

Elastic Service Availability: Utility Framework and Optimal Provisioning

Dahai Xu, *Member, IEEE*, Ying Li, *Student Member, IEEE*, Mung Chiang, *Member, IEEE*,
and A. Robert Calderbank, *Fellow, IEEE*

Abstract—Service availability is one of the most closely scrutinized metrics in offering network services. It is important to cost-effectively provision a managed and differentiated network with various service availability guarantees under a unified platform. In particular, demands for availability may be elastic and such elasticity can be leveraged to improve cost-effectiveness. In this paper, we establish the framework of provisioning elastic service availability through network utility maximization, and propose an optimal and distributed solution using differentiated failure recovery schemes.

First, we develop a utility function with configurable parameters to represent the satisfaction perceived by a user upon service availability as well as its allowed source rate. Second, adopting Quality of Protection [1] and shared path protection, we transform optimal provisioning of elastic service availability into a convex optimization problem. The desirable service availability and source rate for each user can be achieved using a price-based distributed algorithm. Finally, we numerically show the tradeoff between the throughput and the service availability obtained by users in various network topologies. This investigation quantifies several engineering implications. For example, indiscriminately provisioning service availabilities for different kinds of users within one network leads to noteworthy sub-optimality in total network utility. The profile of bandwidth usage also illustrates that provisioning high service availability exclusively for critical applications leads to significant waste in bandwidth resource.

Index Terms—Service availability, shared protection, network utility maximization, resource allocation.

I. INTRODUCTION

A. Motivation

IN the selection of a network service, availability of the service is one of the most closely scrutinized metrics, sometimes even more important than other Quality of Service (QoS) parameters such as latency, jitter, packet loss, etc. A typical market analysis [2] shows that 50% of subscribers to network transmission services expect at least the 99.99% service availability, while very few expect less than 99.7% service availability. A network vendor provides several types of services with different levels of service availability guarantee. For example, AT&T guarantees the availability of 99.9% for Toll-free Service, 99.99% for US Packet Network, and 99.999% for Managed Internet Service ([3], 2006).

Usually, a stronger availability guarantee means higher service price. Since different users (e.g., home user, small

company/education institution, financial business, etc.) have different sensitivity to the service availability, a customer with a limited budget need decide to buy what level of service availability from which provider. On the other hand, for a network provider, the higher service availability can generate more revenue from the customers, but at the cost of higher capital expenditure. Provisioning high service availability usually requires redundant network resources, such as spare routers, switches, links, etc., and fast failure recovery/repair schemes. Accordingly, it is important to cost-effectively provision a managed and differentiated network with various service availability guarantees under a unified framework.

Similar to Quality of Service, various concepts of Reliability of Service (RoS) (or differentiated failure recovery schemes) [1], [4]–[6] have been proposed to distinguish service availability. For example, in [1], with a continuous grade on Quality of Protection (QoP) to represent the probability of initiating recovery in case of failure, the connections are categorized as the connection with guaranteed protection, the connection with best effort protection, unprotected connection, and preemptable connection. In Differentiated Reliability (DiR) [5], a service availability (independent of the recovery scheme) is chosen for each application, and the appropriate failure recovery will be triggered in case of failure to meet the specified service availability. However, none of these schemes address the scenario where the users have *elastic* demands for service availability.

For many users, demands for availability is indeed elastic: the user would prefer a higher availability but can also tolerate a lower one. When demands are elastic, the satisfaction of a user can be represented by a utility function of service availability achieved, as well as the source rate obtained when the service is available. Therefore, the optimal provisioning of elastic service availability can be realized by solving an appropriately formulated Network Utility Maximization (NUM) problem [7], [8]. The notion of network utility was first advocated in the seminal paper [9] in 1995 for bandwidth allocation among elastic demands on *source rate*. This paper demonstrates that utility maximization can also provide a quantitative approach to satisfy elastic demands on *service availability* and to quantify the tradeoff between rate and availability.

B. Related Work

The framework of NUM has many applications in network resource allocation such as Internet congestion control (e.g., [10], [11]) and protocol stack design. There is also a useful

Manuscript received May 14, 2007. A preliminary version entitled “Optimal Provisioning of Elastic Service Availability” appeared in INFOCOM 2007.

The authors are with the Department of Electrical Engineering, Princeton University, Princeton, NJ 08544, USA (e-mail: {dahaiXu, yingli, chiangm, calderbk}@princeton.edu).

Digital Object Identifier 10.1109/JSACOCN.2008.030607.

economics interpretation of the dual-based distributed algorithm for NUM, in which the Lagrange dual variables can be interpreted as shadow prices for resource allocation, and each end user and the network maximize his/her net utility and net revenue, respectively. Primal and dual-based distributed algorithms have been proposed to solve for the global optimum of NUM problems (e.g., [12]–[14]).

Consider a communication network with L logical links, each with a fixed capacity of c_l bps, and S sources (i.e., end users), each transmitting at a source rate of x_s bps. Each source s emits one flow, using a fixed set $L(s)$ of links in its path, and has a utility function $U_s(x_s)$. Each link l is shared by a set $S(l)$ of sources. NUM, in its basic version, is the following problem of maximizing the network utility $\sum_s U_s(x_s)$, over the source rates \mathbf{x} , subject to linear flow constraints $\sum_{s \in S(l)} x_s \leq c_l$ for all links l :

$$\begin{aligned} & \text{maximize} && \sum_s U_s(x_s) \\ & \text{subject to} && \sum_{s \in S(l)} x_s \leq c_l, \forall l, \\ & \text{variables} && \mathbf{x} \geq 0. \end{aligned} \quad (1)$$

Making the standard assumption on concavity of the utility functions, problem (1) is a simple concave maximization of decoupled terms under linear constraints, which has long been studied in optimization theory as a monotropic program [15].

The basic NUM (1) has been extended to include other layers to understand network architecture [16], as well to achieve fair resource allocation in the network provisioning QoS and Differentiated Service (DiffServ) [17]. Thus the utility function is not solely decided by the transmission rate. Instead, it depends on the QoS (such as end-to-end delay, jitter, packet loss, etc.) guaranteed for the transmission as well as the transmission rate [18]–[21]. However, among the extensive literature on NUM and its generalizations, most works treat utility as a function of throughput or throughput per unit of energy, with a few publications examining utility as a function of communication reliability or delay. In contrast, the question of how to optimally provision the network for elastic *service availability* has not been tackled through the utility formulation. Throughout this paper, we will encounter several new challenges in tackling this new question, from the introduction of both primary and backup paths for each source to the nonconvexity in the problem formulation.

C. Communication Reliability vs. Service Availability

In [22], [23], the QoS of end-to-end communication reliability is incorporated into the framework of NUM. Due to channel noise or fading, not all the signals can be successfully decoded at the receiver. On some communication links, the physical layer's adaptive error correction mechanisms can change the link capacity and decoding error probability, e.g., through adaptive channel coding in Digital Subscriber Loop (DSL) broadband access networks or adaptive diversity-multiplexing control in Multiple-Input-Multiple-Output (MIMO) wireless systems. Lee et al. investigate the intrinsic tradeoff between rate and communication reliability (end-to-end signal quality) [22]. Marbukh proposed a method of integrating diverse routing and retransmission as an alternative to single path routing for each flow [23].

Communication reliability and *service availability* are two different concepts. Some of their differences can be demonstrated by a simple example. Assuming a customer requests 1Mbps connection service from a carrier, but the carrier grants 1.01Mbps because either the communication reliability or service availability is only 99%. For the former scenario, on average, one bit of every 100 bits is lost during the transmission. Such loss can be compensated by retransmissions or appropriate coding in a fast timescale [22]. In contrast, in the latter scenario, the connection is available except for an *unpredictable* 7 hours of downtime every month. In general, the customer does not have the same satisfaction/utility in the two scenarios.

In addition to timescale difference, communication reliability and service availability also rely on different mitigation methods: channel coding and lost packet retransmission are used to ensure communication reliability, and backup bandwidth provisioning for path restoration and protection are used to enhance service availability.

D. Summary of Contributions

In this paper, we address the resource allocation when elastic service availability is considered. Service will be temporarily unavailable because of the failures due to human mistakes (e.g., mis-configuration), software bugs, hardware defects, natural disasters (e.g., flooding or earthquakes), or even perpetrators (e.g., terrorists or hackers). Such failures in general cannot be repaired immediately or compensated by retransmission as in the cases of [22], [23]. To ensure the high availability required by some critical applications, failure recovery has to be implemented where the affected traffic is rerouted in case of failure. An effective failure recovery scheme usually consists of three components: establishing backup paths disjoint from the primary paths, provisioning network resource (e.g. bandwidth) prior to failure, and real-time failure detection and signaling to reroute traffic [24]. The first component has been extensively studied with graph theoretic methods. The third has been investigated by the system research community. In this work, we focus on the second component: bandwidth provisioning to achieve the optimal service availability through NUM.

This work is the first to investigate elastic service availability provisioning using differentiated failure recovery:

- *Framework*: We develop the NUM framework for elastic service availability, and present a utility function with configurable parameters to represent the satisfaction perceived by different users upon service availability and source rate.
- *Centralized solution*: With Quality of Protection (QoP) and shared path protection, we transform the problem into a convex optimization, thus efficiently solvable for global optimality through standard centralized algorithms.
- *Distributed solution*: With regular updates of backup path provisioning, we also propose a price-based distributed algorithm to optimally provision elastic service availability and source rate.
- *Simulation*: We carry out numerical experiments over realistic network topologies, and present the optimal

TABLE I
SUMMARY OF KEY NOTATION

Notation	Meaning
q_s	Service availability provided for source s
ρ_s	Number of 9's in service availability q_s
η_s	Probability of initiating failure recovery for source s
$U_s(\cdot)$	Utility function of source rate and service availability
$V_s(\cdot)$	Normalized utility function of service availability, which may take in arguments of q_s , ρ_s or η_s
x_s	Data rate of source s
y_s	Expected backup rate of source s
w_s	Adjusted rate of source s
a_s	Criticality parameter (in service availability) of U_s
b_s	Elasticity parameter (in service availability) of U_s
$L(s)$	Primary/working path of source s
$M(s)$	Backup path of source s
c_l	Capacity of link l
z_l	Backup bandwidth reserved on link l
$S(l)$	Set of connections using link l on primary path
$T(l)$	Set of connections using link l on backup path

tradeoff between the throughput and the service availability.

Engineering implications of this work quantify several intuitions on elastic service availability. For example, we show that indiscriminately provisioning service availabilities for different kinds of users within one network leads to noteworthy sub-optimality in terms of maximizing network utility. By profiling bandwidth usage, we illustrate that provisioning high service availability exclusively for critical users/applications leads to significant waste in bandwidth resource.

The rest of the paper is organized as follows. In Sec. II, we incorporate the elastic service availability into the framework of NUM with differentiated failure recovery. In Sec. III, a price-based distributed algorithm is proposed to determine desirable service availability and source rate for each user. Then we present results from extensive numerical experiments in Sec. IV. We conclude and discuss future work on provisioning of elastic service availability in Sec. V. The key notation used throughout this paper is summarized in Table I.

II. SYSTEM MODEL

Consider a similar setup as that for problem (1), but now source s has a utility function $U_s(x_s, q_s)$, where x_s is a source rate and q_s is the service availability provided for source s . We assume that the utility function is a continuous, strictly increasing function of x_s and q_s .

To provision high network availability, spare bandwidth has to be reserved in advance. Such bandwidth is usually not used under normal situation except by some preemptable connections. Denote z_l as the backup bandwidth reserved on link l . Let Γ be the recovery scheme to be used in case of failure, and denote also by $\Gamma(\mathbf{x}, \mathbf{z})$ the function mapping the source rates \mathbf{x} and backup bandwidth reservation \mathbf{z} to

the service availabilities achieved under the failure recovery scheme Γ . Then the resulting formulation is as follows:

$$\begin{aligned} & \text{maximize} && \sum_s U_s(x_s, q_s) \\ & \text{subject to} && \sum_{s \in S(l)} x_s + z_l \leq c_l, \forall l, \\ & && \mathbf{q} = \Gamma(\mathbf{x}, \mathbf{z}) \\ & \text{variables} && \mathbf{x}, \mathbf{z}, \mathbf{q} \succeq 0. \end{aligned} \quad (2)$$

This problem formulation is the starting point of the development in this section. Next, we need to specify function U_s and function Γ .

A. Utility Function of Service Availability

We need an appropriate utility function $U_s(x_s, q_s)$ to measure the satisfaction perceived by a user from both rate $x_s \geq 0$ and service availability $q_s \in [0, 1]$. In this work, we choose

$$U_s(x_s, q_s) = U_s(x_s \cdot V_s(q_s)),$$

where $V_s(q_s) \in [0, 1]$ denotes the normalized utility function of service availability. Let $w_s \triangleq x_s \cdot V_s(q_s)$ be the *adjusted rate*, and $U_s(w_s)$ be a strictly concave function. Obviously, if $q_s = 1$, then $V_s(q_s) = 1$ and $U_s(x_s, q_s) = U_s(x_s)$.

Note that service availability, q_s , is generally measured in the number of 9's. E.g. 99.99% has four 9's. We use ρ_s to represent the number of nines for service availability q_s as follows:

$$\rho_s \triangleq -\log_{10}(1 - q_s). \quad (3)$$

With a slight abuse of notation V_s to denote the function of ρ_s as well: $V_s(\rho_s) = V_s(1 - 10^{-\rho_s})$, and $V_s(\rho_s)$ is the normalized utility function of ρ_s , the number of 9's in service availability q_s .

Note that $V_s(\rho_s)$ should be a strictly increasing function of ρ_s and bounded within $[0, 1]$. Moreover, each user has a threshold of acceptable service availability, $\rho_{s,min}$. Failing to provision such service availability will result in near-zero utility no matter what source rate can be achieved. Base on the above observations, we propose the following utility function:

$$V_s(\rho_s) = 1 - 10^{a_s - b_s \rho_s}, \quad a_s \geq 0, b_s \geq 0 \quad (4)$$

The proposed utility function (4), depicted in Fig. 1, is a concave function of ρ_s with two parameters: *criticality* parameter a_s and *elasticity* parameter b_s . In Fig. 1, three typical kinds of users (normal, important and critical users) with different sensitivities to service availability are illustrated. For example, home users can be categorized as normal users, school users are important users, and financial business are critical users. The larger value of criticality parameter a_s means higher service availability requirement and the larger value of elasticity parameter b_s means steeper curve and suggests less elasticity in service availability requirement. Note that, to ensure $V_s(\rho) \in [0, 1]$, we have $\rho_{s,min} = \frac{a_s}{b_s}$ and $q_{s,min} = 1 - 10^{-\frac{a_s}{b_s}}$.

We now demonstrate how parameters a_s and b_s can be set from customer's requirements stated in two other parameters of direct engineering implications. Given a customer s who has a minimum requirement on service availability, $q_{s,min}$,

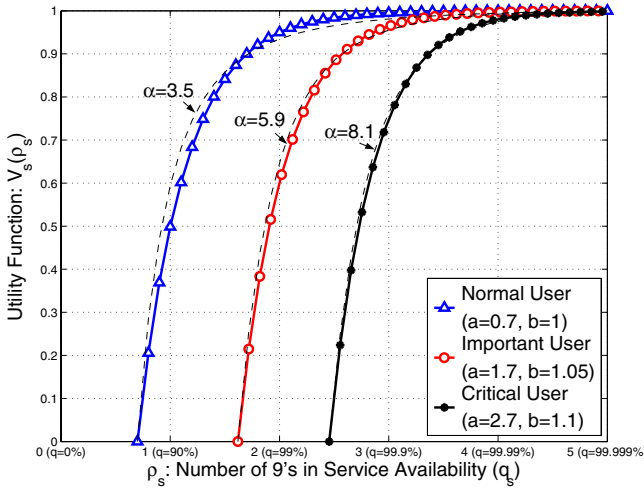


Fig. 1. Utility $V_s(\rho_s)$ as a function of ρ_s (number of 9's in service availability q_s) for three typical kinds of users.

and achieves ϕ percent of satisfaction at service availability $q_{s,\phi}$. Then we have

$$\begin{aligned} q_{s,\min} &= 1 - 10^{-\frac{a_s}{b_s}} \\ \phi/100 &= 1 - 10^{a_s - b_s(-\log_{10}(1 - q_{s,\phi}))}. \end{aligned}$$

Solving the equations above for a_s and b_s , we can choose the following parameters, thus specifying the utility function accordingly:

$$\begin{aligned} b_s &= \frac{\log_{10}(1 - \phi/100)}{\log_{10} \frac{1 - q_{s,\phi}}{1 - q_{s,\min}}} \\ a_s &= b_s (-\log_{10}(1 - q_{s,\min})). \end{aligned}$$

A family of utility functions widely used in NUM for resource allocation formulations are the α -fair utilities [25], which can be normalized such that $V_s(\rho_{s,\min}) = 0$ and $V_s(\rho_{s,\max}) = 1$:

$$V_s(\rho_s) = \begin{cases} \frac{\rho_s^{1-\alpha} - \rho_{s,\min}^{1-\alpha}}{\rho_{s,\max}^{1-\alpha} - \rho_{s,\min}^{1-\alpha}}, & \text{if } \alpha \neq 1 \\ \frac{\log \rho_s - \log \rho_{s,\min}}{\log \rho_{s,\max} - \log \rho_{s,\min}}, & \text{if } \alpha = 1 \end{cases}. \quad (5)$$

It turns out that such utility functions will result in a *non-convex* optimization problem for service availability provisioning, thus losing the desirable properties of efficient solutions (centralized or distributed) for global optimality. Fortunately, the curves of the utility function we proposed are very close to those of normalized utility function (5) with appropriate α parameters, which are shown as the dashed lines in Fig. 1. Therefore, a different parametrization of utility curves whose shapes are very close to the standard α -fair curves lead to a much more tractable convexity structure in the problem formulation to be shown later this section.

B. Enhance Service Availability with Shared Path Protection

Failure recovery is usually required for provisioning high service availability. There are two main failure recovery schemes: *protection* and *restoration*. The major difference between the two is that, in protection, a detour around a possible failure is determined at the time of connection setup and the spare capacity is allocated and updated periodically

along the detour prior to the failure, whereas in restoration, the detour is dynamically determined after the failure occurs. Accordingly, protection schemes can in general recover from a failure quicker (as long as the detour is not affected by any other failures), but are less bandwidth efficient than restoration schemes. On the other hand, restoration schemes can survive one or multiple failures (as long as the destination is still reachable, with sufficient connectivity and bandwidth), but they cannot guarantee the recovery time, or the amount of information loss for real-time applications, making them unsuitable for mission-critical applications. In this paper, we will focus on improving service availability with various *protection* schemes.

In many applications, we mainly consider the scenario of single failure. Then we can use *shared protection* [26]–[28] to reduce bandwidth usage in a mesh network since the backup bandwidth reserved by multiple connections on a same link can be shared as long as no single failure can affect them simultaneously.

C. Quality of Protection under Shared Path Protection

With shared path protection, some schemes with reliability of service [1], [5], [6] can be implemented to differentiate service availability. For example, in Quality of Protection (QoP), each connection is associated with a continuous QoP grade $\eta_s \in [0, 1]$, which is equivalent to the *probability* that connection will be restored immediately in case of failure. Such probabilistic model can be implemented in a deterministic way by reserving $\eta_s x_s$, the expected backup bandwidth, along its backup path [1].

In this work, we adopt shared path protection and QoP to provision elastic service availability. Therefore, for each connection s , besides its fixed working/primary path $L(s)$, it has a pre-planned disjoint backup path $M(s)$, i.e., $L(s) \cap M(s) = \emptyset$. In case of failure at link m , the connection s using m on its primary path will reroute the traffic along its backup path with a probability of η_s .

Let r_l denote the link availability of link l and all link failures are assumed to be statistically independent. Then the availability of the primary path and backup path of the connection s are

$$\begin{aligned} q_{0,s} &= \prod_{l \in L(s)} r_l, \\ q_{1,s} &= \prod_{l \in M(s)} r_l \end{aligned} \quad (6)$$

respectively. When at most one link failure is considered, the availability for connection s ,

$$q_s \approx q_{0,s} + (1 - q_{0,s})q_{1,s}\eta_s \quad (7)$$

To search for the optimal solution to problem (2) with the failure recovery scheme (Γ) mentioned above, we need to optimally determine the source rates (\mathbf{x}) and service availabilities (\mathbf{q}) for all the connections simultaneously. We refer to such optimization procedure as Γ_{OPT} .

For comparison purpose, we also introduce two extreme cases of recovery scheme as follows.

- Γ_{NR} (*No Recovery*): No failure recovery scheme is implemented, i.e. $\eta_s = 0$. Thus the service availability

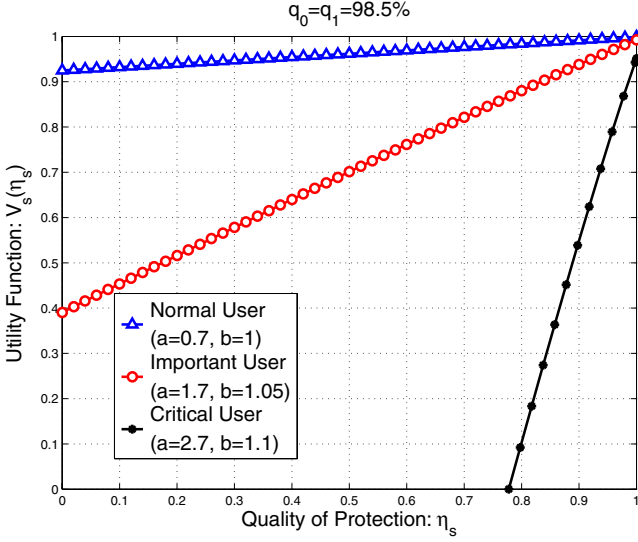


Fig. 2. Utility $V_s(\eta_s)$ as a function of η_s (Quality of Protection) for three typical kinds of users.

achieved by any connection is just the availability of its primary path, $q_{0,s}$.

- Γ_{SA} (Sufficient Availability): As a conservative approach, sufficient backup bandwidth will be reserved along backup path (i.e. $\eta_s = 1$) regardless of users' elastic demands for service availability.

The methods of solving problem Γ_{OPT} can be easily extended to the other two recovery schemes Γ_{NR} and Γ_{SA} by specifying the value η_s in advance as the additional constraints.

Obviously, from (7), we have $1 - q_s = (1 - q_{0,s})(1 - q_{1,s}\eta_s)$. Then $V_s(q_s) = V_s(1 - (1 - q_{0,s})(1 - q_{1,s}\eta_s))$, and with a slight abuse of notation, we also use $V_s(\eta_s)$ to represent the normalized utility function of η_s (QoP). From (3) and (4),

$$\begin{aligned} V_s(\eta_s) &= 1 - 10^{a_s + b_s \log_{10}(1 - q_{0,s})(1 - q_{1,s}\eta_s)} \\ &= 1 - 10^{a_s} \left((1 - q_{0,s})(1 - q_{1,s}\eta_s) \right)^{b_s}. \end{aligned} \quad (8)$$

Fig. 2 shows the curves of utility $V_s(\eta_s)$ as the function of (η_s) for the three representative classes of users with elastic service availability demands (same parameters as those in Fig. 1) assuming the availabilities for both primary path and backup path are 98.5%. Note that, the curve will not be a straight line if b_s is not equal to 1. In addition, to ensure $V_s(\eta_s) \in [0, 1]$, the lower bound of QoP acceptable for user s is

$$\eta_{s,min} = \max \left\{ \frac{1 - \frac{10^{-a_s}}{10^{b_s}(1 - q_{0,s})}}{q_{1,s}}, 0 \right\}.$$

Confirming our intuition, the more important users may have a strictly positive $\eta_{s,min}$, as illustrated by the 'critical user' curve in Fig. 2.

D. Optimal Provisioning of Elastic Service Availability

Denote $y_s = \eta_s x_s$ as the expected backup bandwidth reserved for connection s along its backup path. Then the

objective of (2) is

$$\max \sum_s U_s(x_s \cdot V_s(\eta_s)) = \max \sum_s U_s(x_s \cdot V_s(\frac{y_s}{x_s})),$$

which is equal to $\max \sum_s U_s(w_s)$, where

$$w_s = x_s \cdot V_s(\frac{y_s}{x_s}). \quad (9)$$

Since $U_s(w_s)$ is a strictly increasing function, the equality constraint (9) can be replaced by $w_s \leq x_s \cdot V_s(\frac{y_s}{x_s})$ as the constraint is always tight at optimality.

As we only consider single link failure, we have (10)

$$z_l = \max_{\forall m \neq l} \sum_{s \in T(l) \cap S(m)} y_s, \forall l, \quad (10)$$

where $T(l)$ denotes the set of connections using l on their backup paths and $S(m)$ denotes the set of connections using m on their primary paths. Eq. (10) means the backup bandwidth reserved on link l just need to be sufficient to recover the worst failure scenario. Therefore, the formulation for problem Γ_{OPT} is summarized as follows:

$$\begin{aligned} \max \quad & \sum_s U_s(w_s) \\ \text{s.t.} \quad & \sum_{s \in S(l)} x_s + \sum_{s \in T(l) \cap S(m)} y_s \leq c_l, \forall l, \forall m \neq l, \\ & w_s \leq x_s \cdot V_s(\frac{y_s}{x_s}), \forall s \in S, \\ & \mathbf{x} \succeq \mathbf{y}, \\ \text{vars.} \quad & \mathbf{w}, \mathbf{x}, \mathbf{y} \succeq 0. \end{aligned} \quad (11)$$

Note that η_s can be recovered from (x_s, y_s) , and constants (a_s, b_s) are implicitly represented in the function V_s .

The rest of this paper examines the solution methods and engineering implications of the above problem.

E. Analysis for a Simple Scenario

Before discussing solution methods in general, we first illustrate some of the interesting aspects of the problem formulation through a simple example. Suppose we have only one source-destination pair, and two single-link paths, l_0 and l_1 , for primary and backup paths respectively. Assume their path availabilities are q_0 and q_1 and their link capacities are c_0 and c_1 , respectively. For the single user, its criticality parameter is a and elasticity parameter is b .

For this simple scenario, the optimization problem is (12),

$$\begin{aligned} \text{maximize} \quad & U(w) \\ \text{subject to} \quad & x \leq c_0 \\ & y \leq c_1 \\ & x \geq y \\ & w \leq x \cdot V(\frac{y}{x}) \\ \text{variables} \quad & x, y, \end{aligned} \quad (12)$$

where $V(\frac{y}{x}) = 1 - 10^a(1 - q_0)^b(1 - q_1\frac{y}{x})^b$.

If $c_0 \leq c_1$, the optimal solution is $x^* = y^* = c_0$, which is trivial. If $c_0 \geq c_1$, it is easy to see that the optimal solution satisfies $c_1 \leq x^* \leq c_0$ and $y^* = c_1$. With $y^* = c_1$, the problem is equivalent to maximizing $w(x) = x \left[1 - 10^a(1 - q_0)^b(1 - q_1\frac{c_1}{x})^b \right]$ over $x \in [c_1, c_0]$. We take the derivative of $w(x)$ with respect to x ,

$$\frac{dw}{dx} = 1 - 10^a(1 - q_0)^b(1 - q_1\frac{y}{x})^b \left[1 + \frac{b(1 + q_1\frac{y}{x})}{1 - q_1\frac{y}{x}} \right]. \quad (13)$$




Fig. 3. Optimal flow on primary path (x) for single-user with different criticality (a) and elasticity (b) parameters.

It is easy to verify that $x^* = c_0$ if $\frac{dw}{dx}|_{x=c_0} \geq 0$, otherwise $x^* < c_0$, which means that link l_0 cannot be fully utilized. Since $\frac{dw}{dx}|_{x=c_0}$ is a function of the parameters a , b , c_0 and c_1 , the values of the parameters may affect whether link l_0 is fully utilized or not.

In Fig. 3 below, we show how the parameters may affect whether link l_0 is fully utilized or not. We set $q_0 = q_1 = 98.5\%$, $c_0 = 2$ units and $c_1 = 1$ unit. We can see that for small a , which corresponds to non-critical requirement on availability, link l_0 can be fully utilized. When a becomes large or user has higher availability requirement, link l_0 become not fully utilized, especially for a smaller elasticity parameter, b .

The important thing illustrated here is that given a and b based on users' sensitivity to availability, there may be some capacity not fully utilized if the capacities for primary use and backup do not match well. We can save capacity by matching the primary use and backup use. For instance, for this simple scenario above with the given a , b and c_0 , by solving $\frac{dw}{dx}|_{x=c_0} = 0$ for c_1 , we can get the smallest capacity for backup such that c_0 can be fully utilized. When there are multiple users sharing links, the saved capacity by matching the rates on primary path and backup path for one user, can be used to support other users. Our interest is to do the matching for all the users with elastic demand on service availability in a systematic way in a general topology.

F. Algorithms: Centralized or Distributed

It can be easily verified that when $a_s > 0, b_s \geq 1$, $x_s \cdot V_s(\frac{y_s}{x_s})$ is a strictly concave function of (x_s, y_s) and the final problem (11) is a convex optimization problem. Highly efficient primal-dual interior point algorithms [29] can thus be used to solve for the unique global optimum of the problem. Such centralized computation is suitable for off-line provisioning of elastic service availability through a centralized network management, which is the most probable application scenario in practice.

In a different scenario, when the users change their preference or utility function over time, in order to enable regular updates through distributed message passing within the network, we need to develop distributed algorithms to solve (11) for the jointly optimal source rates and service availabilities. Most likely, such distributed updates of service availability provisioning is only needed once over a long time. This is the subject of Sec. III.

III. DISTRIBUTED ALGORITHM

In this section, we use a dual decomposition approach to distributively solve problem (11). Using both l and m to index links, and denoting the stacked vector of dual variables (or pricing variables) as λ , we first write the Lagrangian associated with problem (11) as

$$\begin{aligned} L(\mathbf{w}, \mathbf{x}, \mathbf{y}, \lambda) &= \sum_s U_s(w_s) + \sum_l \sum_{m:m \neq l} \lambda_{l,m} (c_l - \sum_{s \in S(l)} x_s \\ &\quad - \sum_{s \in T(l) \cap S(m)} y_s) \\ &= \sum_s \{ U_s(w_s) - \sum_{l \in L(s)} \sum_{m:m \neq l} \lambda_{l,m} x_s \\ &\quad - \sum_{l \in M(s)} \sum_{m \in L(s)} \lambda_{l,m} y_s \} + \sum_l \sum_{m:m \neq l} \lambda_{l,m} c_l. \end{aligned} \quad (14)$$

The Lagrange dual function is

$$Q(\lambda) = \max_{\substack{w_s \leq x_s \cdot V_s(\frac{y_s}{x_s}), \forall s, \\ \mathbf{x} \succeq \mathbf{y}, \\ \mathbf{w}, \mathbf{x}, \mathbf{y} \succeq \mathbf{0}},} L(\mathbf{w}, \mathbf{x}, \mathbf{y}, \lambda), \quad (15)$$

where $\mathbf{0}$ is the vector whose elements are all zeros.

The dual problem is formulated as

$$\begin{aligned} \min \quad & Q(\lambda) \\ \text{s.t.} \quad & \lambda \succeq \mathbf{0}. \end{aligned} \quad (16)$$

To solve the dual problem, we first consider problem (15). Since the Lagrangian is separable, this maximization of the Lagrangian over $\mathbf{w}, \mathbf{x}, \mathbf{y}$ can be conducted in parallel at each source s :

$$\begin{aligned} \max \quad & U_s(w_s) - \sum_{l \in L(s)} \sum_{m:m \neq l} \lambda_{l,m} x_s \\ & - \sum_{l \in M(s)} \sum_{m \in L(s)} \lambda_{l,m} y_s \\ \text{s.t.} \quad & w_s \leq x_s \cdot V_s(\frac{y_s}{x_s}), \\ & x_s \geq y_s, \\ \text{vars.} \quad & w_s, x_s, y_s \geq 0. \end{aligned} \quad (17)$$

Then, dual problem (16) can be solved by using the gradient projection algorithm as

$$\begin{aligned} \lambda_{l,m}(t+1) &= \left[\lambda_{l,m}(t) - \alpha(t) \left(c_l - \sum_{s \in S(l)} x_s(t) - \sum_{s \in T(l) \cap S(m)} y_s(t) \right) \right]^+, \\ &\quad \forall l, \forall m \neq l, \end{aligned} \quad (18)$$

where $\alpha(t)$ is the step size and $x_s(t)$ and $y_s(t)$ are solutions of problem (17) for a given $\lambda(t)$.

We now propose the following distributed algorithm for Γ_{OPT} where each source solves its own problem with only local information. There is an *important difference* of the following algorithm compared with the standard NUM for rate allocation only: each link l maintains a set of congestion prices $\lambda_{l,m}$ for all $m \neq l$.

Distributed Algorithm for Γ_{OPT}

In each iteration t , by solving the following problem (19) over (w_s, x_s, y_s) , each source s determines its adjusted rate $(w_s(t))$, rate on primary path $x_s(t)$ and rate on backup path $(y_s(t))$ that maximize its net utility based on the prices $(\lambda_x^s(t), \lambda_y^s(t))$ in the current iteration.

Source problem at source s :

$$\begin{aligned} \max \quad & U_s(w_s) - \lambda_x^s(t)x_s - \lambda_y^s(t)y_s \\ \text{s.t.} \quad & w_s \leq x_s \cdot V_s\left(\frac{y_s}{x_s}\right), \\ & x_s \geq y_s, \\ \text{vars.} \quad & w_s, x_s, y_s \geq 0. \end{aligned} \quad (19)$$

where $\lambda_x^s(t) = \sum_{l \in L(s)} \sum_{m: m \neq l} \lambda_{l,m}(t)$ is the end-to-end (primary path) congestion price at iteration t , and $\lambda_y^s(t) = \sum_{l \in M(s)} \sum_{m \in L(s)} \lambda_{l,m}(t)$ is the end-to-end (backup path) congestion price at iteration t .

In addition, by price update equation (20), the link adjusts its congestion prices for the next iteration.

Update of the set of congestion prices at link l :

$$\lambda_{l,m}(t+1) = \left[\lambda_{l,m}(t) - \alpha(t) \left(c_l - x^l(t) - y^{l,m}(t) \right) \right]^+, \forall m \neq l, \quad (20)$$

where $x^l(t) = \sum_{s \in S(l)} x_s(t)$ is the aggregate rate of those connections using l on their primary paths at iteration t , and $y^{l,m}(t) = \sum_{s \in T(l) \cap S(m)} y_s(t)$ is the aggregate rate of those connections using l and m on their backup paths and primary paths, respectively.

Message passing in the above algorithm only needs to be carried out between each source and the links on its (primary and backup) paths. To decide the congestion price according to (20), link l needs to know $x^l(t)$ and $y^{l,m}(t)$ for all $m \neq l$. Therefore, if source s changes its rate x_s , it can send an update message (containing the new value of x_s) to the links ($l \in L(s)$) along its primary path. In contrast, if source s changes its rate y_s , it has to send an update message (containing the new value of y_s and the route of its primary path, i.e. $m \in L(s)$) to notify the links ($l \in M(s)$) along its backup path.

On the other hand, to solve source problem (19), source s needs to know the end-to-end primary and backup congestion prices, $\lambda_x^s(t)$ and $\lambda_y^s(t)$. First, $\lambda_x^s(t)$ can be obtained by a notification message originated from the destination that summarizes the congestion price $\sum_{m: m \neq l} \lambda_{l,m}(t)$ of each link ($l \in L(s)$) along its primary path. Second, $\lambda_y^s(t)$ can be obtained by the notification message originated from the destination to sum up the congestion price $\sum_{m \in L(s)} \lambda_{l,m}(t)$ of each link ($l \in M(s)$) along its backup path.

In addition, to calculate its utility on service availability according to (6) and (8), the source also needs to know the availabilities of the links (r_l) of its primary and backup paths, which are usually static in wired networks.

After the above dual decomposition, the following result can be proved using standard techniques in distributed gradient algorithm's convergence analysis:

Theorem 1: Assume $a_s > 0, b_s \geq 1$, by Distributed Algorithm for Γ_{OPT} , dual variables $\lambda(t)$ converge to the optimal dual solutions λ^* and the corresponding primal variables w^*, x^* and y^* are the globally optimal primal solutions of (11).

Outline of the Proof:

Since strong duality holds for problem (11) and its Lagrange dual problem (16), we solve the dual problem through distributed gradient method and recover the primal optimizers from the dual optimizers. By Danskin's Theorem [30],

$$\frac{\partial Q(\lambda(t))}{\partial \lambda_{l,m}(t)} = c_l - x^l(t) - y^{l,m}(t), \forall l, \forall m \neq l.$$

Hence, the algorithm in (20) is a gradient projection algorithm for dual problem (16). Since the dual objective function $Q(\lambda)$ is a convex function, there exists a step size $\alpha(t)$ that guarantees $\lambda(t)$ to converge to the optimal dual solutions λ^* [30]. Also, if $\nabla Q(\lambda)$ satisfies a Lipschitz continuity condition, i.e., there exists a constant $H > 0$ such that

$$\|\nabla Q(\lambda_1) - \nabla Q(\lambda_2)\| \leq H \|\lambda_1 - \lambda_2\|, \forall \lambda_1, \lambda_2 \succeq \mathbf{0},$$

then $\lambda(t)$ converges to the optimal dual solution λ^* with a sufficiently small constant step size $\alpha(t) = \alpha, 0 < \alpha < 2/H$ [30]. The Lipschitz continuity condition is satisfied if the curvatures of the utility functions are bounded away from zero, see [14] for further details.

Furthermore, since problem (11) is a strictly convex optimization problem and problem (19) have unique solutions, w^*, x^* and y^* are the globally optimal primal solutions of (11) [31]. ■

IV. PERFORMANCE EVALUATION AND ENGINEERING IMPLICATIONS

In this section, we present the numerical results in provisioning elastic service availabilities. Recall that inputs to our problem are: a set of given primary and backup paths, a set of specified protection schemes, and the (a_s, b_s) parameters for each user s .

We consider two network topologies. The first is shown in Fig. 4, which is the same as that used in [26] and has 15 nodes and 28 bi-directed edges (for a total of 56 links). The capacity of each dark (bold) link is 4 times as large as that of the other (thin) links. The second is a large network called USnet shown in Fig. 5 [32] (with 46 nodes and 76 bi-directed edges of uniform capacity) is also considered. Without being stated explicitly, the capacities of the thin links in the 15-node network and all the links in the 46-node network are assumed to be 1000 units.

For all test scenarios, there is an elastic demand between each node pair. The utility function of user s is $U_s(x_s, q_s) = U_s(x_s \cdot V_s(q_s))$ as discussed in Sec. II. We use $U_s(w_s) = \log(w_s)$ as the utility function of adjusted rate $w_s = x_s \cdot V_s(q_s)$. Each user also has its own criticality parameter a_s and elasticity parameter b_s for elastic service availability demand.

A link-disjoint pair of primary path and backup path are chosen for each demand. Since the traffic is carried on the primary path most of the time and the backup bandwidth can be shared by several connections, we use the shorter one of the disjoint path pair as the primary path [33]. In the simulation, we use Dijkstra (shortest path) algorithm to find a primary path first followed by finding a backup path after removing the links along the primary path. For the case of no utility function for elastic service availability, previous studies have shown that the above primary-path-first heuristic can achieve




Fig. 4. A 15-node network with heterogeneous link capacities.




Fig. 5. 46-node USnet network with uniform link capacities.

near-optimal efficiency in bandwidth usage compared to other counterparts with complicated routing [33].

For the first and simple scenario to be investigated, all the users have the same service availability parameters a_s and b_s , and only optimal recovery scheme Γ_{OPT} is tested. Then in the second scenario, the users could have various service availability parameters and the other two recovery schemes, Γ_{NR} (No Recovery) and Γ_{SA} (Sufficient Availability), are also tested and compared.

A. Scenario with Uniform Service Availability Parameters

In this set of tests, the availabilities of links are all 99% and all the demands are assumed to have the same settings on criticality, a_s , and elasticity, b_s , parameters. The resulting service availabilities achieved by the demands could still be different since their primary/backup paths use different number of hops and thus have different path availabilities. In addition, the congestion price of using a link could also be different.

Impacts of elastic service availability on throughput-availability tradeoff. We trace the globally optimal tradeoff curve between network throughput and service availability with optimal recovery scheme Γ_{OPT} on the 15-node network. At first, the value of b_s is fixed at 1, and the value of

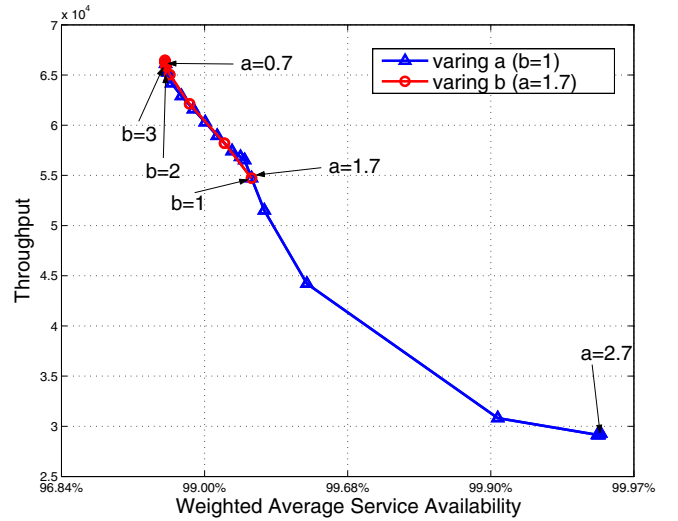


Fig. 6. Optimal tradeoff curves between network throughput ($\sum x_s$) and weighted average service availability ($\frac{\sum (x_s \rho_s)}{\sum x_s}$) with various criticality, a_s (producing a much larger dynamic range), and elasticity, b_s , parameters for a 15-node network.

a_s varies from 0.7 to 2.7 at step size of 0.1. Then we fix the value of a_s as 1.7 and vary the value of b_s from 1 to 3 at step size of 0.2. The two resulting curves are shown in Fig. 6, which demonstrate the tradeoff between the network throughput ($\sum_s x_s$) and the weighted average service availability ($\frac{\sum_s x_s q_s}{\sum x_s}$). Quantifying our intuition, a larger criticality parameter a_s (i.e. more sensitive to service availability) leads to higher service availability at the expense of lower throughput, since more bandwidth have to be used for backup purpose. From the least sensitivity ($a_s = 0.7$) to the highest sensitivity ($a_s = 2.7$), the weighted average service availability increases from 98.63% to 99.96% while the network throughput decreases by 55.7%. The results with various elasticity parameters, b_s , also confirm the above observation on the tradeoff between throughput and service availability. Note that, when b_s is large enough (≥ 2), the utility functions $V_s(\rho_s)$ shown in Fig. 1 will be very close to step functions. Such lack of elasticity in service availability leads to little variation in rate allocation, thus the points with $b_s \in [2, 3]$ are very close to each other in Fig. 6.

Impacts of elastic service availability on bandwidth usage.

Fig. 7 shows the percentages of the total bandwidth used by all primary paths and backup paths when the value of b_s is fixed at 1 and the value of a_s varies from 0.7 to 2.7. It turns out that if the demands are not sensitive to service availability, almost all the bandwidth can be used by primary paths, and thus an optimal recovery scheme would have the similar performance as no recovery scheme. With the increasing sensitivity to service availability for the demands, more bandwidth have to be used for backup purpose. Another interesting observation of Fig. 7 is that when the users are very sensitive to service availability (i.e., a large value of a_s) a significant fraction of bandwidth is wasted since a user cannot increase its rate along its primary path if it cannot simultaneously increase its backup bandwidth reservation to maintain the appropriate service availability. For example, the percentage of the unused

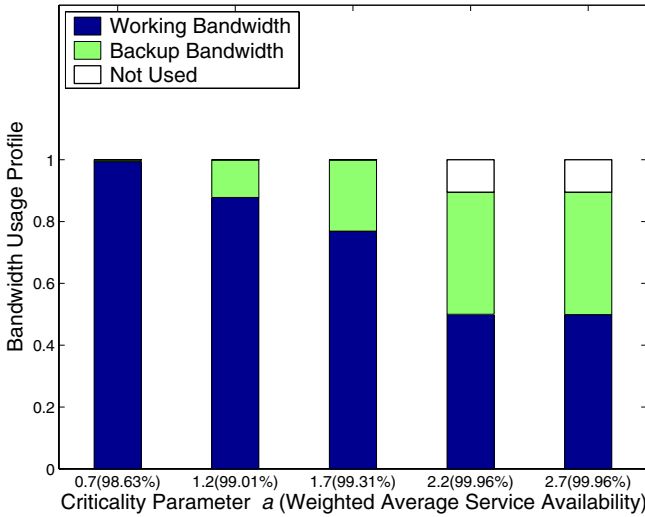


Fig. 7. Distribution (percentage) of bandwidth usage with various criticality parameters, a_s , and the corresponding weighted average service availabilities in a 15-node network.




Fig. 8. Comparison of the achieved network utility ($\sum_s U_s$) among optimal recovery scheme (Γ_{OPT}), no recovery scheme (Γ_{NR}) and sufficient availability scheme (Γ_{SA}) for a 15-node network where all the links have the same availability r .

bandwidth is as high as 10.5% when $a_s \geq 2.2$ (as shown by the 4th and 5th bars in Fig. 7), i.e., provisioning high service availability exclusively for critical users/applications leads to significant waste in bandwidth resource. There are two possible ways to reduce the bandwidth waste: 1) employing dynamic (and possibly complicated) routing for primary paths and backup paths, and 2) allowing for demands with various sensitivities to service availability, which is discussed next.

B. Scenario with Diverse Service Availability Parameters

For this test scenario, all source nodes are categorized as normal user, important user and critical user, with a population ratio of 9:3:1. Their criticality, a_s , and elasticity, b_s , parameters are same as those illustrated in Fig. 2. All the links have the same availability, r . Fig. 8 shows the network utility achieved in a 15-node network when the link availability varies from 0.975 to 0.999. Obviously all curves are monotonically

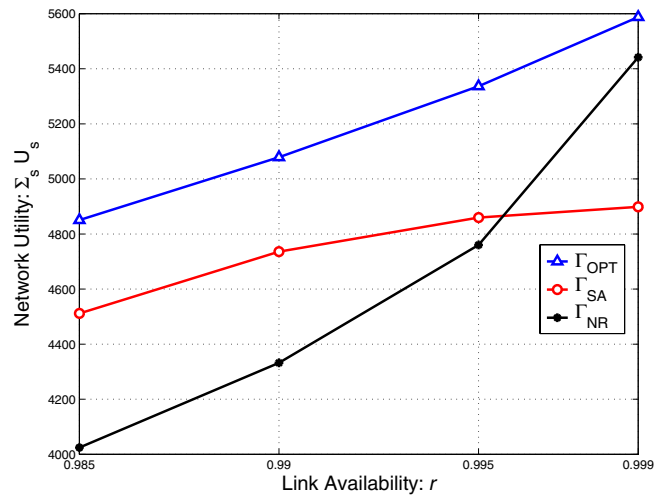


Fig. 9. Comparison of the achieved network utility for the 46-node USnet.




Fig. 10. Network utility ($\sum_s U_s$) achieved by optimal recovery scheme for the 15-node network with various link capacity (c) and link availability (r).

increasing since network-wide availability increases as each link's availability increases. It is clear that the proposed optimal recovery scheme (Γ_{OPT}) is consistently (and indeed provably) better than the other two regular schemes: no recovery scheme (Γ_{NR}) and sufficient availability scheme (Γ_{SA}). When the link is not reliable, by selectively provisioning failure recovery, Γ_{OPT} achieves 26.9% more utility than Γ_{NR} . Moreover, sufficient availability scheme could be even worse than no recovery scheme in terms of total utility when links are very reliable. Fig. 9 shows the network utility achieved in the 46-node USnet as link availability varies, where similar observations can be made.

Fig. 10 shows that the network utility achieved by optimal recovery scheme in the 15-node network will increase when we uniformly raise link capacity (c) or improve link availability (r). Given an operating point of link capacities, link availabilities, and the achieved network utility, to reach a higher network utility, it will be interesting to investigate which way of increasing link capacity and improving link availability is more cost-effective. In this test sample, if the current link capacity, link availability and network utility

achieved are 750 units, 0.98 and 898.2 respectively, to increase network utility by 6.5%, we can either increase link capacity by 33% or improve link availability to 0.995. Which network upgrade (link capacity increase or availability enhancement) is more cost-effective depends on the detailed capacity and equipment cost models. For example, most likely in well-deployed optical networks, it is less expensive to increase link capacity by lighting dark fibers than to enhance link level availability through more advanced transceiver optical components. Graphs such as the one in Figure 10 will allow operators to choose between alternative modes of service availability enhancement to best satisfy the overall elastic demands.

V. CONCLUSION

We establish the framework of provisioning elastic service availability through network utility maximization. This work complements the existing literature on either bandwidth allocation for elastic demands but no availability concerns, or bandwidth allocation for availability provisioning but ignoring demand elasticity. By developing a utility function of service availability in addition to source rate, we transform optimal provisioning into a convex optimization problem using differentiated failure recovery. The desirable service availability and source rate for each user can be achieved using a price-based distributed algorithm, where each link maintains multiple prices. We carry out numerical experiments over realistic network topologies, and present the optimal tradeoff between the throughput and the service availability. Engineering implications of this work quantify several intuitions on elastic service availability.

We initiate a utility-based study of network resilience by addressing optimal provisioning for elastic service availability through quality of protection and shared path protection. It would be interesting to investigate the elastic service availability provisioning for other schemes, e.g. employing other differentiated failure recovery schemes, using restoration instead of protection, recovering from multiple failures, etc. Combined with detailed cost models, this work will also lead to a quantification of the minimum-cost tradeoff between adding capacity and improving link availability for maximization of utility of service availability.

ACKNOWLEDGMENT

We would like to acknowledge the support of NSF CNS-0430487 and CNS-0519880, DARPA Grant HR0011-06-1-0008, and AFOSR Grant FA9550-06-1-0297. We also appreciate the helpful discussions with Ruby B. Lee, Xin Wang, and Jiaping Liu.

REFERENCES

- [1] O. Gerstel and G. Sasaki, "Quality of protection (QoP): a quantitative unifying paradigm to protection service grades," in *Proc. Optical Networking and Communications Conference (OPTICOMM'03)*, Denver, CO, 2001.
- [2] G. Holland, "Carrier class metro networking," Riverstone Networks, Technology white paper no. 135, July 2002.
- [3] AT&T Business Service Guide, <http://new.serviceguide.att.com>.
- [4] L. Z. C. Saradhi and M. Gurusamy, "Differentiated QoS for survivable WDM optical networks," *IEEE Commun. Mag.*, vol. 42, no. 5, pp. 8–14, May 2004.
- [5] A. Fumagalli and M. Tacca, "Differentiated reliability (DiR) in WDM rings without wavelength converters," in *Proc. IEEE International Conference on Communications (ICC'01)*, Helsinki, Finland, pp. 2887–2891.
- [6] S. Arakawa, J. Katou, and M. Murata, "Design method of logical topologies with quality of reliability in WDM networks," in *Photonic Network Commun.*, vol. 5, no. 2, pp. 107–121, Mar. 2003.
- [7] F. Kelly, A. Maulloo, and D. Tan, "Rate control in communication networks: shadow prices, proportional fairness and stability," *J. Operational Research Society*, vol. 49, no. 3, pp. 237–252, Mar. 1998.
- [8] H. Yäiche, R. R. Mazumdar, and C. Rosenberg, "A game theoretic framework for bandwidth allocation and pricing in broadband networks," *IEEE/ACM Trans. Networking*, vol. 8, no. 5, pp. 667–678, 2000.
- [9] S. Shenker, "Fundamental design issues for the future Internet," *IEEE J. Select. Areas Commun.*, vol. 13, no. 7, pp. 1176–1188, Sept. 1995.
- [10] S. H. Low, "A duality model of TCP and queue management algorithms," *IEEE/ACM Trans. Networking*, vol. 11, no. 4, pp. 525–536, 2003.
- [11] R. Srikant, *The Mathematics of Internet Congestion Control (Systems and Control: Foundations and Applications)*. Springer Verlag, 2004.
- [12] S. S. Kunniyur and R. Srikant, "End-to-end congestion control schemes: utility functions, random losses and ECN marks," *IEEE/ACM Trans. Networking*, vol. 11, no. 5, pp. 689–702, 2003.
- [13] R. J. La and V. Anantharam, "Utility-based rate control in the Internet for elastic traffic," *IEEE/ACM Trans. Networking*, vol. 10, no. 2, pp. 272–286, 2002.
- [14] S. H. Low and D. E. Lapsley, "Optimization flow control—part I: basic algorithm and convergence," *IEEE/ACM Trans. Networking*, vol. 7, no. 6, pp. 861–874, 1999.
- [15] R. T. Rockafellar, *Network Flows and Monotropic Programming*. Athena Scientific, 1998.
- [16] M. Chiang, S. H. Low, A. R. Calderbank, and J. C. Doyle, "Layering as optimization decomposition: a mathematical theory of network architectures," *Proc. IEEE*, Jan. 2007.
- [17] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss, "An architecture for differentiated service," IETF, RFC 2475, 1998.
- [18] C. Bouras and A. Sevasti, "SLA-based QoS pricing in DiffServ networks," *Computer Commun.*, vol. 27, no. 18, pp. 1868–1880, 2004.
- [19] X. Wang and H. Schulzrinne, "Pricing network resources for adaptive applications," in *IEEE/ACM Trans. Networking*, Apr. 2006.
- [20] L. DaSilva, D. Petr, and N. Akar, "Equilibrium pricing in multiservice priority based networks," in *Proc. IEEE Global Communications Conference '97*, Phoenix, AZ, Nov. 1997.
- [21] —, "Static pricing and quality of service in multiple service networks," in *Proc. 5th Conference on Computer Science and Informatics*, Atlantic City, NJ, 2000, pp. 355–358.
- [22] J. W. Lee, M. Chiang, and R. A. Calderbank, "Pricing-based distributed algorithms for rate-reliability tradeoff in network utility maximization," in *IEEE Conference on Computer Communications (INFOCOM'06)*, Barcelona, Spain, Apr. 2006.
- [23] V. Marbukh, "On aggregate utility maximization based network management: challenges and possible approaches," in *Proc. IEEE International Conference on Communications (ICC'04)*, Paris, France, 2004.
- [24] V. Sharma *et al.*, "Framework for MPLS-based recovery," Internet draft, draft-ietf-mpls-recovery-frmwk-08.txt (work in progress), Oct. 2002.
- [25] J. Mo and J. C. Walrand, "Fair end-to-end window-based congestion control," *IEEE/ACM Trans. Networking*, vol. 8, no. 5, pp. 556–567, 2000.
- [26] M. Kodialam and T. V. Lakshman, "Dynamic routing of bandwidth guaranteed tunnels with restoration," in *Proc. IEEE Conference on Computer Communications (INFOCOM'00)*, Tel Aviv, Israel, pp. 902–911.
- [27] G. Li *et al.*, "Efficient distributed path selection for shared restoration connections," in *Proc. IEEE Conference on Computer Communications (INFOCOM'02)*, New York, NY, pp. 140–149.
- [28] C. Qiao and D. Xu, "Distributed partial information management (DPIM) schemes for survivable networks—part I," in *Proc. IEEE Conference on Computer Communications (INFOCOM'02)*, New York, NY, pp. 302–311.
- [29] S. Boyd and L. Vandenberghe, *Convex Optimization*. New York: Cambridge University Press, 2004.
- [30] D. P. Bertsekas, *Nonlinear Programming*, 2nd ed. Athena Scientific, 1999.
- [31] M. Minoux, *Mathematical Programming: Theory and Algorithms*. Wiley, 1986.

- [32] S. Baroni, P. Bayvel, and R.J.Gibbens, "On the number of wavelength in arbitrarily-connected wavelength-routed optical networks," University of Cambridge, Statistical Laboratory Research Report 1998-7, available at <http://www.statslab.cam.ac.uk/reports/1998/1998-7.pdf>, 1998.
- [33] Y. Liu, D. Tipper, and P. Siripongwutikorn, "Approximating optimal spare capacity allocation by successive survivable routing," in *Proc. IEEE Conference on Computer Communications (INFOCOM'01)*, Anchorage, AK, pp. 699-708.



Dahai Xu (S'01-M'05) received the B.Eng. degree in Applied Electronics in 1996 and M.Eng. in Computer Science and Engineering in 1999 from Shanghai Jiao Tong University, China, and got his Ph.D. degree in Computer Science and Engineering at State University of New York at Buffalo in 2005.

Currently, he is a Postdoctoral Research Associate in the Department of Electrical Engineering at Princeton University. His research interests include survivability and restoration in IP/MPLS, optical networks, network design and protocol development

for next generation Internet and performance evaluation (modeling, simulation and measurements).



Ying Li (S'05) has been working towards her Ph.D. degree in Electrical Engineering at Princeton University, Princeton, NJ, USA, since September 2003, supervised by Professor A. Robert Calderbank and Professor Mung Chiang. She received the B.E. degree (with honor) and the M.E. degree in Electrical Engineering from Xi'an Jiaotong University, Xi'an, China, in 1997 and 2000 respectively, and the M.A. degree in Electrical Engineering at Princeton University, Princeton, NJ, USA, in 2005. She was a visiting Ph.D. student in Swiss Federal Institute

of Technology (EPFL), Switzerland, in summer 2007, and in Motorola Multimedia Research Lab, Schaumburg, IL, USA, in fall 2007, respectively. She worked as a faculty member of research and teaching assistant in Dept. of Information and Communication Engineering at Xi'an Jiaotong University, China, from 2000 to 2003, and as a visiting scholar in Fuji Xerox Co. Ltd, Japan, from 2000 to 2001. Her research interests include communications, networking, optimization, information theory, and signal processing.



Mung Chiang (S'00-M'03) is an Assistant Professor of Electrical Engineering and an affiliated faculty of Applied and Computational Mathematics and of Computer Science at Princeton University. He received the B.S. (Honors) in Electrical Engineering and Mathematics, M.S. and Ph.D. degrees in Electrical Engineering from Stanford University in 1999, 2000, and 2003, respectively.

He conducts research in the areas of nonlinear optimization of communication systems, theoretical foundation of network architectures, algorithms for broadband access networks, and stochastic analysis of communications and networking. He received CAREER Award from the National Science Foundation, Young Investigator Award from the Office of Naval Research, Howard B. Wentz Junior Faculty Award and Engineering Teaching Commendation from Princeton University, School of Engineering Terman Award from Stanford University, New Technology Introduction Award from SBC Communications, and was a Hertz Foundation Fellow and Stanford Graduate Fellow. For his work on broadband access networks and Internet traffic engineering, he was selected for the TR35 Young Technologist Award in 2007, a list of top 35 innovators in the world under the age of 35. His work on Geometric Programming was selected by Mathematical Programming Society as one of the top 3 papers by young authors in the area of continuous optimization during 2004-2007. His work on Layering As Optimization Decomposition became a Fast Breaking Paper in Computer Science by ISI citation. He also co-authored papers that were IEEE Infocom best paper finalist and IEEE Globecom best student paper.

He has served as an associate editor for *IEEE Transactions on Wireless Communications* and *Springer Journal of Optimization and Engineering*, a lead guest editor for *IEEE Journal of Selected Areas in Communications*, a guest editor for *IEEE/ACM Transactions on Networking* and *IEEE Transactions on Information Theory*, a Program Co-Chair of the 38th Conference on Information Sciences and Systems, and a co-editor of the new Springer book series on "Optimization and Control of Communication Systems."



A. R. Calderbank (M'89 - SM'97 - F'98) received the BSc degree in 1975 from Warwick University, England, the MSc degree in 1976 from Oxford University, England, and the PhD degree in 1980 from the California Institute of Technology, all in mathematics.

Dr. Calderbank is Professor of Electrical Engineering and Mathematics at Princeton University where he directs the Program in Applied and Computational Mathematics. He joined Bell Telephone Laboratories as a Member of Technical Staff in 1980, and retired from AT&T in 2003 as Vice President of Research. Dr. Calderbank has made significant contributions to a wide range of research areas, from algebraic coding theory and quantum computing to wireless communication and active sensing.

Dr. Calderbank served as Editor in Chief of the *IEEE Transactions on Information Theory* from 1995 to 1998, and as Associate Editor for Coding Techniques from 1986 to 1989. He was a member of the Board of Governors of the IEEE Information Theory Society from 1991 to 1996 and began a second term in 2006. Dr. Calderbank was honored by the IEEE Information Theory Prize Paper Award in 1995 for his work on the Z4 linearity of Kerdock and Preparata Codes (joint with A.R. Hammons Jr., P.V. Kumar, N.J.A. Sloane, and P. Sole), and again in 1999 for the invention of space-time codes (joint with V. Tarokh and N. Seshadri). He received the 2006 IEEE Donald G. Fink Prize Paper Award and the IEEE Millennium Medal, and was elected to the US National Academy of Engineering in 2005.