

Simplification of Network Dynamics in Large Systems

Xiaojun Lin, *Student Member, IEEE*, and Ness B. Shroff, *Senior Member, IEEE*

Abstract—We show that when networks are large significant simplicity can be achieved for pricing-based control. 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. 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. These results deepen our understanding of pricing-based network control. In particular, they provide us with the insight that, when the system is large, an appropriate pricing strategy based on the average network conditions (hence, slowly changing) can approach optimality.

Index Terms—Large systems, network pricing, routing, simplicity.

I. INTRODUCTION

IN this work, we use pricing as the network control mechanism for achieving 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 [3], [4], [5], [6], [7] 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 [8], 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 link). 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. In Section II, we extend the result of [8] from the single-link case to a general loss network with arbitrary holding time distributions. (Independently of our work that was first reported in [1], Paschalidis and Liu have also extended the work in [8] from a single-link case to a network case [9]. However, their result still uses the exponential holding-time assumption.) Our technical contributions in this paper are three-fold:

- 1) We generalize the results of [8] and [9] to *non-exponential* holding time distributions. 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 holding 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 holding 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 U.S. may be an economically good pricing mechanism.

- 2) We also investigate whether more sophisticated schemes can improve network performance (e.g., schemes that have prior knowledge of the duration of individual flows, schemes that predict future congestion levels, etc.). We find that the performance gains using such schemes become increasingly marginal as the system size grows.

- 3) We study two types of scaling to model large networks. We study the original scaling in [8] and [9], which requires

Manuscript received xxxxxxxx xx, 200x; revised xxxxxxxx xx, 200x.

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. Earlier versions of this paper have appeared in the Proceedings of 2001 Tyrrhenian International Workshop on Digital Communications [1] and ACM/IEEE IWQoS'02 [2].

The authors are with the School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN 47907, USA (email: {linx,shroff}@ecn.purdue.edu).

that the number of users between each source-destination pair is scaled proportionally with the capacity of the network. We also study a new type of scaling that is suitable for the case when the capacity of the network is large but when the number of users between each source-destination pair is small. Our new scaling is more appropriate for modeling certain Internet scenarios where the topology is complex and the routing is diverse. We show that appropriate static schemes are asymptotically optimal under both scaling models.

We then weaken the assumption on fixed routing and study a dynamic routing model in Section III where flows can choose among several alternative routes based on the current network congestion level. We show that our *invariance* type of result still holds in this more general model, i.e., when the system is large, 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 the *average network condition* 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.

Our work is also related to the work in [10], [11], and the references therein. However, in their work, the price is set *a priori*, and the focus is on how to admit and route each flow. Our work (as well as [8], [9]) explicitly models the users' price-elasticity, and considers the problem of how to set the price. The impact of large systems has also been studied for flow control problems in [12], [13], however, these works focus on the single-link case with a fixed number of flows.

Before we proceed with the details of our results, we would like to comment on what "average network conditions" and "static schemes" mean in real networks. In practice, over a long enough period of time, the network condition is usually non-stationary, and even the average network statistics could vary. Hence, in real networks, when we say a "static scheme", it does not necessarily mean that the price is fixed over the entire time. Rather, the "static scheme" should be interpreted as *prices being static over the time period for which the network statistics do not change* (which is typically still fairly long in real networks) and the *average network condition* should be interpreted as the average of the dominant network condition *over such a time period*. Over longer time scales, the prices should still adapt to the changes in the dominant network condition, although relatively slowly.

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 , and $C_i^l = 0$ 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 .

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. Therefore $\lambda_i(u_i)$ can be viewed as the "demand function" and it represents the *price-elasticity* of class i . We 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. We also assume that $\lambda_i(u_i)$ is a continuous and strictly decreasing function for $u_i \in (0, u_{\max,i})$. 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 times¹ are *i.i.d.* with mean $1/\mu_i$. The service time distribution is general².

The bandwidth requirement determines the set of feasible states $\Omega = \{\vec{n} : \sum_{i=1}^I n_i r_i C_i^l \leq R^l \quad \forall l\}$. 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 in which 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 that we study in this section is more sophisticated than the one in [8]. 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, the network provider can potentially exploit this correlation, i.e., it can use the past to predict the future, and use such prediction to determine price.

¹We use the notions of *service time* and *holding time* interchangeably.

²By assuming that the demand function $\lambda_i(u_i)$, the mean holding time $1/\mu_i$ and the capacity R^l are fixed, we have assumed that the network condition is stationary. We will comment in Section II-G how the results of this paper should be interpreted when this assumption does not hold.

Secondly, we allow the network provider to exploit prior knowledge of the parameters of the incoming flows. In particular, the network provider knows the holding time of the incoming flows, and can charge a different price accordingly. In order to achieve a higher revenue, the network provider 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 their holding times.

For convenience of exposition, we restrict ourselves to the case when the range of the service time can be partitioned into a collection 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. Specifically, let $\{a_k\}, k = 1, 2, \dots$ be an increasing sequence of positive numbers that approach $+\infty$ as $k \rightarrow +\infty$. Let $a_0 = 0$. The service time of a flow is a non-negative real number and hence must fall into one of the segments $[a_{k-1}, a_k], k = 1, 2, \dots$. 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

$$u_i(t, T_i) = u_{ik}(t) = g_{ik}(\vec{n}(s), s \in [t-d, t]), \quad (1)$$

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. One such condition is stated in the following proposition.

Proposition 1: Assume that for all classes i , the arrival rates $\lambda_i(u)$ are bounded from above by some constant λ_0 . Further, for each class i , the holding times are *i.i.d.* with finite mean and independent of the holding times of other classes and all arrivals. If the price is only dependent on the current state of the system, and/or a finite amount of past history (i.e., prediction based on past history), and/or the holding times of the incoming flows, then the stochastic process $\vec{n}(t)$ (i.e. the system state) is asymptotically stationary as $t \rightarrow \infty$ and the stationary version is ergodic.

Proposition 1 can be proved by using Borovkov's Ergodic Theorem [14]. We can construct a so-called "regenerative event" for the system. Once such an event occurs, the system will evolve independently from the past. By showing that such an event occurs with positive probability, we can conclude that any stochastic process that is only a function of the system state is asymptotically stationary as $t \rightarrow \infty$, and the stationary version is ergodic. The details of the proof are omitted here due to space constraints. They were reported in [1] and are also available in our online technical report [15].

We are now ready to define the performance objective function. For each class i , let \tilde{T}_{ik} be the mean service time for flows of class i whose service time T_i falls into segment $[a_{k-1}, a_k)$, i.e., $\tilde{T}_{ik} = \mathbf{E}\{T_i | T_i \in [a_{k-1}, a_k)\}$. 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

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

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). The performance of the *optimal* dynamic policy is thus given by:

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

Finally, we construct the performance objective for *static pricing schemes*. In a static pricing 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:

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

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

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

When the exponential holding time assumption is removed, we can no longer use the MDP approach as in [8] to find the optimal dynamic pricing scheme. We will instead study the behavior of the optimal dynamic pricing scheme by bounding its performance. These bounds will then help us establish the main result that, when the network is large, an appropriately chosen static pricing scheme can achieve almost the same performance as that of the optimal dynamic scheme. By definition, the performance of any static pricing scheme becomes a lower bound for the performance of the optimal dynamic pricing scheme. An upper bound is presented next.

B. An Upper Bound

The upper bound is of a similar form to that in [8]. Let $\lambda_{\max,i} = \lambda_i(0)$ be the maximal value of λ_i . For convenience, we write u_i as a function of λ_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:

$$\begin{aligned} \max_{\lambda_i, i=1, \dots, I} \quad & \sum_{i=1}^I F_i(\lambda_i) \frac{1}{\mu_i} \\ \text{subject to} \quad & \sum_{i=1}^I \frac{\lambda_i}{\mu_i} r_i C_i^l \leq R^l \quad \text{for all } l, \end{aligned} \quad (2)$$

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

Proof: Consider an optimal dynamic pricing policy. Let $n_{ik}(t)$ be the number of flows of substream k of class i in the system at time t , hence $n_i(t) = \sum_{k=1}^{\infty} n_{ik}(t)$. Let $\lambda_{ik}(t) = \lambda_i(u_{ik}(t))$. From Little's Law, we have

$$\mathbf{E}[n_{ik}(t)] = \mathbf{E}[\lambda_{ik}(t) p_{ik}] \tilde{T}_{ik}.$$

The expectation is taken with respect to the steady state distribution.

Now let

$$\lambda_i^* = \frac{\sum_{k=1}^{\infty} \mathbf{E}[\lambda_{ik}(t)] p_{ik} \tilde{T}_{ik}}{\sum_{k=1}^{\infty} p_{ik} \tilde{T}_{ik}}.$$

Note that $\sum_{k=1}^{\infty} p_{ik} \tilde{T}_{ik} = 1/\mu_i$, therefore

$$\frac{\lambda_i^*}{\mu_i} = \sum_{k=1}^{\infty} \mathbf{E}[\lambda_{ik}(t)] p_{ik} \tilde{T}_{ik} = \sum_{k=1}^{\infty} \mathbf{E}[n_{ik}(t)] = \mathbf{E}[n_i(t)].$$

At any time t , $\sum_{i=1}^I n_i(t) r_i C_i^l \leq R^l$ for all l . Therefore

$$\sum_{i=1}^I \frac{\lambda_i^*}{\mu_i} r_i C_i^l \leq R^l \quad \text{for all } l.$$

Since the functions F_i are concave, using Jensen's inequality, we have,

$$\begin{aligned} J_{ub} &\geq \sum_{i=1}^I F_i(\lambda_i^*) \frac{1}{\mu_i} \\ &\geq \sum_{i=1}^I \sum_{k=1}^{\infty} F_i(\mathbf{E}[\lambda_{ik}(t)]) p_{ik} \tilde{T}_{ik} \\ &\geq \sum_{i=1}^I \sum_{k=1}^{\infty} \mathbf{E}[F_i(\lambda_{ik}(t))] p_{ik} \tilde{T}_{ik} = J^*. \end{aligned}$$

The upper bound (2) has a simple and intuitive form. Its objective function can be viewed as an approximation of the

average revenue without taking into account blocking, while the constraints simply keep the load at all links to be no greater than 1 (where the load at a link l is defined by $\frac{1}{R^l} \sum_{i=1}^I \frac{\lambda_i}{\mu_i} r_i C_i^l$).

The maximizer $[\lambda_1, \dots, \lambda_I]$ of the upper bound (2) also induces a set of optimal prices $u_i^{ub} = u_i(\lambda_i)$. 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 (2) and its induced optimal prices are indifferent to these additional mechanisms.

Remark: The concavity assumption on F_i is essential in the proof of the upper bound and the result that follows. A linear $\lambda_i(u_i)$, as used in many applications, guarantees that F_i is concave. If F_i is not concave, some sort of ‘‘convexification’’ procedure needs to be invoked. Readers can refer to [8] for discussions on the implications of non-concave F_i .

C. Asymptotic Optimality of the Static Scheme

In this paper, we are interested in studying *large systems*, i.e., when the capacity and the number of users in the network are large. We will show that, as the network grows large, the relative difference between the revenue of the optimal dynamic scheme, the revenue of the optimal static scheme, and the upper bound, will approach zero. We first study the following scaling **(S1)**, which is also used in [8], for modeling large systems. A different scaling will be studied in Section II-F.

(S1) Let $c \geq 1$ be a scaling factor. We consider a series of systems scaled by c . All systems have the same topology. 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 optimal dynamic revenue, optimal static revenue, and the upper bound, respectively, for the c -scaled system.

We are interested in the performance difference of the dynamic pricing schemes and the static pricing schemes when $c \uparrow \infty$. We first note that, under the scaling **(S1)**, the normalized upper bound J_{ub}^c/c is a constant independent of c , because J_{ub}^c is obtained by maximizing $\sum_{i=1}^I c\lambda_i u_i(\lambda_i)/\mu_i$,

subject to the constraints $\sum_{i=1}^I c\lambda_i r_i C_i^l / \mu_i \leq cR^l$, for all l . Therefore the optimal price induced by the upper bound is also independent of c .

Fix a set of static prices \vec{u} . Let $\mathbf{P}_{loss,i}^c[\vec{u}]$ denote the blocking probability of class i in the c -scaled system. The following lemma illustrates the behavior of $\mathbf{P}_{loss,i}^c[\vec{u}]$ as $c \rightarrow \infty$ under the scaling **(S1)**. Recall that the load at a link l is defined by $\frac{1}{R^l} \sum_{i=1}^I \frac{\lambda_i}{\mu_i} r_i C_i^l$.

Lemma 3: Let λ_i be the arrival rate of flows from class i at a given set of static prices \vec{u} . 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.,

$$\sum_{i=1}^I \frac{\lambda_i}{\mu_i} r_i C_i^l \leq R^l \quad \text{for all } l,$$

then under scaling **(S1)**, as $c \rightarrow \infty$, the blocking probability $\mathbf{P}_{loss,i}^c[\vec{u}]$ of each class i goes to 0, and the speed of convergence is at least $1/\sqrt{c}$, i.e., $\mathbf{P}_{loss,i}^c[\vec{u}] = O(\frac{1}{\sqrt{c}})$. Further, when the load at all links that class i traverses is strictly less than 1, the speed of convergence is exponential, i.e.,

$$\limsup_{c \rightarrow \infty} \frac{1}{c} \log \mathbf{P}_{loss,i}^c[\vec{u}] \leq \max_{l: C_i^l=1} \inf_{w>0} \Lambda_l(w) < 0,$$

where

$$\Lambda_l(w) = \sum_{j=1}^I \frac{\lambda_j}{\mu_j} (e^{r_j w} - 1) C_j^l - w R^l. \quad (3)$$

The proof of this lemma can be found in the Appendix. We now use this lemma to show the following main result:

Proposition 4: If the function F_i is concave in $(0, \lambda_{\max,i})$ for all i , then under the scaling **(S1)**,

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

Instead of proving Proposition 4, we will prove the following stronger result that the static prices induced by the upper bound are in fact asymptotically optimal.

Proposition 5: Let $\vec{u}^{ub} = [u_1^{ub}, \dots, u_I^{ub}]$ denote the set of static prices induced by the upper bound, i.e., let $u_i^{ub} = u_i(\lambda_i)$ for all i , where $[\lambda_i, i = 1, \dots, I]$ is the maximizer of the upper bound (2). Let \tilde{J}_s^c be the revenue for the c -scaled system under this static price. If the function F_i is concave in $(0, \lambda_{\max,i})$ for all i , then under the scaling **(S1)**,

$$\lim_{c \rightarrow \infty} \frac{1}{c} \tilde{J}_s^c = \lim_{c \rightarrow \infty} \frac{1}{c} J_{ub}^c = J_{ub}.$$

Proof: Since $\tilde{J}_s^c \leq J_s^c \leq J^{*,c} \leq J_{ub}^c = c J_{ub}$, in order to prove the above two propositions, we only need to show that $\lim_{c \rightarrow \infty} \tilde{J}_s^c / c = J_{ub}$.

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

$$\sum_{i=1}^I \frac{c \lambda_i(u_i) r_i C_i^l}{\mu_i} \leq c R^l \quad \text{for all } l, \quad (4)$$

let J_0^c denote the revenue under this static price. Since (4) guarantees that the condition of Lemma 3 is met, we have $\mathbf{P}_{loss,i}^c[\vec{u}] \rightarrow 0$, as $c \rightarrow \infty$. Therefore

$$\begin{aligned} \lim_{c \rightarrow \infty} \frac{J_0^c}{c} &= \lim_{c \rightarrow \infty} \sum_{i=1}^I \lambda_i(u_i) u_i \frac{1}{\mu_i} (1 - \mathbf{P}_{loss,i}^c[\vec{u}]) \\ &= \sum_{i=1}^I \lambda_i(u_i) u_i \frac{1}{\mu_i}. \end{aligned} \quad (5)$$

If we take the static price u_i^{ub} induced by the upper bound as our static price, then inequality (4) is satisfied. By definition, $J_0^c = \tilde{J}_s^c$, and the right hand side of (5) is exactly the upper bound. Therefore,

$$\lim_{c \rightarrow \infty} \frac{\tilde{J}_s^c}{c} = J_{ub},$$

and Propositions 4 and 5 then follow. \blacksquare

Proposition 2 and 4 are parallel to Theorem 6 and 7, respectively, in [8]. In [8], they are shown under the assumption of a single link and exponential holding time distribution. Proposition 2 and 4 tell us that extending the results of [8] 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 exist 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 (2) turns out to be indifferent to these additional mechanisms. Therefore, these extra mechanisms only have a minimal effect on the long term revenue when the system is large.

To see how fast the gap between the performance of the optimal dynamic scheme and that of the static schemes diminishes to zero as $c \rightarrow \infty$, note that

$$\begin{aligned} \frac{J^{*,c} - J_s^c}{J^{*,c}} &\leq \frac{J_{ub} - \frac{1}{c} J_0^c(\vec{u}^{ub})}{J_{ub}} \\ &= \frac{\sum_{i=1}^I \lambda_i(u_i^{ub}) u_i^{ub} \frac{1}{\mu_i} \mathbf{P}_{loss,i}^c[\vec{u}^{ub}]}{\sum_{i=1}^I \lambda_i(u_i^{ub}) u_i^{ub} \frac{1}{\mu_i}} \leq \max_i \mathbf{P}_{loss,i}^c[\vec{u}^{ub}]. \end{aligned}$$

Therefore, the speed of convergence of Propositions 4 and 5 can be determined by the $\mathbf{P}_{loss,i}^c[\vec{u}^{ub}]$ that has the slowest speed of convergence to zero. If at the price u_i^{ub} induced by the upper bound, the load of all links is strictly less than 1, then the convergence of Propositions 4 and 5 is exponential:

$$\limsup_{c \rightarrow \infty} \frac{1}{c} \log \frac{J^{*,c} - J_s^c}{J^{*,c}} \leq \max_l \inf_{w>0} \Lambda_l(w) < 0.$$

On the other hand, if at the price u_i^{ub} the load of some links is equal to 1, then the convergence is $1/\sqrt{c}$, i.e.,

$$\frac{J^{*,c} - J_s^c}{J^{*,c}} = O\left(\frac{1}{\sqrt{c}}\right).$$

Our result is different from that of [8] and [9] in the following aspects: Firstly, we remove the assumption on the exponential holding time distributions. Secondly, we characterize two different regions for the speed with which the performance of the optimal static scheme approaches that of the optimal dynamic scheme. When the load on some links is equal to 1, the convergence is $1/\sqrt{c}$. When the load of all links is strictly less than 1, the convergence is exponential. In [8] and [9], only the latter region is characterized.

Thirdly, and more importantly, we show in Proposition 5 that the static price induced by the upper bound is by itself asymptotically optimal. Hence, this result permits us to use the price induced by the upper bound directly in the static schemes. In [8] and [9], although the authors show the asymptotic optimality of the *optimal static scheme*, it is not clear whether the price induced by the upper bound can serve as a viable alternative, because with this price the load at some links could be equal to 1, and the convergence in this region is not characterized in their work.

D. Distance Neutral Pricing

When prices are congestion-dependent, two classes that traverse different routes will typically be charged different




Fig. 1. The network topology




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.

prices, even if they have the same price-elasticity function. A class that traverses a longer distance, or one that traverses a more congested route, will likely be charged a higher price.

However, if the following condition is satisfied, this distance-dependence can be eliminated when we adopt an asymptotically optimal static pricing scheme: We say a network has no *significant constraint* of resources if the *unconstrained* maximizer of $\sum_{i=1}^I F_i(\lambda_i)$ satisfies the constraint in (2). If there is no *significant constraint* of resources, there exists an asymptotically optimal static scheme whose prices depend only on the price-elasticity of each class, and are independent of their routes. To see this, we go back to the formulation of the upper bound (2). If the *unconstrained* maximizer of $\sum_{i=1}^I F_i(\lambda_i)$ satisfies the constraint, then it is also the maximizer of the *constrained* problem. In this case, if we use the prices induced by the upper bound as the asymptotically optimal static prices, the static prices will depend only on the function F_i , which represents the price-elasticity of the users.

This result suggests that the use of flat pricing, as in interstate long distance telephone service in the United States, can also be economically near-optimal under appropriate conditions. Assuming that there is no significant constraint of resources in the domestic telephone network in the U.S., and all consumers have the same price-elasticity, then our result indicates that a flat (i.e., independent of both time and distance) pricing scheme will suffice, given that the capacity of the network is very large.

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 μ_i	0.002	0.001	0.002	0.001
Bandwidth r_i	2	1	1	2

TABLE II
SOLUTION OF THE UPPER BOUND (2) 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

E. Numerical Results

We report a few numerical results here. 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

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

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 μ_i , and bandwidth requirement r_i for each class are given in Table I.

We first consider a base system where the capacity of the five links are 10, 10, 5, 15, and 15 bandwidth units, respectively. The solution of the upper bound (2) 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

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

difference in performance between the static pricing scheme and the optimal dynamic scheme is further reduced. 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 next 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 $J_{ub} = 137.5$. Simulation using the static prices induced by the upper bound shows similar curves as in Fig. 2. This latter example also demonstrates distance-neutral pricing. For example, classes 1 and 2, and classes 3 and 4 have different routes but have the same price (and price-elasticity). Readers can verify that, in this example, if we lift the constraints in (2), and solve the upper bound again, we will get the same maximizer. Thus there is no *significant constraint* of resources, and hence the optimal price 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.

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. 3 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 in Fig. 2, 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.

F. Scaling with Topological Changes

In Section II-C, we have studied a large network under the scaling (S1), where the network topology is fixed and we increase both the demand function $\lambda_i(u_i)$ and the capacity R^l proportionally. Hence, when the network is large, the number of users of each class i (i.e., between each source-destination pair) is also assumed to be large. This scaling is suitable for a *large capacity* network with a *simple* topology, e.g., network backbones. In this section, we will consider large networks where the number of users over each source-destination pair is *much smaller than the capacity of the network*. Large networks of this type arise naturally when the network topology becomes increasingly complex, for instance, when we include




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

the access links into the topology. To illustrate a real world example, note that even though the links between Purdue University, Indiana and Columbia University, New York could have large capacities, the number of users communicating *between these two institutions* at any time is usually quite low. The high-capacity access link that connects Purdue University to the Internet is to accommodate the large aggregate traffic between Purdue University and *all* destinations on the Internet, while the amount of traffic to *any single* destination is much smaller.

The question we attempt to answer in this section is: will similar simplicity results as in Section II-C hold in this type of large networks? We will first present a different scaling model that allows the topology of the network to become more complex as we scale its capacity. We will then show that, under some reasonable assumptions, the performance of appropriately chosen static schemes will still approach that of the optimal dynamic schemes, as long as *the capacity of the network is large*.

We consider a series of networks indexed by the scaling factor c . We use $L(c)$, $I(c)$, and $C(c)$ to denote the number of links, the number of classes, and the routing matrix, respectively, in the c -th network. Note that in the scaling (S1) in Section II-C, these quantities are assumed to be fixed for all c . In this section, in order to accommodate topological changes, we allow them to vary with c . We still use $\lambda_i^c(u)$ and $R^{l,c}$ to denote the demand function for class i and the capacity of link l , respectively, in the c -th system. However, unlike the case for the scaling (S1), $\lambda_i^c(u)$ and $R^{l,c}$ do not need to follow the linear scaling rule.

The *scale* of the c -th network is defined as

$$S(c) = \min_{l=1,\dots,L(c)} \frac{R^{l,c}}{\max_{i:C_i^c=1} r_i^c}.$$

We will consider the scaling (S2) that satisfies the following set of assumptions:

- (A) $S(c) \rightarrow \infty$, as $c \rightarrow \infty$.
- (B) For all networks, the maximum number of links on any route is bounded from above by a number M .

The first assumption simply states that, with large c , the capacity of each link in the network will be large compared to

the bandwidth requirement of each flow that goes through the link. The second assumption limits the maximum number of hops each flow can traverse. This is a reasonable assumption: for example, in TCP/IP, the TTL (time-to-live) field in the IP header occupies only 8 bits. This effectively put an upper bound of 255 on the number of hops a flow can traverse within a network. Real Internet topology exhibits the small world phenomenon [16], where the average number of hops of a source-destination pair is typically small. Any series of networks under scaling **(S1)** trivially satisfy scaling **(S2)**.

Finally, let $\lambda_{\max,i}^c$ be the maximal value of $\lambda_i^c(u)$. Let $u_i^c(\lambda_i)$ be the inverse function of $\lambda_i^c(u_i)$. We assume that:

- (C) The function $F_i^c(\lambda_i) = \lambda_i u_i^c(\lambda_i)$ is concave in $(0, \lambda_{\max,i}^c)$ for all c and i .

Let $J^{*,c}$, J_s^c , and J_{ub}^c be the optimal dynamic revenue, the optimal static revenue, and the upper bound, respectively, for the c -th network. By Proposition 2, under Assumption C, the optimal dynamic revenue $J^{*,c}$ is no greater than J_{ub}^c for all c . In general, J_{ub}^c does not follow the simple linear rule as in scaling **(S1)**. Hence, Proposition 4 will not hold under scaling **(S2)**. Instead, we can show the following:

Proposition 6: Under Scaling **(S2)** and Assumption C, as $c \rightarrow \infty$, the relative difference among the optimal dynamic revenue $J^{*,c}$, the optimal static revenue J_s^c , and the upper bound J_{ub}^c will converges to zero. That is,

$$\lim_{k \rightarrow \infty} \frac{J_{ub}^c - J_s^c}{J_{ub}^c} = \lim_{k \rightarrow \infty} \frac{J_{ub}^c - J^{*,c}}{J_{ub}^c} = 0.$$

The proof is provided in the Appendix. Proposition 6 tells us that no matter how complex the topology of the network is, as long as the capacity of the network is large, the performance of the optimal static scheme will be close to that of the optimal dynamic scheme. We defer relevant numerical results until Section III. The definition of $S(c)$ allows us to apply this result to certain networks with heterogeneous capacity. The network could have both large-capacity links and small-capacity links, and both large-bandwidth flows and small-bandwidth flows. As long as the capacity of each link is large compared to the bandwidth requirement of the flows *that traverse the link*, static schemes will suffice. On the other hand, there may exist network scenarios where $S(c)$ is not large, which means that the capacity of some links can only accommodate a small number of flows that go through them. Proposition 6 will not hold in such scenarios.

It is easy to check that Proposition 4 and Proposition 6 are equivalent under scaling **(S1)**. There is yet a difference between the results we can obtain under the two different types of scaling. Under scaling **(S1)**, we can show that the price \bar{u}^{ub} induced by the upper bound suffices to be the near-optimal static price (see Proposition 5). Such conclusion cannot be drawn under scaling **(S2)**. The difficulty is that, under scaling **(S2)**, we cannot show that the blocking probability $\mathbf{P}_{loss,i}^c[\bar{u}^{ub}]$ at the static price \bar{u}^{ub} goes to zero as $c \rightarrow \infty$ *when the constraints in the upper bound (2) are satisfied with equality*. In spite of this technical difficulty, we expect that in most cases the price induced by the upper bound would still suffice under scaling **(S2)**. We can argue as follows using the familiar *independent blocking assumption* underlying

the Erlang Fixed Point approximation [17], i.e., we assume that blocking occurs independently at each link. Under the independent blocking assumption, the blocking probability B^l at each link l can be calculated *as if the traffic offered to the link l comprises independent Poisson streams with arrival rates that are “thinned” by other links in the network*. At the price induced by the upper bound, the offered load at any link l *after “thinning”* will be no greater than the offered load *before “thinning”*, i.e., $\sum_i \frac{\lambda_i^c}{\mu_i} r_i$, which is no greater than the capacity $R^{l,c}$ of the link l . Applying the techniques in the proof of Lemma 3 to a single link, we can show that, *if the independent blocking assumption holds*, the blocking probability at each link will go to zero as $S(c) \rightarrow \infty$, and the convergence is uniform over all links. Since the number of hops each route can traverse is upper bounded by M , we can then infer that the blocking probability of all classes will go to zero uniformly as $S(c) \rightarrow \infty$. Therefore, using an argument similar to that of Proposition 4, we can infer that the price induced by the upper bound will suffice. The validity of the above argument relies on the independent blocking assumptions. A rigorous characterization of this convergence is an interesting problem for future work.

G. Remarks on Non-stationary Scenarios

We have shown under two different types of scaling that the performance of an appropriately chosen static pricing scheme will approach that of the optimal dynamic pricing scheme when the capacity of the network is large. We conclude this section with some discussion on the assumption we have made. In our model, we assume that the demand function $\lambda_i(u_i)$, the mean service time $1/\mu_i$ and the capacity R^l are fixed. Hence, we have assumed that the network condition is stationary. However, in practice the network condition is usually non-stationary and even the average network statistics can change over time. If we assume that the changes of these average network conditions are at a much slower time scale than that of the flow arrivals and departures, we can still use the result in this section with the following interpretation: the static scheme should be interpreted as *prices being fixed over a time period for which the network statistics do not change* (which is typically a fairly long time in real networks) and the average network condition should be interpreted as *the average of the dominant network condition over such a time period*. Similarly, the invariance result regarding prediction should also be interpreted under this context: if the dynamic pricing scheme uses a prediction window (i.e., d in Equation (1)) that is smaller than the time scale of the changes of the average network condition, then the dynamic scheme will not significantly outperform the static pricing scheme. Over the longer time scale, the prices should still adapt to the changes in the dominant network condition and prediction on the changes of the average network statistics will help.

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 for dynamic routing, the network is typically required to first collect link states (such as available bandwidth, 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” routing policies 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 this type of routing policies 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 quality of the routing decision.

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 types of routing policies. We will first establish an upper bound on the performance of all *dynamic schemes*, which are schemes that can compute both the prices and the routing decisions based on the instantaneous congestion level of the network. We will then show that the performance of an appropriate chosen *static scheme*, which uses static prices and 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 the 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 dynamic schemes to “pack” more flows into the system. Yet, we will show that an appropriately chosen static 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

$$\begin{aligned} \text{There exists } P_{ij}^k \text{ such that } & \sum_{j=1}^{\theta(i)} P_{ij}^k = 1, \forall i, k, \\ \text{and } & \sum_{i=1}^I \sum_{j=1}^{\theta(i)} r_i H_{ij}^l \sum_{k=1}^{n_i} P_{ij}^k \leq R^l \quad \text{for all } l. \end{aligned} \quad (6)$$

The set of feasible states is $\Omega = \{\vec{n} \text{ such that (6) is satisfied}\}$.

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 Ω . Once the flow is admitted, its route (i.e., P_{ij}^k) is assigned based on (6), 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 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 scheme* is given by:

$$J^* \triangleq \max_{\vec{g}} \mathbf{E} \left\{ \sum_{i=1}^I \lambda_i(u_i(t)) u_i(t) \frac{1}{\mu_i} \right\} \quad (7)$$

subject to (6).

The expectation is taken with respect to the steady state distribution. Note that (7) 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 do not advocate 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 (and idealized) dynamic scheme. The static schemes do not require the aforementioned 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 2, we can derive the following upper bound on the optimal revenue in (7). The proof is a natural extension of that of Proposition 2 and are available online in [15].

Proposition 7: 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:

$$\begin{aligned} J_{ub} \triangleq \max_{\lambda_{ij}} & \sum_{i=1}^I F_i \left(\sum_{j=1}^{\theta(i)} \lambda_{ij} \right) \frac{1}{\mu_i} \\ \text{subject to} & \sum_{i=1}^I \sum_{j=1}^{\theta(i)} \frac{\lambda_{ij}}{\mu_i} H_{ij}^l r_i \leq R^l \quad \forall l. \end{aligned} \quad (8)$$

We next construct our static scheme *using probabilistic routing* 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 scheme 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 scheme, 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 . The performance of the *optimal static*

scheme is then given by:

$$J_s \triangleq \max_{u_i^s, P_{ij}^s} \sum_{i=1}^I \sum_{j=1}^{\theta(i)} \lambda_i(u_i^s) u_i^s P_{ij}^s \frac{1}{\mu_i} [1 - \mathbf{P}_{Loss,ij}], \quad (9)$$

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

We consider a special static scheme derived from the solution of the upper bound in Proposition 7. If λ_{ij}^{ub} is the maximal solution to the upper bound, we let $u_i^s = u_i(\sum_{j=1}^{\theta(i)} \lambda_{ij}^{ub})$,

and $P_{ij}^s = \frac{\lambda_{ij}^{ub}}{\sum_{j=1}^{\theta(i)} \lambda_{ij}^{ub}}$. The revenue with this static scheme 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 (S1), we can show that, as $c \rightarrow \infty$, $\mathbf{P}_{loss,ij} \rightarrow 0$. Therefore, we have our invariance result (stated next).

Proposition 8: In the dynamic routing model, if the function F_i is concave in $(0, \lambda_{\max,i})$ for all i , then under the scaling (S1),

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

Proof: Analogous to that of Propositions 4 and 5. Details are available online at [15]. ■

A similar result under scaling (S2) can be stated analogous to Proposition 6.

When the routing is fixed, by replacing λ_{ij} with λ_i , and H_{ij}^l with C_i^l , we recover Propositions 2 and 4 from the results in this section. When there are multiple alternate routes, the upper bound in Proposition 7 is typically larger than that of Proposition 2. Therefore one can indeed improve revenue by employing dynamic routing. However, Proposition 8 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 asymptotically optimal static scheme also reveals the macroscopic structure of the optimal dynamic 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.




Fig. 4. Dynamic Routing Problem: There are 3 classes of flows, AB , BC , CA . For each class, there are two alternate routes. For example, for class AB , the direct one-link path is $A \rightarrow B$, while the indirect two-link path is $A \rightarrow C \rightarrow B$.

We use the following examples to illustrate the results in this section. We first consider a triangular network (Fig. 4). There are three classes of flows, AB , BC , CA . There are two possible routes for each class of calls, i.e., a direct one-link path (route 1), and an indirect two-link path (route 2). Each call consumes one bandwidth unit along the link(s) and holds the link(s) for a mean time of 1 unit. Let the capacity of all links be R .

Let $\vec{\lambda} = \{\lambda_{AB,1}, \lambda_{AB,2}, \lambda_{BC,1}, \lambda_{BC,2}, \lambda_{CA,1}, \lambda_{CA,2}\}$. By Proposition 7, we can formulate the upper bound as (based on (8)):

$$J_{ub} = \max_{\vec{\lambda}} \sum_{i=AB,BC,CA} \lambda_i u_i(\lambda_i) \quad (10)$$

$$\lambda_i = \lambda_{i,1} + \lambda_{i,2}, \quad i = AB, BC, CA,$$

subject to the following resource constraints:

$$\begin{aligned} \lambda_{AB,1} + \lambda_{BC,2} + \lambda_{CA,2} &\leq R \\ \lambda_{BC,1} + \lambda_{AB,2} + \lambda_{CA,2} &\leq R \\ \lambda_{CA,1} + \lambda_{AB,2} + \lambda_{BC,2} &\leq R \\ \lambda_{i,j} &\geq 0, \quad i = AB, BC, CA, j = 1, 2. \end{aligned}$$

Once the upper bound is solved, we can find the near-optimal static scheme using the one induced by the upper bound, i.e., the price charged to class i , $i = AB, BC, CA$, is

$$u_i^s = u_i(\lambda_{i,1} + \lambda_{i,2}),$$

and the bifurcation probabilities are

$$P_{i,1}^s = \frac{\lambda_{i,1}}{\lambda_{i,1} + \lambda_{i,2}}$$

$$P_{i,2}^s = \frac{\lambda_{i,2}}{\lambda_{i,1} + \lambda_{i,2}}.$$

Let $R = 100$. We consider the following examples:

1) When the price-elasticity function of all classes are $\lambda_{AB}(u) = \lambda_{BC}(u) = \lambda_{CA}(u) = 100(1 - u)$, the solution of (10) gives $\lambda_{AB} = \lambda_{BC} = \lambda_{CA} = 50$, and the prices for each class are $u_{AB}^s = u_{BC}^s = u_{CA}^s = 0.5$. It also coincides with the solution of the unconstrained version of (10). This corresponds to the case of light traffic load. The price is only determined by the price-elasticity of each class.

TABLE IV

SOLUTION OF THE UPPER BOUND IN THE DYNAMIC ROUTING PROBLEM:
WHEN THE PRICE-ELASTICITY OF CLASS AB IS $\lambda_{AB}(u) = 500(1 - u)$

	Class AB	Class BC	Class CA
$\lambda_{i,1}$	100	40.91	40.91
$\lambda_{i,2}$	59.09	0	0
λ_i	159.09	40.91	40.91
Price u_i^s	0.682	0.591	0.591
$P_{i,1}^s$	62.86%	100%	100%
$P_{i,2}^s$	37.14%	0	0

There are multiple solutions for the bifurcation. One example is $\lambda_{i,1} = 50$, and $\lambda_{i,2} = 0$, $i = AB, BC, CA$, i.e., all calls use the direct link.

2) When we change the price-elasticity of class AB to $\lambda_{AB}(u) = 500(1 - u)$, the solution of (10) is shown in Table IV. This corresponds to the case of heavy traffic load. The price are raised from that of the unconstrained problem in order to limit incoming traffic. All constraints are binding. Note that here in order to maximize the revenue, class AB has a higher arrival rate λ_i than that of class BC and CA , and the network should allow flows from class AB to use indirect two-link path, while flows from classes BC and CA should not be allowed to use the indirect routes.

We next use a larger network example to demonstrate the optimality of static schemes with probabilistic routing. We use the BRITE topology generation tool [18] and the Barabasi-Albert model [19] to generate a random network with 100 nodes and 197 links. The Barabasi-Albert topology model is able to capture the power-law of node-connectivity in real Internet topologies. There are a total of 9900 source-destination (s-d) pairs. For each s-d pair, we use the set of minimum-hop paths as the alternate paths. The demand function for each source-destination pair is the same. We formulate the upper bound for the randomly-generated network, and use standard convex optimization methods to solve the static prices and the bifurcation probabilities. We then use simulation to obtain the static revenue under the prices and the bifurcation probabilities induced by the upper bound.

We present the following result from a typical simulation. The bandwidth for each link is 1000 units. The bandwidth requirement of each flow is 1 unit and the mean holding time is 1 unit. The demand function for each s-d pair is $\lambda(u) = 10(1 - u)$. At this level of demand, around 22% of the links experience congestion, i.e., their respective constraints in (8) are binding. The upper bound is found to be 2.45×10^4 , while the static revenue obtained from the simulation is 2.40×10^4 . The relative difference is just 2%. This validates our result that the performance of static schemes is close to that of the optimal dynamic scheme and the upper bound. Among the 9900 s-d pairs, around 49% have multiple alternate routes. However, among those with multiple routes, only 30% actually use multiple routes. Hence, a large number of s-d pairs does not benefit from multiple alternative paths. The average number of routes between a s-d pair is 2.15. Finally, note that in this example, the end-to-end demand of each source-destination pair is at most $10(1 - 0)/1 = 10$, which is much

smaller than the capacity of the links (1000 units). Hence, this simulation serves to validate our result under scaling (S2) where the number of users of each source-destination pair is much smaller than the capacity of the network. In all of our simulations, the prices induced by the upper bound turn out to be near-optimal, even although we have not been able to establish this optimality rigorously under scaling (S2). We have also run simulations using other topology models and find similar results.

IV. CONCLUSION, DISCUSSION, AND FUTURE WORK

In this work, we have studied pricing as a mechanism to control large networks. We have shown under general settings that the performance of an appropriately chosen *static scheme* can approach that of the *optimal dynamic scheme* when the system is large. These results have important implications for the design and control of large-capacity networks. Compared with the optimal dynamic scheme, the static scheme has several desirable features. The static schemes are much easier to obtain because of their simple structures. They are also much easier to execute since they do not require the collection of instantaneous load information. Instead, they only depend on some average parameters, such as the average load. Hence, they introduce less computation and communication overhead, and they are less sensitive to feedback delay. These advantages make the static scheme an attractive alternative for controlling large networks.

However, one should keep in mind that static schemes also have their disadvantage, namely, their lack of adaptivity. Static schemes could be more sensitive to modeling errors than dynamic schemes [20], [21]. If the parameters of the model are estimated incorrectly, the resulting static scheme may lead to bad performance. Further, as we discussed at the end of Section II, the network condition may be non-stationary and even the average network parameters may change over time. The static prices that are good for one time may not be good for the next moment. Therefore, we are not advocating that in practice *purely* static schemes (i.e., prices being fixed for all time) be used.

Nonetheless, we believe that the results in this paper can be exploited to develop practical network control algorithms that are both *simple* and *adaptive*. One direction is to develop efficient algorithms that can compute the static prices based on the current dominant network condition and allow the prices to *adaptively* track the changes in the average network parameters. Here we briefly discuss one possible approach. Note that although the static prices are calculated by solving a *global* optimization problem, i.e., the upper bound (2) or (8), it is possible to develop a distributed solution. Indeed, we can associate a non-negative Lagrange multiplier p^l for the constraint 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 implicit 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 update these implicit costs based on the congestion level at each link and have the ingress router serve as “brokers” to probe these implicit costs and determine the price offered to users of each class i . The idea of this decomposition has been used in [6] and [7] to develop distributed algorithms for optimization flow control, and it is also mentioned in [8] for computing the static prices in the single-link case.

The distributed algorithm described above can achieve adaptivity in several ways. Firstly, the edge router can use the *online measurement* of flow arrivals at different price levels to update its estimate of the demand function. Secondly, the core router can use the *online measurement* of the congestion level at each link to update the implicit costs. We have recently developed such *quasi-static* adaptive schemes for QoS routing [22].

It is instructive to compare such a *quasi-static* distributed algorithm with typical dynamic and static schemes. Note that since the distributed algorithm updates the implicit costs based on online measurements of the congestion level at the link, it can also be viewed as a *dynamic scheme*. However, the distributed algorithm is based on the asymptotic optimality of the static schemes. It attempts to solve for the static prices according to the current dominant network condition. Hence, we refer to the distributed algorithm as being *quasi-static*. On one hand, the distributed algorithm exploits the simplicity of the static scheme, thus has a simple form and is easier to implement than the optimal dynamic scheme. When the network condition is stationary, the prices computed by the distributed algorithm will converge to that of the near-optimal static scheme. On the other hand, the distributed algorithm is by definition also dynamic in that, when the network condition is non-stationary, the prices computed by the distributed algorithm will *track* the long term changes. Hence, the distributed algorithm is more robust than *purely* static schemes.

As a final remark, we note that there are also possibilities of extremal changes in network conditions, such as failures of network components. When such situations occur, an appropriate *immediate* response is usually more important than an optimal but *slower* one. For such situations, other levels of network control, such as failure detection and fault recovery, are more appropriate than the pricing-based control studied in this paper.

APPENDIX

A. Proof of Lemma 3

The key idea is to use an insensitivity result from [23]. In [23], Burman *et. al.* investigate a blocking network model, where a call instantaneously seizes channels along a route between the originating and terminating node, holds the channels for a randomly distributed length of time, and frees them instantaneously at the end of the call. If no channels are available, the call is blocked. When the arrivals are Poisson and the holding time distributions are general, the authors in [23] show that the blocking probabilities are still in product form, and are insensitive to the call holding-time distributions.

This means that they depend on the call duration only through its mean.

When the static prices are given, our system is a special case of [23]. Hence, we can reuse results for loss networks with Poisson arrivals and exponential holding-times. If we assume that the bandwidth requirements r_i are integers with greatest common divisor being 1, an upper bound on the blocking probability for calls of class i is given by [24, Proposition 2.1]:

$$\mathbf{P}_{loss,i}^c \leq 2^L \sum_{l:C_i^l=1} \frac{e^{-cI_l(R^l)} (1 - e^{-r_i y_l})}{\sqrt{2\pi c} \Gamma_l^2} \frac{1 - e^{-r_i y_l}}{1 - e^{-y_l}} \left[1 + O\left(\frac{1}{\sqrt{c}}\right) \right], \quad (11)$$

where y_l is the unique solution of

$$R^l = \sum_{i=1}^I \frac{\lambda_i}{\mu_i} r_i e^{-y_l r_i} C_i^l, \quad (12)$$

and

$$\begin{aligned} \Gamma_l^2 &= \sum_{i=1}^I \frac{\lambda_i}{\mu_i} r_i^2 e^{-y_l r_i} C_i^l \\ I_l(R^l) &= \sum_{i=1}^I \frac{\lambda_i}{\mu_i} (1 - e^{-y_l r_i}) C_i^l - y_l R^l. \end{aligned}$$

From (12), it is easy to verify that $y_l < 0$ if $\sum_{i=1}^I \frac{\lambda_i}{\mu_i} r_i C_i^l < R^l$, and $y_l = 0$ if $\sum_{i=1}^I \frac{\lambda_i}{\mu_i} r_i C_i^l = R^l$. Note that $-y_l$ is also the minimizer of $\Lambda_l(w)$ in (3) over $w \geq 0$, and thus

$$I_l(R^l) = - \inf_{w>0} \Lambda_l(w).$$

Hence, if the load at each resource is less than or equal to 1, i.e.,

$$\sum_{i=1}^I \frac{\lambda_i}{\mu_i} r_i C_i^l \leq R^l \text{ for all } l,$$

then $y_l \leq 0$ and $I_l(R^l) \geq 0$ for all l . We have,

$$\mathbf{P}_{loss,i}^c \leq 2^L \sum_{l:C_i^l=1} \frac{1}{\sqrt{2\pi c} \Gamma_l^2} \frac{1 - e^{-r_i y_l}}{1 - e^{-y_l}} \left[1 + O\left(\frac{1}{\sqrt{c}}\right) \right],$$

i.e., $\mathbf{P}_{loss,i}^c = O\left(\frac{1}{\sqrt{c}}\right)$. On the other hand, when the load of all the links that class i traverses is strictly less than 1, then $y_l < 1$ and $I_l(R^l) > 0$ for all link l that class i traverses. Hence, the exponential terms in (11) dominate. We thus have,

$$\begin{aligned} \limsup_{c \rightarrow \infty} \frac{1}{c} \log \mathbf{P}_{loss,i}^c &\leq \max_{l:C_i^l=1} -I_l(R^l) \\ &= \max_{l:C_i^l=1} \inf_{w>0} \Lambda_l(w) < 0. \end{aligned}$$

The above techniques (and that of [24]) can easily be generalized to the case when the bandwidth requirements r_i are positive real numbers. For a more elementary proof that does not use the result of [24], see our work in [1] and [15].

B. Proof of Proposition 6

We first focus on the c -th network. To simplify notation, we will drop the index c when there is no source of confusion. Let $\lambda_i, i = 1, \dots, I$ denote the solution of the upper bound (2). Let ϵ be a positive real number smaller than 1. Let $\lambda_i^\epsilon = (1 - \epsilon)\lambda_i$ and $u_i^\epsilon = u_i(\lambda_i^\epsilon)$ for all i . Then the static revenue at static prices $u_i^\epsilon, i = 1, \dots, I$ is

$$\begin{aligned} J_s^{c,\epsilon} &= \sum_{i=1}^I u_i(\lambda_i^\epsilon) \lambda_i^\epsilon \frac{1}{\mu_i} (1 - \mathbb{P}_{loss,i}^\epsilon) \\ &\geq \sum_{i=1}^I (1 - \epsilon) u_i(\lambda_i) \lambda_i \frac{1}{\mu_i} (1 - \mathbb{P}_{loss,i}^\epsilon), \end{aligned}$$

where $\mathbb{P}_{loss,i}^\epsilon$ is the blocking probability of users of class i at static price $u_i^\epsilon, i = 1, \dots, I$, and we have used the property that $u_i(\cdot)$ is decreasing. Thus the relative difference between $J_s^{c,\epsilon}$ and J_{ub}^c is

$$\frac{J_{ub}^c - J_s^{c,\epsilon}}{J_{ub}^c} \leq \epsilon + \max_{i=1,\dots,I} \mathbb{P}_{loss,i}^\epsilon. \quad (13)$$

Next we estimate $\mathbb{P}_{loss,i}^\epsilon$. Let n_j be the random variable that represents the number of flows of class j that are in the system. We now consider another network with the same topology and the same demand λ_i^ϵ . However, each link in the new network has infinite capacity. Let n_j^∞ be the random variable that represents the number of flows of class j that are in the infinite capacity system. By a sample path argument, $n_j \leq n_j^\infty$. Therefore,

$$\begin{aligned} \mathbb{P}_{loss,i}^\epsilon &= \mathbb{P}\{\text{There exists } l \text{ such that } C_i^l = 1 \\ &\quad \text{and } \sum_{j=1}^I n_j r_j C_j^l \geq R^l - r_i\} \\ &\leq \sum_{l:C_i^l=1} \mathbb{P}\{\sum_{j=1}^I n_j r_j C_j^l \geq R^l - r_i\} \\ &\leq \sum_{l:C_i^l=1} \mathbb{P}\{\sum_{j=1}^I n_j^\infty r_j C_j^l \geq R^l - r_i\}. \quad (14) \end{aligned}$$

In the new system with infinite capacity, $n_j^\infty, j = 1, \dots, I$ are independent Poisson random variables (by well known $M/G/\infty$ results). We can calculate their moment generating functions as

$$\mathbb{E}[\exp(\theta n_j^\infty)] = \exp\left[\frac{\lambda_j^\epsilon}{\mu_j} (e^\theta - 1)\right] \text{ for } \theta > 0.$$

Fix i and l such that $C_i^l = 1$. By invoking Markov Inequality, we have,

$$\begin{aligned} \mathbb{P}\left\{\sum_{j=1}^I n_j^\infty r_j C_j^l \geq R^l - r_i\right\} &\leq \frac{\mathbb{E}[\exp(\sum_{j=1}^I \theta r_j n_j^\infty C_j^l)]}{\exp[\theta(R^l - r_i)]} \\ &= \exp\left[\sum_{j=1}^I \frac{\lambda_j^\epsilon}{\mu_j} C_j^l (e^{\theta r_j} - 1) - \theta(R^l - r_i)\right] \\ &\leq \exp\left[\sum_{j=1}^I \frac{\lambda_j^\epsilon}{\mu_j} r_j C_j^l \frac{S(c)}{R^l} (e^{\theta R^l/S(c)} - 1) - \theta(R^l - r_i)\right] \end{aligned}$$

$$\leq \exp[(1 - \epsilon)S(c)(e^{\frac{\theta R^l}{S(c)}} - 1) - \frac{\theta R^l}{S(c)}(S(c) - 1)], \quad (15)$$

where in the last two inequalities we have used the definition of the *scale* $S(c)$ such that, for all j with $C_j^l = 1$, we have

$$r_j \leq \frac{R^l}{S(c)},$$

and thus,

$$\frac{e^{\theta r_j} - 1}{r_j} \leq \frac{e^{\theta R^l/S(c)} - 1}{R^l/S(c)};$$

and we have used the assumption that,

$$\sum_{j=1}^I \frac{\lambda_j^\epsilon}{\mu_j} r_j C_j^l = \sum_{j=1}^I (1 - \epsilon) \frac{\lambda_j}{\mu_j} r_j C_j^l \leq (1 - \epsilon)R^l.$$

Taking infimum of (15) over all $\theta > 0$, we have

$$\mathbb{P}\left\{\sum_{j=1}^I n_j^\infty r_j C_j^l \geq R^l - r_i\right\} \leq \exp[f(S(c))], \quad (16)$$

where

$$f(s) = \inf_{\theta > 0} [(1 - \epsilon)s(e^\theta - 1) - \theta(s - 1)]. \quad (17)$$

Note that the inequality (16) holds for all i and l such that $C_i^l = 1$. Substituting (16) into (14), we have,

$$\mathbb{P}_{loss,i}^\epsilon \leq \sum_{l:C_i^l=1} \exp[f(S(c))] \leq M \exp[f(S(c))], \quad (18)$$

where M is the maximum number of hops for all routes by Assumption B of scaling **(S2)**. Note that the right hand side is uniform for all class i . Substituting (18) into (13), we have,

$$\frac{J_{ub}^c - J_s^{c,\epsilon}}{J_{ub}^c} \leq \epsilon + M \exp[f(S(c))].$$

The function $f(s)$ can be evaluated analytically. We can easily show that

$$\lim_{s \rightarrow +\infty} \frac{f(s)}{s} = \epsilon + \ln(1 - \epsilon).$$

Since $\epsilon \in (0, 1)$, $\epsilon + \ln(1 - \epsilon) < 0$. Hence, $\lim_{s \rightarrow \infty} f(s) = -\infty$. Now by Assumption A of scaling **(S2)**, $S(c) \rightarrow \infty$ as $c \rightarrow \infty$. Fix ϵ and let $c \rightarrow \infty$, we have

$$\lim_{c \rightarrow \infty} \frac{J_{ub}^c - J_s^{c,\epsilon}}{J_{ub}^c} \leq \lim_{c \rightarrow \infty} \{\epsilon + M \exp[f(S(c))]\} = \epsilon.$$

Note that $J_s^{c,\epsilon}$ is always no greater than the optimal static revenue J_s^c . Hence,

$$\lim_{c \rightarrow \infty} \frac{J_{ub}^c - J_s^c}{J_{ub}^c} \leq \epsilon.$$

This holds for any $\epsilon \in (0, 1)$. Letting $\epsilon \rightarrow 0$ and noting that $J_s^c \leq J^{*,c} \leq J_{ub}^c$, the result then follows.

ACKNOWLEDGMENT

We thank the anonymous reviewers for their insightful comments and suggestions.

REFERENCES

- [1] X. Lin and N. B. Shroff, "Pricing-Based Control of Large Networks," in *Evolutionary Trends of the Internet, Proceedings of Tyrrhenian International Workshop on Digital Communications (IWDC 2001)*, Taormina, Italy, September 2001, pp. 212–231.
- [2] X. Lin and N. B. Shroff, "Simplification of Network Dynamics in Large Systems," in *Tenth International Workshop on Quality of Service (IWQoS 2002)*, Miami Beach, FL, May 2002.
- [3] 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.
- [4] 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.
- [5] F. P. Kelly, A. Maulloo, and D. Tan, "Rate Control in Communication Networks: Shadow Prices, Proportional Fairness and Stability," *Journal of the Operational Research Society*, vol. 49, pp. 237–252, 1998.
- [6] 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, December 1999.
- [7] 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.
- [8] 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.
- [9] I. Ch. Paschalidis and Y. Liu, "Pricing in Multiservice Loss Networks: Static Pricing, Asymptotic Optimality, and Demand Substitution Effects," *IEEE/ACM Transactions on Networking*, vol. 10, no. 3, pp. 425–438, June 2002.
- [10] 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.
- [11] P. B. Key, "Optimal Control and Trunk Reservation in Loss Networks," *Probability in the Engineering and Informational Sciences*, vol. 4, pp. 203–242, 1990.
- [12] T. Basar and R. Srikant, "Revenue-Maximizing Pricing and Capacity Expansion in a Many-Users Regime," in *Proceedings of IEEE INFOCOM*, New York, New York, June 2002, pp. 1556–1563.
- [13] P. Tinnakornsrisuphap and A. M. Makowski, "TCP Traffic Modeling via Limit Theorems," *Technical Report, Institute for Systems Research, University of Maryland*, 2002.
- [14] A. A. Borovkov, "Ergodicity and Stability Theorems for a Class of Stochastic Equations and Their Applications," *Theory of Probability and Its Applications*, vol. 23, no. 2, pp. 227–247, June 1978.
- [15] X. Lin and N. B. Shroff, "Simplification of Network Dynamics in Large Systems," *Technical Report, Purdue University*, <http://min.ecn.purdue.edu/~linx/papers.html>, 2001.
- [16] S. Jin and A. Bestavros, "Small-World Internet Topologies: Possible Causes and Implications on Scalability of End-System Multicast," *Technical Report BUCS-TR-2002-004, Boston University, Boston, MA*, 2002.
- [17] F. P. Kelly, "Loss Networks," *Annals of Applied Probability*, vol. 1, pp. 319–378, 1991.
- [18] A. Medina, A. Lakhina, I. Matta, and J. Byers, "Brite: Universal Topology Generation From a User's Perspective," in *Proceedings of the International Workshop on Modeling, Analysis and Simulation of Computer and Telecommunications Systems (MASCOTS '01)*, Cincinnati, Ohio, August 2001.
- [19] A. Barabasi and R. Albert, "Emergence of Scaling in Random Networks," *Science*, vol. 286, pp. 509–512, 1999.
- [20] P. Marbach, O. Mihatsch, and J. N. Tsitsiklis, "Call Admission Control and Routing in Integrated Service Networks Using Neuro-Dynamic Programming," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 2, pp. 197–208, February 2000.
- [21] P. Whittle, "Approximation in Large-Scale Circuit-Switched Networks," *Prob. Engrn. Inf. Sci.*, vol. 2, pp. 279–291, 1988.
- [22] X. Lin and N. B. Shroff, "An Optimization Based Approach for Quality of Service Routing in High-Bandwidth Networks," in *Proceedings of IEEE INFOCOM*, Hong Kong, China, March 2004.
- [23] D. Y. Burman, J. P. Lehoczy, and Y. Lim, "Insensitivity of Blocking Probabilities in a Circuit-Switching Network," *Journal of Applied Probability*, vol. 21, pp. 850–859, 1984.
- [24] F. Theberge, A. Simonian, and R. R. Mazumdar, "Upper Bounds for Blocking Probabilities in Large Multi-rate Loss Networks," *Telecommunication Systems*, vol. 9, pp. 23–39, 1998.



Xiaojun Lin (S'02) received his B.S. in Electronics and Information Systems from Zhongshan University, Guangzhou, China, in 1994, and his M.S. in Electrical and Computer Engineering from Purdue University, West Lafayette, IN, in 2000. From 1995 to 1998, he worked as an I/T Specialist in IBM China. In 2002, he spent three months in IBM T.J. Watson Research Labs. Since 2000, he has been working towards the Ph.D. degree in Electrical and Computer Engineering at Purdue University, West Lafayette, IN.

His research interests are resource allocation, optimization, network pricing, routing, congestion control, network as a large system, cross-layer design in wireless networks, mobile ad hoc and sensor networks. He received the 2005 best paper of the year award from *Journal of Communications and Networks*. His paper was also one of the two runner-up papers for the best-paper award at IEEE INFOCOM 2005.



Ness B. Shroff (S'91 / M'93 / SM'01) received his Ph.D. degree from Columbia University, NY in 1994. He has been at Purdue University since 1994, and is currently a Professor of Electrical and Computer Engineering at Purdue.

His research interests span the areas of wireless and wireline communication networks. He is especially interested in fundamental problems in the design, performance, control, security, and pricing of these networks.

Dr. Shroff is an editor for the IEEE/ACM Trans. on Networking and the Computer Networks Journal, and past editor of IEEE Communications Letters. He was the Technical Program co-chair for IEEE INFOCOM'03 (San Francisco, CA), the panel co-chair for ACM Mobicom'02 (Atlanta, GA), program co-chair for the symposium on high-speed networks, Globecom 2001 (San Francisco, CA), and conference chair for the 14th Annual IEEE Computer Communications Workshop (Estes Park, CO). He was the co-organizer of the NSF Workshop on "Fundamental Research in Networking" in April 2003. He has received numerous awards including the NSF Career award in 1996, the *Computer Network Journal's* best paper award of 2003, and the 2005 best paper of the year award from *Journal of Communications and Networks*.