

# Simplification of Network Dynamics in Large Systems

Xiaojun Lin and Ness B. Shroff  
School of Electrical and Computer Engineering  
Purdue University  
West Lafayette, IN 47906, U.S.A.  
Email: {linx, shroff}@ecn.purdue.edu.

## *Abstract—*

In this paper we show that significant simplicity can be exploited for pricing-based control of large networks. We first consider a general loss network with Poisson arrivals and arbitrary holding time distributions. In dynamic pricing schemes, the network provider can charge different prices to the user according to the current utilization level of the network and also other factors. We show that, when the system becomes large, the performance (in terms of expected revenue) of an appropriately chosen static pricing scheme, whose price is independent of the current network utilization, will approach that of the optimal dynamic pricing scheme. Further, we show that under certain conditions, this static price is independent of the route that the flows take. This indicates that we can use the static scheme, which has a much simpler structure than the optimal dynamic scheme, to control large communication networks. We then extend the result to the case of dynamic routing, and show that the performance of an appropriately chosen static pricing scheme with bifurcation probability determined by average parameters can also approach that of the optimal dynamic routing scheme when the system is large. Finally, we study the control of elastic flows and show that there exist schemes with static parameters whose performance can approach that of the optimal dynamic resource allocation scheme (in the large system limit). We also identify the applications of our results for QoS routing and rate control for real-time streaming.

## I. INTRODUCTION

In this work, we use pricing as the mechanism of controlling a network to achieve certain performance objectives. The performance objectives can be modeled by some revenue- or utility-functions. Such a framework has received significant interest in the literature (e.g., see [1], [2], [3], [4], [5] and the references therein) wherein price provides a good control signal because it carries monetary incentives. The network can use the current price of a resource as a feedback signal to coerce the users into modifying their actions (e.g., changing the rate or route).

In [6], Paschalidis and Tsitsiklis have shown that the performance (in terms of expected revenue or welfare) of an appropriately chosen *static pricing scheme* approaches the performance

of the optimal *dynamic pricing scheme* when the number of users and the network capacity become very large. Note that a *dynamic pricing scheme*, is one where the network provider can charge different prices to the user according to the varying levels of congestion in the network, while a *static pricing scheme* is one where the price only depends on the average levels of congestion in the network (and is hence invariant to the instantaneous levels of congestion). The result is obtained under the assumption of Poisson flow arrivals, exponential flow holding times, and a single resource (single node). This elegant result is an example of the type of simplicity that one can obtain when the system becomes large. In this paper, we find that simple static network control can also approach the optimal dynamic network control under more general assumptions and a variety of other network problems.

For simplicity of exposition, we structure the paper as follows:

We first extend the result of [6] from the single-link case to a general loss network with arbitrary holding time distributions.<sup>1</sup> Note that while the assumption of Poisson arrivals for flows in the network is usually considered reasonable, the assumption of exponential holding time distribution is not. For example, much of the traffic generated on the Internet is expected to occur from large file transfers which do not conform to exponential modeling. By weakening the exponential service time assumption we can extend our results to more realistic systems. We show that a static pricing scheme is still asymptotically optimal, and that the correct static price depends on the service time distribution only through its mean. A nice observation that stems from this result is that under certain conditions, the static price depends only on the price elasticity of the user, and not on the specific route or distance. This indicates, for example, that the flat pricing scheme used in the domestic long distance telephone service in the US may be a sufficiently good pricing mechanism.

We then investigate whether more sophisticated schemes can improve network performance (e.g., schemes that have prior

This work has been partially supported by the National Science Foundation through the NSF award ANI-0099137, and the Indiana 21st Century Research and Technology Award 1220000634.

<sup>1</sup>Independently, in [7], Pachalidis and Liu have extended the work in [6] from a single-link case to a Markovian loss network. But the authors still have the exponential holding time assumption.

knowledge of the duration of individual flows, schemes that predict the future congestion levels, etc.). We find that the performance gains using such schemes become increasingly marginal as the system size grows.

We then weaken the assumptions of fixed routing and fixed bandwidth flows. In our dynamic routing model, flows can choose among several alternative routes based on the current network congestion level. In our elastic flow model, users are allowed to modify their rates when facing different prices, similar to the way in which TCP and some elastic multimedia traffic react to changing network conditions. In these more general models, when the system is large, we show that the invariance result still holds, i.e., there still exists a static pricing scheme whose performance can approach that of the optimal dynamic scheme.

In networks of today and in the future, the capacity will be very large, and the network will be able to support a large number of users. The work reported in this paper demonstrates under general assumptions and different network problem settings that, when a network is large, significant simplicity can be exploited for pricing based network control. Our result also shows the importance of *average information* when the system is large, since the parameters of the static schemes are determined by average conditions rather than instantaneous conditions. These results will help us develop more efficient and realistic algorithms for controlling large networks. We have identified the applications of our results in QoS routing and rate control for real-time streaming.

Our work also has similarities to the work in [8], [9], and the reference therein. However, in their work, the price is fixed, and the focus is on how to admit and route each flow. Our work (as well as [6]) explicitly models the users' price-elasticity, and consider the optimality of the pricing schemes. Our model of elastic flows is also similar to the optimization flow control model in [3], [10], [4], [5]. However, their models assume that the number of users in the system is fixed. Hence their optimization is done for a snapshot in time, while we explicitly consider the *dynamics* of the network by taking into account the flow arrivals and departures.

Due to space limitations, we leave out most of the proofs. Details of these proofs are available online in [11].

## II. PRICING IN A GENERAL MULTI-CLASS LOSS NETWORK

### A. Model

The basic model that we consider in this section is that of a multi-class loss network with Poisson arrivals and arbitrary service time distributions. There are  $L$  links in the network. Each link  $l \in \{1, \dots, L\}$  has capacity  $R^l$ . There are  $I$  classes of users. We assume that flows generated by users from each class have a fixed route through the network. The routes are characterized by a matrix  $\{C_i^l, i = 1, \dots, I, l = 1, \dots, L\}$ , where  $C_i^l = 1$  if the route of class  $i$  traverses link  $l$ ,  $C_i^l = 0$  otherwise.

Let  $\vec{n} = \{n_1, n_2, \dots, n_I\}$  denote the state of the system, where  $n_i$  is the number of flows of class  $i$  currently in the network. We assume that each flow of class  $i$  requires a fixed amount of bandwidth  $r_i$ . *The fixed routing and fixed bandwidth assumption will be weakened in Sections III and IV, respectively.*

Flows of class  $i$  arrive to the network according to a Poisson process with rate  $\lambda_i(u_i)$ . The rate  $\lambda_i(u_i)$  is a function of the price  $u_i$  charged to users of class  $i$ . Here  $u_i$  is defined as the price per unit time of connection. We assume that  $\lambda_i(u_i)$  is a non-increasing function of  $u_i$ . Therefore  $\lambda_i(u_i)$  represents the *price-elasticity* of class  $i$ . We also assume that for each class  $i$ , there is a "maximal price"  $u_{\max,i}$  such that  $\lambda_i(u_i) = 0$  when  $u_i \geq u_{\max,i}$ . Therefore by setting a high enough price  $u_i$  the network can prevent users of class  $i$  from entering the network. Once admitted, a flow of class  $i$  will hold  $r_i$  amount of resource in the network and pay a cost of  $u_i$  per unit time, until it completes service, where  $u_i$  is the price set by the network at the time of the flow arrival. The service times are i.i.d. with mean  $1/\mu_i$ . The service time distribution is general.

The bandwidth requirement determines the set of feasible states  $\Omega = \{\vec{n} : \sum_i n_i r_i C_i^l \leq R^l \quad \forall l\}$ . A flow will be blocked if the system becomes infeasible after accommodating it. Other than this feasibility constraint, the network provider can charge a different price to each flow, and by doing so, the network provider strives to maximize the revenue collected from the users. The way price is determined can range from the simplest *static pricing schemes* to more complicated *dynamic pricing schemes*. In a *dynamic pricing scheme*, the price at time  $t$  can depend on many factors at the moment  $t$ , such as the current congestion level of the network, etc. On the other hand, in a *static pricing scheme*, the price is fixed over all time  $t$ , and does not depend on these factors. Intuitively, the more factors a pricing scheme can be based on, the more information it can exploit, and hence the higher the performance (i.e., revenue) it can achieve.

*The dynamic pricing scheme we study in this section is more sophisticated than the one in [6].* Firstly, we allow the network provider to exploit the knowledge of the immediate past history of states up to length  $d$ . Note that when the exponential holding time assumption is removed, the system is no longer Markovian. There will typically be correlations between the past and the future given the current state. In order to achieve a higher revenue, we can potentially take advantage of this correlation, i.e., we can use the past to predict the future, and use such prediction to determine the price.

Secondly we allow the network provider to exploit prior knowledge of the parameters of the incoming flows. In particular, the network knows the holding time of the incoming flows, and can charge a different price accordingly. In order to achieve a higher revenue, the network can thus use pricing to control the composition of flows entering the network, for example, short flows may be favored under certain network conditions, while long flows are favored under others. We assume that the price-elasticity of flows is independent of these parameters.

For convenience of exposition, we restrict ourselves to the case when the range of the service time can be partitioned into a series of disjoint segments, and the price is the same for flows that are from the same class and whose service times fall into the same segment. In particular, let  $\{a_k\}, k = 1, 2, \dots$  be an increasing series of positive numbers, i.e.,  $0 < a_1 < a_2 < \dots$  and let  $a_0 = 0$ . We assume that at any time  $t$ , for all flows of class  $i$  whose service times  $T_i$  fall into segment  $[a_{k-1}, a_k)$ , we charge the same price  $u_{ik}(t)$ , i.e. we do not care about the exact value of  $T_i$  as long as  $T_i \in [a_{k-1}, a_k)$ .

The dynamic pricing scheme can thus be written as

$$u_i(t, T_i) = u_{ik}(t) = g_{ik}(\vec{n}(s), s \in [t-d, t]),$$

$$\text{for } T_i \in [a_{k-1}, a_k),$$

where  $\vec{n}(s), s \in [t-d, t]$  reflects the immediate past history of length  $d$ ,  $T_i$  is the holding time of the incoming flow of class  $i$ , and  $g_{ik}$  are functions from  $\Omega^{[-d, 0]}$  to the set of real numbers  $\mathbf{R}$ . By incorporating the past history in the functions  $g_{ik}$ , we can study the effect of prediction on the performance of the dynamic pricing scheme without specifying the details of how to predict. Let  $\vec{g} = \{g_{ik}, i = 1, \dots, I, k = 1, 2, \dots\}$ .

The system under such a dynamic pricing scheme can be shown to be stationary and ergodic under very general conditions. For example, when the arrival rates  $\lambda_i(u)$  are bounded above by some constant  $\lambda_0$ , one can construct a so-called ‘‘regenerative event’’ (due to the Poisson nature of the arrivals), which is the event that the system is empty in the time interval  $[t-d, t]$ . One can show that such an event is a stationary event and occurs with positive probability. This ensures that any stochastic process that is only a function of the system state is asymptotically stationary and the stationary version is ergodic. For details, please see [11], [12].

We are now ready to define the performance objective function. For each class  $i$ , let  $\tilde{T}_{ik} = \mathbf{E}\{T_i | T_i \in [a_{k-1}, a_k)\}$  be the mean service time for flows of class  $i$  whose service time  $T_i$  falls into segment  $[a_{k-1}, a_k)$ . The expectation is taken with respect to the service time distribution of class  $i$ . Let  $p_{ik} = \mathbf{P}\{T_i \in [a_{k-1}, a_k)\}$  be the probability that the service time  $T_i$  of an incoming flow of class  $i$  falls into segment  $[a_{k-1}, a_k)$ . We can decompose the original arrivals of each class into a spectrum of substreams. Substream  $k$  of class  $i$  has service time in  $[a_{k-1}, a_k)$ . Its arrival is thus Poisson with rate  $\lambda_i(u)p_{ik}$ , since we assume that the price-elasticity of flows is independent of  $T_i$ .

For any dynamic pricing scheme  $\vec{g}$ , the expected revenue achieved per unit time is given by

$$\lim_{\zeta \rightarrow \infty} \sum_{i=1}^I \frac{1}{\zeta} \mathbf{E} \left[ \int_0^\zeta \sum_{k=1}^{\infty} \lambda_i(u_{ik}(t)) u_{ik}(t) \tilde{T}_{ik} p_{ik} dt \right]$$

$$= \sum_{i=1}^I \sum_{k=1}^{\infty} \mathbf{E} \left[ \lambda_i(u_{ik}(t)) u_{ik}(t) \tilde{T}_{ik} p_{ik} \right],$$

where the expectation is taken with respect to the steady state distribution. The limit on the left hand side as the time  $\zeta \rightarrow \infty$

exists and equals to the right hand side due to stationarity and ergodicity. Note that the right hand side is independent of  $t$  (from stationarity).

Therefore, the performance of the *optimal* dynamic policy is

$$J^* \triangleq \max_{\vec{g}} \sum_{i=1}^I \sum_{k=1}^{\infty} \mathbf{E} \left[ \lambda_i(u_{ik}(t)) u_{ik}(t) \tilde{T}_{ik} p_{ik} \right].$$

When the exponential holding time assumption is removed, we can no longer use the MDP approach as in [6] to find the optimal dynamic pricing scheme. We will instead study the behaviour of the dynamic pricing scheme and its relationship with the static pricing scheme when the system is large. In particular, we will establish an upper bound for the performance of dynamic pricing schemes and show that the performance of an appropriately chosen static pricing scheme can approach this upper bound as the system is large. We will then conclude that, when the system is large, the performance of an appropriately chosen static pricing scheme can approach that of the *optimal* dynamic pricing scheme. Further, we show that the performance gains of schemes that use such sophisticated mechanisms as prediction and charging based on prior knowledge of the holding times are minimal when the system is large.

## B. An Upper Bound

We find that the upper bound of the form in [6] is also an upper bound for our case. Let  $\lambda_{\max, i} = \lambda_i(0)$  be the maximal value of  $\lambda_i$ . For convenience, we write  $u_i$  as a function of  $\lambda_i$ . Let  $F_i(\lambda_i) = \lambda_i u_i(\lambda_i)$ ,  $\lambda_i \in [0, \lambda_{\max, i}]$ . Further, let  $J_{ub}$  be the optimal value of the following nonlinear programming problem:

$$\max_{\lambda_i, i=1, \dots, I} \sum_i F_i(\lambda_i) \frac{1}{\mu_i} \quad (1)$$

$$\text{subject to} \quad \sum_i \frac{\lambda_i}{\mu_i} r_i C_i^l \leq R^l \quad \text{for all } l, \quad (2)$$

where  $1/\mu_i, r_i$  are the mean holding time and the bandwidth requirement, respectively, for flows from class  $i$ ,  $C_i^l$  is the routing matrix and  $R^l$  is the capacity of link  $l$ .

*Proposition 1:* If the function  $F_i$  is concave in  $(0, \lambda_{\max, i})$  for all  $i$ , then  $J^* \leq J_{ub}$ .

*Proof:* Details of this proof are available online in [11]. ■

The maximizer of the upper bound (1) induces a set of optimal prices  $u_i = u_i(\lambda_i)$ . It is interesting to note that although the dynamic pricing scheme can use prediction and exploit prior knowledge of the parameters of the incoming flows, the upper bound (1) and its induced optimal prices are indifferent to these additional mechanisms.

## C. Static Policy

We now consider the static pricing scheme. In this scheme, the price for each class is fixed, i.e., it does not depend on the

current state of the network, nor does it depend on the individual holding time of the flow. Let  $u_i$  be the static price for class  $i$ . Let  $\vec{u} = [u_1, \dots, u_I]$ . Under this static pricing scheme  $\vec{u}$ , the expected revenue per unit time is:

$$J_0 = \sum_{i=1}^I \lambda_i(u_i) u_i \frac{1}{\mu_i} (1 - \mathbf{P}_{loss,i}[\vec{u}]),$$

where  $\mathbf{P}_{loss,i}[\vec{u}]$  is the blocking probability for class  $i$ . Therefore the performance of the *optimal* static policy is

$$J_s \triangleq \max_{\vec{u}} \sum_{i=1}^I \lambda_i(u_i) u_i \frac{1}{\mu_i} (1 - \mathbf{P}_{loss,i}[\vec{u}]).$$

By definition  $J_s \leq J^*$ .

Throughout this paper we will focus on large systems with many small users. To be specific, we consider the following scaling (S):

(S) Let  $c \geq 1$  be a scaling factor. We consider a series of systems scaled by  $c$ . The scaled system has capacity  $R^{l,c} = cR^l$  at each link  $l$ , and the arrivals of each class  $i$  has rate  $\lambda_i^c(u) = c\lambda_i(u)$ . Let  $J^{*,c}$ ,  $J_s^c$  and  $J_{ub}^c$  be the dynamic revenue, static revenue, and upper bound, respectively, for the  $c$ -scaled system.

We are interested in the performance of the dynamic pricing scheme and the static pricing scheme when  $c \uparrow \infty$ , i.e., when both the capacity and the number of users in the system become very large. We first note that *the normalized upper bound  $J_{ub}^c/c$  is fixed over all  $c$* , since  $J_{ub}^c$  is obtained by maximizing  $\sum_i c\lambda_i u_i (\lambda_i)/\mu_i$ , subject to the constraints  $\sum_i c\lambda_i r_i C_i^l/\mu_i \leq cR^l$ , for all  $l$ . Therefore the optimal price induced by the upper bound is also independent of  $c$ .

The following lemma illustrates the behaviour of the blocking probability  $\mathbf{P}_{loss,i}$  as  $c \rightarrow \infty$  under scaling (S).

*Lemma 2:* Let  $\lambda_i$  be the arrival rate of flows from class  $i$  and let  $1/\mu_i$  be the mean holding time. Under the assumptions of Poisson arrivals and general holding time distributions, if the load at each resource is less than or equal to 1, i.e.,

$$\sum_i \frac{\lambda_i}{\mu_i} r_i C_i^l \leq R^l \text{ for all } l,$$

then under scaling (S), as  $c \rightarrow \infty$ , the blocking probability of each class goes to 0, and the speed of convergence is at least  $1/\sqrt{c}$ .

*Proof:* The proof of this lemma is analogous to the proof of Proposition 3 in [12]. Details of this proof are available online in [11]. ■

We will use this lemma to show the following main result:

*Proposition 3:* If the function  $F_i$  is concave in  $(0, \lambda_{\max,i})$  for all  $i$ , then

$$\lim_{c \rightarrow \infty} \frac{1}{c} J_s^c = \lim_{c \rightarrow \infty} \frac{1}{c} J^{*,c} = \lim_{c \rightarrow \infty} \frac{1}{c} J_{ub}^c = J_{ub}.$$

*Proof:* Since  $J_s^c \leq J^{*,c} \leq J_{ub}^c = cJ_{ub}$ , we only need to show that  $\lim_{c \rightarrow \infty} J_s^c/c \geq J_{ub}$ .

Now consider  $J_s^c$ . For every static price  $\vec{u} = [u_1, \dots, u_I]$  falling into the constraint of  $J_{ub}$ , i.e.,

$$\sum_i \frac{c\lambda_i(u_i) r_i C_i^l}{\mu_i} \leq cR^l \text{ for all } l, \quad (3)$$

let  $J_0^c$  denote the revenue under this static price. Since (3) guarantees that the condition of Lemma 2 is met, we have  $\mathbf{P}_{loss,i}[\vec{u}] \rightarrow 0$ , as  $c \rightarrow \infty$ . Therefore

$$\begin{aligned} \lim_{c \rightarrow \infty} \frac{J_0^c}{c} &= \lim_{c \rightarrow \infty} \sum_i \lambda_i(u_i) u_i \frac{1}{\mu_i} (1 - \mathbf{P}_{loss,i}[\vec{u}]) \\ &= \sum_i \lambda_i(u_i) u_i \frac{1}{\mu_i}. \end{aligned} \quad (4)$$

If we take the optimal price induced by the upper bound as our static price, then the right hand side of (4) is exactly the upper bound. Therefore,

$$\lim_{c \rightarrow \infty} \frac{J_s^c}{c} \geq \lim_{c \rightarrow \infty} \frac{J_0^c}{c} \geq J_{ub}$$

and the result follows. ■

Proposition 3 can be seen as a network version (with also general holding times) of Theorem 6 in [6]. It tells us that extending the result of [6] from a single link to a network of links and from exponential holding time distributions to arbitrary holding time distributions does not change the invariance result. In other words, there still exists static pricing schemes whose performance can approach that of the optimal dynamic pricing scheme when the system is large. Further, even though the dynamic pricing scheme can use prediction and exploit prior knowledge of the parameters of the incoming flows, the upper bound (1) turns out to be indifferent to these additional mechanisms. This shows that these extra mechanisms have a minimal effect on the long term revenue when the system is large.

The static schemes are much easier to implement because they do not require the collection of instantaneous load information. Instead, they only depend on some average parameters, such as the average load, etc. The static schemes are also much easier to obtain because of their simple structure. Hence, they introduce less communication and computation overhead and they are insensitive to feedback delays. In future work we intend to develop efficient distributed algorithms that can find these static prices. We will discuss this briefly in Section V.

Here we report a few numerical results. Consider the network in Fig. 1. There are 4 classes of flows. Their routes are shown in the figure. Their arrivals are Poisson. The function  $\lambda_i(u)$  for each class  $i$  is of the form

$$\lambda_i(u) = \left[ \lambda_{\max,i} \left( 1 - \frac{u}{u_{\max,i}} \right) \right]^+,$$




Fig. 1. The network topology

TABLE I  
TRAFFIC AND PRICE PARAMETERS OF 4 CLASSES

	Class 1	Class 2	Class 3	Class 4
$\lambda_{\max,i}$	0.01	0.01	0.02	0.01
$u_{\max,i}$	10	10	20	20
Service Rate $\mu_i$	0.002	0.001	0.002	0.001
Bandwidth $r_i$	2	1	1	2

i.e.,  $\lambda_i(0) = \lambda_{\max,i}$  and  $\lambda_i(u_{\max,i}) = 0$  for some constants  $\lambda_{\max,i}$  and  $u_{\max,i}$ . The price elasticity is then

$$-\frac{\lambda'_i(u)}{\lambda_i(u)} = \frac{1/u_{\max,i}}{1 - u/u_{\max,i}}, \text{ for } 0 < u < u_{\max,i}.$$

The function  $F_i$  is thus

$$F_i(\lambda_i) = \lambda_i \left(1 - \frac{\lambda_i}{\lambda_{\max,i}}\right) u_{\max,i},$$

which is concave in  $(0, \lambda_{\max,i})$ . The holding time is exponential with mean  $1/\mu_i$ . The parameters  $\lambda_{\max,i}$ ,  $u_{\max,i}$ , service rates  $\mu_i$ , and bandwidth requirement  $r_i$  for each class are shown in Table I.

First, we consider a base system where the 5 links have capacity 10, 10, 5, 15, and 15 respectively. The solution of the upper bound (1) is shown in Table II. The upper bound is  $J_{ub} = 127.5$ . We then use simulations to verify how tight this upper bound is and how close the performance of the static pricing policy can approach this upper bound when the system is large. We use the price induced by the upper bound calculated above as our static price. We first simulate the case when the holding time distributions are exponential. We simulate  $c$ -scaled versions of the base network where  $c$  ranges from 1 to 1000. For each scaled system, we simulate the static pricing scheme, and report the revenue generated. In Fig. 2 we show the normalized revenue  $J_0/c$  as a function of  $c$ .

As we can see, when the system grows large, the difference in performance between the static pricing scheme and the upper bound decreases. Although we do not know what the optimal dynamic scheme is, its normalized revenue  $J^*/c$  must lie somewhere between that of the static scheme and the upper bound. Therefore the difference in performance between the static pricing scheme and the optimal dynamic scheme is further reduced.

TABLE II  
SOLUTION OF THE UPPER BOUND (1) WHEN THE CAPACITY OF LINK 3 IS 5 BANDWIDTH UNITS. THE UPPER BOUND IS  $J_{ub} = 127.5$

	Class 1	Class 2	Class 3	Class 4
$u_i$	9.00	5.00	12.00	10.00
$\lambda_i(u_i)$	0.00100	0.00500	0.00800	0.00500
$\lambda_i(u_i)/\mu_i$	0.500	5.00	4.00	5.00

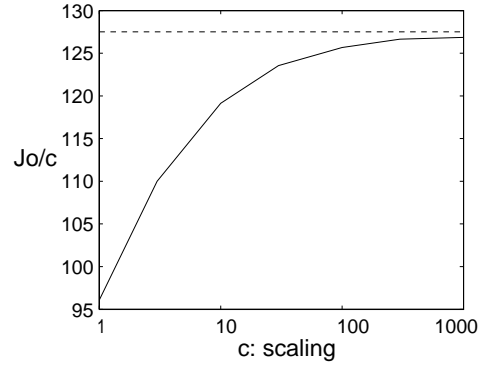


Fig. 2. The static pricing policy compared with the upper bound: when the capacity of link 3 is 5 bandwidth units. The dotted line is the upper bound.

For example, when  $c = 10$ , which corresponds to the case when the link capacity can accommodate around 100 flows, the performance gap between the static policy and the upper bound is less than 7%. The gap decreases as  $1/\sqrt{c}$ .

We now change the capacity of link 3 from 5 bandwidth units to 15 bandwidth units. The solution of the upper bound is shown in Table III. The upper bound is  $J^* = 137.5$ . The simulation result (Fig. 3) confirms again that the performance of the static policy approaches the upper bound when the system is large. At  $c = 10$ , the performance gap between the static policy and the upper bound is around 10%. Note that in this latter example, the static price is the same for users with the same price-elasticity even if they traverse different routes. For example, classes 1 & 2 and classes 3 & 4 have different routes but have the same price (and price-elasticity). In general, *if there is no significant constraint of resources, the maximizing price structure will be independent of the route of the connection.* (A network has no significant constraint of resources if the *unconstrained* maximizer of  $\sum_i F_i(\lambda_i)$  satisfies the constraint (2).) To see this, we go back to the formulation of the upper bound (1). If the *unconstrained* maximizer of  $\sum_i F_i(\lambda_i)$  satisfies the constraint, then it is also the maximizer of the *constrained* problem. In this case the price only depends on the function  $F_i$ , which is determined by the price elasticity of the users. Readers can verify that, in our second example, when the capacity of link 3 is 15 bandwidth units. if we lift the constraints in (2), and solve the upper bound again, we will get the same result. Therefore in our example, *the optimal price*

TABLE III

SOLUTION OF THE UPPER BOUND WHEN THE CAPACITY OF LINK 3 IS 15 BANDWIDTH UNITS. THE UPPER BOUND IS  $J_{ub} = 137.5$

	Class 1	Class 2	Class 3	Class 4
$u_i$	5.00	5.00	10.00	10.00
$\lambda_i(u_i)$	0.00500	0.00500	0.0100	0.00500
$\lambda_i(u_i)/\mu_i$	2.50	5.00	5.00	5.00

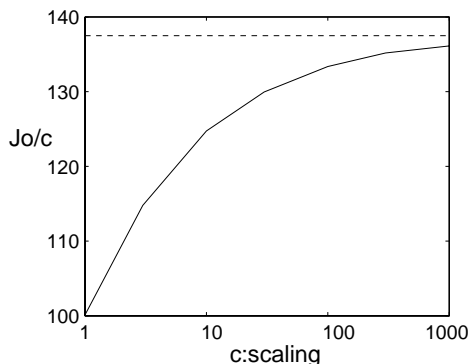


Fig. 3. The static pricing policy compared with the upper bound: when the capacity of link 3 is 15 bandwidth units. The dotted line is the upper bound.

will only depend on the price elasticity of each class and not on the specific route. Since class 1 has the same price elasticity as class 2, its price is also the same as that of class 2, even though it traverses a longer route through the network. This result perhaps justifies the use of flat pricing in inter-state long distance telephone service in the United States.

We also simulate the case when the holding time distribution is deterministic. The result is the same as that of the exponential holding time distribution. The simulation result with heavy tail holding time distribution also shows the same trend except that the sample path convergence (i.e., convergence in time) becomes very slow, especially when the system is large. For example, Fig. 4 is obtained when the holding time distribution is Pareto, i.e., the cumulative distribution function is  $1 - 1/x^a$ , with  $a = 1.5$ . We use the same set of parameters as the constrained case above, and let the Pareto distribution have the same mean as that of the exponential distribution. Note that this distribution has finite mean but infinite variance. This demonstrates that our result is indeed invariant of the holding time distribution.

### III. DYNAMIC ROUTING

We next consider a system with dynamic routing. Many results in the QoS routing literature focus on finding the “best” route for each individual flow based on the instantaneous network conditions. When these QoS routing algorithms are used in a dynamic routing setting, the network is typically required to first collect link information (such as available bandwidth,




Fig. 4. The static pricing policy compared with the upper bound: when the capacity of link 3 is 5 bandwidth units and the service time distribution is Pareto.

delay, etc.) on a regular basis. Then, when a request for a new flow arrives, the QoS routing algorithms are invoked to find a route that can accommodate the flow. When there are multiple routes that can satisfy the request, certain heuristics are used to pick one of the routes. However, such “greedy” schemes may be sub-optimal system wide, because a greedy selection may result in an unfavorable configuration such that more future flows are blocked. Further, an obstacle to the implementation of these dynamic schemes is that it consumes a significant amount of resources to propagate link states throughout the network. Propagation delay and stale information will also degrade the performance of the dynamic routing schemes.

In this section, we will formulate a dynamic routing problem that directly optimizes the total system revenue. Although our model is simplified, it reveals important insight on the performance tradeoff among different dynamic routing schemes. We will establish an upper bound on the performance of the dynamic schemes, and show that the performance of an appropriate chosen static pricing scheme, which selects routes based on some pre-determined probabilities, can approach the performance of the optimal dynamic scheme when the system is large. The static scheme only requires some average parameters. It consumes less communication and computation resources, and is insensitive to network delay. Thus the static scheme is an attractive alternative for control of routing in large networks.

The network model is the same as in the last section, except that now a user of class  $i$  has  $\theta(i)$  alternative routes that are represented by matrix  $\{H_{ij}^l\}$  such that  $H_{ij}^l = 1$ , if route  $j$  of class  $i$  uses resource  $l$  and  $H_{ij}^l = 0$ , otherwise. The dynamic schemes we consider have the following idealized properties: the routes of existing flows can be changed during their connection; and the traffic of a given flow can be transmitted on multiple routes at the same time. Thus our model captures the packet-level dynamic routing capability in the current Internet. These idealized capabilities allow the dynamic schemes to “pack” more flows into the system. Yet, we will show that an appropriately chosen static routing scheme will have comparable performance to the optimal dynamic scheme.

Let  $n_i$  be the number of flows of class  $i$  currently in the network. Consider the  $k$ -th flow of class  $i$ ,  $k = 1, \dots, n_i$ . Let  $P_{ij}^k$  denote the proportion of traffic of flow  $k$  assigned to route  $j$ ,  $j = 1, \dots, \theta(i)$ . Then, state  $\vec{n} = \{n_1, \dots, n_I\}$  is feasible if and only if

$$\begin{aligned} &\text{There exists } P_{ij}^k \text{ such that } \sum_j P_{ij}^k = 1, \forall i, k, \\ &\text{and } \sum_{i,j} r_i H_{ij}^l \sum_{k=1}^{n_i} P_{ij}^k \leq R^l \text{ for all } l. \end{aligned} \quad (5)$$

The set of feasible states is  $\Omega = \{\vec{n} \text{ such that (5) is satisfied}\}$ .

A dynamic scheme can charge prices based on the current state of the network, or a finite amount of past history, i.e., prediction based on past history. (For simplicity we consider pricing schemes that are insensitive to the individual holding times.) An incoming flow will be admitted if the resulting state is in  $\Omega$ . Once the flow is admitted, its route (i.e.,  $P_{ij}^k$ ) is assigned based on (5), involving (in an idealized dynamic scheme) possible rearrangement of routes of all existing flows. We assume that such rearrangement can be carried out instantaneously. Thus a dynamic pricing scheme can be modeled by  $u_i(t) = g_i(\vec{n}(s), s \in [t-d, t])$ , where  $g_i$  is a function from  $\Omega^{[-d,0]}$  to  $\mathbf{R}$ . Let  $\vec{g} = \{g_1, \dots, g_I\}$ .

The performance objective is again the expected revenue per unit time generated by the incoming flows admitted into the system. The performance of the *optimal dynamic routing scheme* is given by:

$$\begin{aligned} J^* \triangleq & \max_{\vec{g}} \mathbf{E} \left\{ \sum_i \lambda_i(u_i(t)) u_i(t) \frac{1}{\mu_i} \right\} \\ & \text{subject to (5)}. \end{aligned} \quad (6)$$

The expectation is taken with respect to the steady state distribution. Note that (6) is independent of  $t$  because of stationarity and ergodicity.

The set of dynamic schemes we have described may require complex capabilities (e.g., rearrangements of routes and transmitting traffic of a single flow over multiple routes) and hence may not be suitable for actual implementation. We make clear here that we are not advocating implementing such schemes but instead advocate implementing static schemes. In fact, we will show that, as the system scales, our static scheme will approach the performance of the optimal idealized dynamic scheme. The static schemes do not require the afore-mentioned complex capabilities and could be an attractive alternative for network routing.

Let  $u_i = u_i(\lambda_i)$  and  $F_i(\lambda_i) = u_i(\lambda_i)\lambda_i$ . Analogous to Proposition 1, we can derive the following upper bound on the optimal revenue in (6). Details are available online in [11].

*Proposition 4:* If the function  $F_i$  is concave in  $(0, \lambda_{\max,i})$  for all  $i$ , then  $J^* \leq J_{ub}$ , where  $J_{ub}$  is defined as the solution for the following optimization problem:

$$J_{ub} \triangleq \max_{\lambda_{ij}} \sum_i F_i \left( \sum_j \lambda_{ij} \right) \frac{1}{\mu_i} \quad (7)$$

$$\text{subject to } \sum_{ij} \frac{\lambda_{ij}}{\mu_i} H_{ij}^l r_i \leq R^l \quad \forall l.$$

We next construct our static routing policy as follows: The network charges a static price to all incoming flows, and the incoming flows are directed to alternative routes based on pre-determined probabilities. *Note that the static policy does not have the idealized capabilities prescribed for the dynamic schemes, i.e., all traffic of a flow has to follow the same path, and rearrangement of routes of existing flows is not allowed.* Let  $\{u_i^s, P_{ij}^s\}$  denote such a static policy, where  $u_i^s$  is the price for class  $i$ , and  $P_{ij}^s$  is the bifurcation probability that an incoming flow from class  $i$  is directed to route  $j$ .

Then the optimal static policy can be found by solving:

$$J_s \triangleq \max_{u_i^s, P_{ij}^s, \sum_j P_{ij}^s = 1} \sum_{ij} \lambda_i(u_i^s) u_i^s P_{ij}^s \frac{1}{\mu_i} [1 - \mathbf{P}_{Loss,ij}], \quad (8)$$

where  $\mathbf{P}_{Loss,ij}$  is the blocking probability experienced by users of class  $i$  routed to  $j$ .

We consider a special static policy derived from the solution of the upper bound in Proposition 4. If  $\lambda_{ij}^{ub}$  is the maximal solution to the upper bound, we let  $u_i^s = u_i(\sum_j \lambda_{ij}^{ub})$ , and  $P_{ij}^s = \frac{\lambda_{ij}^{ub}}{\sum_j \lambda_{ij}^{ub}}$ . The revenue with this static policy differs from the upper bound only by the term  $(1 - \mathbf{P}_{Loss,ij})$ , and this revenue will be less than  $J_s$ . However, under scaling (S), we can show that, as  $c \rightarrow \infty$ ,  $\mathbf{P}_{Loss,ij} \rightarrow 0$ . Therefore, we have our invariance result (stated next).

*Proposition 5:* In the dynamic routing model, if the function  $F_i$  is concave in  $(0, \lambda_{\max,i})$  for all  $i$ , then

$$\lim_{c \rightarrow \infty} J_s^c / c = \lim_{c \rightarrow \infty} J^{*,c} / c = \lim_{c \rightarrow \infty} J_{ub}^c / c = J_{ub}$$

*Proof:* Analogous to that of Proposition 3. Details are available online at [11].  $\blacksquare$

When the routing is fixed, by replacing  $\lambda_{ij}$  with  $\lambda_i$ , and  $H_{ij}^l$  with  $C_i^l$ , we recover Propositions 1 and 3 from the results in this Section. When there are multiple available routes, the upper bound in Proposition 4 is typically larger than that of Proposition 1. Therefore one can indeed improve revenue by employing dynamic routing. However, Proposition 5 shows that, when the system is large, most of the performance gain can also be obtained by simpler static schemes that routes incoming flows based on pre-determined probabilities. Further, what we learn is that for large systems the capability to rearrange routes and to transmit traffic of a single flow on multiple routes does not lead to significant performance gains.

Not only can the static schemes be asymptotically optimal, they also have a very simple structure. Their parameters are determined by average conditions rather than instantaneous conditions. Collecting average information introduces less communication and processing overhead, and it is also insensitive to network delay. Hence the static schemes are much easier to implement in practice.

The optimal static scheme also reveals the macroscopic structure of the optimal dynamic routing scheme. For example,

the static price  $u_i^s$  shows the preference of some classes than the others, and the static bifurcation probability  $P_{ij}^s$  reveals the preference on certain routes than the other. While a “greedy” routing scheme tries to accommodate each individual flow, the optimal static scheme may reveal that one should indeed prevent some flows from entering the network, or prevent some routes from being used. For our future work we plan to study efficient distributed algorithms to derive these optimal static parameters.

#### IV. ELASTIC FLOWS

In previous sections we have restricted ourselves to the case when the bandwidth requirements of flows are fixed. In this section we will extend the model to the case when users can change their bandwidth requirements according to the current price. For ease of exposition we assume that there is only one route for each class  $i$ . The routes are again represented by the matrix  $\{C_i^l\}$  as in Section II. Flows of class  $i$  enter the network according to a Poisson process with rate  $\lambda_i$ . The service times of flows of class  $i$  are i.i.d. with mean  $1/\mu_i$ . The service time distribution is general. Let  $U_i(x_i)$  be the utility function for each class  $i$ , where  $x_i$  is the amount of resource assigned to a class  $i$  flow along its route. We assume that  $U_i$  is a continuous differentiable and strictly concave function of  $x_i$ , and  $U_i(0) = 0$ . This model is appropriate for real-time streaming applications that can change the transmission rate according to the network congestion level. For example, the utility function  $U_i(x_i)$  can be taken as the index of reception quality when the real-time stream is transmitting at rate  $x_i$ .

The network tries to allocate resources to the flows so that the total utility of all flows supported by the network is maximized. For each flow, the resource allocation may vary over time. In this section, we will first establish the optimal dynamic scheme. We will then show, as before, that there exists a static scheme whose performance will approach that of the optimal dynamic scheme when the system is large. Surprisingly, this near-optimal solution is in a “fixed-bandwidth” and “loss-network” form as in Section II.

##### A. The Optimal Dynamic Scheme

Let  $n_i(t)$  be the number of flows from class  $i$  that are in the network at time  $t$ . Let  $\vec{n}(t) = \{n_1(t), n_2(t), \dots, n_I(t)\}$ . The optimal resource assignment is then given by the solution to the following problem:

$$J^*(\vec{n}(t)) \triangleq \max_{x_i} \sum_{i=1}^I n_i(t) U_i(x_i) \quad (9)$$

subject to  $\sum_{i=1}^I n_i(t) x_i C_i^l \leq R^l,$

where  $J^*(\vec{n}(t))$  can be interpreted as the maximal total utility achieved by the system at time  $t$ . For each  $t$  we can solve (9)

and obtain the optimal assignment  $x_i(t)$ . Over time, this policy will optimize the total utility.

Remark: in the optimal assignment (9), each flow of class  $i$  will consume same amount of resource  $x_i$ . This is a consequence of the concavity of  $U_i$ . For details, see [11].

In the past (e.g., [3], [4], [10]) this model has been used to study the behavior of TCP congestion control when the number of flows in the system is fixed. It has been shown that there exist distributed algorithms that can drive the flows to the optimal resource assignment. The notion of “price” arises naturally as Lagrange multipliers for the constraints. Some examples of such distributed algorithms resemble the control of TCP in the Internet. Therefore, TCP congestion control can be seen to maximize the total utility of a group of users with concave utility functions. Our model is different from theirs because we consider the dynamics caused by the arrivals and departures of flows. We are interested in finding alternative forms of resource assignment schemes that can also achieve near optimal total utility when the system is large. These schemes can then be used in cases when TCP does not work as well.

##### B. An Upper Bound

Let  $\mathbf{E}[n_i]$  be the stationary mean of  $n_i(t)$ , i.e.,  $\mathbf{E}[n_i] = \lambda_i/\mu_i$ . We formulate another optimization problem:

$$J_{ub} \triangleq \max_{x_i} \sum_{i=1}^I \mathbf{E}[n_i] U_i(x_i) \quad (10)$$

subject to  $\sum_{i=1}^I \mathbf{E}[n_i] x_i C_i^l \leq R^l.$

*Proposition 6:* The expected total utility of the optimal dynamic scheme is upper bounded by  $J_{ub}$ , i.e.  $\mathbf{E}[J^*] \leq J_{ub}$ , where the expectation is taken with respect to the steady state distribution of  $n_i(t)$ .

*Proof:* Note that  $J^*$  is a function of  $\vec{n}(t) = \{n_i(t), i = 1, \dots, I\}$ . Then  $J_{ub} = J^*(\mathbf{E}[\vec{n}])$ .

To show that  $\mathbf{E}[J^*(\vec{n})] \leq J^*(\mathbf{E}[\vec{n}])$ , it is sufficient to show that  $J^*(\vec{n})$  is a concave function of  $\vec{n}$ , which is a consequence of the concavity of  $U_i(x_i)$ . For details, please refer to [11]. ■

##### C. Static Policy

Let  $x^0 = \{x_1^0, x_2^0, \dots, x_I^0\}$  be the corresponding maximizing parameter of (10). Now consider the following control algorithm with a static rate assignment: when a new flow from class  $i$  arrives to the network, it will be assigned a rate  $x_i^0$  if there is enough capacity available along its route, otherwise it will either be blocked, or, equivalently, be assigned a rate 0. Therefore, the flow is still elastic except that the rate is chosen according to the average condition as in (10) rather than the instantaneous condition as in (9). The flow will hold the same amount of resource  $x_i^0$  until it leaves the system.



In such a system, the expected total utility will be

$$J_s \triangleq \sum_{i=1}^I \frac{\lambda_i}{\mu_i} U_i(x_i^0)(1 - \mathbf{P}_{loss,i}),$$

where  $\mathbf{P}_{loss,i}$  is the blocking probability of class  $i$ . Under scaling ( $\mathbf{S}$ ), we have the following proposition.

*Proposition 7:* In the elastic flow model,

$$\lim_{c \rightarrow \infty} \frac{1}{c} J_s^c = \lim_{c \rightarrow \infty} \frac{1}{c} \mathbf{E}[J^{*,c}] = \lim_{c \rightarrow \infty} \frac{1}{c} J_{ub}^c = J_{ub}.$$

*Proof:* Analogous to that of Proposition 3. Details are available online at [11]. ■

An application of this result is on the rate control of real-time flows (e.g. audio and video streaming) on the Internet. A central question in congestion control of streaming traffic is its fairness with respect to TCP. When real-time flows and TCP flows coexist in the same network, they should consume comparable bandwidth, and neither flows should be starved by the other. Among the existing congestion control schemes for real-time flows, some use the same AIMD (Additive Increase Multiplicative Decrease) idea as TCP [13]. They are usually fair with TCP if timeouts occur infrequently. However, these schemes typically produce a TCP-like saw-tooth type of trajectory, which leads to rapid changes in reception quality. Such rapid changes in quality are disconcerting for the viewer of multimedia flows [14]. Equation-based congestion control does not use AIMD and produces smoother rates at small time-scales. However, simulation results show that at time-scales around 10 seconds, the fluctuation is still quite significant [15]. There are yet other schemes, such as some binomial algorithms [16], which change the rate slower than TCP. However they are also slower in adapting to changing network conditions.

Note that fairness objectives are very closely related to the utility maximization objectives. For example, proportional fairness is equivalent to maximizing the total utility of a group of users with log-utility functions. If we adopt utility maximization as a substitute for the fairness requirement, we can use the result above to obtain a new class of congestion-control algorithms for real-time traffic. For example, consider the special case when  $\alpha$  portion of the flows are real-time flows, and the rest are TCP flows. To be precise, let  $n_i^{\text{RT}}(t)$  and  $n_i^{\text{TCP}}(t)$  denote the number of real-time flows and TCP flows, respectively, at time  $t$ . Then their stationary means are  $\mathbf{E}[n_i^{\text{RT}}] = \alpha \mathbf{E}[n_i]$  and  $\mathbf{E}[n_i^{\text{TCP}}] = (1 - \alpha) \mathbf{E}[n_i]$ . Let us assign the fixed bandwidth  $x_i^0$  to real-time flows, and allow them to use the same amount of bandwidth throughout the connection. Such fixed bandwidth allocation is beneficial to streaming applications because it ensures a stable reception quality for the viewer. Therefore the expected total utility achieved by real-time flows is given by

$$J^{\text{RT}} = \mathbf{E}[n_i^{\text{RT}}] U_i(x_i^0)(1 - \mathbf{P}_{loss,i}^{\text{RT}}) = \alpha J_{ub}(1 - \mathbf{P}_{loss,i}^{\text{RT}}),$$

where  $\mathbf{P}_{loss,i}^{\text{RT}}$  is the blocking probability experienced by the real-time flows. The total utility achieved by TCP flows at

time  $t$  is given by the following optimization problem:

$$J^{\text{TCP}} \triangleq \max_{x_i} \sum_{i=1}^I n_i^{\text{TCP}}(t) U_i(x_i), \quad (11)$$

subject to 
$$\sum_{i=1}^I n_i^{\text{TCP}}(t) x_i C_i^l \leq R^l - \sum_{i=1}^I n_i^{\text{RT}}(t) x_i C_i^l.$$

The expected total utility achieved by both the real-time flows and the TCP flows,  $J^{\text{RT}} + \mathbf{E}[J^{\text{TCP}}]$ , is bounded from above by  $J_{ub}$  and bounded from below by  $J_s$ . Therefore, by Proposition 7,

$$\lim_{c \rightarrow \infty} \frac{J^{\text{RT},c} + \mathbf{E}[J^{\text{TCP},c}]}{c} = \lim_{c \rightarrow \infty} \frac{J_{ub}^c}{c} = J_{ub},$$

where  $J^{\text{RT},c}$ ,  $J^{\text{TCP},c}$  and  $J_{ub}^c$  are the respective utility when the system is scaled by  $c$ . Now by Lemma 2,

$$\lim_{c \rightarrow \infty} \mathbf{P}_{loss,i}^{\text{RT}} = 0.$$

Therefore

$$\lim_{c \rightarrow \infty} \frac{J^{\text{RT},c}}{c} = \alpha J_{ub},$$

and we conclude that

$$\lim_{c \rightarrow \infty} \frac{\mathbf{E}[J^{\text{TCP},c}]}{c} = (1 - \alpha) J_{ub}.$$

Note that by Proposition 7,  $(1 - \alpha) J_{ub}$  is also the limit of the normalized expected total utility achieved by the TCP flows as  $c \rightarrow \infty$ , when the remaining portion  $\alpha$  of the flows are also TCP flows. This shows that when the same utility functions are used for both the real-time flows and TCP flows, assigning the fixed bandwidth  $x_i^0$  to real-time flows does not degrade the performance of the TCP flows when the system is large.

It is interesting to compare existing congestion-control schemes with our scheme above. In existing schemes, flows start from an arbitrary initial condition, and congestion control is exercised *during* the connection. In our scheme, congestion control is exercised *at the beginning* of the connection. The congestion controller reacts to changing network condition by choosing the correct initial bandwidth assignment for incoming flows. Although our scheme does not modify the bandwidth assignment for on-going flows, the difference between the total utility of our scheme and the optimal utility is minimal (when the system is large). Therefore, in the long run, the real-time flows and TCP flows will receive fair share of the bandwidth. In future work we plan to investigate the problem of efficiently distributing our congestion controller over the network.

## V. CONCLUSION AND FUTURE WORK

In this work we study pricing as a mechanism to control large networks. We show under very general settings that an appropriately chosen static pricing scheme is asymptotically optimal

when the system is large. We have established these results for admission control, dynamic routing, and control of elastic flows.

The above results have important implications in the networks of today and in the future. Compared with dynamic pricing schemes, static pricing schemes have some desirable properties. They are less computationally intensive, and introduce less communication overhead. Their performance will not degrade as the network delay grows. Our results show that when the system is large, as in broadband networks, the difference between static pricing schemes and dynamic pricing schemes is minimal.

Having said this, it should be noted that the parameters of the static schemes are obtained from some *global* optimization problem (e.g., (1), (7) and (10)) which requires coordination among possibly all elements of the network. To make the static schemes implementable, it is important to develop efficient *distributed* algorithms that can find these parameters. Here we briefly discuss one possible approach (that is further elaborated in [11]). We can associate a non-negative Lagrange multiplier  $p^l$  for the constraint (e.g. (2)) at each resource  $l$ . The Lagrange multiplier  $p^l$  can be viewed as the implicit cost that summarizes the congestion information at link  $l$ . Given  $p^l$ , in order to determine the price for class  $i$ , one only needs to know the price-elasticity of class  $i$  (i.e., the function  $F_i$ ) and the sum of the implied costs along the path that flows of class  $i$  traverse. *Therefore we can decompose the global optimization problem into several subproblems for each class.* We can have the core routers generate these implicit costs and have the ingress router for each class  $i$  probe these implicit costs and determine the price for class  $i$ . The idea of this decomposition has been used in [4] to develop distributed algorithms for optimization flow control, and it is also mentioned in [6] for the static pricing scheme in the single-link case. Following this path, we are currently studying distributed algorithms for static pricing schemes under more general settings.

Note that in the distributed algorithms outlined above for static pricing schemes, the network still needs to generate and communicate the implicit costs between different network elements. However, the computation and communication involved will be much smaller than in the optimal dynamic pricing scheme. In the optimal dynamic pricing scheme, the network has to acquire the instantaneous global state  $\vec{n}(t)$ , and then compute the right price for each network state. While in the static schemes, only one set of static prices needs to be found given the network topology and the functions  $F_i$ . Therefore,

the computation and propagation of the implicit costs  $p^l$  can be much slower than the evolution of the network state. Once the distributed algorithm converges, the prices  $u_i$  (and the implicit costs  $p^l$ ) stay unchanged until the network topology or the load condition  $F_i$  change.

## REFERENCES

- [1] J. K. Mackie-Mason and H. R. Varian, "Pricing the Internet," in *Public Access to the Internet (B. Kahin and J. Keller, Eds.)*, pp. 269–314. The MIT Press, Cambridge, MA, 1995.
- [2] A. A. Lazar and N. Semret, "The Progressive Second Price Auction Mechanism for Network Resource Sharing," in *8th International Symposium on Dynamic Games and Applications, Maastricht, the Netherlands*, July 1998, pp. 359–365.
- [3] F. P. Kelly, A. Maulloo, and D. Tan, "Rate Control in Communication Networks: Shadow Prices, Proportional Fairness and Stability," *Journal of Operational Research Society*, vol. 49, pp. 237–252, 1998.
- [4] S. H. Low and D. E. Lapsley, "Optimization Flow Control—I: Basic Algorithm and Convergence," *IEEE/ACM Transactions on Networking*, vol. 7, no. 6, pp. 861–874, 1999.
- [5] H. Yaiche, R. Mazumdar, and C. Rosenberg, "A Game Theoretic Framework for Bandwidth Allocation and Pricing in Broadband Networks," *IEEE/ACM Transactions on Networking*, vol. 8, no. 5, pp. 667–678, Oct. 2000.
- [6] I. Ch. Paschalidis and J. N. Tsitsiklis, "Congestion-Dependent Pricing of Network Services," *IEEE/ACM Transactions on Networking*, vol. 8, no. 2, pp. 171–184, April 2000.
- [7] I. Ch. Paschalidis and Y. Liu, "Pricing in Multiservice Loss Networks: Static Pricing, Asymptotic Optimality, and Demand Substitution Effects," *IEEE/ACM Transactions on Networking*, 2002, in print, available at <http://ionia.bu.edu/>.
- [8] R. J. McEliece and K. N. Sivarajan, "Maximizing Marginal Revenue In Generalized Blocking Service Networks," in *Proc. 30th Annual Allerton Conference on Communication, Control, and Computing*, 1992, pp. 455–464.
- [9] P. B. Key, "Optimal Control and Trunk Reservation in Loss Networks," *Probability in the Engineering and Informational Sciences*, vol. 4, pp. 203–242, 1990.
- [10] S. Kunniyur and R. Srikant, "End-to-End Congestion Control Schemes: Utility Functions, Random Losses and ECN Marks," in *Proceedings of IEEE INFOCOM*, Tel-Aviv, Israel, March 2000.
- [11] X. Lin and N. B. Shroff, "Simplification of Network Dynamics in Large Systems," *Technical Report, Purdue University*, <http://shay.ecn.purdue.edu/~linx/papers.html>, 2001.
- [12] X. Lin and N. B. Shroff, "Pricing-Based Control of Large Networks," in *Evolutionary Trends of the Internet, Proceedings of Thyrrenian International Workshop on Digital Communications (IWDC 2001)*, Taormina, Italy, September 2001, pp. 212–231.
- [13] R. Rejaie, M. Handley, and D. Estrin, "RAP: an End-to-End Rate-Based Congestion Control Mechanism for Realtime Streams in the Internet," in *Proceedings of IEEE INFOCOM*, New York, NY, March 1999.
- [14] B. Girod, "Psychovisual Aspects of Image Communications," *Signal Processing*, vol. 28, no. 3, pp. 239–251, September 1992.
- [15] S. Floyd, M. Handley, and J. Padhye, "A Comparison of Equation-Based and AIMD Congestion Control," in <http://www.aciri.org/tfrc/aimd.ps>, 2000.
- [16] D. Bansal and H. Balakrishnan, "Binomial Congestion Control Algorithms," in *Proceedings of IEEE INFOCOM*, Anchorage, Alaska, April 2001.