Service Overlay Networks: SLAs, QoS, and Bandwidth Provisioning

Zhenhai Duan, *Member, IEEE*, Zhi-Li Zhang, *Member, IEEE*, and Yiwei Thomas Hou, *Member, IEEE*

***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 quality of service (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 SON purchases bandwidth with certain QoS guarantees from the individual network domains via bilateral service level agreement (SLA) to build a logical end-to-end service delivery infrastructure on top of the existing data transport networks. Via a service contract, users directly pay the SON for using the value-added services provided by the SON.**

In this paper, we study the bandwidth provisioning problem for an SON which buys bandwidth from the underlying network domains to provide end-to-end value-added QoS sensitive services such as VoIP and Video-on-Demand. A key problem in the SON deployment is the problem of bandwidth provisioning, which is critical to cost recovery in deploying and operating the value-added services over the SON. The paper is devoted to the study of this problem. We formulate the bandwidth provisioning problem mathematically, taking 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 SON bandwidth provisioning.

***Index Terms*—Bandwidth provisioning, overlay networks, service level agreements.**

I. 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, a 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 multilateral business relationship involving multiple domains, 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 multiservice network infrastructure with end-to-end quality of service (QoS) support.

We propose and advocate the notion of the 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 SON network architecture relies on well-defined business relationships between the SONs, the underlying network domains and users of the SONs to provide support for end-to-end QoS: the SON purchases bandwidth with certain QoS guarantees from the individual network domains via a bilateral service level agreement (SLA) to build a logical end-to-end service delivery infrastructure on top of the existing data transport networks; via a service contract (e.g., a usage-based or fixed price service plan), users[1] directly pay the SON for using the value-added services provided by the SON.

Fig. 1 illustrates the SON architecture. The 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, for example, bypasses the peering points among the network domains, and thus avoids the potential performance problems associated with them. Relying on the bilateral SLAs the SON 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

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

Manuscript received January 6, 2003; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor E. Knightly. This work was supported in part by the National Science Foundation under Grants CAREER Award NCR-9734428, EIA-9818338, and ITR ANI-0085824. Any opinions, findings, conclusions, or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the National Science Foundation. An earlier, abridged version of this paper appeared in the Proceedings of the 10th IEEE International Conference on Network Protocols, Paris, France, November, 2002.

Z. Duan is with the Computer Science Department, Florida State University, Tallahassee, FL 32306 USA (e-mail: duan@cs.fsu.edu).

Z.-L. Zhang is with the Department of Computer Science and Engineering, University of Minnesota, Minneapolis, MN 55455 USA (e-mail: zhzhang@cs.umn.edu).

Y. T. Hou is with the Bradley Department of Electrical and Computer Engineering, Virginia Polytechnic and State University, Blacksburg, VA 24061 USA (e-mail: thou@vt.edu).

Digital Object Identifier 10.1109/TNET.2003.820436

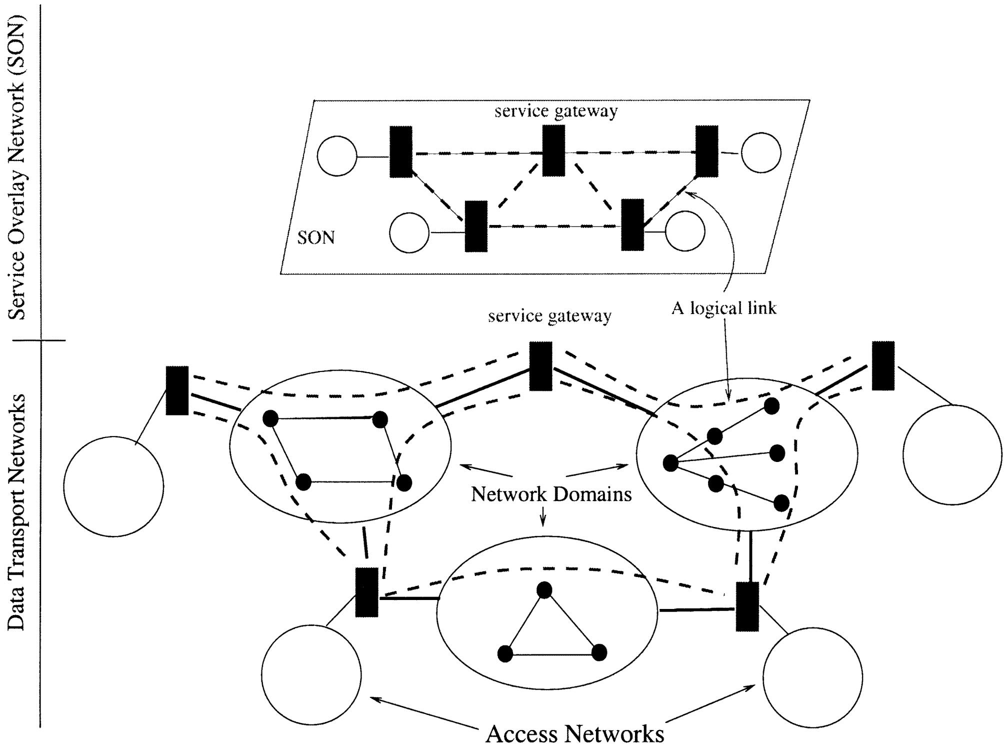

Fig. 1. Illustration of a service overlay network.

of network service management and control, especially in terms of QoS management and control. The 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—the *service aggregate*; the underlying network domains can aggregate traffic based on the SONs they belong to and perform traffic and QoS control accordingly based on the corresponding SLAs. Under the SON architecture, an SON is responsible for ensuring end-to-end QoS for its services. Because of its service awareness, an 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 the network QoS management and makes it more scalable, but also enables flexible creation and deployment of new (value-added) services.

Obviously, deployment of SON is a capital-intensive investment. It is, therefore, imperative to consider the cost recovery issue for the SON. Among the many costs the SON deployment incurs (e.g., equipment such as service gateways), a dominant recurring cost is the cost of bandwidth that the SON must purchase from the underlying network domains to support its services. The 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*, which is the focus of this 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, which due to space limitations we do not consider in this paper. In Section II, we describe how the SON logical topology can be represented under the pipe SLA model and present the model assumptions. Using the pipe SLA model, we present a basic static SON bandwidth provisioning solution in Section III, and study the problems of the more general static and dynamic SON bandwidth provisioning, respectively, in Sections IV and V. 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 industry [4], [5] as a means to deliver diverse QoS-sensitive services over the Internet. The service overlay networks we propose here is simply a generalization of these ideas. Perhaps what is particularly interesting is the use of SONs to address 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., [8], [10]) in that we explicitly take into account various factors such as SLAs, QoS, and traffic demand distributions.

II. Service Overlay Networks: Assumptions and Bandwidth Provisioning Problems

In this section, we first describe a logical topology representation of the 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 describing the traffic demand model and a few notations regarding service revenue and bandwidth cost that will be used later in this paper.

A. SON and Service QoS

The pipe SLA model is a common SLA model used in today's Internet. Under the pipe model, the 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 the service gateways and the underlying network domains, we denote the logical (unidirectional) 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 the SON and access networks where traffic to the SON originate and terminate, 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 $l = \langle u, v; D\rangle$, the 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 SON bandwidth provisioning problem 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 an 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 some specified condition. In other words, for the purpose of bandwidth provisioning, we assume that it is possible to map the service QoS guarantee requirements to a link utilization threshold.[2] 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 provisioned for link l must be such that the (average) link utilization stays below η_l.

B. Bandwidth Provisioning Modes

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

C. Traffic Demand, Service Revenue, and Bandwidth Cost

We now describe the traffic demand model for the 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 the SON can be represented by the traffic demands over these routes: for each $r \in R$, let ρ_r 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 several minutes, compared to the time scale of static bandwidth provisioning, denoted by T, which could be in several hours or days. 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 also that the traffic demands are always measured in units of bandwidth.

[2] This 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], [9], [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.

To capture the traffic demand fluctuations over time, we assume that the traffic demand ρ_r along each route r varies according to some distribution[3]. 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 the traffic demand distributions along the different routes are independent. In this paper, we will study the bandwidth provisioning problem by considering two different traffic demand models. The first one takes into account the widely observed self-similar property of the Internet traffic by employing the $M/G/\infty$ input model [12], [13]; the second is based on the measurements of real Internet traffic.

For each route r, we assume that the 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, the 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 nondecreasing function.

III. Basic Static Bandwidth Provisioning Model

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. In this section, we present a basic static bandwidth provisioning model and analyze its properties. This basic model will serve as the basis for other bandwidth provisioning models that we will consider in this paper.

In the basic model, the SON provisions bandwidth on each link based on the long-term average traffic demand matrix $\{\bar{\rho}_r\}$, and attempts to maximize the expected net income. 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 prespecified threshold η_l. 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

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

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

Given that c_l amount of bandwidth is provisioned on each link l, the expected net income of the SON is $\bar{W} = \sum_{r\in R} e_r \bar{\rho}_r - \sum_{l\in L} \Phi_l(c_l)$. Hence, the basic bandwidth provisioning problem can be formulated as the following optimization problem:

$$\max_{c_l:l\in L} \bar{W} \text{ subject to (1).}$$

Since Φ_l's are nondecreasing, it is easy to see that the optimal solution to the optimization problem is given by

$$c_l^* = \frac{(1+\epsilon_l)\bar{\rho}_l}{\eta_l} \quad \forall\, l \in L. \qquad (2)$$

Hence, under the basic bandwidth provisioning model, once we fix the overprovisioning parameters, the optimal amount of bandwidth to be provisioned for each link can be derived using (2).

Assuming that Φ_l's are sub-additive, we see that a sufficient condition for the SON to have positive expected net income is to ensure that

$$e_r > \frac{\sum_{l\in r} \Phi_l(c_l^*)}{\bar{\rho}_r} = \frac{\sum_{l\in r} \Phi_l\left(\frac{\bar{\rho}_r(1+\epsilon_l)}{\eta_l}\right)}{\bar{\rho}_r}. \qquad (3)$$

The relationship (3) provides a useful guideline for the SON to determine how it should set its price structure for charging users of its services to recover its cost of bandwidth provisioning. It has a simple interpretation: we can regard $\Phi_l(\bar{\rho}_r(1+\epsilon_l)/\eta_l)/\bar{\rho}_r$ as the average cost of carrying one unit of traffic demand per unit of time along route r on link l. Then the right-hand side of (3) is the total cost of carrying one unit of traffic demand per unit of time along route r. To recover its cost, the SON must then charge users of its services more than this amount. If Φ_l's are strictly concave (i.e., nonlinear), in other words, the per-unit bandwidth cost decreases as the amount of reserved bandwidth increases, the economy of scale will benefit the SON: the higher the average long-term traffic demands, the lower the average cost of providing its services, yielding higher net income. In the case Φ_l's are linear, i.e., $\Phi_l(c_l) = \phi_l c_l$, then (3) becomes $e_r > \sum_{l\in r} \phi_l(1+\epsilon_l)/\eta_l$ which is independent of the traffic demands.

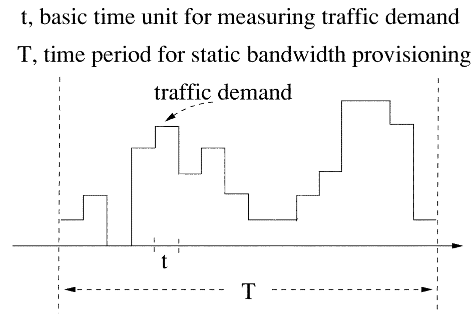

Fig. 2. Traffic demands.

IV. Static Bandwidth Provisioning With Penalty

In the basic static bandwidth provisioning model, we assume that the overprovisioning parameters are given. We now consider the problem of 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 violation when the actual traffic demands exceed the target link utilization. For this purpose, we assume that *the SON*

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

may suffer a penalty when the target utilization on a link is exceeded, and, therefore, service QoS may potentially be violated. For example, it is quite likely that the service contract between the SON and its user is such that when the service QoS is poor (e.g., due to network congestion), a lower rate is charged, or the user may demand a refund. In the case that some form of admission control is used by the SON to guide against possible QoS violation, the penalty can be used to reflect the lost revenue due to declined user service requests. We will refer to this model as the static bandwidth provisioning with penalty model, or simply, the 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 route r is potentially violated.* Given the 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 the SON for servicing the given traffic demand matrix $\{\rho_r\}$ can be expressed as

$$W(\{\rho_r\}) = \sum_{r\in R} e_r\rho_r - \sum_{l\in L}\Phi_l(c_l) - \sum_{r\in R}\pi_r\rho_r B_r(\{\rho_r\}) \quad (4)$$

where 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 will drop $\{\rho_r\}$ from the notation.

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

$$E(W) = \int\cdot\int_{\{\rho_r\}} W(\{\rho_r\})d\{\rho_r\}, \quad (5)$$

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.,

$$\max_{\{\epsilon_l\}} E(W) \text{ subject to (1)}. \quad (6)$$

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. To circumvent this difficulty, in the following, we shall derive an approximate solution (a lower bound) based on the so-called link independence assumption: the link overload events (i.e., exceeding the target utilization threshold) occur on different links independently. Clearly, this assumption does not hold in reality, but it enables us to express B_r in terms of $B_l(\rho_l, c_l)$, the probability that the target utilization level η_l on link l is exceeded, where $\rho_l = \sum_{r:l\in r}\rho_r$. (Again, we may drop the variables ρ_l and c_l in $B_l(\rho_l, c_l)$ if there is no confusion.) Such link independence assumption has been used extensively in teletraffic analysis and capacity planning in the telephone networks (see, e.g., [10]). Under the link independence assumption, the service QoS violation probability B_r, i.e., at least one of the links on route r is overloaded, is given by

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

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

$$\int_{\hat{\rho}_r}^{\infty} \rho_r d\rho_r \le \delta. \quad (8)$$

Since $\int_{\hat{\rho}_r}^{\infty}\rho_r d\rho_r \ge \hat{\rho}_r\int_{\hat{\rho}_r}^{\infty} d\rho_r = \hat{\rho}_r Pr\{\rho_r \ge \hat{\rho}_r\}$, we have $Pr\{\rho_r \ge \hat{\rho}_r\} \le \delta/\hat{\rho}_r$. In other words, (8) 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 Appendix I for the detailed derivation):

$$E(W) \ge \sum_{r\in R} e_r\bar{\rho}_r - \sum_{l\in L}\Phi(c_l) - \sum_{r\in R}\pi_r\bar{\rho}_r B_r(\{\hat{\rho}_r\}) - \sum_{r\in R}\pi_r\delta\left(1 + \sum_{r'\ne r}\frac{\bar{\rho}_r}{\hat{\rho}_{r'}}\right). \quad (9)$$

Denote the right-hand side of the above equation by V, then $E(W) \ge V$. Comparing the lower bound V with the expected net income $\bar{W} = \sum_{r\in R} e_r\bar{\rho}_r - \sum_{l\in L}\Phi_l(c_l)$ without taking penalty into account, we see that ignoring the extremal traffic demands (i.e., when $\rho_r \ge \hat{\rho}_r$), we pay at most a penalty of $\pi_r B_r(\{\hat{\rho}_r\})$ per unit of traffic demand on route r for potential service QoS violations. For given $\delta > 0$, the penalty incurred due to extremal traffic demands is upper bounded by $\sum_{r\in R}\pi_r\delta(1 + \sum_{r'\ne r}\bar{\rho}_r/\hat{\rho}_{r'})$. Note also that $B_r(\{\hat{\rho}_r\})$ is the probability of service QoS violation along route r when the long-term average traffic demands are assumed to be $\hat{\rho}_r$. Thus, in using V as an approximation to $E(W)$, we are being conservative by overestimating the probability of potential QoS violations.

From $E(W) \ge V$, we have $\max_{\{\epsilon_r\}} E(W) \ge \max_{\{\epsilon_r\}} V$. 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 (6). Using the solution to the basic bandwidth provisioning problem (2), we assume $c_l = (1+\epsilon_l)\bar{\rho}_l/\eta_l$ for a given set of $\{\epsilon_l\}$, i.e., the target utilization constraints (1) hold with equality. Under this assumption, let $\{\epsilon_l^*\}$ be the solution to the optimization problem $\max_{\{\epsilon_r\}} V$, and refer to them as the approximate optimal overprovisioning parameters. In the following, we demonstrate how $\{\epsilon_l^*\}$ can be derived.

Using (7), we can rewrite V as

$$V = \sum_{r\in R}(e_r - \pi_r)\bar{\rho}_r - \sum_{l\in L}\Phi_l(c_l) + \sum_{r\in R}\pi_r\bar{\rho}_r\prod_{l\in r}(1 - B_l(\hat{\rho}_l, c_l)) - \sum_{r\in R}\pi_r\delta\left(1 + \sum_{r'\ne r}\frac{\bar{\rho}_r}{\hat{\rho}_{r'}}\right) \quad (10)$$

where $\hat{\rho}_l = \sum_{r:l\in r} \hat{\rho}_r$.

Assume B_l is a continuous and everywhere differentiable function of c_l. (See Section V for a discrete case.) For each link l, define

$$\hat{s}_l = \sum_{r:l\in r} \pi_r \bar{\rho}_r \prod_{k\in r, k\neq l} [1 - B_k(\hat{\rho}_k, c_k)]\, \zeta_l \qquad (11)$$

where $\zeta_l = -dB_l(\hat{\rho}_l, c_l)/dc_i$.

Through some simple algebraic manipulation, it is not too difficult to show that

$$\frac{\partial V}{\partial \epsilon_l} = \frac{\partial V}{\partial c_l}\frac{\partial c_l}{\partial \epsilon_l} = \left(-\frac{\partial \Phi_l(c_l)}{\partial c_l} + \hat{s}_l\right)\frac{\bar{\rho}_l}{\eta_l}. \qquad (12)$$

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. Thus, from (12) we must have

$$\frac{\partial \Phi_l(c_l)}{\partial c_l} = \hat{s}_l \quad \forall\, l \in L. \qquad (13)$$

Intuitively, $\hat{s}_l$ measures the sensitivity of potential penalty reduction to bandwidth increase on link l, whereas $\partial\Phi_l(c_l)/\partial c_l$ measures the sensitivity of bandwidth cost to bandwidth increase on link l. Hence, the optimal (or rather, the approximate optimal) overprovisioning parameter ϵ_l^* should be chosen such that the two values coincide. In the following discussion, we will loosely refer to $\hat{s}_l$ as the per-unit bandwidth gain in potential penalty reduction and to $\partial\Phi_l(c_l)/\partial c_l$ as the increase in per-unit bandwidth cost.

In the above derivation of the approximate optimal solution to the static bandwidth provisioning problem, we have simply assumed the existence of B_l, the probability that the target utilization level η_l on link l is exceeded. The particular form of it depends on the distribution of (average) traffic demands on the link. In Subsections IV.A and IV.B, we consider two different traffic demand models—a self-similar traffic demand model and a traffic demand model based on real Internet traffic measurements to demonstrate the static bandwidth provisioning problem.

A. $M/G/\infty$ *Traffic Demand Model*

Since the pioneering work of [11], the self-similar (or long-range dependent) property has been observed in Ethernet local-area network [11], wide-area network [13], and World Wide Web traffic [7]. The observed self-similar property of the Internet traffic has important implications on the dimensioning and provisioning of the IP networks. In this section, we consider a traffic demand model, $M/G/\infty$, that captures the (asymptotically) self-similar property of the Internet traffic [12], [13].

Consider an $M/G/\infty$ queue, 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 $t = 0, 1, 2, \ldots$. Then the count process $\{X_t\}_{t=0,1,2,\ldots}$ is asymptotically self-similar. Let ρ denote the customer arrival rate to the $M/G/\infty$ queue 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 on each route. Consider an arbitrary route r. We assume that the traffic demand (i.e., the average traffic arrival rate per unit time) is governed by the count process $\{X_t\}_{t=0,1,2,\ldots}$ of an $M/G/\infty$ queue. Let ρ_r denote the mean traffic demand on the route. It is easy to see that $\rho_r = \rho\mu$, where ρ and μ are the customer arrival rate and the mean service time, respectively, of the $M/G/\infty$ queue. As traffic demands along all the routes are assumed to be independent, the average overall traffic load on a link l is $\rho_l = \sum_{r:l\in r} \rho_r$.

Given the average overall load ρ_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 $B_l(\rho_l, c_l) = (\sum_{i=(\bar{c}_l+1)}^{\infty} \rho_l^i/i!)e^{-\rho_l}$. Extending the definition of $B_l(\rho_l, c_l)$ to noninteger values of c_l by linear interpolation, at integer values of c_l define the derivative of $B_l(\rho_l, c_l)$ with respect to c_l to be the left derivative. Then $dB_l(\rho_l, c_l)/dc_l = B_l(\rho_l, c_l) - B_l(\rho_l, c_l - 1)$. Therefore

$$\begin{aligned}\zeta_l &= -\frac{d}{dc_l} B_l(\hat{\rho}_l, c_l)\\ &= \eta_l\{B_l(\hat{\rho}_l, (\eta_l c_l - 1))\\ &\qquad - B_l(\hat{\rho}_l, \eta_l c_l)\}\\ &= \eta_l \frac{\hat{\rho}_l^{\lceil \eta_l c_l \rceil}}{\lceil \eta_l c_l \rceil!} e^{-\hat{\rho}_l}.\end{aligned}$$

By this definition of B_l, we are able to obtain the (approximate) optimal overprovisioning parameters ϵ_l^*'s by solving (13).

We now discuss 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 particular, there exists $\hat{\epsilon}_l$ such that $\hat{s}_l$ 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 toward $\hat{\epsilon}_l$, 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 $\partial\Phi_l(c_l)/\partial c_l = \phi_l$. Hence, (13) becomes $\phi_l = \hat{s}_l$. Suppose $\phi_l = \hat{s}_l$ holds for some $\epsilon_l \geq 0$. Because of the shape of $\hat{s}_l$, there potentially exists two solutions $\epsilon_{l,1}$ and $\epsilon_{l,2}, 0 \leq \epsilon_{l,1} \leq \hat{\epsilon}_l \leq \epsilon_{l,2}$ such that $\phi_l = \hat{s}_l$. In particular, as $\hat{s}_l$ is a decreasing function in the range $[\hat{\epsilon}_l, \infty)$, $\epsilon_{l,2}$ always exists. As $\partial V/\partial c_l$ 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 $\epsilon_{l,1}$). Intuitively, when only a small amount of bandwidth is overprovisioned on link l, the per-unit bandwidth gain in potential penalty reduction is too small to offset the per-unit bandwidth cost, hence, V decreases. However, as we increases the amount of bandwidth overprovisioned, the per-unit bandwidth gain in potential penalty reduction becomes sufficiently large and offsets the per-unit bandwidth cost, hence, V increases

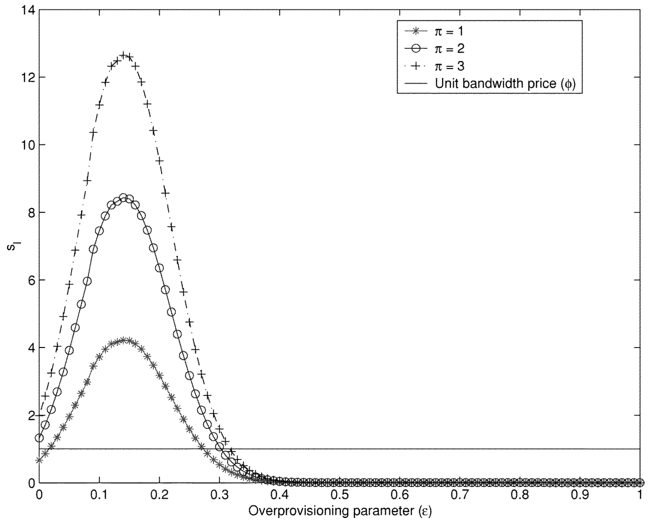

Fig. 3. Relationship between $\hat{s}_l$, ϵ, & ϕ_l.

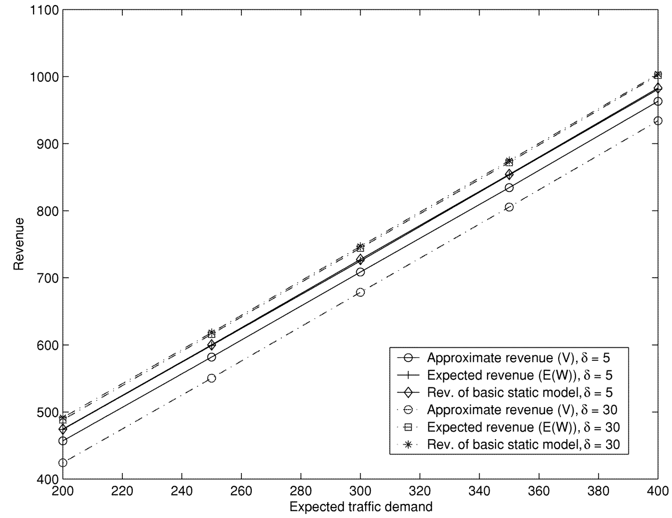

Fig. 4. Comparison of V and $E(W)$.

until it reaches a maximum. Due to the diminished return in the per-unit bandwidth gain in potential penalty reduction, V decreases again when too much bandwidth is overprovisioned on link c. In the special case that ϕ_l is such that $\phi_l > \hat{s}_l$ for all $\epsilon_l \geq 0$, then as $\partial V/\partial c_l < 0$, V attains its maximum at $\epsilon_l^* = 0$ with respect to link l. Intuitively it says that when the per-unit bandwidth cost on link l is higher than the per-unit bandwidth gain in potential penalty reduction, there is no benefit in overprovisioning any bandwidth on link l to guide against any potential QoS violation penalty. These observations can be extended to other bandwidth cost functions such as concave or convex cost functions. In general, we see that the tradeoff between the bandwidth cost and overprovisioning bandwidth to guide against service QoS violations is critical to the problem of SON bandwidth provisioning. It is also clear from the above discussion that as the per-unit bandwidth cost decreases, there is more benefit in overprovisioning. Finally, we comment that from (11) and (13) and the above observations, we can compute the approximate optimal overprovisioning parameters ϵ_l^*'s using fixed-point approximation.

1) Numerical Examples: 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. Numerical studies in more complex settings will be performed in a later section.

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 $\hat{s}_l$ as a function of ϵ_l with three different values of π_r, namely, $\pi_r = 1, 2, 3$. 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 $\hat{s}_l = \phi_l$. Recall from Section IV 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 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.

In Fig. 4, we compare the lower bound V with the actual expected net income $E(W)$ for two given values of δ ($\delta = 5$ and 30). For comparison, we also include the expected net income $\bar{W}$ under the basic static model, where the overprovisioning parameter ϵ_l^* is obtained from the static-penalty model. From the figure, we see that for both values of δ, the lower bound V provides a reasonable approximation to $E(W)$. Note also that the difference between the actual expected net income $E(W)$ under the static-penalty model and the expected net income $\bar{W}$ under the basic static is almost invisible. This is likely due to the fact that the additional revenue generated when the traffic demand exceeds its long-term average (the first term in $E(W)$) and the potential penalty incurred due to service QoS violations (the third term in $E(W)$) cancel each other out on average. From Fig. 4, it is clear that the lower bound depends on the choice of δ. The smaller the δ, the closer the approximate revenue V is to the expected revenue $E(W)$. To further explore the relation between δ and V, in Fig. 5 we plot V as a function of δ (upper plot). In the figure, we also include the overprovisioning parameter ϵ_l^* as a function of δ (lower plot). We see that V is a concave function of δ, and thus there is a unique δ that maximizes V. On the other hand, ϵ_l^* is a nonincreasing function of δ.

To highlight the relationship between bandwidth cost and overprovisioning, in Fig. 6 we plot the overprovisioning parameter ϵ_l^* as a function of the per-unit bandwidth cost ϕ_l. 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.

B. Measurement-Based Traffic Demand Model

A key property of the presented approximate optimal solution to the static bandwidth provisioning problem is that it only relies on the marginal distribution of the traffic demand on each link. In this section, we will study the static bandwidth provisioning problem based on the measurements of real Internet traffic. That

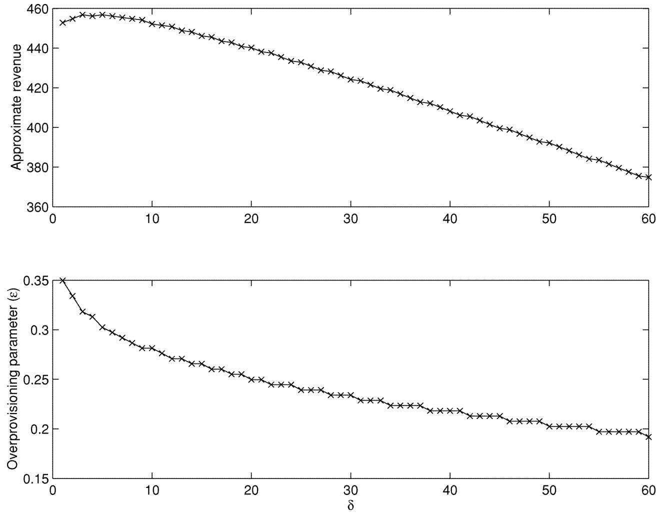

Fig. 5. Impact of δ on V and ϵ^*.

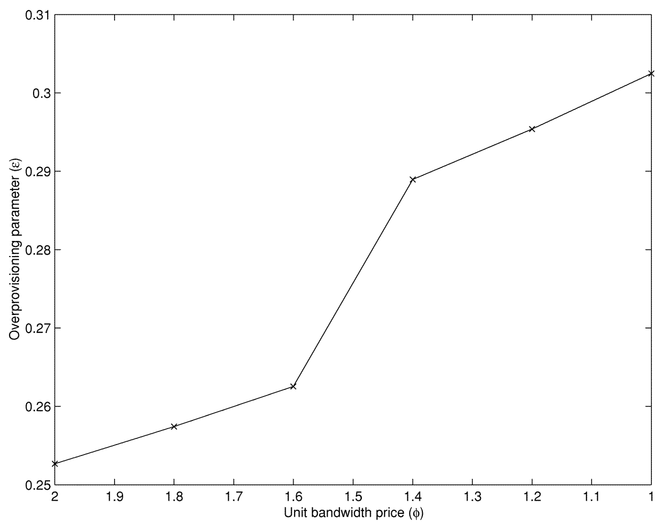

Fig. 6. Impact of unit bandwidth price on ϵ^*.

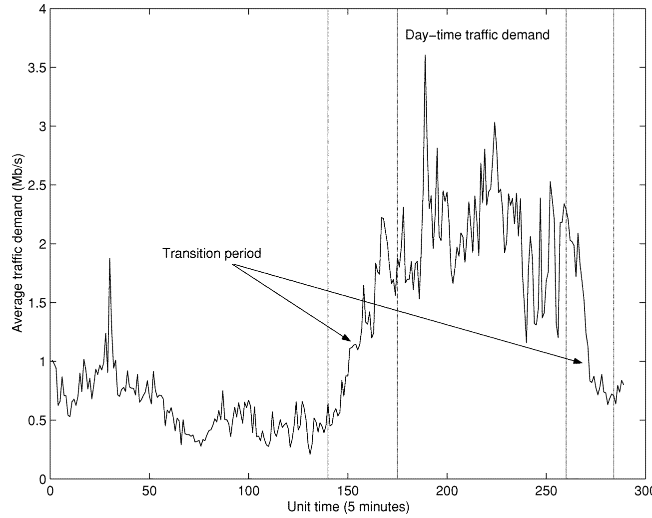

Fig. 7. Traffic demands of the Auckland data trace.

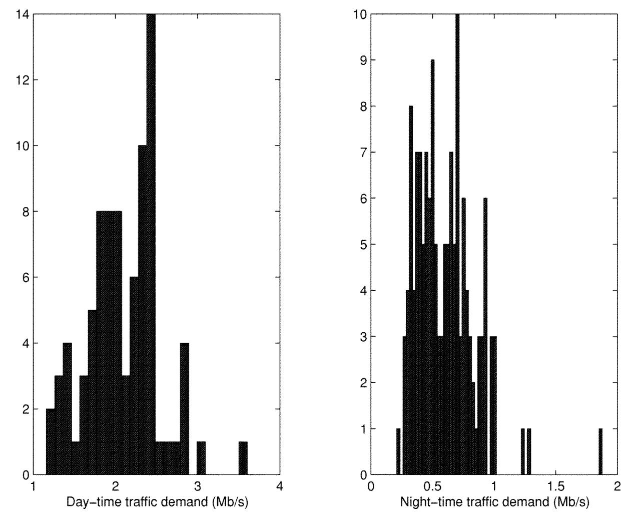

Fig. 8. Histogram of the Auckland data trace's traffic demands.

is, we estimate the marginal distributions of the traffic demands on the links by the long-term measurements of the traffic data, and then apply the estimated marginal distributions of the traffic demands to our static bandwidth provisioning problem.

The data trace we will use was collected at the University of Auckland Internet access link on December 1, 1999, and lasted roughly for 24 hours (referred to as the *Auckland data trace*) [15]. In the Auckland data trace, there are in total 32 628 004 packet arrivals. Fig. 7 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-min time interval (which is also used as the base unit of time, i.e., $t = 5$ min; see Fig. 2). Given the largely different traffic arrival patterns during the daytime and nighttime, we will accordingly provision bandwidth differently for them, where the daytime is defined to be from 10:00 AM to 5:00 PM and nighttime from 7:00 PM to 7:00 AM. We will refer to the traffic demands during the daytime and nighttime as *daytime traffic demand* and *nighttime traffic demand*, respectively. All other times are considered to be transition times. The bandwidth provisioned for the daytime and nighttime should be switched during the transition times based on certain criteria, which is not considered in this paper.

Now let us consider the properties of the daytime and the nighttime traffic demands. The mean traffic arrival rate is 2.1 Mb/s over the whole daytime duration, and 0.6 Mb/s over the nighttime duration. Fig. 8 plots the histograms of the traffic demands for the daytime (left-hand side) and nighttime (right-hand side) separately, where the bin sizes for the daytime traffic demands and the nighttime traffic demands are 100 and 50 kb/s, respectively. From the plots, we see that the daytime traffic demands are relatively symmetrically centered at its mean arrival rate, while the nighttime traffic demands are more skewed. In the following studies, we will model the daytime traffic demands by a *Normal* distribution and the nighttime traffic demands by a *Lognormal* distribution to retain the different traffic characteristics during the daytime and the nighttime. Table I presents the mean traffic demands and the standard deviations (STD) of the daytime and nighttime traffic demands, where the base unit of bandwidth (traffic demand) is 1 kb/s.

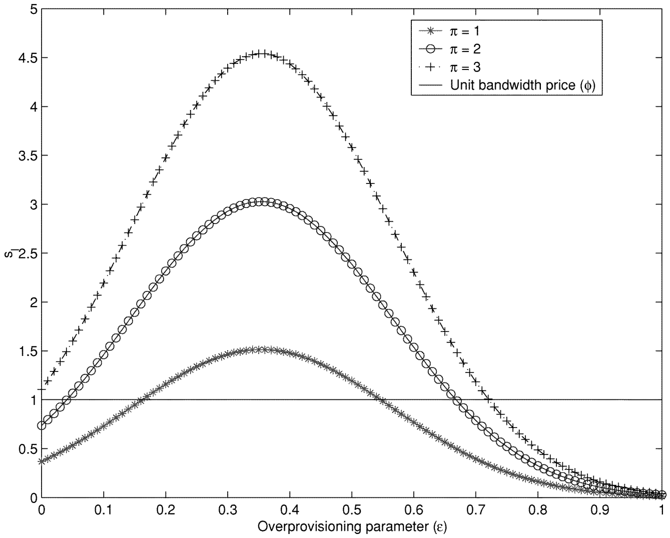

Fig. 9. Relationship between $\hat{s}_l$, ϵ, and ϕ_l for daytime traffic.

In the following, we will conduct numerical studies to illustrate the static bandwidth provisioning using the Auckland data trace. In all these studies, we again consider the simple setting: a single route over a single link. The per-unit bandwidth per-unit time earning $e_r = 4$, and $\phi_l = 1$, $\pi_r = 2$. We set the target utilization threshold $\eta_l = 0.8$.

Similar to the numerical example for the $M/G/\infty$ traffic demand model, in Fig. 9, we show $\hat{s}_l$ as a function of ϵ_l with three different values of π_r, namely, $\pi_r = 1, 2, 3$, for the daytime traffic demands. The value of δ used is 140. 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 $\hat{s}_l = \phi_l$ [see (13)]. Following a similar argument as that in Section IV-A, there potentially exists two solutions $\epsilon_{l,1}$ and $\epsilon_{l,2}$, $0 \leq \epsilon_{l,1} \leq \epsilon_{l,2}$ such that $\phi_l = \hat{s}_l$. Moreover, with respect to link l, V is maximized at either $\epsilon_l^* = \epsilon_{l,2}$ or at $\epsilon_l^* = 0$. From Fig. 9, we can draw the similar conclusions as that in the $M/G/\infty$ traffic demand model. In particular, we see that as the penalty π_r increases, ϵ_l^* also increases. Hence, for 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. However, compared with the result in Fig. 3, we see that we obtain larger overprovisioning parameters here. This is caused by the high traffic fluctuation in the Auckland data trace. Table I gives the *coefficient of variance* (C.O.V.) for the daytime and the nighttime traffic demands. This value (0.21) is much higher than that in Fig. 3, which is 0.07.

To compare the different provisioning behaviors during the daytime and nighttime, we present the overprovisioning parameters for both the daytime and nighttime traffic demands in Table I. To obtain these results, we have searched for the best δ's that yield the maximal V's, respectively. In the table, we also include the approximate revenue V's (per-unit time) for the daytime and nighttime traffic demands. From the table, we see that for the daytime traffic demands, the overprovisioning parameter $\epsilon_l^* = 0.67$, while for the nighttime traffic demands $\epsilon_l^* = 0$. The reason is that even though the average traffic demands during nighttime are much lower than that during daytime, we observe a much higher traffic demand fluctuation during the nighttime than during the daytime (see Table I for their corresponding coefficients of variance). It is too expensive to accommodate this high traffic demand variance during the nighttime ($\epsilon_{l,2}$ is dramatically large), therefore, no overprovisioning is provided in this case. During daytime, the (per-unit time) approximate revenue is 3446, which is higher than that during the nighttime (1672). This is not unexpected.

TABLE I
PROVISIONING FOR THE AUCKLAND TRAFFIC DEMANDS

	Mean	STD	C.O.V.	ϵ_l^*	V
Day-time	2096	442	0.21	0.67	3446
Night-time	609	240	0.39	0	1672

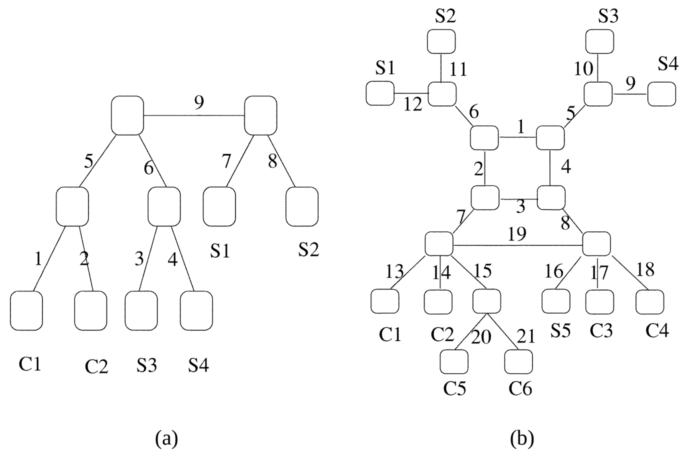

Fig. 10. SON topologies. (a) Tree. (b) Mesh tree.

C. Performance Evaluation

We now use two SON topologies—the *tree* [Fig. 10(a)] and the mesh-tree [Fig. 10(b)] topologies—to illustrate the effect of traffic load distribution among various routes of an SON on static bandwidth provisioning. 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 case there are two such paths, only one is chosen. In the numerical studies below, we will use the $M/G/\infty$ traffic demand model. We set $e_r = 10$, $\pi_r = 2$ for all the routes, and $\phi_l = 1$ for all the links. The value of δ is chosen in such a way that $\delta_r = 1/40\rho_r$.

In the tree topology, four routes are used: $R1 = S3 \rightarrow C1$, $R2 = S1 \rightarrow C1$, $R3 = S4 \rightarrow C2$, and $R4 = S2 \rightarrow C2$. To investigate the effects of different traffic loads on bandwidth provisioning, we consider two types of traffic load distribution among the routes: the *balanced* load where the expected traffic demand for all routes is 200, and the *unbalanced* load where the expected traffic demands on routes $R1$, $R2$, $R3$, and $R4$ are 300, 100, 250, and 150, respectively. Table II presents the resulting overprovisioning parameter ϵ_l^* and provisioned bandwidth c_l for six representative links: link 1, 4, 5, 7, 8, and 9. The corresponding average traffic loads $\bar{\rho}_l$'s on these four links

TABLE II
TREE TOPOLOGY

Link ID		1	4	5	7	8	9
Balanced	ρ_l	400	200	800	200	200	400
	ϵ_l^*	0.26	0.3	0.23	0.3	0.3	0.26
	c_l	630	325	1230	325	325	630
Unbalanced	ρ_l	400	250	800	100	150	250
	ϵ_l^*	0.26	0.27	0.23	0.41	0.34	0.33
	c_l	630	397	1230	176	251	416

TABLE III
MESH-TREE TOPOLOGY

Link ID		2	6	11	18	19	21
Balanced	ρ_l	1200	800	400	400	400	200
	ϵ_l^*	0.22	0.23	0.26	0.26	0.26	0.3
	c_l	1830	1230	630	630	630	325
Unbalanced	ρ_l	1350	1100	500	400	400	100
	ϵ_l^*	0.21	0.2	0.24	0.26	0.26	0.41
	c_l	2042	1650	775	630	630	176

are also given in the table. From the results, we see that under the balanced load, links with a higher average traffic load have a smaller overprovisioning parameter. This is due to statistical multiplexing gains for carrying a higher load on a link. In the unbalanced case, similar results can be observed. Note that even though links 4 and 9 have the same traffic demand load, they are overprovisioned differently. This is because there are two routes traversing link 9 while there is only one on link 4.

We now consider the mesh-tree topology. In this case, there are ten routes: $R1 = S1 \rightarrow C1$, $R2 = S2 \rightarrow C2$, $R3 = S3 \rightarrow C1$ (1), $R4 = S4 \rightarrow C2$ (1), $R5 = S1 \rightarrow C3$ (3), $R6 = S2 \rightarrow C4$ (3), $R7 = S3 \rightarrow C3$, $R8 = S4 \rightarrow C4$, $R9 = S5 \rightarrow C5$, and $R10 = S5 \rightarrow C6$. The number in the parentheses following a route shows a link that the route traverses in case there are multiple paths between the source and destination with the same path length. Again for the balanced load case, all the routes have an average traffic demand of 200, while for the unbalanced load case, the average demands for routes $R1$ to $R10$ are 300, 250, 100, 150, 300, 250, 100, 150, 300, and 100, respectively. Table III shows the results for six representative links: link 2, 6, 11, 18, 19, and 21. From the table, we can see that similar observations also hold for the mesh-tree topology.

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 Section V, we will study the dynamic bandwidth provisioning mode, where the link bandwidth can be dynamically adjusted according to the traffic demand fluctuations in relatively shorter time intervals.

V. DYNAMIC BANDWIDTH PROVISIONING

In this section, we study the dynamic bandwidth provisioning problem. As pointed out in Section II, to account for the potential higher cost in supporting dynamic bandwidth provisioning, it is likely that the underlying network domains will charge the 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 the SON is to determine how much bandwidth should be reserved statically on each link l a priori to meet certain base traffic demands, while dynamically allocating bandwidth to meet the additional traffic demands as needed*. The objective is again to maximize the overall long-term expected net income of the SON.

To focus on the dynamic bandwidth problem, we assume that the underlying network domains possess abundant bandwidth that the dynamic requests for additional bandwidth from the 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 Section IV. 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 the overlay network, and then present an adaptive heuristic algorithm for dynamic bandwidth provisioning based on online traffic measurements.

A. 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 reserved to meet the traffic demands by the following expression:

$$\Delta c_l = \left\{ \frac{\rho_l}{\eta_l} - c_l \right\}^+ \tag{14}$$

where $\rho_l = \sum_{l \in r} \rho_r$. Then $\Delta c_l > 0$ if and only if $\rho_l > \eta_l c_l$.

Using (14), we can write the approximate overall net income the SON generates for the given traffic demand matrix $\{\rho_r\}$:

$$\tilde{W}(\{\rho_r\}) = \sum_{r \in R} e_r \rho_r - \sum_{l \in L} \Phi_l(c_l) - \sum_{l \in L} \Phi'_l(\Delta c_l). \tag{15}$$

Integrating on both sides of (15) over the (joint) distribution of $d\{\rho_r\}$, we have

$$E(\tilde{W}) = \sum_{r \in R} e_r \bar{\rho}_r - \sum_{l \in L} \Phi_l(c_l) - \sum_{l \in L} \int \cdot \int \Phi'_l(\Delta c_l) d\{\rho_r\}. \tag{16}$$

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

$$\max_{\{c_l\}} E(\tilde{W}). \tag{17}$$

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 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 (17) 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, please refer to the technical report version of this paper [18].

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 of c'_l's is an (approximate) solution to the dynamic bandwidth provisioning problem, i.e., c'_l is the amount of bandwidth to be statically provisioned, while the portion to be dynamically allocated on link l is given by (14), 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 l, 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 capacity needs to be reserved.

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 provisioned statically *a priori* c_l 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. 11 presents the bandwidth provisioned statically c_l (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 bandwidth price for dynamically allocated bandwidth increases, more bandwidth needs to be provisioned statically *a priori*. However, the increase in the amount of static bandwidth is not dramatic 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 partly due to the fact that an SON needs to statically provision more bandwidth *a priori* on each link, as well as the fact that the SON needs to pay more for the 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 the approximate revenues obtained, using the tree network topology [Fig. 10(a)], with a similar setting as that in Section IV-C. In particular, we use the balanced

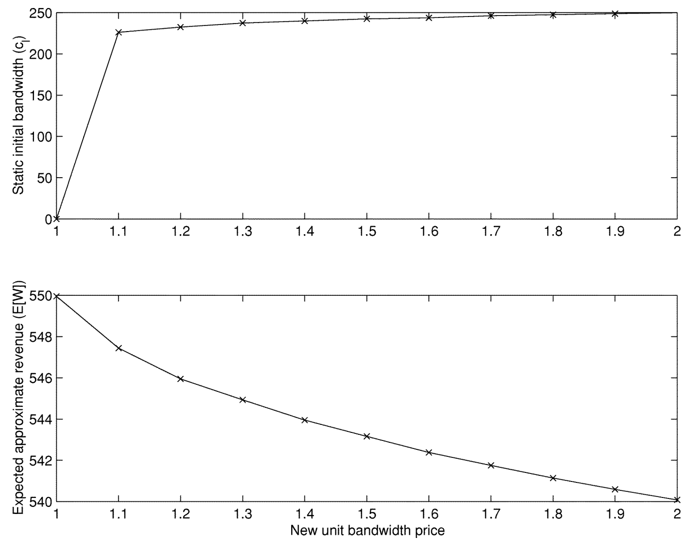

Fig. 11. Effects of ϕ'_l on c_l and $E(\tilde{W})$.

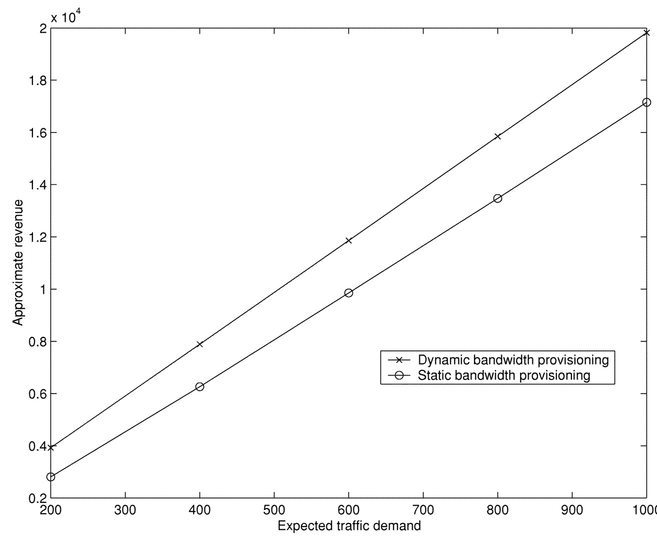

Fig. 12. Dynamic versus static bandwidth provisioning.

traffic load model and assume the traffic demand on each route is governed by the $M/G/\infty$ model. For static bandwidth provisioning, $\pi_r = 2$. Fig. 12 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, and the dynamic bandwidth provisioning has a higher approximate revenue than that of the static bandwidth provisioning. Moreover, as the average traffic demand increases, the difference between the approximate revenues of the dynamic bandwidth provisioning and the 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 in this case.

B. 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 [see (14)]. In this section, we present an adaptive online bandwidth provisioning model (or, simply, online dynamic model) that dynamically adjust the allocated bandwidth on a link according to the measurement of the traffic demands on the links of the network.

As before, let $\bar{\rho}_r$ denote the long-term average traffic demand on route r, and $\bar{\rho}_l = \sum_{r:l\in r} \bar{\rho}_r$ denote the long-term average traffic demand on link l. Based on the measurement of the 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 the network.

Let t denote a fixed time interval. In the online dynamic model, the average traffic demand ρ_r 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 resulted 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 following, 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 *a priori*. In the online dynamic model, c_l is chosen in such a manner that, if the average traffic demand on a link l does not exceed $\bar{\rho}_l$, the service QoS will be honored, i.e.,

$$c_l = \lceil \frac{\bar{\rho}_l}{\eta_l \Theta} \rceil \Theta. \tag{18}$$

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

Next, we discuss the allocation of additional bandwidth and deallocation of extra bandwidth on an arbitrary link l. To reduce the possibility that the service QoS is violated, the online dynamic model will allocate the additional bandwidth (a new quota) when 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 if $\rho_l > C_l \eta_l - \iota^f$, where ι^f is the forward threshold for allocating a new quota. Similarly, a backward threshold for deallocating an extra quota is defined as (denoted by ι^b (a positive number)): an extra quota is released from link l only if $\rho_l < (C_l - \Theta)\eta_l - \iota^b$.

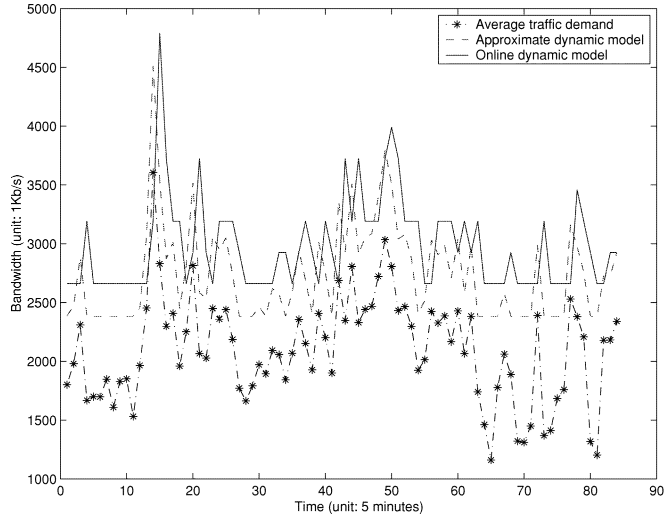

Fig. 13. Dynamic bandwidth provisioning with approximate model and online model.

Because the online dynamic model only adjusts bandwidth on the links at the end of the each measurement interval, it is possible that the service QoS is violated during the course of the measurement time interval. As in static bandwidth provisioning with penalty in Section IV, 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 route* r *is violated*. Then the revenue of the online dynamic model for a measurement time interval is

$$\bar{V} = \sum_{r\in R} e_r \rho_r - \sum_{l\in L} \Phi_l(c_l) - \sum_{l\in L} \Phi'_l(\Delta c_l) - \sum_{r\in R} \pi_r \rho_r \mathbf{1}_{\{\rho_l/C_l > \eta_l : l\in r\}} \tag{19}$$

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. The studies are carried out in the simple network setting using the daytime traffic demands of the Auckland data trace (see Fig. 7). The following parameters are used. The base unit of bandwidth for the Auckland data trace is 1 kb/s. The measurement time interval (i.e., unit time) is 5 min. 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 quota $\Theta = 0.6\sigma$, where σ is the standard deviation of the daytime traffic demands of the Auckland data trace (see Table II). The forward and backward threshold $\iota^f = \iota^b = 0.3\Theta$.

Fig. 13 presents the average traffic demands (per 5 min) and the corresponding provisioned 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

TABLE IV
PER-UNIT TIME AVERAGE REVENUE

	Approximate model	Online model
Average revenue	5468	4152

of the traffic demands on the link and, meanwhile, remains insensitive to small shorttime fluctuations in traffic demands (for example, see the provisioned bandwidth at times 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.) Note also that, under this parameter setting, the approximate dynamic model has a smaller initial static bandwidth than the online dynamic model. Moreover, the approximate dynamic model is more sensitive to the fluctuations in traffic demands than the online dynamic model.

Table IV gives the mean revenues (per-unit time) of the approximate dynamic model and the online dynamic model, averaged over the whole duration of the daytime traffic demands of the Auckland data trace. From the table, we see that the approximate dynamic model has a higher per-unit time average revenue than the online dynamic model. There are two possible reasons. First, under this parameter setting, the amount of initial static bandwidth is larger than the approximate dynamic model, therefore, it results in more cost on the overlay. Second, 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 could be violated during a time interval and incurs penalty on the overlay. However, the online dynamic model has the advantage that it does not make any assumption about the (average) traffic demands (except the long-term average traffic demand and its standard deviation).

VI. CONCLUSIONS AND FUTURE WORK

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

The approximate optimal solution we presented 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 environment, the dynamic bandwidth provisioning model outperforms the static bandwidth provisioning model, albeit with more complex and frequent network resource managements. We investigated the effects of various parameters like static and dynamic bandwidth costs on the revenue that an SON can obtain, which provides useful guidelines on how an SON should be provisioned to stay profitable.

In this paper, we have assumed the route between a source gateway and a destination gateway is predetermined. Currently, we are investigating the functionalities of the service gateways in support of service-aware (multipath) routing, which will have great impact on how an SON should be provisioned.

APPENDIX I
A LOWER BOUND ON $E(W)$ OF THE STATIC BANDWIDTH PROVISIONING WITH PENALTY

From (4) and (5), it is easy to see that

$$E[W] = \sum_{r\in R} e_r\bar{\rho}_r - \sum_{l\in L}\Phi_l(c_l) - \sum_{r'\in R}\int\cdot\int_{\{\rho_r\}} \pi_{r'}\rho_{r'}B_{r'}(\{\rho_r\})d\{\rho_r\}. \quad (20)$$

Moreover

$$\begin{aligned}\sum_{r'\in R}\int\cdot\int_{\{\rho_r\}} &\pi_{r'}\rho_{r'}B_{r'}(\{\rho_r\})d\{\rho_r\}\\ &\le \sum_{r'\in R}\int\cdot\int_{\{0\}}^{\{\hat{\rho}_r\}} \pi_{r'}\rho_{r'}B_{r'}(\{\rho_r\})d\{\rho_r\}\\ &\quad + \sum_{r'\in R}\sum_{r''\in R, r''\ne r'}\int_{\hat{\rho}_{r''}}^{\infty}\left(\int\cdot\int_{\{0\}_{R-r''}}^{\{\infty\}}\right)\\ &\quad \times \pi_{r'}\rho_{r'}B_{r'}(\{\rho_r\})d\{\rho_r\}\\ &\quad + \sum_{r'\in R}\int_{\hat{\rho}_{r'}}^{\infty}\left(\int\cdot\int_{\{0\}_{R-r'}}^{\{\infty\}}\right)\\ &\quad \times \pi_{r'}\rho_{r'}B_{r'}(\{\rho_r\})d\{\rho_r\}. \end{aligned} \quad (21)$$

As $B_{r'}(\{\rho_r\}) \le B_{r'}(\{\hat{\rho}_r\})$ when $\rho_r \le \hat{\rho}_r, \forall\, r$

$$\begin{aligned}\int\cdot\int_{\{0\}}^{\{\hat{\rho}_r\}} &\pi_{r'}\rho_{r'}B_{r'}(\{\rho_r\})d\{\rho_r\}\\ &\le \int\cdot\int_{\{0\}}^{\{\hat{\rho}_r\}} \pi_{r'}\rho_{r'}B_{r'}(\{\hat{\rho}_r\})d\{\rho_r\}\\ &\le \pi_{r'}B_{r'}(\{\hat{\rho}_r\})\bar{\rho}_{r'}. \end{aligned} \quad (22)$$

Notice $B_{r'}(\{\rho_r\}) \le 1$ and the definition of δ, we have (note $r'' \ne r'$)

$$\begin{aligned}\int_{\hat{\rho}_{r''}}^{\infty}\left(\int\cdot\int_{\{0\}_{R-r''}}^{\{\infty\}}\right)&\pi_{r'}\rho_{r'}B_{r'}(\{\rho_r\})d\{\rho_r\}\\ &\le \int_{\hat{\rho}_{r''}}^{\infty}\left(\int_0^{\infty}\pi_{r'}\rho_{r'}d\rho_{r'}\right)d\rho_{r''}\\ &\le \pi_{r'}\delta\frac{\bar{\rho}_{r'}}{\hat{\rho}_{r''}}. \end{aligned} \quad (23)$$

Similarly

$$\begin{aligned}\int_{\hat{\rho}_{r'}}^{\infty}\left(\int\cdot\int_{\{0\}_{R-r'}}^{\{\infty\}}\right)&\pi_{r'}\rho_{r'}B_{r'}(\{\rho_r\})d\{\rho_r\}\\ &\le \int_{\hat{\rho}_{r'}}^{\infty}\pi_{r'}\rho_{r'}d\rho_{r'} \le \pi_{r'}\delta. \end{aligned} \quad (24)$$

Substituting (22), (23), and (24) into (21), and then recursively into (20), we have

$$E(W) \geq \sum_{r\in R} e_r\bar{\rho}_r - \sum_{l\in L}\Phi(c_l) - \sum_{r\in R}\pi_r\bar{\rho}_r B_r(\{\hat{\rho}_r\}) - \sum_{r\in R}\pi_r\delta\left(1+\sum_{r'\in R, r'\neq r}\frac{\bar{\rho}_r}{\hat{\rho}_{r'}}\right). \quad (25)$$

ACKNOWLEDGMENT

The authors would like to thank the anonymous reviewers of IEEE ICNP 2002 and this TRANSACTIONS for many valuable comments.

REFERENCES

[1] D. G. Andersen, H. Balakrishnan, M. Kaashoek, and R. Morris, "Resilient overlay networks," in *Proc. 18th ACM SOSP*, Banff, Canada, Oct. 2001, pp. 131–145.

[2] A. Charny and J.-Y. Le Boudec, "Delay bounds in a network with aggregate scheduling," in *Proc. QoFIS*, Berlin, Germany, Oct. 2000, pp. 1–13.

[3] Y. Chawathe, S. Fink, S. McCanne, and E. Brewer, "A proxy architecture for reliable multicast in heterogeneous environments," in *Proc. ACM Multimedia*, Bristol, U.K, Sept. 1998, pp. 151–159.

[4] Virtela Communications. [Online]. Available: http://www.virtele.com

[5] Internap Network Services Corp. [Online]. Available: http://www.internap.com

[6] D. Cox and V. Isham, *Point Processes*. London, U.K.: Chapman Hall, 1980.

[7] M. Crovella and A. Bestavros, "Self-similarity in world wide web traffic: Evidence and causes," in *Proc. ACM SIGMETRICS*, Philadelphia, PA, May 1996, pp. 160–169.

[8] A. Girard, *Routing and Dimentioning in Circuit-Switched Networks*. Boston, MA: Addison-Wesley, 1990.

[9] S. Jamin, P. Danzig, S. Shenker, and L. Zhang, "A measurement-based call admission control for integrated services packet networks," in *Proc. ACM SIGCOMM*, Cambridge, MA, August 1995, pp. 2–13.

[10] F. P. Kelly, "Routing in circuit-switched networks: Optimization, shadow prices and decentralization," *Adv. Appl. Probabil.*, vol. 20, pp. 112–144, 1988.

[11] W. E. Leland, M. S. Taqqu, W. Willinger, and D. Wilson, "On the self-similar nature of Ethernet traffic," *IEEE/ACM Trans. Networking*, vol. 2, pp. 1–15, Feb. 1994.

[12] M. Parulekar and A. M. Markowski, "M/G/∞ input processes: A versatile class of models for network traffic," in *Proc. IEEE INFOCOM*, Kobe, Japan, Apr. 1997, pp. 419–426.

[13] V. Paxon and S. Floyd, "Wide area traffic: The failure of Poisson modeling," in *Proc. ACM SIGCOMM*, Aug. 1994, pp. 257–268.

[14] S. Savage *et al.*, "Detour: A case for informed internet routing and transport," *IEEE Micro*, vol. 19, pp. 50–59, Jan. 1999.

[15] Auckland Data Trace [Online]. Available: 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 *Proc. IEEE Int. Conf. Network Protocols (ICNP)*, Riverside, CA, Nov. 2001, pp. 129–137.

[17] ——, "On scalable design of bandwidth brokers," *IEICE Trans. Commun.*, vol. E84-B, no. 8, Aug. 2001.

[18] Z. Duan, Z.-L. Zhang, and Y. T. Hou, "Service overlay networks: SLAs, QoS, and bandwidth provisioning," Dept. Comput. Sci., Florida State Univ., Tallahassee, FL, Tech. Rep., http://www.cs.fsu.edu/~duan/publications/SON_TR.pdf.

Zhenhai Duan (S'97–M'03) received the B.S. degree from Shandong University, China, in 1994, the M.S. degree from Beijing University, China, in 1997, and the Ph.D. degree from the University of Minnesota, Minneapolis, in 2003, all in computer science.

He is currently an Assistant Professor in the Computer Science Department, Florida State University, Tallahassee. His research interests include computer networks and multimedia communications, especially scalable QoS control and management in the Internet, Internet routing protocols and service architectures, and networking security.

Dr. Duan was a corecipient of the 2002 IEEE International Conference on Network Protocols Best Paper Award. He has been a member of the ACM since 2003.

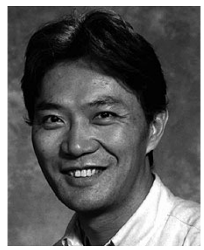

Zhi-Li Zhang (M'97) received the B.S. degree in computer science from Nanjing University, China, in 1986 and the M.S. and Ph.D. degrees in computer science from the University of Massachusetts, Amherst, in 1992 and 1997, respectively.

In 1997 he joined the computer science and engineering faculty at the University of Minnesota, Minneapolis, where he is currently an Associate Professor. From 1987 to 1990, he conducted research in the Computer Science Department, Århus University, Denmark, under a fellowship from the Chinese National Committee for Education. He has held visiting positions at Sprint Advanced Technology Laboratories, IBM T. J. Watson Research Center, Fujitsu Laboratories of America, Microsoft Research China, and INRIA, Sophia Antipolis, France. His research interests include computer communication and networks, especially the QoS guarantee issues in high-speed networks, multimedia and realtime systems, and modeling and performance evaluation of computer and communication systems.

Dr. Zhang currently serves on the editorial board of IEEE/ACM TRANSACTIONS ON NETWORKING and *Computer Network*, an international journal. He was a Program Co-chair of the SPIE ITCOM 2002 Conference on Scalability and Traffic Control in IP Networks, served on the Executive Committee for IEEE INFOCOM 2001 and INFOCOM 2003, and on the Technical Program Committees of various conferences and workshops, including IEEE INFOCOM and ACM SIGMETRICS. He received the National Science Foundation CAREER Award in 1997. He was also awarded the prestigious McKnight Land-Grant Professorship by the University of Minnesota. He was a co-recipient of an ACM SIGMETRICS Best Paper Award and an IEEE ICNP Best Paper Award. He is a member of INFORMS Telecommunication Section and has been a member of the ACM since 1997.

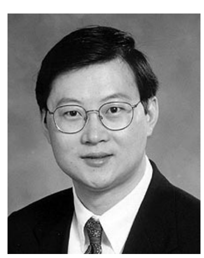

Yiwei Thomas Hou (S'91–M'98) received the B.E. degree (*summa cum laude*) from the City College of New York in 1991, the M.S. degree from Columbia University, New York, NY, in 1993, and the Ph.D. degree from Polytechnic University, Brooklyn, NY, in 1998, all in electrical engineering.

From 1997 to 2002, he was a Research Scientist and Project Leader with the IP Networking Research Department, Fujitsu Laboratories of America, Sunnyvale, CA. He is currently an Assistant Professor with the Bradley Department of Electrical and Computer Engineering, Virginia Polytechnic and State University (Virginia Tech), Blacksburg, VA. His research interests include wireless video sensor networks, multimedia delivery over wireless networks, scalable architectures, protocols, and implementations for differentiated services Internet, and service overlay networking. He has published extensively in the above areas.

Dr. Hou was a co-recipient of the 2002 IEEE International Conference on Network Protocols Best Paper Award and the 2001 IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY Best Paper Award. He has been a member of the ACM since 1995.