods to charge for the usage of the network services. Charging for campus network access means a big cultural change and new policies need to be in place. In most campus, users are not used to think about costs every time they use the network. It is important to note that this research is quite general and can be used to charge for effective usage at many levels of network access, ranging from individual users to large organizations. It can be applied to any packet switching technology and can be used under both deterministic and statistical multiplexing. The extension to our approach to networks consisting of more than one link raises several further issues. Important choices concern whether a user sees a single charge from its immediate service provider, or whether a user might see several charges arising from various intermediate networks. The simulations showed that it is possible to gain both network efficiency and economic efficiency by using pricing scheme. Our pricing scheme, time-volume pricing, can be classified as a value pricing strategy which means offering the target market a high-quality product at a fair price and with good service.

6. References

1 F. P. Kelly, "Tariffs and Effective Bandwidths in Multiservice Networks", Proceedings of the ITC '94

2 C. Courcoubetis, R. Weber, "Pricing Communication Networks", John Wiley, 2003

3 MacKie-Mason, J.K., and H. Varian, "Pricing the Internet." *Public Access to the Internet*, Prentice-Hall, 1994

4 Market, Economics & Finance Unit, "IP-based Networks: Pricing of telecommunications services", Telecommunication Development Bureau, ITU, 2002

5 Altmann, J. H., H. Oliver, and A. Sanchez-Beato Suarez, "How to market-manage a QoS network", IEEE InfoCom 2002, New York.

Biographies

Nazli Mollah
Millersville University, PA
Degree History:
[BSc: Chemical Engineering, Clarkson University, Potsdam, NY] [MBA: Media Management and IT, Metropolitan College of New York, NY] [PhD: Computer Science, City University of New York, NY]
Areas of Interest:
Networks Security & Forensics, Intelligent Networks, ISPs, Network Feasibility Engineering, E-Commerce, Decision Support Systems, Database Management & Intelligent Information Retrieval, Wireless Security, Manchester United Football Club

Marcos Pinto
Assistant Professor, CityTech, CUNY
Degree History:
BS: Physics, Univ. Fed. Rio de Janeiro, Brasil; BA: Computer Science, Hunter College; MBA: Finance, Baruch College; PhD: Computer Science, Graduate Center, CUNY.
Areas of Interest:
Networks Security, Intelligent Network, Human Computer Interaction, etc.