

Service Overlay Networks: SLAs, QoS and Bandwidth Provisioning

Zhenhai Duan[†], Zhi-Li Zhang[†], and Yiwei Thomas Hou[‡]

[†] Dept. of Computer Science & Engineering
University of Minnesota
Minneapolis, MN 55455
{duan, zhzhang}@cs.umn.edu

[‡] Bradley Dept. of Electrical & Computer Engineering
Virginia Tech
Blacksburg, VA 24061
thou@vt.edu

Abstract

We advocate the notion of service overlay network (SON) as an effective means to address some of the issues, in particular, end-to-end QoS, plaguing the current Internet, and to facilitate the creation and deployment of value-added Internet services such as VoIP, Video-on-Demand, and other emerging QoS-sensitive services. A SON purchases bandwidth with certain QoS guarantees from individual network domains via bilateral service level agreement (SLA) to build a logical end-to-end service delivery infrastructure on top of existing data transport networks. Via a service contract, users directly pay the SON provider for using the value-added services provided by the SON.

In this paper we study the bandwidth provisioning problem for a service overlay network which is critical to the cost recovery in deploying and operating value-added services over the SON. We mathematically formulate the bandwidth provisioning problem, taking into account various factors such as SLA, service QoS, traffic demand distributions, and bandwidth costs. Analytical models and approximate solutions are developed for both static and dynamic bandwidth provisioning. Numerical studies are also performed to illustrate the properties of the proposed solutions and demonstrate the effect of traffic demand distributions and bandwidth costs on the bandwidth provisioning of a SON.

1 Introduction

Today's Internet infrastructure supports primarily *best-effort connectivity* service. Due to historical reasons, the Internet consists of a collection of network domains (i.e., autonomous systems owned by various administrative entities). Traffic from one user to another user typically traverses multiple domains; network domains enter various bilateral business relationships (e.g., provider-customer, or peering) for traffic exchange to achieve global connectivity. Due to the

nature of their business relationships, each network domain is only concerned with the network performance of its own domain and responsible for providing service guarantees for its customers. As it is difficult to establish multi-lateral business relationship involving multiple domains, the deployment of end-to-end services beyond the best-effort connectivity that requires support from multiple network domains is still far from reality. Such problems have hindered the transformation of the current Internet into a truly multi-service network infrastructure with end-to-end QoS support.

We propose and advocate the notion of *service overlay network* (SON) as an effective means to address some of the issues, in particular, end-to-end QoS, plaguing the current Internet, and to facilitate the creation and deployment of *value-added Internet services* such as VoIP, Video-on-Demand, and other emerging QoS-sensitive services. The network architecture of a SON relies on well-defined business relationships between the SON, the underlying network domains and users of the SON to provide support for end-to-end QoS: the SON purchases bandwidth with certain QoS guarantees from individual network domains via *bilateral service level agreement* (SLA) to build a logical end-to-end service delivery infrastructure on top of existing data transport networks; via a service contract (e.g., a usage-based or fixed price service plan), users¹ directly pay a SON provider for using the value-added services provided by the SON.

Figure 1 illustrates the SON architecture. A SON is pieced together via *service gateways* which perform service-specific data forwarding and control functions. The *logical* connection between two service gateways is provided by the underlying network domain with certain bandwidth and other QoS guarantees. These guarantees are specified in a bilateral SLA between the SON and the network domain. This architecture bypasses the peering points among the network domains, and thus avoids potential performance problems associated with them. Relying on the bilateral SLAs a SON

¹Users may also need to pay (i.e., a monthly fee) the access networks for their right to access the Internet.

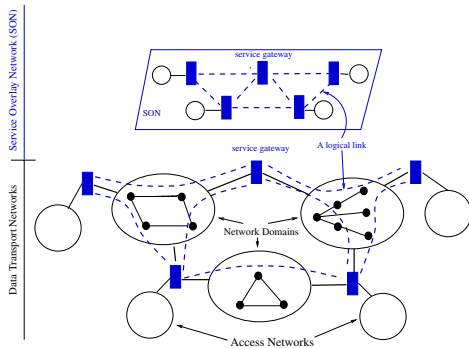


Figure 1. An illustration of a service overlay network.

can deliver end-to-end QoS sensitive services to its users via appropriate provisioning and service-specific resource management.

In addition to its ability to deliver end-to-end QoS sensitive services, the SON architecture also has a number of other important advantages. For example, it decouples application services from network services, thereby reducing the complexity of network service management and control, especially in terms of QoS management and control. Network domains are now concerned primarily with provisioning of data transport services with associated bandwidth management, traffic engineering and QoS guarantees on a much coarser granularity (per SON). In particular, the notion of SON also introduces a new level of traffic aggregation—*service aggregate*: underlying network domains can aggregate traffic based on the SON to which it belongs and perform traffic and QoS control accordingly based on the corresponding SLA. Under this architecture, a SON is responsible for ensuring the end-to-end QoS of its services. Because of its service awareness, a SON can deploy service-specific provisioning, resource management and QoS control mechanisms (e.g., at service gateways) to optimize its operations for its services. Hence the SON architecture not only simplifies network QoS management and makes it more scalable, but also enables the flexible creation and deployment of new (value-added) services.

Obviously the deployment of a SON is a capital-intensive investment. It is therefore imperative to consider the *cost recovery* issue for a SON. Among many costs incurred in the deployment of a SON (e.g., equipment such as service gateways), a dominant *recurring* cost is the cost of bandwidth that a SON must purchase from underlying network domains to support its services. A SON must provision adequate bandwidth to support its end-to-end QoS-sensitive services and meet traffic demands while minimizing the bandwidth cost so that it can generate sufficient revenue to recover its service deployment cost and stay profitable. *The bandwidth*

provisioning problem is therefore a critical issue in the deployment of the SON architecture. This study is devoted to this issue. The design and implementation of the SON architecture will be left to another paper.

We develop analytical models to study the problem of SON bandwidth provisioning and investigate the impact of various factors on SON bandwidth provisioning: SLAs, service QoS, bandwidth costs and traffic demands. We consider the so-called *pipe* SLA model as an example to illustrate how the SON bandwidth provisioning problem can be formally defined. The analyses and solutions can be adapted to the so-called *hose* SLA model [9], which due to space limitation we do not consider in this paper. In Section 2 we describe how the SON logical topology can be represented under the pipe SLA model and present the assumptions of our model. We study the *static* and *dynamic* SON bandwidth provisioning problems in Section 3 and Section 4, respectively. Analytical models and approximate solutions are developed for both static and dynamic bandwidth provisioning. Numerical studies are also performed to illustrate the properties of the proposed solutions and demonstrate the effect of traffic demand distributions and bandwidth costs on SON bandwidth provisioning.

The notion of overlay networks has been used widely in telecommunication and data networks. For example, more recently content distribution networks and application layer multicast networks have been used for multimedia streaming [3]; *Detour* [14] and *Resilient Overlay Network* (RON) [1] employ the overlay technique to provide better routing support. Moreover, the overlay technique has attracted a lot of attention from industries [4, 5] as a means to deliver diverse QoS-sensitive services over the Internet. The service overlay network we propose here is simply a generalization of these ideas. Perhaps what is particularly interesting is the use of SONs to address the end-to-end QoS deployment issue. The major contribution of our paper however lies in the study of the SON bandwidth provisioning problem. Our approach and formulation also differ from the traditional capacity planning in telephone networks (e.g. [10, 11]) in that we explicitly take into account various factors such as SLAs, QoS, traffic demand distributions.

2 Service Overlay Networks: Assumptions and Bandwidth Provisioning Problems

In this section we first describe a logical topology representation of a SON under the pipe SLA model and a simplifying assumption on service QoS. Two modes of bandwidth provisioning—*static* and *dynamic* bandwidth provisioning—are then introduced. We conclude this section by presenting a traffic demand model and a few notations regarding the service revenue and bandwidth cost for formulating the band-

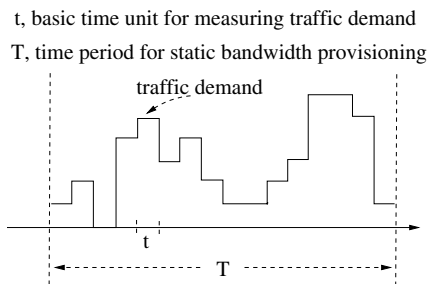


Figure 2. Traffic demands.

width provisioning problem.

2.1 SON and Service QoS

The pipe SLA model is a common SLA model used in today's Internet. Under the pipe model, a SON can request bandwidth guarantees between any two service gateways across a network domain (see Fig. 1); in other words a "pipe" with certain bandwidth guarantee is provisioned between the two service gateways across the network domain. To emphasize the relationship between service gateways and underlying network domains, we denote the *logical* (uni-directional) connection from a service gateway u to a neighboring service gateway v across a network domain D by $\langle u, v; D \rangle$, and refer to it as a *logical link* (or simply a *link*) between u and v across D . Note that between a SON and the access networks where traffic to the SON originates and terminates, the *hose* SLA model is assumed to be used where certain amount of bandwidth is reserved for traffic *entering* or *exiting* the SON. We can treat each access network A as a *fictitious* service gateway u_A . Then we can talk about "connection" between u_A and a neighboring service gateway v across A and the corresponding "logical link" $\langle u_A, v; A \rangle$.

Given a logical link $l = \langle u, v; D \rangle$, a SON provider will contract with the network domain D to provide a certain amount of bandwidth guarantee c_l between the service gateways u and v across D . The bandwidth provisioning problem of the SON is then to determine how much bandwidth to be provisioned for each link $l = \langle u, v; D \rangle$ so that: 1) the end-to-end QoS required by its services can be supported adequately; and 2) its overall revenue or net income can be maximized.

Although the QoS that a SON must support for its services can be quite diverse (e.g., bandwidth, delay or delay jitter guarantees), in almost all cases a key component in providing such guarantees is to exert some form of control on the link utilization level, i.e., to ensure the overall load on a link does not exceed certain specified condition. Consequently, for the purpose of bandwidth provisioning, we assume that it is possible to map service QoS guarantee requirements to a

link utilization threshold². To state this assumption formally, we assume that a link utilization threshold η_l is specified for each link l ; and to ensure service QoS, the bandwidth c_l on link l must be provisioned in such a way that the average link utilization stays below η_l (averaged over the basic unit of time, see Section 2.3).

2.2 Bandwidth Provisioning Modes

We consider two modes of bandwidth provisioning under the pipe model: *static* bandwidth provisioning and *dynamic* bandwidth provisioning. In the *static* bandwidth provisioning mode, a SON contracts and purchases a fixed amount of bandwidth *a priori* for each link connecting the service gateways from underlying network domains. In other words, the bandwidth is provisioned for a (relatively) long period of time without changing. In the *dynamic* bandwidth provisioning mode, in addition to the ability to contract and purchase bandwidth for each link *a priori*, a SON can also *dynamically* request for additional bandwidth from underlying network domains to meet its traffic demands, and pay for the dynamically allocated bandwidth accordingly. To account for the potential higher cost in supporting dynamic bandwidth provisioning, it is likely that underlying network domains will charge a SON different prices for statically provisioned and dynamically allocated bandwidth. Hence in either mode the key question in bandwidth provisioning for a SON is to determine the appropriate amount of bandwidth to be purchased *a priori* so that the total net income of the SON is maximized while maintaining the service QoS to meet the traffic demands.

2.3 Traffic Demand, Service Revenue and Bandwidth Cost

We now describe the traffic demand model for a SON. Recall that we assume that traffic always originates from and terminates at access networks. Given a source node s and destination node d , for simplicity we assume that a fixed route r consisting of a series of links connecting s and d is used to forward traffic from s to d . Let R denote the collection of routes between the source and destination nodes. Then the traffic demands over a SON can be represented by the traffic demands over these routes: for each $r \in R$, let ρ_r

²This particularly will be the case if the underlying network domain employs aggregate packet scheduling mechanisms such as FIFO or priority queues. For example, it has been shown [2, 16] that in order to provide end-to-end delay guarantees, link utilization must be controlled at a certain level. Hence from the bandwidth provisioning perspective we believe that this assumption on service QoS is not unreasonable in practice. In fact it is said that many of today's network service providers use a similar utilization based rule (e.g., an average utilization threshold of 60% or 70%) to provision their Internet backbones.

denote the (average) traffic demand (also referred to as traffic load) along route r measured over some period of time t (see Fig. 2). The period t is relatively short, for example in seconds or a few minutes, compared to the time scale of static bandwidth provisioning, denoted by T , which could be in several hours or days (or longer). The period t is considered as the basic unit of time. The set $\{\rho_r : r \in R\}$ then represents the traffic demands over the SON during the time unit they are measured, and is referred to as the traffic demand matrix of the SON. Note that traffic demands are always measured in units of bandwidth.

To capture traffic demand fluctuations over time, we assume that the traffic demand ρ_r along a route r varies according to some distribution³. We denote the probability density function of the traffic demand distribution of ρ_r by $d\rho_r$. Then the probability that the traffic demand ρ_r exceeds x units of bandwidth is given by $\int_x^\infty d\rho_r$. Let $\bar{\rho}_r = \int_0^\infty \rho_r d\rho_r$, i.e., $\bar{\rho}_r$ is the (long-term) average traffic demand along route r over the time period for static bandwidth provisioning. Furthermore, we assume that traffic demand distributions along different routes are *independent*. In this paper, we will study the bandwidth provisioning problem by considering a traffic demand model based on the $M/G/\infty$ input process [12, 13], which takes into account the widely observed self-similar property of the Internet traffic. (See [7, 8] for studies based on other traffic demand models.)

For each route r , we assume that a SON receives e_r amount of revenue for carrying one unit of traffic demand per unit of time along route r . On the other hand, for each logical link or pipe l connecting two service gateways, a SON must pay a cost of $\Phi_l(c_l)$ per unit of time for reserving c_l amount of bandwidth from the underlying network domain. We refer to Φ_l as the bandwidth cost function of link l . Without loss of generality, we assume that Φ_l is a *non-decreasing* function.

3 Static Bandwidth Provisioning with Penalty

In static bandwidth provisioning, a certain amount of bandwidth *overprovisioning* is needed to accommodate some degree of fluctuation in traffic demands. The key challenge in static bandwidth provisioning is therefore to decide the *optimal* amount of bandwidth overprovisioning. To accommodate some degree of fluctuation from the long-term average traffic demands, we introduce an *overprovisioning parameter* ϵ_l on each link l , $\epsilon_l \geq 0$. The meaning of the overprovisioning parameter ϵ_l is given as follows: we will provision c_l amount of bandwidth on link l such that as long as the overall traffic load on link l does not exceed its long-term average load by ϵ_l , the service QoS can be maintained, i.e., the link utilization is kept below the pre-specified threshold η_l . To

³This traffic demand distribution can be obtained, for example, through long-term observation and measurement.

put it formally, define $\bar{\rho}_l = \sum_{r:l \in r} \bar{\rho}_r$, where $l \in r$ denotes that link l lies on route r . Then

$$\bar{\rho}_l(1 + \epsilon_l) = (1 + \epsilon_l) \sum_{r:l \in r} \bar{\rho}_r \leq \eta_l c_l, \forall l \in L \quad (1)$$

where L is the set of all links of the SON.

We now consider how to obtain the *optimal* overprovisioning parameters under given traffic demand distributions. We study this problem by taking into account the consequence of potential QoS violations when actual traffic demands exceed the target link utilization. For this purpose, we assume that *a SON may suffer a penalty when the target utilization on a link is exceeded, and therefore service QoS may potentially be violated*. We refer to this model as the *static bandwidth provisioning with penalty model*, or in short, *static-penalty model*.

For each route r , let π_r denote the average penalty suffered by per unit of traffic demand per unit of time along route r when the service QoS along the route is potentially violated. Given a traffic demand matrix $\{\rho_r\}$, let $B_r(\{\rho_r\})$ denote the probability that the service QoS along route r is potentially violated, more specifically, *the target utilization on one of its links is exceeded*. Then the total net income of a SON for servicing a given traffic demand matrix $\{\rho_r\}$ can be expressed as follows:

$$W(\{\rho_r\}) = \sum_{r \in R} e_r \rho_r - \sum_{l \in L} \Phi_l(c_l) - \sum_{r \in R} \pi_r \rho_r B_r(\{\rho_r\}), \quad (2)$$

where in the above we use $W(\{\rho_r\})$ to emphasize the dependence of the total net income on the traffic demand matrix $\{\rho_r\}$. When there is no confusion, we may drop $\{\rho_r\}$ from the notation.

Let $d\{\rho_r\}$ denote the joint probability density function of a traffic demand matrix $\{\rho_r\}$, where recall that $d\rho_r$ is the probability density function of a traffic demand ρ_r along route r . Then the expected net income of a SON under the traffic demand distributions $\{d\rho_r\}$ is given by

$$E(W) = \int \cdot \int_{\{\rho_r\}} W(\{\rho_r\}) d\{\rho_r\}, \quad (3)$$

where $\int \cdot \int_{\{\rho_r\}}$ denotes multiple integration under the joint traffic demand distribution $\{d\rho_r\}$.

Now we can state the problem of static bandwidth provisioning with penalty as the following optimization problem: finding the optimal overprovisioning parameters $\{\epsilon_l\}$ to maximize the expected net income, i.e.,

$$\max_{\{\epsilon_l\}} E(W) \text{ subject to (1)}. \quad (4)$$

Unfortunately, the exact solution to this optimization problem is in general difficult to obtain. It depends on both

the particular forms of the traffic demand distributions $\{d\rho_r\}$ and the service QoS violation probabilities B_r . In the following, instead of the exact solution, we shall derive an approximate solution based on a lower bound on $E(W)$. (Due to page limitations, we only sketch the derivation. We refer interested readers to [8].) Before we present the approximate optimal solution, we need to introduce one more set of notations. Define a small real number $\delta > 0$. For each route r , let $\hat{\rho}_r > \bar{\rho}_r$ be such that

$$\int_{\hat{\rho}_r}^{\infty} \rho_r d\rho_r \leq \delta. \quad (5)$$

Since $\int_{\hat{\rho}_r}^{\infty} \rho_r d\rho_r \geq \hat{\rho}_r \int_{\hat{\rho}_r}^{\infty} d\rho_r = \hat{\rho}_r Pr\{\rho_r \geq \hat{\rho}_r\}$, we have $Pr\{\rho_r \geq \hat{\rho}_r\} \leq \delta/\hat{\rho}_r$. In other words, (5) basically says that $\hat{\rho}_r$ is such that the probability the traffic demand along route r exceeds $\hat{\rho}_r$ is very small, and thus negligible.

With these notations in place, we now present a lower bound on $E(W)$ as follows (see [8] for the detailed derivation).

$$\begin{aligned} E(W) \geq & \sum_{r \in R} e_r \bar{\rho}_r - \sum_{l \in L} \Phi(c_l) - \sum_{r \in R} \pi_r \bar{\rho}_r B_r(\{\hat{\rho}_r\}) \\ & - \sum_{r \in R} \pi_r \delta \left(1 + \sum_{r' \neq r} \frac{\bar{\rho}_r}{\hat{\rho}_{r'}}\right), \end{aligned} \quad (6)$$

where B_r is the service QoS violation probability, i.e., at least one of the links on route r is overloaded:

$$B_r = 1 - \prod_{l \in r} (1 - B_l). \quad (7)$$

Denote the right-hand side of (6) by V , then $E(W) \geq V$. From $E(W) \geq V$, we have $\max_{\{\epsilon_r\}} E(W) \geq \max_{\{\epsilon_r\}} V$. Therefore we can obtain the *best* overprovisioning parameters that maximize V instead of the expected net income $E(W)$ as an approximate solution to the original optimization problem (4). Let $\{\epsilon_l^*\}$ be the solution to the optimization problem $\max_{\{\epsilon_r\}} V$, and refer to them as the *approximate optimal overprovisioning parameters*. Suppose that $\{\epsilon_l^*\}$ are strictly positive, then a necessary condition for them to be an optimal solution is that the gradient ∇V (with respect to $\{\epsilon_l\}$) must vanish at ϵ_l^* 's. Based on this observation and through some simple algebraic manipulation, it is not too hard to show that, $\{\epsilon_l^*\}$ can be obtained by solving the following equations

$$\frac{\partial \Phi_l(c_l)}{\partial c_l} = \hat{s}_l, \quad \forall l \in L, \quad (8)$$

in the above equation, \hat{s}_l is defined as

$$\hat{s}_l = \sum_{r: l \in r} \pi_r \bar{\rho}_r \prod_{k \in r, k \neq l} [1 - B_k(\hat{\rho}_k, c_k)] \zeta_l, \quad (9)$$

where $\zeta_l = -\frac{d}{dc_l} B_l(\hat{\rho}_l, c_l)$.

In the above derivation of the approximate optimal solution to the static bandwidth provisioning problem, we have simply assumed the existence of B_l but not its form. Its particular form depends on the distribution of (average) traffic demands on link l . In the following subsection, we consider a traffic demand model based on the $M/G/\infty$ input process to demonstrate the approximate optimal solution to the static bandwidth provisioning problem.

3.1 $M/G/\infty$ Traffic Demand Model

Consider an $M/G/\infty$ input process, where the service time has a heavy-tailed distribution. We assume that the distribution of the service time has a finite mean. Let X_t denote the number of customers in the system at time t , for $t = 0, 1, 2, \dots$. Then the count process $\{X_t\}_{t=0,1,2,\dots}$ is asymptotically self-similar. Let ρ denote the customer arrival rate of the $M/G/\infty$ input process and μ the mean service time, then X_t has a Poisson marginal distribution with mean $\rho\mu$ [6].

Now we are ready to present the $M/G/\infty$ traffic demand model. Consider an arbitrary route r . In the $M/G/\infty$ traffic demand model, the (average) traffic demand (i.e., the average traffic arrival rate in each unit of time) on the route is governed by the count process $\{X_t\}_{t=0,1,2,\dots}$ of an $M/G/\infty$ input process. For example, let $\rho_{r,i}$ denote the average traffic demand in the i th unit of time, then we have $\rho_{r,i} = X_i$. Let $\bar{\rho}_r$ denote the long-term average traffic demand on the route. It is easy to see that $\bar{\rho}_r = \rho\mu$, where ρ and μ are the customer arrival rate and the mean service time, respectively, of the $M/G/\infty$ input process. As traffic demands along all the routes are assumed to be independent, the average overall traffic load on a link l is $\bar{\rho}_l = \sum_{r: l \in r} \bar{\rho}_r$.

Given the average overall load $\bar{\rho}_l$ and the link capacity c_l , it can be shown that the probability that the total load on link l exceeds $\bar{c}_l = \eta_l c_l$ during any given unit of time is given by $B_l(\bar{\rho}_l, c_l) = (\sum_{i=(\bar{c}_l+1)}^{\infty} \frac{\bar{\rho}_l^i}{i!}) e^{-\bar{\rho}_l}$. We extend the definition of $B_l(\bar{\rho}_l, c_l)$ to the non-integer values of c_l by linear interpolation. Moreover, at the integer values of c_l we define the derivative of $B_l(\bar{\rho}_l, c_l)$ with respect to c_l to be the left derivative. Then $\frac{d}{dc_l} B_l(\bar{\rho}_l, c_l) = B_l(\bar{\rho}_l, c_l) - B_l(\bar{\rho}_l, c_l - 1)$. Therefore, $\zeta_l = -\frac{d}{dc_l} B_l(\hat{\rho}_l, c_l) = \eta_l \{B_l(\hat{\rho}_l, (\eta_l c_l - 1)) - B_l(\hat{\rho}_l, \eta_l c_l)\} = \eta_l \frac{\hat{\rho}_l^{[\eta_l c_l]}}{[\eta_l c_l]!} e^{-\hat{\rho}_l}$. By this definition of B_l , we can obtain the (approximate) optimal overprovisioning parameters ϵ_l^* 's by solving (8).

We now discuss the effect of the *shapes* of \hat{s}_l ' and Φ_l on (approximate) optimal overprovisioning parameters ϵ_l^* 's as well as their implication in static bandwidth provisioning. Note first that the shape of \hat{s}_l is determined by ζ_l , which has a shape of (skewed) bell-shape with a center approximately at $\hat{\rho}_l$ (it is essentially a Poisson probability density function). Hence \hat{s}_l is a concave function of $\epsilon_l \geq 0$. In

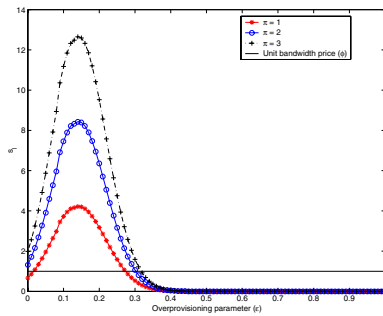


Figure 3. Relationship between $\hat{\epsilon}_l$, ϵ_l , & ϕ_l .

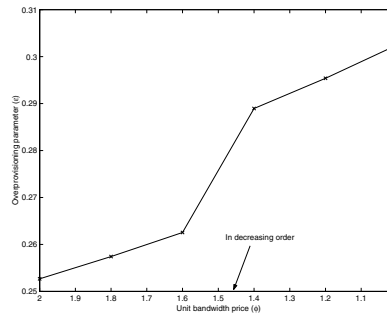


Figure 4. Impact of unit bandwidth price on ϵ_l^* .

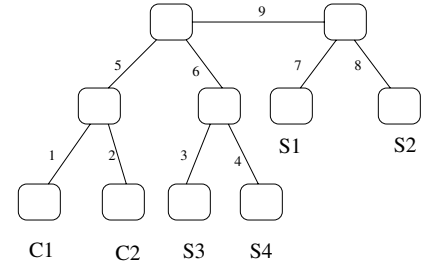


Figure 5. Tree topology.

particular, there exists $\hat{\epsilon}_l$ such that $\hat{\epsilon}_l$ is an increasing function in the range $[0, \hat{\epsilon}_l]$ and a decreasing function in the range $[\hat{\epsilon}_l, \infty)$ (see Fig. 3). Intuitively, this means that as ϵ_l moves from 0 towards $\hat{\epsilon}_l$, there is an increasing benefit in bandwidth overprovisioning in terms of *reducing potential QoS violation penalty*. However, as ϵ_l moves beyond $\hat{\epsilon}_l$, there is a *diminished return* in overprovisioning in terms of reducing potential QoS violation penalty.

Suppose that Φ_l is a linear function, i.e., $\Phi_l(c_l) = \phi_l c_l$. Then $\frac{\partial \Phi_l(c_l)}{\partial c_l} = \phi_l$. Hence (8) becomes $\phi_l = \hat{\epsilon}_l$. Suppose $\phi_l = \hat{\epsilon}_l$ holds for some $\epsilon_l \geq 0$. Because of the shape of $\hat{\epsilon}_l$, there potentially exist two solutions $\epsilon_{l,1}$ and $\epsilon_{l,2}$, $0 \leq \epsilon_{l,1} \leq \hat{\epsilon}_l \leq \epsilon_{l,2}$ such that $\phi_l = \hat{\epsilon}_l$. In particular, as $\hat{\epsilon}_l$ is a decreasing function in the range $[\hat{\epsilon}_l, \infty)$, $\epsilon_{l,2}$ always exists. As $\frac{\partial V}{\partial c_l}$ is positive in the range $(\epsilon_{l,1}, \epsilon_{l,2})$, and is negative in the ranges $[0, \epsilon_{l,1})$ and $(\epsilon_{l,2}, \infty)$, we see that with respect to link l , V is maximized at either $\epsilon_l^* = \epsilon_{l,2}$ or at $\epsilon_l^* = 0$ (whereas it is minimized at $\epsilon_{l,1}$).

In the following, we conduct numerical studies to illustrate the properties of the analytic results we obtained and demonstrate the effects of various parameters on static bandwidth provisioning. For this purpose, we consider a simple setting: a single route over a single link. (See [8] for numerical studies in more complex settings.)

Unless otherwise stated, the following parameters will be used in the numerical studies: the long-term average traffic demand on the route is 200 (measured in unit of bandwidth per unit of time), i.e., $\bar{\rho}_r (= \bar{\rho}_l) = 200$, and $e_r = 4$, $\phi_l = 1$, $\pi_r = 2$. We set $\delta = 5$ and the target utilization threshold $\eta_l = 0.8$.

Fig. 3 shows $\hat{\epsilon}_l$ as a function of ϵ_l with three different values of π_r : $\pi_r = 1, 2, 3$. In the figure we also include a line corresponding to $\phi_l = 1$ to illustrate how ϵ_l^* can be obtained as the solution to $\hat{\epsilon}_l = \phi_l$. Recall that, $\epsilon_l^* = \epsilon_{l,2}$ (the right intersecting point). From Fig. 3 we see that as the penalty π_r increases, ϵ_l^* also increases. Hence for a higher penalty it is necessary to overprovision more bandwidth to guide against

potential QoS violations. Likewise, as we increase the per-unit bandwidth cost ϕ_l (i.e., moving up the line of ϕ_l), ϵ_l^* decreases. In other words, as the bandwidth cost increases, it is beneficial to reduce overprovisioned bandwidth so as to maximize the net income.

To highlight the relationship between bandwidth cost and overprovisioning in Fig. 4 we plot the overprovisioning parameter ϵ_l^* as a function of the per-unit bandwidth cost ϕ_l (note the decreasing order of ϕ_l on x-axis). We see that as the per-unit bandwidth cost ϕ_l decreases (from 2 to 1), the overprovisioning parameter ϵ_l^* increases, i.e., it is more beneficial to overprovision more bandwidth. This is not surprising.

In this section, we have studied the static bandwidth provisioning mode, where during a relatively long period, the provisioned bandwidth on a link will not be changed. The static bandwidth provisioning mode is simple in bandwidth management, but may result in inefficient bandwidth usage facing traffic demand fluctuations. In the next section, we will study the dynamic bandwidth provisioning mode, where the link bandwidth could be dynamically adjusted according to the traffic demand fluctuations in relatively shorter time intervals.

4 Dynamic Bandwidth Provisioning

In this section we study the dynamic bandwidth provisioning problem. As pointed out in Section 2, to account for potential higher cost in supporting dynamic bandwidth provisioning, it is likely that underlying network domains will charge a SON different prices for statically provisioned and dynamically allocated bandwidth. Hence we assume that for each link l , the cost for reserving c_l amount of bandwidth *statically* is, as before, $\Phi_l(c_l)$; while the cost of reserving the same amount of bandwidth *dynamically* is $\Phi'_l(c_l)$, where $\Phi'_l(c_l) \geq \Phi_l(c_l)$. Given this price differential, a *key question for a SON is to determine how much bandwidth should be reserved statically on each link l a priori to meet certain base*

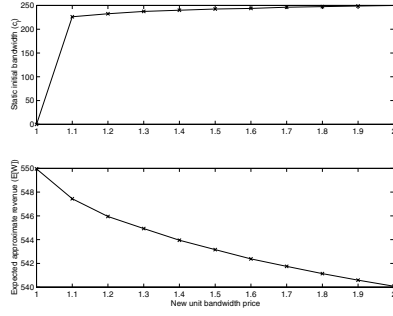


Figure 6. Effects of ϕ'_l on c_l and $E(\tilde{W})$.

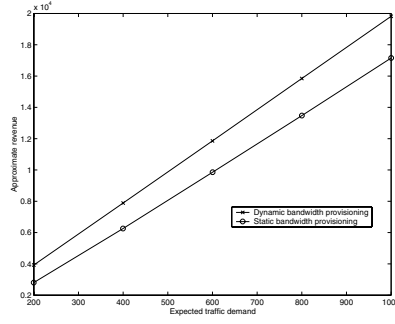


Figure 7. Dynamic vs. static bandwidth provisioning.

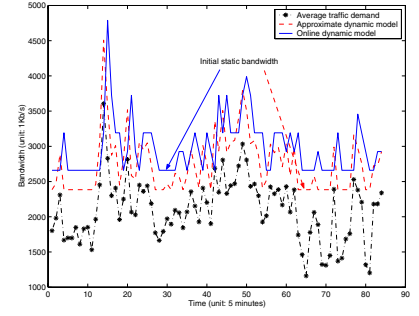


Figure 8. Dynamic bandwidth provisioning.

traffic demands, while dynamically allocating bandwidth to meet additional traffic demands as needed. The objective is again to maximize the overall long-term expected net income of a SON.

To focus on the dynamic bandwidth provisioning problem, we assume that underlying network domains possess abundant bandwidth that dynamic requests for additional bandwidth from a SON are always satisfied. In other words, no request is blocked. Under this assumption, for a given traffic demand matrix $\{\rho_r\}$, it is possible to compute the expected additional bandwidth that needs to be dynamically allocated to meet the traffic demands. This can be done, for example, using the $M/G/\infty$ traffic demand model introduced in the previous section. However such precise formulation is extremely complicated, and consequently the corresponding optimization problem is unlikely to be tractable. In the following, we will first describe an approximate model based on the marginal distributions of the traffic demands on the links of a SON; and then present an adaptive heuristic algorithm for dynamic bandwidth provisioning based on *online* traffic measurements.

4.1 Approximate Model

Suppose for each link $l \in L$, c_l amount of bandwidth has been provisioned statically *a priori*. Given a traffic demand matrix $\{\rho_r\}$, we approximate the *expected* additional bandwidth that must be dynamically allocated to meet the traffic demands by the following expression:

$$\Delta c_l = \left\{ \frac{\rho_l}{\eta_l} - c_l \right\}^+, \quad (10)$$

where $\rho_l = \sum_{r \in R} \rho_r$. Then $\Delta c_l > 0$ if and only if $\rho_l > \eta_l c_l$. Using (10) we can write down the *approximate* overall net income a SON generates for a given traffic demand matrix

$\{\rho_r\}$:

$$\tilde{W}(\{\rho_r\}) = \sum_{r \in R} e_r \rho_r - \sum_{l \in L} \Phi_l(c_l) - \sum_{l \in L} \Phi'_l(\Delta c_l). \quad (11)$$

Integrating on both sides of (11) over the (joint) distribution of $d\{\rho_r\}$, we have

$$E(\tilde{W}) = \sum_{r \in R} e_r \bar{\rho}_r - \sum_{l \in L} \Phi_l(c_l) - \sum_{l \in L} \int \int \Phi'_l(\Delta c_l) d\{\rho_r\}. \quad (12)$$

The dynamic bandwidth provisioning problem can now be formulated as the following optimization problem:

$$\max_{\{c_l\}} E(\tilde{W}). \quad (13)$$

Note that unlike the static bandwidth provisioning problem, here we do not have any explicit QoS or target utilization constraints. This is because we implicitly assume that whenever the target utilization threshold of a link is about to be exceeded, additional bandwidth is dynamically allocated on the link to meet the service QoS. We will refer to the optimization problem (13) as the *approximate model* for dynamic bandwidth provisioning. In the following, we will present an (approximate) solution to the approximate model of the dynamic bandwidth provisioning problem. For the detailed analysis, we refer interested readers to [8].

Assume both bandwidth cost functions are linear, i.e., for any $l \in L$, $\Phi_l(c_l) = \phi_l c_l$ and $\Phi'_l(\Delta c_l) = \phi'_l \Delta c_l$, where $\phi_l \leq \phi'_l$ for any l . Let c'_l be such that $Pr\{\rho_l > \eta_l c'_l\} = \phi_l / \phi'_l$. Then the set $\{c'_l\}$ is an (approximate) solution to the dynamic bandwidth provisioning problem. That is, c'_l is the amount of bandwidth to be statically provisioned on link l , while the portion to be dynamically allocated on the link is given by (10) with c_l replaced by c'_l , for a given traffic demand matrix $\{\rho_r\}$.

An intuitive interpretation of the above results is that under the dynamic bandwidth allocation model, we need to statically reserve at most c'_l amount of bandwidth on each link

l , where the probability that the (average) aggregate load on link l exceeds the statically reserved link bandwidth c'_l equals the ratio of the two prices on the link, ϕ_l/ϕ'_l . In the special case that $\phi_l = \phi'_l$, i.e., the unit price of dynamically allocated bandwidth is the same as that of the statically reserved one, we have $c'_l = 0$. Hence in this case, no static bandwidth needs to be reserved.

4.1.1 Numerical Examples

In this section we perform numerical studies to illustrate the properties of the dynamic bandwidth provisioning model, and compare it with the static bandwidth provisioning model. Unless otherwise stated, the per-unit bandwidth per-unit time earning $e_r = 4$, and $\phi_l = 1$, $\phi'_l = 1.5$. The target link utilization threshold η_l is 0.8.

In the first set of studies, we examine the effects of the per-unit bandwidth price ϕ'_l for dynamically allocated bandwidth on the amount of bandwidth c_l provisioned statically *a priori* and the approximate revenue $E(\tilde{W})$. In these studies, we use the simple network setting: a single route over a single link. The traffic demand model is $M/G/\infty$ and the long term average traffic demand on the route is 200. Fig. 6 presents the bandwidth c_l provisioned statically (upper plot) and the approximate revenue $E(\tilde{W})$ (lower plot) as functions of ϕ'_l , respectively. From the figure we see that as the per-unit dynamic-bandwidth price increases, more bandwidth needs to be provisioned statically *a priori*. However, the increase in the amount of static bandwidth is not significant as ϕ'_l increases from $\phi'_l = 1.1$ to $\phi'_l = 2$. On the other hand, as we increase the price for dynamically allocated bandwidth, the approximate revenue $E(\tilde{W})$ decreases. This is due to the fact that a SON needs to statically provision more bandwidth *a priori* on each link, in addition to having to pay more for dynamically allocated bandwidth.

In the next set of numerical studies, we compare the dynamic bandwidth provisioning model with the static bandwidth provisioning model in terms of obtained approximate revenues. We use the *tree* network topology (see Fig. 5). In the following $a \rightarrow b$ denotes a route from service gateway a to service gateway b . The path with *minimum* “hop-count” (i.e., service gateways) is used as the route between two service gateways. In the tree topology, four routes are used: $R1 = S3 \rightarrow C1$, $R2 = S1 \rightarrow C1$, $R3 = S4 \rightarrow C2$, and $R4 = S2 \rightarrow C2$. In the numerical studies below, we use the $M/G/\infty$ traffic demand model. Moreover, the expected traffic demand for all routes is 200. We set $e_r = 10$, $\pi_r = 2$ for all routes, and $\phi_l = 1$ for all links. The value of δ is chosen in such a way that $\delta_r = \frac{1}{40}\rho_r$. Fig. 7 presents the approximate revenue as a function of the (long-term) average traffic demands for dynamic and static bandwidth provisioning, respectively. From the figure we see that, for both dynamic and static bandwidth provisioning models, the approximate

revenue increases as the average traffic demand increases. Moreover, the dynamic bandwidth provisioning model has a higher approximate revenue than that of the static bandwidth provisioning model. Note also that as the average traffic demand increases, the difference between the approximate revenues of dynamic bandwidth provisioning and static bandwidth provisioning becomes larger. This is possibly due to the fact that, as the average traffic demand on a route increases, traffic along the route becomes more bursty (recall that the marginal distribution of traffic demand on a route is Poisson), and the dynamic bandwidth provisioning model works better than the static bandwidth provisioning model in this case.

4.2 Adaptive Online Bandwidth Provisioning Algorithm

In developing the approximate dynamic bandwidth provisioning model, we have assumed that the (average) traffic demands are known *a priori* for determining the additional bandwidth that must be dynamically allocated to meet the traffic demands (10). The approximate model has very nice computation and performance properties but in general the traffic demand matrix may not be available *a priori*. In this section, we present a heuristic *online* bandwidth allocation algorithm (for short *online dynamic model*) that emulates the approximate dynamic bandwidth provisioning model. The online dynamic model dynamically adjusts the allocated bandwidth on a link according to the *measurement* of the traffic demands on the links of a SON.

As before, let $\bar{\rho}_r$ denote the long-term average traffic demand on route r , and $\bar{\rho}_l = \sum_{r:l \in r} \bar{\rho}_r$, the long-term average traffic demand on link l . Based on the measurement of traffic demands on the links, our target in this section is to determine the amount of bandwidth c_l that should be statically provisioned *a priori* to meet certain base traffic demands, and the amount of bandwidth Δc_l that should be allocated dynamically to accommodate the traffic demand dynamics in a SON.

Let t denote a fixed time interval. In the online dynamic model, the average traffic demand ρ_l during each such time interval is calculated at the end of the time interval. Based on the measured average traffic demands and the contracted service QoS, the bandwidth allocated on each link will be adjusted accordingly at the end of the time interval. Moreover, the allocated bandwidth will be kept constant during the next measurement time interval. In other words, the allocated bandwidth is only adjusted at the end of each measurement time interval. To reduce the frequency of allocating additional bandwidth or de-allocating extra bandwidth caused by short-term traffic fluctuations, bandwidth will be allocated in units of quota, which is a chunk of bandwidth [17] and normally much larger than one unit of bandwidth. In the fol-

lowing, we will denote the size of a quota by Θ (in unit of bandwidth).

Let c_l denote the amount of bandwidth that has been provisioned statically on a link l *a priori*. In the online dynamic model, c_l is chosen in such a manner that, if the average traffic demand on the link does not exceed $\bar{\rho}_l$, the service QoS will be honored, i.e.,

$$c_l = \lceil \frac{\bar{\rho}_l}{\eta_l \Theta} \rceil \Theta, \quad (14)$$

note that, the initial static bandwidth is allocated in units of quota.

Next, we discuss the allocation of additional bandwidth and de-allocation of extra bandwidth on an arbitrary link l . To reduce the possibility that the service QoS is violated, the online dynamic model will allocate additional bandwidth (a new quota) as soon as the average traffic demand is *approaching* the target link utilization level threshold, instead of until the threshold is exceeded. Let ι_f denote a positive number, and C_l the current total bandwidth on link l , i.e., $C_l = c_l + \Delta c_l$. Then an additional quota will be allocated onto link l as soon as $\rho_l > C_l \eta_l - \iota_f$. ι_f is called the forward threshold for allocating a new quota. Similarly, a backward threshold for de-allocating an extra quota is defined as (denoted by ι_b (a positive number)): an extra quota is released from link l only if $\rho_l < (C_l - \Theta) \eta_l - \iota_b$. As a summary, we present the online dynamic model in pseudo-code in Fig. 9.

1. Set initial static bandwidth $c_l = \lceil \frac{\bar{\rho}_l}{\eta_l \Theta} \rceil \Theta$.
2. At the end of each measurement interval:
3. **for** each link l :
4. /* C_l : current total bandwidth on link l */
5. Calculate the average traffic demand ρ_l .
6. **if** $\rho_l > C_l \eta_l - \iota_f$
7. $C_l = C_l + \Theta$
8. **else if** $\rho_l < (C_l - \Theta) \eta_l - \iota_b$
9. $C_l = \max\{c_l, C_l - \Theta\}$

Figure 9. Online dynamic model.

Because the online dynamic model only adjusts bandwidth on links at the end of a measurement interval, it is possible that the service QoS is violated during the course of the interval. As in static bandwidth provisioning with penalty in Section 3, certain penalty will apply in this case. Let π_r denote the average penalty suffered by per unit of traffic demand per unit of time (the measurement time interval) along route r when the service QoS along the route is violated. Then the revenue of the online dynamic model for a measurement time interval is,

$$\bar{V} = \sum_{r \in R} e_r \rho_r - \sum_{l \in L} \Phi_l(c_l) - \sum_{l \in L} \Phi'_l(\Delta c_l) - \sum_{r \in R} \pi_r \rho_r \mathbf{1}_{\{\rho_l / C_l > \eta_l : l \in r\}}, \quad (15)$$

where the indicator function $\mathbf{1}_{\{\rho_l / C_l > \eta_l : l \in r\}} = 1$ if $\rho_l / C_l > \eta_l$ holds for any link l on route r , 0 otherwise.

In the following, we perform numerical studies to illustrate the bandwidth allocation behavior of the online dynamic model based on the measurements of real Internet traffic. The data trace we use was collected at the University of Auckland Internet access link on December 1, 1999, lasted roughly for 24 hours [15]. In this study, we only use the portion of measurement from 10:00AM to 5:00PM and refer to it as *Auckland data trace*. Fig. 8 presents the average traffic arrival rates (i.e. traffic demands) of the Auckland data trace, where each point represents the average traffic demand for a 5 minute time interval (which is also used as the basic unit of time, i.e., $t = 5$ minutes, see Fig. 2). Let the basic unit of bandwidth (traffic demand) be 1 Kb/s , then the mean traffic demand and the standard deviation of the Auckland data trace traffic demand are 2096 and 442, respectively.

The following studies are carried out in the simple network setting and the following parameters are used. The per-unit bandwidth per-unit time earning $e_r = 4$, and $\phi_l = 1$, $\phi'_l = 1.5$, $\pi_r = 2$. The target utilization threshold $\eta_l = 0.8$. The size of a quota $\Theta = 0.6\sigma$, where σ is the standard deviation of the Auckland data trace. The forward and backward threshold $\iota_f = \iota_b = 0.3\Theta$.

Fig. 8 presents the average traffic demands (per 5 minutes) and the corresponding allocated bandwidth in the online dynamic model. For the purpose of comparison, we also include the bandwidth provisioning behavior of the approximate dynamic model. From the figure we see that the online dynamic model is able to adjust the link bandwidth according to the dynamics of the traffic demands on the link and meanwhile remains insensitive to small short-time fluctuations in traffic demands (for example, see the allocated bandwidth at time 24, 25 and 26). Because of the nature of the online dynamic model, sometimes the bandwidth on a link could be less than the average traffic demand on the link (for example, at time 14), where a penalty will apply. (A penalty may apply in other cases.) Comparing the curve of the approximate dynamic model with that of the online dynamic model, we see that the online dynamic model approaches the approximate dynamic model reasonably well (with a small lag interval), except that the approximate dynamic model has a smaller initial static bandwidth than the online dynamic model. (However, recall that the initial static bandwidth of the online dynamic model is only based on the long-term average traffic demand while of the approximate dynamic model, it relies on the distribution of the average traffic demands.) Note also that the approximate dynamic model is more sensitive to the fluctuations in traffic demands than the online dynamic model.

Table 1 gives the mean revenues (per-unit time) of the approximate dynamic model and the online dynamic model. From the table we see that the approximate dynamic model

Table 1. Per-unit time average revenue.

	Approximate model	Online model
Average revenue	5468	4152

has a higher per-unit time average revenue than the online dynamic model. There are possibly two reasons. Firstly, under this parameter setting, the amount of initial static bandwidth of the online dynamic model is larger than that of the approximate dynamic model; moreover, the bandwidth is allocated in units of quota in the online dynamic model, which also tends to reserve more bandwidth than needed. These two factors cause a higher expense on the overlay with the online dynamic model. Secondly, the online dynamic model is measurement-based and the bandwidth on a link is only adjusted at the end of the measurement time intervals. Consequently, as we discussed before, service QoS may be violated during a time interval and incurs penalty on the overlay. However, given that the approximate dynamic model requires the traffic demand matrix to be known *a priori* while the online dynamic model does not, we believe the latter is a good approximation to the former overall.

5 Conclusions and Future Work

In this paper, we studied the bandwidth provisioning problem for service overlay networks (SONs). We considered both the static and dynamic bandwidth provisioning models, and our formulation of the SON bandwidth provisioning problem took into account various factors such as service QoS, traffic demand distributions, and bandwidth costs.

The approximate optimal solution we developed to the static bandwidth provisioning problem is generic in the sense that it applies to different marginal distributions of the traffic demands on the routes in a network, which makes the solution very attractive facing different traffic arrival behaviors. The static bandwidth provisioning model is simple in terms of network resource management but may result in inefficient network resource usage if the traffic demands are highly variable. In this kind of environments, the dynamic bandwidth provisioning model outperforms the static bandwidth provisioning model, albeit with more expensive network resource management. We investigated the effects of various parameters like static and dynamic bandwidth costs on the revenue that a SON can obtain, which provides useful guidelines on how a SON should be provisioned to stay profitable.

Currently, we are investigating the effects of time granularity for measuring (average) traffic demands on the bandwidth provisioning of a SON and the resulting network performance. We are also interested in exploring the functionalities of service gateways in support of service-aware (multi-path) routing, which may have great impact on how a SON

should be dimensioned and provisioned.

References

- [1] D. G. Andersen, H. Balakrishnan, M. Kaashoek, and R. Morris. Resilient overlay networks. In *Proc. 18th ACM SOSP*, Banff, Canada, October 2001.
- [2] A. Charny and J.-Y. Le Boudec. Delay bounds in a network with aggregate scheduling. In *Proceedings of QoFIS*, Berlin, Germany, October 2000.
- [3] Y. Chawathe, S. Fink, S. McCanne, and E. Brewer. A proxy architecture for reliable multicast in heterogeneous environments. In *Proceedings of ACM Multimedia*, Bristol, U.K., September 1998.
- [4] Virtela Communications. <http://www.virtela.com>.
- [5] Internap Network Services Corporation. <http://www.internap.com>.
- [6] D. Cox and V. Isham. *Point Processes*. Chapman and Hall, 1980.
- [7] Z. Duan, Z.-L. Zhang, and Y. T. Hou. Bandwidth provisioning for service overlay networks. In *Proceedings of SPIE ITCOM (Scalability and Traffic Control in IP Networks) 2002*, Boston, MA, July 29 – August 1 2002.
- [8] Z. Duan, Z.-L. Zhang, and Y. T. Hou. Service overlay networks: SLAs, QoS and bandwidth provisioning. Technical report, Computer Science Department, University of Minnesota, February 2002.
- [9] N. Duffield, P. Goyal, A. Greenberg, P. Mishra, K. Ramakrishnan, and J. Merwe. A flexible model for resource management in virtual private networks. In *Proc. ACM SIGCOMM*, Cambridge, Massachusetts, September 1999.
- [10] A. Girard. *Routing and Dimensioning in Circuit-Switched Networks*. Addison-Wesley, 1990.
- [11] F. P. Kelly. Routing in circuit-switched networks: Optimization, shadow prices and decentralization. *Advances in Applied Probability*, 20:112–144, 1988.
- [12] M. Parulekar and A. M. Markowski. M/G/∞ input processes: A versatile class of models for network traffic. In *Proc. IEEE INFOCOM*, pages 419–426, Kobe, Japan, April 1997.
- [13] V. Paxson and S. Floyd. Wide area traffic: The failure of poisson modeling. In *Proc. ACM SIGCOMM*, pages 257–268, August 1994.
- [14] S. Savage, T. Anderson, and et al. Detour: a case for informed internet routing and transport. *IEEE Micro*, 19(1):50–59, January 1999.
- [15] Auckland Data Trace. <http://pma.nlanr.net/traces/long/auck2.html>.
- [16] Z.-L. Zhang, Z. Duan, and Y. T. Hou. Fundamental trade-offs in aggregate packet scheduling. In *Proceedings of IEEE International Conference on Network Protocols (ICNP)*, Riverside, CA, November 2001.
- [17] Z.-L. Zhang, Z. Duan, and Y. T. Hou. On scalable design of bandwidth brokers. *IEICE Transaction on Communications*, E84-B(8), August 2001.