Equilibrium Allocation of Variable Resources for Elastic Traffics

Steven H. Low
Dept. of Electrical & Electronic Engineering
University of Melbourne, Australia
slow@ee.mu.oz.au

Abstract— **Consider a set of connections sharing a network node under an allocation scheme that provides each connection with a fixed minimum and a random extra amount of bandwidth and buffer. Allocations and prices are adjusted to adapt to resource availability and user demands. We consider two scenarios of user behavior. In the first scenario a connection purchases an allocation to maximize its expected utility in such a way that the resource cost of the new allocation, and hence its connection charge, remains the same as that for the old allocation. In the second scenario this budget constraint is relaxed and a connection tries to maximize its benefit, expected utility minus the resource cost. Equilibrium is achieved when all connections achieve their optimality and demand equals supply for non-free resources. We show that at equilibrium expected return on purchasing variable resources can be higher than that on fixed resources. Thus connections must balance the marginal increase in utility due to higher return on variable resources and the marginal decrease in utility due to their variability. We further show that in equilibrium where such tradeoff is optimized *all* connections hold strictly positive amounts of variable bandwidth and buffer.**

I. Introduction

In order to guarantee quality of service (QoS), one popular approach is to reserve resources for connections. Reservation may take the form of a fixed amount of dedicated bandwidth, as in a circuit switched network, or it may also include a variable component, as in the available bit rate (ABR) service of an ATM network where a connection can receive a minimum cell rate (MCR) plus some random excess bandwidth. There have been a large number of proposed scheduling policies under which a connection is guaranteed a minimum share of resources and gets random extra amounts depending on network condition; some recent examples include, e.g., generalized process sharing of [26], self-clock fair queueing of [11], the GMUL (guaranteed minimum and upper limit) policy of [3], bandwidth allocation scheme of [1], and the virtual partitioning policy of [22], etc. For elastic traffics [29] that can tolerate some degree of delay or loss buffer is also a scarce resource to be traded-off in network resource allocation, e.g., [19], [10], [27], [8]. Again buffer allocation can be implemented by schemes ranging from complete partitioning, in which all connections are guaranteed fixed amounts of buffer, to complete sharing, in which no connection is guaranteed any fixed amounts of buffer [13], [4], [2]. In this paper we describe a model for such bandwidth and buffer allocation schemes and study the equilibrium allocation that would result when a large number of connections interact under such schemes. Equilibrium allocations are appealing because, when utility functions are concave, as we assume here, equilibrium allocations are Pareto optimum, i.e., it is impossible by some reallocation to increase a connection's utility without at the same time reducing the utility of other connections.

The problem is also motivated by resource renegotiation in future networks. Resource reservation is typically made during connection setup. It is difficult however to predict how much resource a connection needs throughout its lifetime. This may be due to multiple time-scale sources such as a movie trace, due to changes in application QoS requirements (e.g., changing from lecturing to question–and–answer session in a teleconferencing application), or due to changes in network conditions. A promising solution is to assign a tentative allocation to a connection during connection setup

The author would like to thank Albert Greenberg of AT&T Laboratories for his hospitality where part of this work was done and to thank Australian Research Council for grant S49813050.

and allow it to renegotiate its allocation from time to time. For example an algorithm is proposed in [12] to optimally renegotiate fixed bandwidth allocation when a video source undergoes slow timescale changes.

Specifically consider an ATM node that offers two types of resources, bandwidth and buffer, in two flavors, guaranteed or fixed resources and variable resources. The amounts of variable resources fluctuate randomly. They may represent unused bandwidth and buffer allocated to other traffic classes that are not under our consideration. The network sets prices for these resources, which are adjusted as their availability or demand changes.

The node is shared by a set of connections. In each period a connection decides how much fixed and variable bandwidth and buffer to purchase in order to maximize its expected utility in that period. In [20] we presented a model to study the equilibrium allocations where the allocations are not restricted to be nonnegative. It is motivated by a possible future scenario of the deregulated telecommunications environment where a large number of resellers may purchase network services (in terms of reserved resources) and retail them to end users. Some of these resellers may themselves own resources and can therefore either purchase services from the network or sell services to other resellers, depending for example on the current prices. In this paper we extend that work by studying two models, both requiring that the allocations be nonnegative.

In the first model a connection's objective is to maximize its expected utility subject to the constraint that the *value* of its new allocation remains the same as its current allocation at current prices, i.e., it finances its new allocation by selling its current allocation. The constraint is motivated by some proposed usage-based payment schemes, e.g., [9],[30, Chapter 8.2], where the connection charge is negotiated during connection setup. This negotiation involves considerable overhead to obtain payment authorization from the user, and therefore may not be reasonable to invoke at each renegotiation instant. With the constraint, however, since the total resource cost, and hence the connection charge, remains unchanged user authorization can be bypassed[1].

The first model excludes the possibility that a connection may pay more to acquire more resources or save in connection charge by requesting less resources[2]. In the second model we remove this constraint and there, a connection's objective is to maximize its expected benefit, i.e., expected utility minus the resource cost.

For the first model we show that, in equilibrium, the "total expected return" in purchasing variable resources can be higher than that in purchasing fixed resources, though variable resources cause a reduction in utility due to their variability. Hence a connection must balance the marginal increase in utility due to higher return of variable resources with the marginal decrease in utility due to the fluctuation in their availability. Such tradeoffs are optimized at equilibrium. It is then interesting that in equilibrium *every* connection holds *strictly positive* shares of variable bandwidth and buffer. This may mean, e.g, for connections whose utility functions satisfy certain conditions, ABR or VBR (variable bit rate) are more desirable than CBR (constant bit rate) services. Finally we show that this interesting property of positive variable resources at equilibrium is also present in the second model where the connections are not restricted to keep their resource costs unchanged.

Pricing network services and resources have received much attention recently; see e.g., [5], [18], [24], [21], [15], [6], [16][30, Chapter 8] and references therein. This paper differs from the previous work in two respects. First we consider *both* bandwidth and buffer in our model. Second, and more importantly, we consider both variable and fixed resources where connections must balance the higher expected return of variable resources and their variability in making resource decisions. A different but interesting pricing model for ABR is studied in [7] where each connection n requests a fixed bandwidth allocation and the network allocates a random extra amount to connection n based on the total available (random) bandwidth capacity and total users' requests.

The paper is structured as follows. In §II we describe and motivate our model and assumptions. We present our main result on equilibrium alloca-

[1]Alternatively we may have the network determine resource allocations in each period with unchanged resource costs, without the involvement of users.

[2]We believe that such changes involve more overhead and should occur infrequently.

tions for the first model in §III, and for the second model in §IV. We conclude in §V with future work. Due to space limitation, all proofs are omitted and can be found in technical reports available from the author.

II. Model and Motivation

We first describe in subsection II-A our model and assumption precisely, and then explain in subsection II-B the motivation behind the model.

A. Model

Consider a set of connections $1, \ldots, N$ sharing an ATM node. Time is slotted. In each period each connection n chooses a bandwidth and buffer allocation. A bandwidth allocation is specified by a pair $x_n = (x_{n0}, x_{n1}) \geq 0$, with the interpretation that a fixed amount $x_{n0}R_0$ and a random amount $x_{n1}R_1$ of bandwidth will be made available for connection n's exclusive use in that period. Similarly a buffer allocation is specified by a pair $y_n = (y_{n0}, y_{n1}) \geq 0$, with the interpretation that n will receive a fixed amount $y_{n0}B_0$ and a random amount $y_{n1}B_1$ of buffer in that period. Here, $R_0 > 0$, $B_0 > 0$ are real numbers and $R_1 \geq 0$, $B_1 \geq 0$ (almost surely) are nonnegative random variables. R_1 and B_1 are generally statistically dependent. R_i (respectively B_i) represents a unit of fixed or variable bandwidth (respectively buffer) that can be sold or purchased.

Fixed bandwidth is priced at p_0 per R_0 amount and variable bandwidth is priced at p_1 per R_1 (random) amount. Fixed buffer is priced at q_0 per B_0 amount and variable buffer is priced at q_1 per B_1 amount. Connection n pays $px_n + qy_n = p_0x_{n0} + p_1x_{n1} + q_0y_{n0} + q_1y_{n1}$ for allocation (x_n, y_n), where $p = (p_0, p_1)$ and $q = (q_0, q_1)$ are price vectors. Hence a connection is priced not on the actual amounts of resources that are made available to it, but by its reservation.

In this paper we only consider a single period. At the beginning of the update period connection n holds an allocation $\bar{x}_n = (\bar{x}_{n0}, \bar{x}_{n1}) > 0$ and $\bar{y}_n = (\bar{y}_{n0}, \bar{y}_{n1}) > 0$ that it inherits from the previous period. Connection n achieves a utility of $U_n(r_n, b_n)$ when it receives $r_n = x_nR$ amount of bandwidth and $b_n = y_nB$ amount of buffer, where $x_nR = x_{n0}R_0 + x_{n1}R_1$ and $y_nB = y_{n0}B_0 + y_{n1}B_1$. As motivated in §I we consider two models with different user behavior.

Model M1:
Connection n's objective is to

$$\max_{x_n, y_n \geq 0} \quad EU_n(x_nR, y_nB)$$
$$\text{subject to} \quad px_n + qy_n = p\bar{x}_n + q\bar{y}_n.$$

i.e., it buys a new allocation (x_n, y_n) to maximize expected utility and finances it by selling its current allocation $(\bar{x}_n, \bar{y}_n)$.

Model M2:
Connection n's objective is to:

$$\max_{x_n, y_n \geq 0} \quad EU_n(x_nR, y_nB) - (px_n + qy_n)$$

i.e., it buys a new allocation to maximize its expected benefit.

We make the following assumptions:

A1: $ER_1 > 0$, $EB_1 > 0$, $\text{cov}(R_1, B_1) \neq 0$ (i.e., there exists no constant a such that $R_1 = aB_1$ almost surely).

A2: $U_n(r, b)$ is concave and strictly increasing in its arguments.

A3: $EU_n(r, b) = u_n(E(r+b), var(r+b))$ where u_n satisfies

$$\frac{\partial u_n}{\partial \mu}(\mu, v) > 0 \quad \text{and} \quad \frac{\partial u_n}{\partial v}(\mu, v) < 0.$$

A4: For all μ, v,

$$\frac{\frac{\partial u_n}{\partial \mu}(\mu, v)}{-\frac{\partial u_n}{\partial v}(\mu, v)} < 2\frac{v}{\mu}.$$

A5: For Model M1 the current allocations are strictly positive, i.e., $\bar{x}_{ni} > 0$ and $\bar{y}_{ni} > 0$ for $i = 0, 1$, $n = 1, \ldots, N$.

By **A3** we assume that connection n's preference depends on the resources only through their first two moments. Moreover increasing the mean allocation increases connection n's utility, while increasing the variance of the allocation decreases its utility. This is motivated by the fact that applications' performance typically improves as they get more network resources and degrades as the availability of these resources fluctuates.[3]

[3]The value of bandwidth and buffer to an application may be different. This effect can be modeled by including in the utility function appropriate weights α, β: instead of **A3**, we can assume that $EU_n(r, b) = u_n(E(\alpha r + \beta b), var(\alpha r + \beta b))$. All results in the paper generalizes in a straightforward manner to the case of non-unity weights.

Assumption **A4** says that the ratio of marginal increase in utility due to increased mean allocation to that due to reduced variance cannot be too high (less than two times the ratio of variance to mean). It imposes a certain structure on the utility function. As an example an expected utility that satisfies **A3** and **A4** is

$$u_n(\mu, v) = \tau_{n1} \log \mu - \frac{v}{\tau_{n2}}$$

such that $\tau_{n1}\tau_{n2} < \min\{\bar{x}_{n1}^2 \sigma_R^2, \ \bar{y}_{n1}^2 \sigma_B^2\}$. This is shown in figure 1 for $\tau_{n1} = 100$ and $\tau_{n2} = 0.5$.

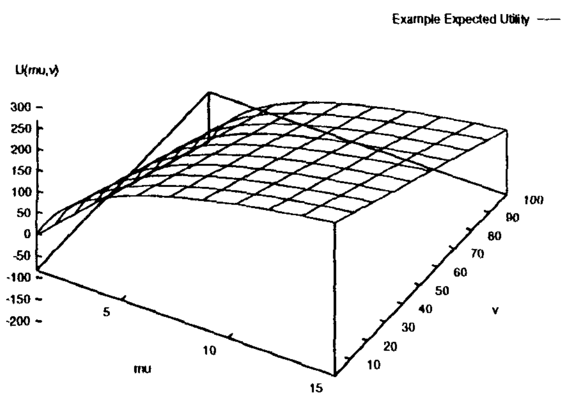

Fig. 1. Example expected utility $u_n(\mu, v)$.

We are concerned with equilibrium situations. Let $X_i = \sum_n x_{ni}$ (respectively, $\bar{X}_i = \sum_n \bar{x}_{ni}$) and $Y_i = \sum_n y_{ni}$ (respectively, $\bar{Y}_i = \sum_n \bar{y}_{ni}$) denote the total demand for (respectively, supply of) bandwidth and buffer i, respectively.

Definition 1: A set of allocation vectors $(x_n, y_n, n = 1, \ldots, N) \geq 0$ and a price vector $(p, q) \geq 0$ form a *competitive equilibrium* (or just *equilibrium*) if

1. (x_n, y_n) solves the maximization problem for all connection n, and
2. for $i = 0, 1$, $X_i \leq \bar{X}_i$, $Y_i \leq \bar{Y}_i$, and $p_i(X_i - \bar{X}_i) = 0$, $q_i(Y_i - \bar{Y}_i) = 0$.

We call (p, q) an *equilibrium price* and $(x_n, y_n, n = 1, \ldots, N)$ an *equilibrium allocation*.

Hence at equilibrium all connections n achieve their optimality and for all *non-free* resources, demand (X_i, Y_i) equals supply $(\bar{X}_i, \bar{Y}_i)$.

B. Motivation

The allocation model seems to fit a large class of proposed resource sharing schemes under which a connection is guaranteed a minimum share of resources and gets random extra amounts depending on the network condition, see e.g., [26], [11], [3], [22], [2] and references therein for recent examples. Here we may think of the random variables R_1 and B_1 as the total amounts of bandwidth and buffer (scaled by $\bar{X}_i$ and $\bar{Y}_i$) that will become available in that period, and the subscription x_{n1} and y_{n1} as the share reserved for connection n.

The connections in our model may represent a set of ABR connections or a set of VBR connections, or a set of VP connections. If the model is for ABR connections, for instance, then the fixed bandwidth allocation x_{n0} may represent the MCR. The variable bandwidth represents the random extra amount a connection gets in excess of its MCR. We allow connections to purchase bigger or smaller shares of the excess bandwidth, by choosing different x_{n1} values, according as how they value the variable bandwidth.

For another example, if the connections under consideration are VPs, a common proposal is to partition the bandwidth of an output link among the VPs sharing that link, and to completely or partially share the *node buffers* among all VPs that go through that node. Effectively each VP will be allocated a fixed amount of bandwidth and a random amount of buffer for its exclusive use. In this case our model would have $x_{n1} = 0$, and $y_{n0} = 0$ if buffers are completely shared and $y_{n0} > 0$ otherwise.

Our model is motivated by resource renegotiation problem. When a connection is admitted, a route is selected, resources are allocated along its route, and a charge per unit of connection time is agreed upon. The charge is determined by the resource allocation and the current resource prices. The route and the charge persist throughout the connection's lifetime, but the resource allocation can be renegotiated from time to time to adapt to changes in network conditions or applications' demand for resources. Changes in network conditions are reflected in changes in the *probability distribution* of R_1 and B_1. Changes in applications' demand for resources are reflected in changes in U_n.

III. Equilibrium Allocations in Model M1

Using standard equilibrium analysis (see, e.g., [25]) we can show that a competitive equilibrium exists.

Proposition 1: There exists a competitive equilibrium.

We next characterize equilibrium prices and allocations.

It is well known that an equilibrium price (p, q) only gives relative prices of the resources in that if $(x_n, y_n, n = 1, \ldots, N)$ is an equilibrium allocation at price (p, q), then it is also an equilibrium allocation at price $(\lambda p, \lambda q)$ for all $\lambda > 0$. The following proposition shows that for an equilibrium price (p, q), $p_0 > 0$ and hence we can fix without loss of generality $p_0 = R_0$ and $q_0 = B_0$. Moreover since prices are positive for fixed resources supply equals demand at equilibrium.

Proposition 2: Let $(p, q, x_n, y_n, n = 1, \ldots, N)$ be an equilibrium. Then $p_0 > 0$, $q_0 > 0$, and $R_0/p_0 = B_0/q_0$. Moreover $X_0 = \bar{X}_0$ and $Y_0 = \bar{Y}_0$.

Our main result for model **M1** is summarized by the following two theorems. Let $\mathbf{P} = \{(p, q) \geq 0 \mid p_0 = R_0, q_0 = B_0,\ p_1 \leq \bar{R}_1, \text{ or } q_1 \leq \bar{B}_1\ \}$. The first result says that we can restrict equilibrium prices to the set $\mathbf{P}$. Indeed if the covariance of R_1 and B_1 $c := cov(R_1, B_1) > 0$ then equilibrium prices can be further restricted to the rectangle $\{(p, q) \geq 0 \mid p_0 = R_0, q_0 = B_0, p_1 \leq \bar{R}_1, q_1 \leq \bar{B}_1\}$.[4] The theorem is illustrated in Figure 2.

Theorem 3: If (p, q) is an equilibrium price then $(p, q) \in \mathbf{P}$, possibly after scaling so that $p_0 = R_0$.

To interpret the theorem we think of $\bar{R}_i/p_i$ and $\bar{B}_i/q_i$ as the *total return* on investing in resource i. Theorem 3 then implies that at equilibrium the total return on fixed resources are equal, and the total expected return on variable resources can be higher (they are indeed higher if the correlation $c = cov(B_1, R_1) > 0$). We note however from assumption **A3** that variable resources causes a reduction in utility because of their variance. At equilibrium since all connections achieve optimality the marginal increase in utility by exchanging one dollar worth of fixed resources for one dollar worth of variable resources must be balanced by the marginal decrease in utility due to increased variability.

[4]The covariance c of R_1 and B_1 may be positive if, for example, R_1 and B_1 model the left over resources not consumed by higher priority traffics where a surge in higher priority traffics leads to reduction in *both* R_1 and B_1.

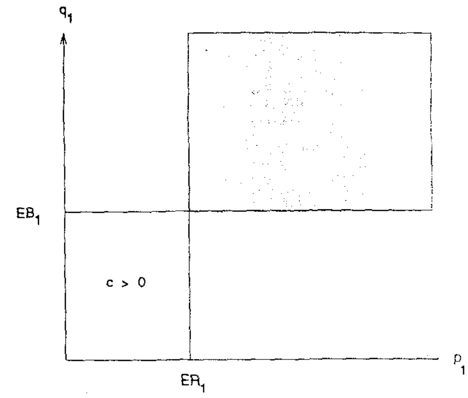

Fig. 2. Region of feasible equilibrium prices. Shaded region is not feasible. If $c = cov(R_1, B_1) > 0$ the feasible region is the lower left rectangle at the origin.

It is then interesting to find that in equilibrium *every* connection holds strictly positive amounts of variable bandwidth and buffer, as the next theorem shows.

Theorem 4: If $(p, q, x_n, y_n, n = 1, \ldots, N)$ is an equilibrium with $p_1 > 0$ or $q_1 > 0$, then $x_{n1} > 0$ and $y_{n1} > 0$ for all n.

From Definition 1 equilibrium price $p_1 = 0$ or $q_1 = 0$ if there are excess variable resources. The qualification that "$p_1 > 0$ or $q_1 > 0$" restricts the theorem to the situation where not both variable bandwidth and buffer are in excess. Indeed if $c = cov(R_1, B_1) < 0$ then the qualification is not necessary.

The theorem has a similar flavor to the well-known fact in the context of investment where it is optimal for every investor to diversify [28], [17], [23], [14]. The security models there however has an important difference: investors are allowed to hold short positions, i.e., (x_n, y_n) can be negative as well as positive. Negative allocations are not meaningful in our context and the nonnegativity constraint in our model complicates greatly the equilibrium analysis.

IV. Equilibrium Allocations in Model M2

In this section we show that the interesting property of strictly positive equilibrium allocation of variable resources persists in model M2 where connections can freely choose allocations that maximize their benefits.

The following result is easy to show.

Proposition 5: There exists a competitive equilibrium. Furthermore the set of equilibrium allocations is convex.

Unlike in model **M1** here the equilibrium prices are not relative. But $(p_0, q_0) > 0$ as before and hence for fixed resources supply equals demand. Moreover the total return on fixed resources are the same.

Theorem 6: Let $(p, q, x_n, y_n, n = 1, \ldots, N)$. Then $p_0 > 0$, $q_0 > 0$, and hence $X_0 = \bar{X}_0$ and $Y_0 = \bar{Y}_0$. Moreover $R_0/p_0 = B_0/q_0$.

Under Assumption **A4** the positive equilibrium allocation of variable resources holds.

Theorem 7: If $(x_n, y_n, n = 1, \ldots, N)$ is an equilibrium allocation then $x_{n1} > 0$ and $y_{n1} > 0$ for all n.

If we replace the strict inequality in assumption **A4** by inequality:

A4': For all μ, v,

$$\frac{\frac{\partial u_n}{\partial \mu}(\mu, v)}{-\frac{\partial u_n}{\partial v}(\mu, v)} \leq 2\frac{v}{\mu}.$$

then the equilibrium situation becomes more complicated. There are still connections that hold strictly positive amounts of both variable bandwidth and variable buffer, but there may be connections that hold zero variable resources. Indeed some connections can be 'priced out' of the market at equilibrium, i.e., $(x_n, y_n) = (0, 0)$. To express precisely the equilibrium situation on the allocation of variable resources, let's define (by slight abuse of notation) $N = \{1, \ldots, N\}$ to be the set of all connections, and partition N into four disjoint subsets according to their equilibrium allocation of *variable* bandwidth and buffer:

$$N = N_0 \cup N_1 \cup N_{11}$$

where N_0 is the set of connections n that hold zero variable bandwidth and zero variable buffer $x_{n1} = y_{n1} = 0$, N_1 is the set of connections n whose allocations are such that exactly one of x_{n1} and y_{n1} is strictly positive, and N_{11} is the set of connections n with $x_{n1} > 0$ and $y_{n1} > 0$.

Theorem 8: Suppose $(p, q, x_n, y_n, n \in N)$ is an equilibrium allocation. If Assumption **A4** is replaced by **A4'** then

(i) $N_{11} \neq \phi$.
(ii) at most one of N_0 and N_1 can be nonempty.
(iii) if N_0 is nonempty then $p_1 > 0$ and $q_1 > 0$.
(iv) if N_1 is nonempty then $p_1 = 0$ or $q_1 = 0$, or both.
(iii) if $x_{n0} > 0$ or $y_{n0} > 0$ then $n \in N_0$ or $n \in N_{11}$.

V. Conclusion

We have described two models for resource allocation schemes that provide a connection with a fixed and a variable share of bandwidth and buffer. In the first model a connection chooses an allocation to maximize its expected utility subject to a budget constraint. In the second model the budget constraint is relaxed and the connection chooses an allocation to maximize its benifit, utility minus resource cost. At equilibrium the marginal increase in utility due to the higher expected return of variable resources is balanced by the marginal decrease in utility due to their variability. We have shown that for both models *all* connections desire strictly positive amounts of variable bandwidth and buffer at equilibrium (under certain assumptions on user utility functions). This implies that ABR or VBR may be more suitable than CBR services for elastic traffics. We are exploiting these interesting equilibrium properties to design renegotiation strategies.

References

[1] Flavio Bonomi, Debasis Mitra, and Judith B. Seery. Adaptive algorithms for feedback-based flow control in high-speed wide-area ATM networks. *IEEE Journal on Selected Areas in Communicaitons*, 13(7):1267–1283, September 1995.

[2] A. K. Choudhury and E. L. Hahne. Dynamic queue length thresholds in a shared memory atm switch. In *IEEE INFOCOM '96*, pages 679–687, San Francisco, California, March 1996.

[3] G. L. Choudhury, K. K. Leung, and W. Whitt. An algorithm to compute blocking probabilities in multi-rate multi-class multi-resource loss models. *Advances in Applied Probability*, 27(4):1104–1143, December 1995.

[4] I. Cidon, L. Giorgiadis, R. Guerin, and A. Khamisy. Optimal buffer sharing. *IEEE Journal on Selected Areas in Communicaitons*, 13(7):1229–1240, September 1995.

[5] R. Cocchi, D. Estrin, S. Shenker, and L. Zhang. Pricing in computer networks: Motivation, formulation and example. *IEEE/ACM Transactions on Networking*, 1(6):614–627, 1993.

[6] Costas Courcoubetis, Vasilios A. Siris, and George D. Stamoulis. Integration of pricing and flow control for ABR services in ATM networks. *Proceedings of Globecom'96*, November 1996.

[7] Costas Courcoubetis, Vasilios A. Siris, George D. Stamoulis, and Richard Weber. ABR pricing model. Preprint, 1997.

[8] N. Duffield and S. Low. The cost of quality in networks of aggregate traffic. In *IEEE Infocom'98*, San Francisco, CA, March 1998.

[9] R. Edell, N. McKeown, and P.P. Varaiya. Billing users and pricing for TCP. *Journal on Selected Areas in Communications*, 13(7):1162–1175, 1995.

[10] A. Elwalid, D. Mitra, and R. Wentworth. A new approach for allocating buffer and bandwidth to heterogeneous, regulated traffic in an atm node. *IEEE Journal on Selected Area in Communications*, 13(6):1115–1127, August 1995.

[11] S. J. Golestani. A self-clocked fair queueing scheme in high speed applications. In *Proceedings of Infocom'94*, pages 636–646, 1994.

[12] M. Grossglauser, S. Keshav, and D. Tse. RCBR: A simple and efficient service for multiple time-scale traffic. In *Proceedings of ACM Sigcomm'95*, pages 219–230, Boston, MA, Aug. 1995.

[13] M. Irland. Buffer management in a packet switch. *IEEE Transactions on Communications*, COM-26:328–337, March 1978.

[14] Robert A. Jarrow. *Finance Theory*. Prentice-Hall, Englewood Cliffs, N.J., 1988.

[15] F. P. Kelly. Charging and accounting for bursty connections. In L. W. McKnight and J. P. Bailey, editors, *Internet Economics*. MIT Press, 1996.

[16] F. P. Kelly. Charging and rate control for elastic traffic. Preprint, 1997.

[17] John Lintner. The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *Review of Economics and Statistics*, 47:13–37, 1965.

[18] S. Low and P. Varaiya. A new approach to service provisioning in ATM networks. *IEEE/ACM Transactions on Networking*, 1(5):547–553, October 1993. For an updated version see http://www.ee.mu.oz.au/staff/slow/research.html.

[19] S. H. Low and P. P. Varaiya. Burst reducing servers in ATM networks. *Queueing Systems*, 20:61–84, 1995.

[20] Steven H. Low. Equilibrium bandwidth and buffer allocations in ATM Networks. In *Proceedings of International Conference on Telecommunications*, Melbourne, Australia, April 1997.

[21] Jeffrey K. MacKie-Mason and Hal R. Varian. Pricing congestible network resources. *IEEE Journal on Selected Areas in Communications*, 13(7):1141–1149, 1995.

[22] Debasis Mitra and Ilze Ziedins. Virtual partitioning by dynamic priorities: fair and efficient resource-sharing by several services. In B. Plattner, editor, *Lecture Notes in Computer Science (Proc. Intl. Zurich Sem. Digital Comm.)*. Spinger, 1996.

[23] J. Mossin. Equilibrium in a capital asset market. *Econometrica*, 34:768–783, 1965.

[24] J. Murphy, L. Murphy, and E. C. Posner. Distributed pricing for embedded ATM networks. In J. Labetoulle and J. W. Roberts, editors, *Proceedings of the 14th International Teletraffic Congress*. Elsevier Science, 1994.

[25] Takashi Negishi. Welfare economics and existence of an equilibrium for a competitive economy. *Metroeconomica*, XII:92–97, 1960.

[26] A. K. Parekh and R. G. Gallager. A generalized processor sharing approach to flow control in integrated services networks - the single node case. *IEEE/ACM Transactions on Networking*, 1(3):344–357, June 1993.

[27] F. L. Presti, Z. Zhang, J. Kurose, and D. Towsley. Source time scale and optimal buffer/bandwidth trade-off for regulated traffic in an atm node. In *Proceedings of Infocom'97*, April 1997.

[28] W. F. Sharpe. Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance*, 19:425–442, 1964.

[29] Scott Shenker. Fundamental deisgn issues for the future internet. *IEEE Journal on Selected Areas in Communications*, 13(7):1176–1188, 1995.

[30] Jean Walrand and Pravin Varaiya. *High-Performance Communication Networks*. Morgan Kaufmann Publisher, San Francisco, CA, 1996.