A New Approach to Service Provisioning in ATM Networks

Steven H. Low, Member, IEEE, and Pravin P. Varaiya, Fellow, IEEE

***Abstract*— We formulate and solve a problem of allocating resources among competing services differentiated by user traffic characteristics and maximum end-to-end delay. The solution leads to an alternative approach to service provisioning in an ATM network, in which the network offers directly for rent its bandwidth and buffers and users purchase freely resources to meet their desired quality. Users make their decisions based on their own traffic parameters and delay requirements and the network sets prices for those resources. The procedure is iterative in that the network periodically adjusts prices based on monitored user demand, and is decentralized in that only local information is needed for individual users to determine resource requests. We derive network's adjustment scheme and users' decision rule and establish their optimality. Since our approach does not require the network to know user traffic and delay parameters, it does not require traffic policing on the part of the network.**

1 Introduction

In this paper, we study a different approach to provisioning services in an ATM network. A *service* is specified by a one-way connection (source, destination, route) and two sets of service parameters. A connection is used to transport a data stream or *message* from the source to the destination; the traffic parameters specify constraint on a user's traffic 'burstiness'; the quality parameters specify maximum end-to-end delay and cell loss rate.[1] We assume that the network offers for sale different types of services, differentiated by the triplet (connection, traffic parameters, quality parameters), and that there is a flow, depending on the service cost, of user requests for these services. A service request is admitted if sufficient resources can be allocated along the connection's route to guarantee service quality. By resources, we mean the bandwidth and buffers in each node along the route. Before transmission, a message is segmented into small, fixed size units called cells. The bandwidth and buffers allocated to a connection can vary over links in its route.

S. Low is with Computing Systems Research Laboratory, AT&T Bell Laboratories, Murray Hill, NJ 07974. (email: slow@research.att.com)

P. Varaiya is with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94720. (email: varaiya@helios.berkeley.edu)

This is the same paper as that appears in the IEEE/ACM Transactions on Networking Volume 1, Number 5, pp. 547–553, October 1993, with minor correction. The correction appears separately in the same journal Volume 2, Number 3, pp. 312, June 1994. Presented in part at Interdisciplinary Workshop on Coordination and Complexity, Berkeley, CA, March 24-25, 1993

[1] We permit no cell loss in this paper though this can be generalized; see §2.

More specifically, a network offers a set S of services. A unit of type s service is provided by a type s connection with the associated traffic and quality parameters, and is sold for a unit price of w_s. The network can produce any amount of type s service provided the required resources do not exceed capacity. Users request services to maximize their individual utilities subject to budget constraints, and the result of this maximization is summarized by an aggregate demand function $D_s(w_s)$.[2] Further assume that the network is regulated to charge its services so as to maximize a welfare function:

$$\max_{w,x} \quad W'(w,x) \tag{1}$$

$$\text{subject to} \quad f(x) \leq 0 \tag{2}$$

$$x_s \leq D_s(w_s), \qquad s \in S$$

where $W'(w,x)$ is the sum of network revenue and user surplus, and (2) is capacity constraint. Standard price-adjustment schemes [21, pp. 188–189] can be used to reach an equilibrium, in which the network sets a price w and supply x, users present their demands $D_s(w_s)$, and the network increases or decreases price according as demand is greater or less than supply, until (w,x) converges.

A key observation is that bandwidth and buffers are 'substitutible resources' to meet a service quality (see below). In this paper, we exploit this tradeoff of different resource combinations to optimize an overall measure of network performance. For each allocation indexed by μ, let $W(\mu)$ denote the welfare achieved in (1) by an equilibrium price – supply vector $(w(\mu), x(\mu))$. Our objective is to derive an iterative and decentralized algorithm that solves for a μ^* that meets service quality and maximizes $W(\mu)$. The algorithm leads to our service provisioning procedure in which the network offers directly for rent its bandwidth and buffers, and the users purchase freely resources to meet its desired quality. A user bases its decision on the knowledge of its *own* traffic and quality parameters, and on the resource price. The network periodically adjusts the prices based on the monitored user request for resources. Unlike the common price adjustment scheme based on the law of demand and supply, our scheme involves a minimization by the network. Furthermore, users of type s service effectively do two optimization in each period: one selects a resource combination along the route to minimize service cost w_s; the other selects a demand $D_s(w_s)$ that maxi-

[2] We assume that demand function can be easily measured by counting the number of service requests, both admitted and rejected.

mizes surplus. It is decentralized in that each user only needs to know the resource price at nodes along its route in addition to its own traffic and quality parameters. The solution makes critical use of the bandwidth-buffer tradeoff described by burstiness curve in [16, 17]: for the network, it determines the resource combination to maximize welfare; for individual users, it provides a simple rule for requesting resources (see (19)).

Our approach differs in three ways from the conventional service provisioning approach in which the network decides beforehand resources to be allocated to the users. First, under our approach, users freely rent resources and package them into services that best meet their needs. Second, since service price is the rent a user pays for the resources it reserves, our procedure ties this price to network performance measured by the welfare function. Third, since the network only guarantees the availability of purchased resources, it is the users' responsibility to shape their traffic in order that the allocated resources can provide the desired quality. Our approach relieves the network of the difficult task of traffic policing and enforcement and can potentially adapt to time-varying user needs expressed by the traffic and quality parameters.

Several sets of previous work are relevant here. In conventional packet-switched networks, bandwidth and buffers are typically not reserved for connections but shared on an on-demand basis. Unrestricted sharing makes it difficult for the network to offer a guaranteed delay to a particular connection, since a more bursty connection can monopolize network resources to the detriment of other connections, possibly leading to congestion. As a consequence of the need to guarantee serviec quality in an ATM network, various bandwidth allocation schemes have recently been designed to guarantee certain amount of bandwidth to a connection despite changes in the number and burstiness of concurrent connections, e.g. [10, 25, 5, 20, 2, 7, 22]. This justifies our approximation of each network node as allocating a fixed bandwidth to a connection. The next question is how much bandwidth to allocate to each connection. To secure service quality, a common proposal is to allocate to each connection enough bandwidth to accommodate its peak rate. This simple proposal leads to inefficiency if the peak rate is much larger than its average rate. This may be unnecessary if the tradeoff between bandwidth and buffers as substitutible resources for service provisioning can be exploited. A more sophisticated proposal is based on the notion of a connection's 'equivalent bandwidth' [13, 6, 14, 3]. [11, 12] study the effect on some overall measure of network performance of varying resource capacities in the network and of varying routing. We vary the resource allocation to the connections rather than the resource capacity, exploiting the bandwidth-buffer tradeoff. Multiple-service-multiple-resource model is also considered in [8]. There the resource allocation is fixed, and the problem is to find a revenue-maximizing admission control policy that may reject a request even when sufficient resources are available in the hope of more profitable requests in the future. In the present formulation resources cannot be reserved in anticipation of future requests. Instead we explore alternative feasible allocations. Pricing has been used previously for control and optimization in communication networks, e.g. [4, 15].

The tradeoff between bandwidth and buffers is easily illustrated. Consider the transfer over a single link of a sequence of video frames; see Figure 1. Suppose a frame contains 512×512 8-bit pixels, and is generated at the video source every 33 msec. The service is to deliver a frame every 33 msec to the display. Suppose the source rate is $m(t)$ as shown in the figure. If we allocate the peak rate 150 Mbps then the received rate is the same as $m(t)$ (except for propagation and processing delay). An alternative is to allocate 65 Mbps and some buffers to the connection; the received rate $m'(t)$ is as shown. The alternative allocation achieves the same service as shown in the figure.

In §2, we introduce our network model and service parameters. In §3, we define optimality and formulate optimal allocation as a game problem. Its solution leads to our iterative and decentralized service provisioning procedure that does not require the network to know user traffic and quality parameters. Concluding remarks are collected in §4, and proofs are relegated to the Appendix.

While our argument is preliminary, it does suggest an alternative to the popular approach in which the network decides which services will best meet user needs. Here, the users decide the resources they need based on their own traffic and quality parameters, and the network coordinates their choices via resource pricing in order to optimize an overall measure of network performance.

2 Model

We consider a network with a set L of links. Link $l \in L$ comprises a transmission capacity of C_l cells per second, or cps, and buffers for B_l cells, see Figure 2. A set R of routes is specified; $r \in R$ also denotes the set of links along that route. The network offers a set S of services. A unit of type s service is provided by a type s connection over route r_s. It is sold at a price of w_s per unit time[3]. Under our service provisioning scheme, this unit price w_s is related to the cost of resources needed to provide the service, and is adjusted periodically to achieve an optimal allocation, as elaborated in §3. We shall assume that user demand, or requests, for type s service is given by the aggregate demand function $D_s(w_s) = \nu_s exp(-w_s), \nu_s := \lambda_s T_s$. $\lambda_s exp(-w_s)$ is the (average) rate type s requests arrive and T_s is the (average) duration of a type s connection. A type s request is admitted and assigned a connection with route r_s provided there is available spare bandwidth of μ_{ls} cps and spare buffers for b_{ls} cells, in each link $l \in r_s$. The routing $s \to r_s$ is fixed, but the allocation

$$s \to \{\mu_{ls}, b_{ls}; l \in r_s\}$$

[3]The price w_s may be some fictitious currency for control purposes that the network can adjust at will.

Figure 1: Tradeoff between bandwidth and buffers

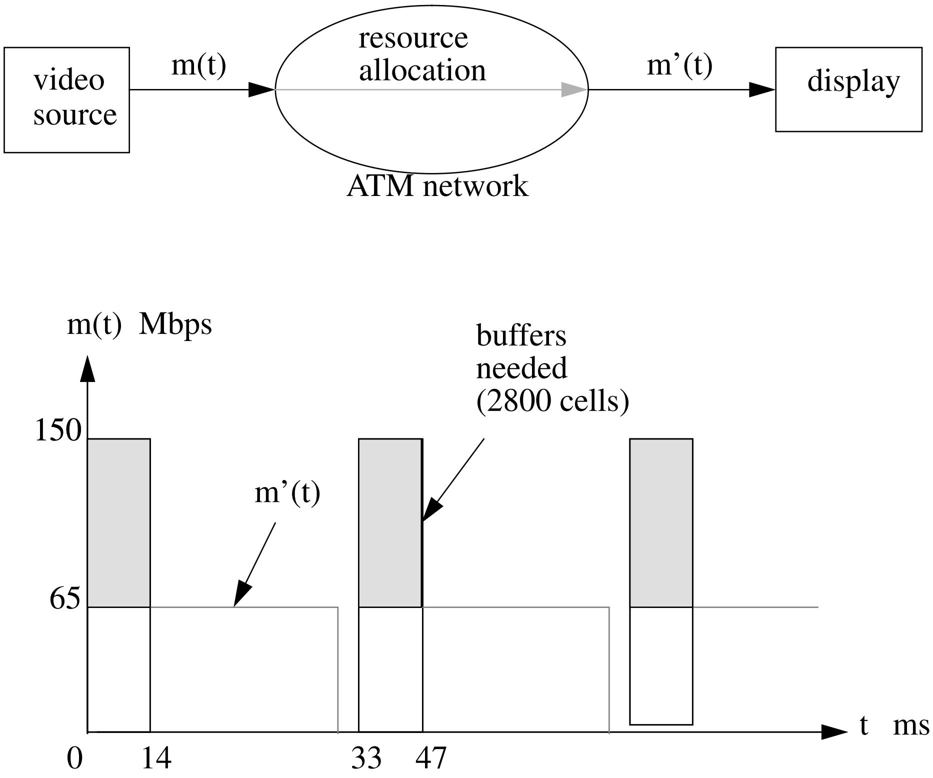

Figure 2: An ATM network

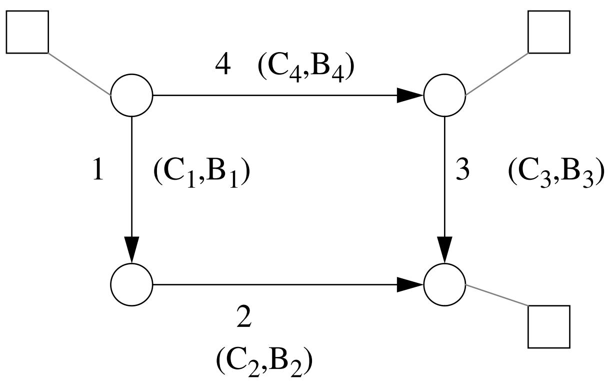

can be freely chosen provided the service quality constraint is met as explained next.

Once a type s request is admitted, a connection is set up, and the user sends a message. A message is a 'fluid flow', $m(t), 0 \leq t \leq T$, where $m(t)$ is the instantaneous rate in cps, and T is the duration. A type s message must satisfy two constraints denoted by $(b_s(\mu), \underline{\mu}_s)$. The parameter $\underline{\mu}_s$ is a positive real number that bounds the average message rate. It is also the minimum bandwidth required for a type s message at each node along its route. The parameter $b_s(\mu)$, $\mu \geq 0$, is a non-negative, decreasing, *strictly* convex function that bounds the message 'burstiness'. A type s message m is said to be *compliant* if

$$\eta := \frac{1}{T} \int_0^T m(\tau) d\tau \leq \underline{\mu}_s \tag{3}$$

and

$$\max_{0 \leq s \leq t \leq T} \int_s^t [m(\tau) - \mu] d\tau \leq b_s(\mu), \qquad \mu \geq \eta \tag{4}$$

Inequality (3) says that the average message rate η cannot exceed $\underline{\mu}_s$. The left-hand side of (4) is the maximum backlog if m is transmitted over a link at a constant speed $\mu \geq \eta$. Hence, inequality (4) says that if m is allocated a bandwidth of μ, then a buffer of size $b_s(\mu)$ is sufficient to prevent cell loss. Note that the larger is μ the smaller is $b_s(\mu)$. Thus the function b_s, called burstiness curve, gives the bandwidth-buffer tradeoff for zero cell loss. To incorporate cell loss, we may relax (4) and let $b_s(\mu)$ be the buffer required to have no more than certain number of lost cells if m is transmitted over a link at a constant speed μ [23].

A word is in order on how this traffic characterization may be used in practice. First, if the duration of a message is very long, e.g. a video program, we divide its duration into disjoint periods T_i, $i = 1, \ldots, k$. A type s message is compliant if the portion of message on every period T_i satisfies (3-4) with T replaced by T_i. (3) then guarantees that no cell backlog carries over to the next period. Second, we do not assume that a user of type s service knows its own message $m(t)$ or the burstiness of $m(t)$. We only assume that they know a bound $b_s(\mu), \mu \geq \eta$, on the burstiness. In fact, condition (4) can be easily enforced by passing an arbitrary user message through a leaky bucket policing device before being admitted into the network. The two parameters of a leaky bucket (and the bound on peak message rate in each period) define a piecewise linear burstiness curve that bounds the burstiness of the output message from the leaky bucket [17, Proposition 3].

For the rest of this paper, we will use the following vector notation. μ_s denotes the vector $\{\mu_{ls},\ l \in r_s\}$ of bandwidth allocation for a type s connection, and μ denotes the vector $\{\mu_s,\ s \in S\}$. Similarly, $b_s(\mu_s)$ denotes the vector $\{b_s(\mu_{ls}),\ l \in r_s\}$ of buffers required for a type s connection, and $b(\mu)$ denotes the vector $\{b_s(\mu_s),\ s \in S\}$. $\underline{\mu}$ denotes the vector $\{\underline{\mu}_s,\ s \in S\}$. We may abuse notation and use '$\mu \geq \underline{\mu}$' to mean '$\{\mu_{ls} \geq \underline{\mu}_s;\ l \in r_s, s \in S\}$'. Finally, $< x, y >$ denotes the inner product of vectors x and y.

To specify the quality of service, the maximum end-to-end delay, we use two results from the theory of burstiness curve in [16, 18, 17]. (See also [1, 20, 24].) Suppose a type s compliant message is transmitted over a connection with route r_s. Suppose that a bandwidth of μ_{ls} cps and buffers of b_{ls} cells are allocated to that connection at each link $l \in r_s$. Suppose that the allocation $\{\mu_{ls}, b_{ls}, l \in r_s\}$ satisfies

$$\mu_{ls} \geq \underline{\mu}_s, \quad b_{ls} \geq b_s(\mu_{ls}), \qquad l \in r_s \tag{5}$$

i.e. at each link, the allocated bandwidth exceeds the minimum bandwidth $\underline{\mu}_s$, and the allocated buffer exceeds the burstiness constraint. Then, (i) no cells will be lost at any link $l \in r_s$ [16, Proposition 4], and (ii) the end-to-end delay is at most [16, Theorem 3]

$$\frac{b_s(\underline{\mu}_s)}{\underline{\mu}_s} + \text{ propagation and processing delay}$$

We shall assume that the 'propagation and processing delay' is constant and omit it from further consideration. Consequently, a maximum end-to-end delay translates into the minimum bandwidth $\underline{\mu}_s$ required at each link along the route.

In summary, there are two service parameters for type s service: the burstiness curve b_s, and the minimum bandwidth $\underline{\mu}_s$ required at each link along the route. A user message is compliant if it satisfies (3-4). An allocation $\{\mu_{ls}, b_{ls}; l \in r_s\}$ is compliant if it satisfies (5). Given service parameters, $\{b_s, \underline{\mu}_s; s \in S\}$, we want to find a compliant allocation $(\mu, b) = \{\mu_{ls}, b_{ls}; l \in r_s, s \in S\}$ that is 'optimal'. We will restrict ourselves to allocations with $b_{ls} = b_s(\mu_{ls})$, since this is sufficient to prevent cell loss. We henceforth represent an allocation by a vector $\mu = \{\mu_{ls}, b_s(\mu_{ls});\ l \in r_s,\ s \in S\}$.

We now formulate the problem and present a solution, which leads to a different approach to service provisioning.

3 Optimal allocation and service provisioning procedure

For the rest of this paper, we relax the constraint of the discrete nature of a type s service and approximate it by a continuous variable. The network can produce any amount x_s of type s service, provided that sufficient resources are available, i.e.

$$x_s \ \leq \ D_s(w), \qquad s \in S \tag{6}$$

$$\textstyle\sum_s x_s \mu_{ls} \leq C_l, \ \sum_s x_s b_s(\mu_{ls}) \leq B_l, \quad l \in L \tag{7}$$

and expects a revenue (rate) of $\sum x_s w_s$. The aggregate demand function summarizes the users' utility such that

$$\sum \int_{w_s}^{\infty} D_s(v) dv$$

is the user surplus [21]. Take as social welfare

$$W'(w, x, \mu) := \sum \int_{w_s}^{\infty} \min(x, D_s(v)) dv + \sum x_s w_s$$

so the problem is to maximize $W'(w,x,\mu)$ subject to (6-7). Consider initially a fixed allocation $\mu \geq \underline{\mu}$.

Definition 1 *A set of prices and amounts of service produced $\{w_s(\mu), x_s(\mu); s \in S\}$ form an equilibrium if, for all $\{x_s\}$ satisfying (6-7),*

$$x_s(\mu) \;=\; \nu_s exp[-w_s(\mu)] \qquad (8)$$

$$\sum x_s(\mu) w_s(\mu) \;\geq\; \sum x_s w_s(\mu) \qquad (9)$$

The conditions say that, in equilibrium, user demand is met and the network maximizes revenue.

Proposition 1 *$(w_s(\mu), x_s(\mu))$ is an equilibrium if and only if there exist $(\alpha(\mu), \beta(\mu)) \geq 0$ such that*

$$x_s(\mu) \;=\; \nu_s exp[-w_s(\mu)] \qquad (10)$$

$$\sum x_s(\mu)\mu_{ls} \leq C_l, \quad \sum x_s(\mu) b_s(\mu_{ls}) \leq B_l \qquad (11)$$

$$w_s(\mu) \;=\; <\alpha(\mu), \mu_s> + <\beta(\mu), b_s(\mu_s)> \qquad (12)$$

$$\sum x_s(\mu) w_s(\mu) \;=\; <\alpha(\mu), C> + <\beta(\mu), B>$$

$$\alpha_l(C_l - \sum x_s(\mu)\mu_{ls}) = 0, \quad \beta_l(B_l - \sum x_s(\mu) b_s(\mu_{ls})) = 0$$

Note that (12) says that the equilibrium price equals the resource cost for providing that service. The cost is estimated by taking as the 'shadow' price or rent per unit time of $\alpha_l(\mu)$ per cps of bandwidth and $\beta_l(\mu)$ per cell of buffer, in link l.

It can be verified that there is a unique equilibrium. Moreover, the unique equilibrium maximizes W' subject to conditions (6–7), as proved next. Let

$$p_s(x_s, w_s) = \begin{cases} D_s^{-1}(x_s) & \text{if } x_s \leq D_s(w_s) \\ w_s & \text{if } x_s > D_s(w_s) \end{cases}$$

Then the welfare function becomes

$$W'(w,x,\mu) := \sum \{x_s p_s(x_s, w_s) + \nu_s exp[-p_s(x_s, w_s)]\}$$

On the set $\{(x,w) | x_s < D_s(w_s)\}$, $p_s(x_s, w_s) = D_s^{-1}(x_s)$. Hence $\frac{\partial W'}{\partial x_s} = D_s^{-1}(x_s) > 0$ for $x_s < \nu_s$ and $\frac{\partial W'}{\partial w_s} = 0$, i.e., any maximizer (x,w) must satisfy $x_s = D_s(w_s)$ provided condition (7) is satisfied. We can restrict our search for a maximizer to the set $\{(x,w) | x_s = D_s(w_s)\}$ and consider the equivalent problem:

$$\max_{x \geq 0} \quad V'(x) := \sum x_s (D_s^{-1}(x_s) + 1)$$
$$\text{subject to} \quad \text{equation (7)}$$

By the Kuhn-Tucker theorem, $x(\mu)$ is a maximizer only if $\exists \alpha(\mu) \geq 0, \beta(\mu) \geq 0$ such that $\frac{\partial V'(x(\mu))}{\partial x_s} - \sum_l (\alpha_l(\mu)\mu_{ls} + \beta_l(\mu) b_s(\mu_{ls})) \leq 0$ with equality if $x_s(\mu) > 0$. Since $x_s(\mu) = D_s(w_s(\mu)) > 0$ and $\frac{\partial V'(x(\mu))}{\partial x_s} = D_s^{-1}(x_s(\mu))$, a necessary condition for $(x(\mu), w(\mu))$ to be a maximizer is

$$w_s(\mu) = D_s^{-1}(x_s(\mu)) = \sum_l (\alpha_l(\mu)\mu_{ls} + \beta_l(\mu) b_s(\mu_{ls}))$$

Since V' is strictly concave ($\frac{\partial^2 V'(x)}{\partial x^2} = -diag(\frac{1}{x_s}) < 0$), the condition is also sufficient. Hence the unique equilibrium maximizes welfare by Proposition 1.

Also from Proposition 1, the maximum welfare is

$$W(\mu) = \sum x_s(\mu) + <\alpha(\mu), C> + <\beta(\mu), B> \qquad (13)$$

Now suppose $\mu \geq \underline{\mu}$ can be freely chosen. We can now formally define an optimal allocation.

Definition 2 *An allocation μ is optimal, or welfare maximizing, if it maximizes the welfare $W(\mu)$ in (13).*

To directly maximizing (13) the network needs to know user traffic and quality parameters $(b_s, \underline{\mu}_s)$. We propose a different approach which does not require such knowledge, and hence does not require any traffic policing and enforcement on the part of the network, though users may still want to shape their messages to comply with $(b_s, \underline{\mu}_s)$ so that the end-to-end delay is met. From (10-11) and the convexity of $G(\mu, \alpha, \beta)$ in α, β, we obtain an alternative expression for the maximum welfare,

$$W(\mu) = \min_{(\alpha,\beta) \geq 0} G(\mu, \alpha, \beta)$$

where

$$G(\mu, \alpha, \beta) = \sum_s \nu_s exp[- <\alpha, \mu_s> - <\beta, b_s(\mu_s)>] + <\alpha, C> + <\beta, B>$$

Hence, μ^* is a welfare-maximizing allocation if

$$W(\mu^*) = \max_{\mu \geq \underline{\mu}} W(\mu) = \max_{\mu \geq \underline{\mu}} \min_{(\alpha,\beta) \geq 0} G(\mu, \alpha, \beta) \qquad (14)$$

The following result is key to our solution. Note that G is convex in (α, β) but not generally concave in μ.

Proposition 2 *There exists a saddle-point $(\mu^*, \alpha^*, \beta^*)$ to the max-min problem (14) that is welfare-maximizing, i.e.*

$$\begin{aligned} & G(\mu^*, \alpha^*, \beta^*) \\ = \; & \max_{\mu \geq \underline{\mu}} \min_{(\alpha,\beta) \geq 0} G(\mu, \alpha, \beta) \\ = \; & \min_{(\alpha,\beta) \geq 0} \max_{\mu \geq \underline{\mu}} G(\mu, \alpha, \beta) \end{aligned}$$

Note that the max-min problem in the proposition is equivalent to the following game: for all $s \in S$,

$$\text{player } U_s: \quad \min_{\mu_s \geq \underline{\mu}_s} \sum_{l \in r_s} (\alpha_l \mu_{ls} + \beta_l b_s(\mu_{ls})) \qquad (15)$$

$$\text{player } N: \quad \min_{(\alpha,\beta) \geq 0} \; G(\mu, \alpha, \beta) \qquad (16)$$

where we recall that $\mu_s = \{\mu_{ls}, \; l \in r_s\}$. The proposition says that if player N chooses the minimizer (α^*, β^*), then player U_s will choose the optimal μ_s^* since $(\mu^*, \alpha^*, \beta^*)$ is a saddle point. Note again that α and β in (15) can be conveniently interpreted as the rent per unit time for one unit of bandwidth and buffer, respectively. Even though the minimizer (α^*, β^*) for $\min_{\alpha,\beta} \; G(\mu^*, \alpha, \beta)$ may not be unique, it can be shown that the unit price $w_s^* = \sum_{l \in r_s} (\alpha_l^* \mu_{ls}^* + \beta_l^* b_s(\mu_{ls}^*))$ for type s service is the same regardless of which minimizer is used as resource price.

This interpretation suggests the following service provisioning procedure to reach $(\mu^*, \alpha^*, \beta^*)$. The procedure is based on an algorithm to solve the following equivalent game problem

$$\begin{array}{lll} \min_{\mu_{ls} \geq \underline{\mu}_s} & \alpha_l \mu_{ls} + \beta_l b_s(\mu_{ls}), & l \in r_s, s \in S \\ \min_{(\alpha,\beta)\geq 0} & G(\mu, \alpha, \beta) & \end{array} \tag{17}$$

Problem (17) is equivalent to (15) since the objective function in (15) is separable in μ_{ls}.

Suppose the network charges each user during connection setup a rent per unit time of α_l per cps of bandwidth and β_l per cell of buffer in link l. The expected cost to a user per request of type s is then[4]

$$T_s \sum_{l \in r_s} (\alpha_l \mu_{ls} + \beta_l b_s(\mu_{ls})) \tag{18}$$

where T_s is the duration of a type s connection. Users may rent any amounts of bandwidth $\mu_{ls} \geq \underline{\mu}_s$ and buffers $b_s(\mu_{ls})$ at each link $l \in r_s$. We *assume* that, given prices (α, β), users will try to minimize their expected service cost (18), subject to their own quality requirement $\mu_s \geq \underline{\mu}_s$. That is, they will take on the role of U_s in (15). We let the network take on the role of N in (16) to calculate the price (α, β). As noted above, if the network charges according to (α^*, β^*), then users will indeed request the optimal allocation μ^*. The following procedure will allow the network to compute α^*, β^* even if it may have no knowledge of user needs expressed by the traffic and quality parameters $\{b_s, \underline{\mu}_s; s \in S\}$.

Service provisioning procedure

Network algorithm:

1. Network initializes the update period $n = 0$.
2. It posts a rent (α^n, β^n) in period n for resources at each link.
3. It monitors the requested bandwidth and buffers (μ^n, b^n) in the entire network, and uses this *observed* (μ^n, b^n) to solve

$$\begin{array}{l} \min_{(\alpha,\beta)\geq 0} \sum_s \nu_s exp[- < \alpha, \mu_s^n > - < \beta, b_s^n >] \\ \qquad + < \alpha, C > + < \beta, B > \end{array}$$

4. It uses any minimizer as rent $(\alpha^{n+1}, \beta^{n+1})$ in the next period. It increments n, and goto step 2.

User algorithm:
In period n,

[4]In charging a user the service cost (18), we implicitly assume that the user does not know the burstiness of its own message but only its bound b_s that can be easily enforced by a leaky bucket. This does not lose generality, since if some of the users know a better bound, we may separate them into a different service type s' without changing the argument.

1. A user of type s service solves (using (19) below)

$$\min_{\mu_{ls} \geq \underline{\mu}_s} \alpha_l^n \mu_{ls} + \beta_l^n b_s(\mu_{ls}), \quad l \in r_s$$

   for minimizer μ_{ls}^n and $b_{ls}^n = b_s(\mu_{ls}^n)$.
2. It requests (μ_s^n, b_s^n), and is admitted if resources are available, and rejected otherwise.
3. If admitted, it pays the rent $T_s \sum_{l \in r_s} (\alpha_l^n \mu_{ls}^n + \beta_l^n b_{ls}^n)$.

The optimality of the procedure is assured by the following theorem.

Theorem 1 *Given (α^0, β^0), construct a sequence $(\mu^n, \alpha^{n+1}, \beta^{n+1})$, $n \geq 0$, by*

$$\begin{array}{rcll} \mu_{ls}^n & = & \arg \min_{\mu_{ls} \geq \underline{\mu}_s} & \alpha_l^n \mu_{ls} + \beta_l^n b_s(\mu_{ls}) \\ & & & l \in r_s, s \in S \\ (\alpha^{n+1}, \beta^{n+1}) & = & \arg \min_{(\alpha,\beta)\geq 0} & G(\mu^n, \alpha, \beta) \end{array}$$

Then any accumulation point of the sequence $(\mu^n, \alpha^{n+1}, \beta^{n+1})$ is a saddle point of G.

The decentralized nature of the procedure is striking: given the price (α, β), the bandwidth and buffer request $(\mu_{ls}, b_s(\mu_{ls}))$ at link l of route r_s that minimizes the service cost (18) depends only on the rent at link l. In fact, given (α, β), the cost minimizing μ_{ls} satisfies, by the Kuhn-Tucker theorem [19],

$$\begin{array}{ll} \mu_{ls}^* = \underline{\mu}_s & \text{if } -\dfrac{\alpha_l^*}{\beta_l^*} < \dfrac{d}{d\mu} b_s(\underline{\mu}_s) \\ \mu_{ls}^* = M_s & \text{if } -\dfrac{\alpha_l^*}{\beta_l^*} > \dfrac{d}{d\mu} b_s(M_s) \\ -\dfrac{\alpha_l^*}{\beta_l^*} = \dfrac{d}{d\mu} b_s(\mu_{ls}^*) & \text{otherwise} \end{array} \tag{19}$$

Here, M_s is the peak message rate at which $b_s(M_s) = 0$. Since $b_s(\mu_{ls})$ is strictly decreasing and convex for $\mu_{ls} < M$, the optimal μ_{ls}^* is unique for each (α^*, β^*). Hence, in each period n, a user minimizes the expected cost (18) by requesting resources according to (19); and this computation can be done locally for each link along the route. By Theorem 1, the service provisioning procedure will achieve a welfare-maximizing allocation. The optimal price (α^*, β^*) determines a minimum-cost allocation μ^*. $(\mu^*, \alpha^*, \beta^*)$ yields the service price w^* and the amount $x_s^* = D_s(w^*)$ of service produced that form an equilibrium.

4 Conclusion

We have studied an alternative approach to service provisioning in an ATM network, which does not require the network to know user needs. In this approach, the network offers directly for rent its bandwidth and buffers and the users purchase them freely to meet their desired quality. The service provisioning procedure is based on a solution of the problem of allocating bandwidth and buffers to meet several types of service requests, differentiated by bounds on the average rate and burstiness of the message and on the end-to-end delay.

The model has one serious deficiency. Once network resources are allocated to a service request, those resources cannot be shared by other connections. Such 'exclusion' is necessary since no cell loss is allowed. However, if one does permit cell loss (as a service quality parameter), then it is possible, indeed desirable, to share resources among concurrent connections, i.e. to permit statistical multiplexing. This will require an understanding of the interaction between resource allocations, cell loss, and delay. A preliminary attempt in this direction is reported in [23].

Appendix

Proof of Proposition 1

By the duality theorem of linear programming, $(x_s(\mu))$ satisfies (9) if and only if there exists $(\alpha(\mu), \beta(\mu)) \geq 0$ such that

$$< \alpha(\mu), \mu_s > + < \beta(\mu), b_s(\mu_s) > \geq \ w_s(\mu)$$

with equality if $x_s(\mu) > 0$. Moreover,

$$\sum x_s(\mu) w_s(\mu) = < \alpha(\mu), C > + < \beta(\mu), B >$$

The result now follows upon noting that, in equilibrium, $x_s(\mu) = \nu_s exp[-w_s(\mu)] > 0$. □

We now proceed to the proof of Proposition 2 and Theorem 1. Since $G(\mu, \alpha, \beta)$ is not concave in μ, we will use a generalization of von Neumann's minimax theorem due to Kakutani [9].

Suppose that M_s is the bound on the peak rate for a type s message, i.e. $b_s(M_s) = 0$. Then, in the zero-sum game of the proposition, we may restrict the strategy set for players $U_s, s \in S$, who maximize $G(\mu, \alpha, \beta)$, to the following compact and convex set

$$S_1 := \{\mu \mid \underline{\mu}_s \leq \mu_{ls} \leq M_s; l \in r_s, s \in S\}$$

Since $G(\mu, 0, 0) = \sum_s \nu_s$, we may restrict the strategy set for player N, who minimizes $G(\mu, \alpha, \beta)$, to the following compact and convex set

$$S_2 := \{(\alpha, \beta) \geq 0 \mid \sum_l (\alpha_l C_l + \beta_l B_l) \leq \sum_s \nu_s\}$$

For each (α, β), player U_s's optimal strategy $R_1(\alpha, \beta)$, given by (19), is unique. For each μ, denote by $R_2(\mu)$ the set of minimizers for (16). Both R_1 and R_2 are nonempty. The convexity of G in (α, β) implies that $R_2(\mu)$ is convex for each μ.

Lemma 1 *R_1 is continuous on S_2. R_2 is upper semi-continuous on S_1, i.e., if $\mu^n \to \mu^*$, $(\alpha^n, \beta^n) \in R_2(\mu^n)$ and $(\alpha^n, \beta^n) \to (\alpha^*, \beta^*)$, then $(\alpha^*, \beta^*) \in R_2(\mu^*)$.*

Proof. The proof for the continuity of R_1 follows from the *strict* convexity of the burstiness curve b_s and the implicit function theorem. We shall prove that R_2 is upper semi-continuous. Suppose in contradiction that $(\alpha^*, \beta^*) \notin R_2(\mu^*)$. Fix $(\hat{\alpha}, \hat{\beta}) \in R_2(\mu^*)$. Then

$$\epsilon := G(\mu^*, \alpha^*, \beta^*) - G(\mu^*, \hat{\alpha}, \hat{\beta}) > 0$$

By continuity of G, for sufficiently large n,

$$G(\mu^n, \alpha^n, \beta^n) > G(\mu^*, \alpha^*, \beta^*) - \frac{\epsilon}{2} \tag{20}$$

Continuity of G at $(\mu^*, \hat{\alpha}, \hat{\beta})$ guarantees a neighborhood $N(\mu^*, \hat{\alpha}, \hat{\beta})$ of $(\mu^*, \hat{\alpha}, \hat{\beta})$ in $S_1 \times S_2$ such that for sufficiently large n, for all $(\mu^n, \alpha, \beta) \in N(\mu^*, \hat{\alpha}, \hat{\beta})$,

$$\begin{aligned} G(\mu^n, \alpha, \beta) &\leq G(\mu^*, \hat{\alpha}, \hat{\beta}) + \frac{\epsilon}{2} \\ &= G(\mu^*, \alpha^*, \beta^*) - \frac{\epsilon}{2} && \text{(By definition of } \epsilon\text{)} \\ &< G(\mu^n, \alpha^n, \beta^n) && \text{(By (20))} \end{aligned}$$

contradicting $(\alpha^n, \beta^n) \in R_2(\mu^n)$. □

Proof of Proposition 2

To prove the existence of a saddle point, define the mapping $\chi : S_1 \times S_2 \to 2^{S_1 \times S_2}$

$$\chi(\mu, \alpha, \beta) \quad = \quad (R_1(\alpha, \beta), R_2(\mu))$$

where $2^{S_1 \times S_2}$ denotes the collection of all subsets of $S_1 \times S_2$. $(\mu^*, \alpha^*, \beta^*)$ is a saddle point if and only if it is a fixed point of χ, because

$$(\mu^*, \alpha^*, \beta^*) \ \in \ \chi(\mu^*, \alpha^*, \beta^*)$$

is equivalent to

$$\begin{aligned} \mu^* &\in R_1(\alpha^*, \beta^*) \\ (\alpha^*, \beta^*) &\in R_2(\mu^*) \end{aligned}$$

which is equivalent to

$$G(\mu, \alpha^*, \beta^*) \ \leq \ G(\mu^*, \alpha^*, \beta^*) \ \leq \ G(\mu^*, \alpha, \beta)$$

for all $(\mu, \alpha, \beta) \in S_1 \times S_2$ We will prove that χ has a fixed point using Kakutani's theorem [9] which states that, given a compact and convex set X, if f is an upper semi-continuous function which assigns to each $x \in X$ a closed and convex subset of X, then there exists some $x \in X$ such that $x \in f(x)$. Now $S_1 \times S_2$ is compact and convex, and by Lemma 1 χ is upper semi-continuous. $\chi(\mu, \alpha, \beta) = R_1(\alpha, \mu) \times R_2(\mu)$ is convex since each component is. We only need to show that $\chi(\mu, \alpha, \beta)$ is closed. $R_1(\alpha, \beta)$ is a singleton and hence closed. The continuity of G implies that $R_2(\mu)$ is closed. Hence $\chi(\mu, \alpha, \beta)$ is closed. □

Proof of Theorem 1

Since $S_1 \times S_2$ is compact, accumulation points of $(\mu^n, \alpha^{n+1}, \beta^{n+1})$ always exist.

Define the mapping $\psi : S_1 \times S_2 \to 2^{S_1 \times S_2}$

$$\psi(\mu, \alpha, \beta) \quad = \quad (R_1(\alpha, \beta), R_2(R_1(\alpha, \beta))$$

Note that if player N starts with the strategy (α, β), then players $U_s, s \in S$, will pick $R_1(\alpha, \beta)$, to which N will respond by picking a strategy in $R_2(R_1(\alpha, \beta))$, and so forth. The following three observations complete the proof:

1. The sequence in the theorem can be constructed, given (α^0, β^0), by
$$(\mu^n, \alpha^{n+1}, \beta^{n+1}) \in \psi(\mu^{n-1}, \alpha^n, \beta^n)$$
2. If $\{\mu^{n(k)}, \alpha^{n(k)+1}, \beta^{n(k)+1}\}$ is a subsequence converging to $(\mu^*, \alpha^*, \beta^*)$, then by the upper semi-continuity of ψ and the first observation, $(\mu^*, \alpha^*, \beta^*)$ is a fixed point of ψ, i.e. $(\mu^*, \alpha^*, \beta^*) \in \psi(\mu^*, \alpha^*, \beta^*)$.
3. As in the proof of Proposition 2, $(\mu^*, \alpha^*, \beta^*)$ is a fixed point of ψ if and only if it is a saddle point.

□

References

[1] Rene L. Cruz. A calculus for network delay Part I:network elements in isolation. *IEEE Transactions on Information Theory*, 37(1):114–131, January 1991.

[2] Alan Demers, Srinivasan Keshav, and Scott Shenker. Analysis and simulation of a fair queueing algorithm. *Proceedings of ACM SIGCOMM'89*, pages 3–12, 1989.

[3] Anwar Elwalid and Debasis Mitra. Effective bandwidth of general Markovian traffic sources and admission control of high speed networks. *IEEE/ACM Transactions on Networking*, June 1993.

[4] D. Ferguson, Y. Yemini, and C. Nikolaou. Microeconomic algorithms for load balancing in distributed computer systems. *Proceedings of the 8th International Conference on Distributed Computing Systems*, pages 491–499, June 1988.

[5] D. Ferrari and D. Verma. A scheme for real-time channel establishment in wide-area networks. *IEEE Journal on Selected Areas in Communications*, 8(3):368–379, April 1990.

[6] R. J. Gibbens and P. J. Hunt. Effective bandwidths for the multi-type UAS channel. *Queueing Systems*, 9(1-2):17–27, October 1991.

[7] S. J. Golestani. Congestion-free communication in high-speed packet networks. *IEEE Transactions on Communications*, 39(12):1802–1812, December 1991.

[8] S. Jordan and P. Varaiya. Throughput in multiple service multiple resource communications networks. *IEEE Transactions on Communications*, 39(8):1216–1222, August 1991.

[9] S. Kakutani. Generalization of Brouwer's fixed point theorem. *Duke Journal of Mathematics*, 8:457–459, 1941.

[10] C. Kalmanek, H. Kanakia, and S. Keshav. Rate controlled servers for very high speed networks. *Proceedings of Globecom'90*, pages 12–20, December 1990.

[11] F. P. Kelly. Routing in circuit-switched networks: optimization, shadow prices and decentralization. *Adv. Appl. Prob.*, pages 112–144, 1988.

[12] F. P. Kelly. Routing and capacity allocation in networks with trunk reservation. *Math. Oper. Res.*, 15:771–793, 1990.

[13] F. P. Kelly. Effective bandwidths at multi-class queues. *Queueing Systems*, 9(1-2):5–15, October 1991.

[14] George Kesidis and Jean Walrand. Effective bandwidths for multiclass Markov fluids and other ATM sources. Preprint, 1992.

[15] J. F. Kurose and R. Simha. A microeconomic approach to optimal resource allocation in distributed computer systems. *IEEE Transactions on Computers*, 38(5):705–717, May 1989.

[16] S. Low and P. Varaiya. A simple theory of traffic and resource allocation in ATM. *Proceedings of Globecom'91*, pages 1633–1637, December 1991.

[17] S. Low and P. Varaiya. Burstiness bounds for some burst reducing servers. *Proceedings of Infocom'93*, pages 2–9, March 1993.

[18] Steven Low. *Traffic Management of ATM Networks: Service Provisioning, Routing, and Traffic Shaping.* PhD thesis, UC Berkeley, May 1992.

[19] David G. Luenberger. *Linear and Nonlinear Programming, 2nd Ed.* Addison-Wesley Publishing Company, 1984.

[20] A. K. Parekh and R. G. Gallager. A generalized processor sharing approach to flow control in integrated services networks - the single node case. *Proceedings of Infocom'92*, pages 915–914, May 1992.

[21] Hal R. Varian. *Microeconomic Analysis.* W. W. Norton & Company Inc., 1978.

[22] Dinesh C. Verma, Hui Zhang, and Domenico Ferrari. Guaranteeing delay jitter bounds in packet switching networks. *Proceedings of TriComm'91*, pages 35–43, April 1991. Chapel Hill, North Carolina.

[23] Michael Wong and Pravin Varaiya. A deterministic fluid model for cell loss in ATM networks. *Proceedings of Infocom'93*, pages 395–400, March 1993.

[24] Opher Yaron and Moshe Sidi. Calculating performance bounds in communication networks. Preprint, 1993.

[25] Lixia Zhang. Virtualclock: A new traffic control algorithm for packet-switched networks. *ACM Transactions on Computer Systems*, 9(2):101–124, May 1991.